Predicting Digging Success for Unmanned Aircraft System Sensor Emplacement⋆

Adam Plowcha[1], Yue Sun[2], Carrick Detweiler[1], and Justin Bradley[1]

[1] Computer Science and Engineering Department, University of Nebraska, Lincoln
[2] Mechanical Engineering Department, University of Nebraska, Lincoln

Abstract. *We have developed an autonomous, digging, Unmanned Aircraft System (UAS) for sensor emplacement. A key challenge is quickly determining whether or not a particular digging activity will lead to successful emplacement, thereby allowing the system to potentially try another location. We have designed a first-of-its-kind decision-making algorithm using a Markov Decision Process to autonomously monitor the activity of a digging UAS activity to quickly decide if success is likely. Further, we demonstrate through many experimental trials that our method outperforms other decision-making methods with an overall success rate of 82.5%.*

Keywords: Sensor emplacement, UAS, Field robotics, Markov decision process

1 Motivation and Problem Statement

Fig. 1: UAS with digging auger.

We have developed a small Unmanned Aircraft System (UAS) that can autonomously fly to a remote location to emplace a sensor underground (shown in Figure 1) [13]. Unlike prior approaches for UAS sensor deployment [1, 5, 8], our approach *digs* into the ground to place the sensor underground for explicit measurements or to conceal the sensor. An aircraft-mounted digging system provides greater flexibility in locations that can be reached and can minimize the time to deploy a sensor. To minimize the weight of the digging system and maximize the UAS range, the system is optimized for a particular type of soil. This optimization may lead to emplacement failure if the system attempts to dig in soil that varies locally from the

⋆ This material is based upon work supported by USSTRATCOM through the 55th Contracting Squadron under Contract No. FA4600-12-D-9000 and was also funded, in part, by NSF-1638099 and USDA-2017-67021-25924.

target type, and since the system is operating far from the operator, autonomy is required to determine if emplacement is likely to succeed.

The key contribution of this work is the development and experimental validation of algorithms to quickly determine whether or not a particular digging activity will lead to successful emplacement. An emplacement may fail, for example, if the soil is too compact or has extraneous obstructions (e.g. roots, rocks), or the sensor may get stuck or partially emplaced, which we consider the worst case failure. If these situations can be accurately detected early, the UAS can extract the sensor and try an alternate location. Setting a time limit on a digging attempt provides an easy litmus test but can lead to either wasting resources on what will be an unsuccessful dig, or aborting a potentially successful dig and subsequently wasting energy flying to a new location. In either case, faster, more reliable prediction will yield improved emplacement results. To illustrate, in our trials we find that 29% of our attempts succeed within 20 seconds, but of those that go longer there is a significant variance in the time it takes to succeed, and only 48% fail to reach the target depth.

Success in emplacement is best characterized as a stochastic event. Although knowing the soil type can reduce uncertainty, the likelihood of unforeseen obstructions slowing or stopping the progress is high and difficult to predict. With this uncertainty in mind, we have designed a first-of-its-kind decision-making algorithm for a digging UAS that enables autonomous monitoring of digging to quickly decide if success is likely or if another digging location should be selected. Markov Decision Processes (MDP) provide optimal decision-making under uncertainty [10], and as such, we develop an MDP to predict the outcome of a single digging event. We also provide a comparison of the performance of our MDP against binary decision trees and a support vector machine, which are commonly used classification techniques. Finally, we compare these methods to the performance of an expert human operator in predicting the outcome during operation of *153 experimental trials involving emplacement of sensors in six different soil types.*

2 Technical Approach

Figure 2 shows a detailed diagram of the UAS auger system, which we describe briefly here for context[3]. The auger is mounted to a DJI Matrice 100 with a custom weight-distributing chassis, a micro-controller, a depth sensor, and a DC motor which turns the auger. The auger is hollow, allowing a sensor of up 25 mm in diameter and 76 mm length to be placed inside. Upon reaching target depth, the auger mechanism (with sensor inside) can be released. Alternately, the auger can be reversed, removing the sensor to try an alternate location.

2.1 Digging System State Characterization

To characterize the digging process we leverage the following key performance metrics that, together, provide insight into the status of the digging operation: motor RPM, motor current, auger depth, and elapsed digging time. For instance, decreasing motor rpm and increasing motor current draw coupled with increasing

[3] Additional details can be found in [13].

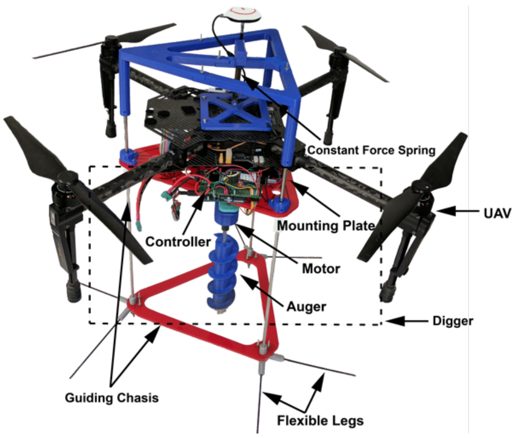

Fig. 2: Details of UAS auger system

auger depth is often indicative of a potentially successful dig. Conversely, high motor rpm with low motor current draw might indicate that the auger has been impeded by something solid and will not be able to achieve a greater depth. In many scenarios the system may initially achieve a certain auger depth, and subsequently fail to advance, indicating a potential failure. This is not captured well by the above metrics, so we add the metric *rate of digging progress* to help identify this scenario.

We impose three discrete measures, "Low," "Medium," and "High," on each metric to create a state space with sufficient fidelity to represent the digging process via an automaton. This results in $3^5 = 243$ system states, a description of which can be seen in Figure 3.

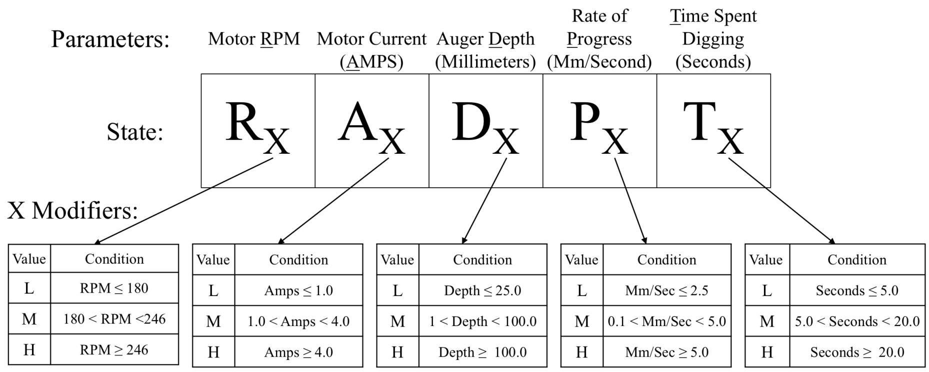

Value	Condition
L	RPM ≤ 180
M	180 < RPM <246
H	RPM ≥ 246

Value	Condition
L	Amps ≤ 1.0
M	1.0 < Amps < 4.0
H	Amps ≥ 4.0

Value	Condition
L	Depth ≤ 25.0
M	1 < Depth < 100.0
H	Depth ≥ 100.0

Value	Condition
L	Mm/Sec ≤ 2.5
M	0.1 < Mm/Sec < 5.0
H	Mm/Sec ≥ 5.0

Value	Condition
L	Seconds ≤ 5.0
M	5.0 < Seconds < 20.0
H	Seconds ≥ 20.0

Fig. 3: State Description - A state with high motor RPM, low amps, low depth, low rate of progress, and low time would be $R_H A_L D_L P_L T_L$.

3 Experimental Setup

An understanding of the nuances of digging helps develop the correct methods and models we use to predict digging success. We analyzed the results of 153 digging trials in a testbed, shown in Figure 4. Soil types were provided to us by our collaborators as described in [13]. To meet these specifications, a testbed having six different soil compositions with a range of compactness and moisture levels representative of the specified target environment was created. During each trial we logged the system state described in Figure 3 at a rate of 10 Hz. We label a dig "successful" when the auger reaches its target depth of 100 mm. If the auger does not reach its target depth, then it is manually stopped by the operator when system parameters indicate that success is unlikely or approximately 120 seconds have passed, and the trial is considered a digging failure.

From the 153 trials, we use eight randomly chosen successful trials and eight randomly chosen failed trials as the training data for our various decision making models. By using approximately 10% of our data for training, we avoid over fitting our data when running those models against the data from the remaining 137 trials. We have chosen four decision-making methods for evaluation and describe each one below.

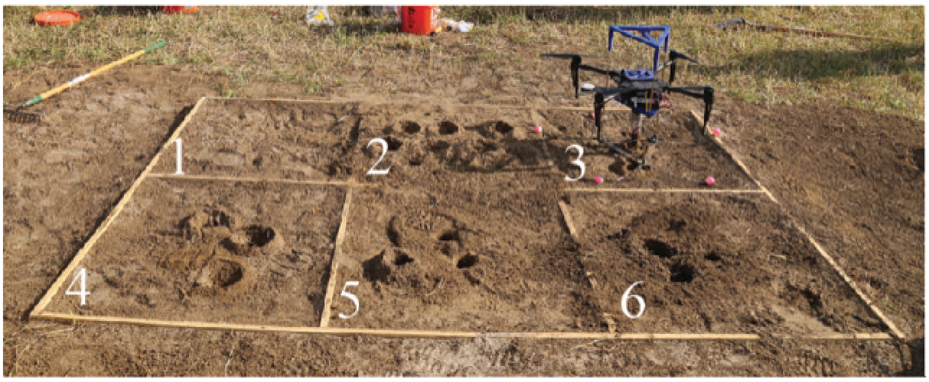

Fig. 4: Test site with different soil types.

3.1 Markov Decision Process

A MDP is a tuple, $S = \{S, A, P(s, a, s'), R(s, a, s'), \gamma\}$, representing states, actions, transition probabilities, rewards, and a discount factor, respectively. The MDP can be solved to generate an optimal policy, π, representing which action should be taken in any state. MDPs in UAS have predominantly been used for path planning [2, 6]. In this case, the stochasticity represented by the MDP captures the dynamic uncertainty from the assumed noise and disturbances potentially preventing the vehicle from attaining the commanded position. MDPs have also been used in UAS for other purposes such as image classification [11], search and rescue decisions [14], and target detection [4]. In each of these cases the stochasticity (i.e. randomness) of their data was suited to the application of a

MDP. Similarly, for our digging success prediction, the transitions between states are highly uncertain, sometimes being equi-probabilistically distributed among 3-4 outgoing transitions from a state. Moreover, we have numerous ending states as seen in Table 1, with numerous paths to those states. This uncertainty stems from the soil, what may be in the soil, the failure conditions of the motor and auger, as well as how well a particular auger performs in a particular soil type. This uncertainty is well represented by a MDP.

Our MDP consists of 243 states, as described in section 2. The actions are either: 1) continue digging or 2) stop digging. We generated transition probabilities from our 16 randomly chosen training trials. We created our reward table after examining our data using the following rationale. Rewards were increased to encourage stopping when a successful outcome seemed unlikely. This includes situations where very little digging depth is achieved over a fixed period of time, or when the system was in a state where high motor current draw and low motor RPM indicated a potential stall situation. States in which the system had achieved its target depth were considered termination states and were assigned the largest reward for the stopping action. To encourage continuation, we assigned higher rewards if a positive rate of digging progress was observed, or the system had been digging only a short time. To further tune our MDP reward table, we ran the initial MDP-generated policy against random successful and unsuccessful digs observing the MDP's "decision" at each time interval within the files. This allowed us to find areas within our reward table that could be adjusted to increase the chance of the policy leading to a successful outcome. For instance, we discovered that we needed to increase the reward for continuing while in a state that had the attributes of medium time with medium progress, i.e. $R_X A_X D_X P_M T_M$. The discount factor, $\gamma = 0.96$, was tuned by comparing success rates at γ values near 1.0 and then decreasing γ until success rates became maximal. We then used version 1.6 of the MDPToolbox [3] for Matlab® to generate our policy which was then run against our remaining data sets.

3.2 Decision Tree

Decision trees are supervised learning methods that use a training set of data to compute a relationship between a set of input attributes and a target attribute. This relationship is a model in the form of a binary decision tree that can be used to predict the target attribute of new data [9]. In our case, the input attributes consist of the digging system's measured parameters (time, motor current draw, motor amps, and depth of auger), and the target attribute is the decision to continue or stop digging. We choose this method as the relatively small number of attributes should generate a model that will be fast and efficient to implement on the hardware of our digging system. Additionally, it has been successfully employed to classify soil conditions from UASs in the past [1].

Using the same 16 randomly chosen digging trials we used as training sets for our MDP, we create a binary classification tree based on the four attributes our system sensors provided using Matlab®'s "fitctree()" method [7].

3.3 Support Vector Machine

A support vector machine (SVM) is a supervised learning method that classifies data from a training set by defining a hyperplane with maximal distance from the respective data points within the set. New data can then be classified by its position relative to that hyperplane. We choose to evaluate this method in our situation for the sake of completeness as there are cases when SVMs are reported to outperform decision trees [12] as well as vice-versa [1].

We create our SVM-based predictive model using Matlab®'s "fitcsvm()" method with the default linear kernel and the same training sets as our previous methods.

3.4 Human Operator Decision

The human operator made predictions about the overall outcome of a digging evolution based on the experience gained from the numerous observations made while creating and testing the digging system. The digger control software used to log the raw sensor data also allowed the user to indicate a prediction of success or failure at any given moment during a dig. Clicking a button labeled "Succeed" or "Fail" would insert a marker at that particular time into the log file (see the upper right corner of Figure 5). On average, the human operator would indicate a prediction after twenty seconds of digging based on observations of the visual and audible clues from the digging system as well as observing the sensor data received from the system on the digging control software. In this way, the human operator had access to the same data used by the machine-based methods of decision making.

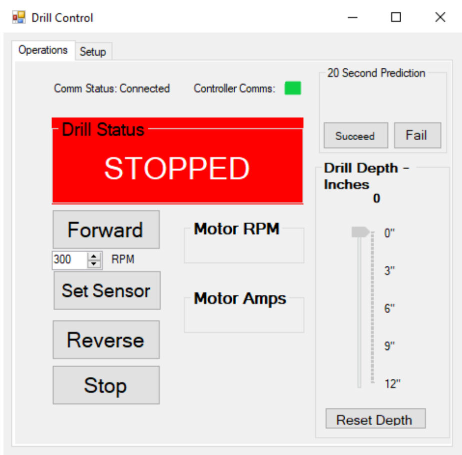

Fig. 5: Digger Control Software.

4 Experimental Results

Figure 6 shows an example profile of a successful digging attempt. It shows an increase in depth and current usage (with a corresponding drop in RPM) as the dig progresses. Figures 7 and 8 show example profiles of failed digging attempts. Figure 7 shows a consistently high motor RPM with low current usage and depth stopping at about 30 millimeters. We call this a failure due to spinning. Figure 8 shows the opposite type

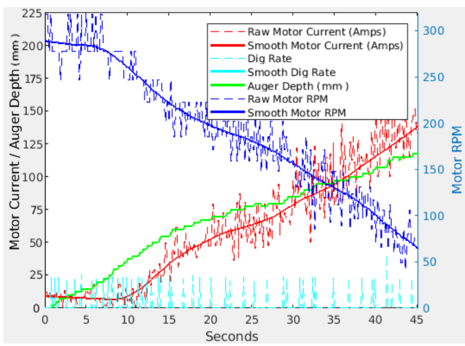

Fig. 6: Raw data from a successful dig.

of failure where something in the soil causes the auger bit to drastically slow down or stop with a commensurate rise in motor current draw. We call this a failure due to stalling.

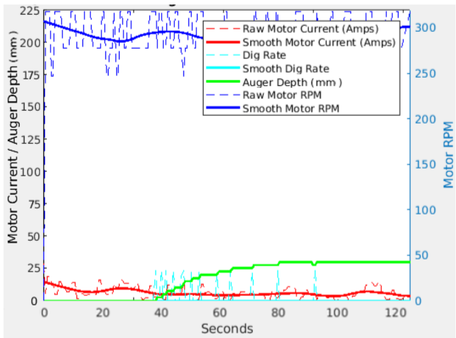

Fig. 7: Raw data from a dig failure due to spinning.

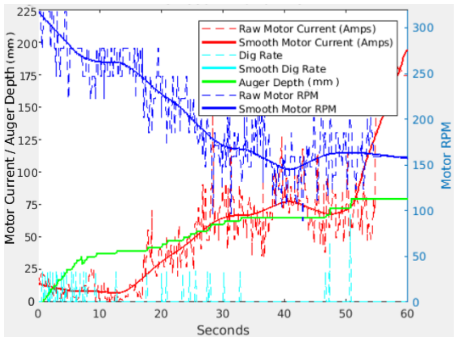

Fig. 8: Raw data from a dig failure due to stalling.

Table 1 shows the final states which define both successful and unsuccessful digs. Of the trials, 80 were successful with the digger reaching its target depth. Those successes were distributed among 19 different end states. There were 73 trials that failed where the digger did not reach its target depth. Those failures were distributed among 19 different end states. The large number of final end states shows the stochasticity of the digging and that a single metric cannot be used to detect failure.

The majority of our successful digs (45 out of 80) occurred within the first 20 seconds as evidenced by the states ending with T_L (low time) and T_M (medium time). We can also separate the majority of the failed digging evolutions into two subgroups based on their RPM and current draw. Trials either failed due to the auger stalling, as indicated by states starting with R_L (low RPM) or R_M (medium RPM) and then A_M (medium amps) or A_H (high amps), or the auger was freely spinning, as indicated by states starting with R_H (high RPM) and then A_L (low amps) or A_M (medium amps).

In some instances there is little difference between the end states of a successful dig and an unsuccessful one. The first row in Table 1 shows that the successful dig and unsuccessful dig differ only in a single separation of modifier for depth (i.e., D_H vs. D_M). This is due to the range of values used to specify medium depth (D_M) vs. high depth (D_H) and medium time (T_M) vs. high time (T_H). D_M spans a range of 75 mm ending just below the D_H mark of 100 mm. T_H is any value greater than 20 seconds. Examination of the data logs for the specific trials in question show that two of the unsuccessful digs were at the lower end of the D_M range, one had been stalled at 90 mm for over 30 seconds, and all three trials had been digging for over 90 seconds. The four successful digs all achieved the target depth of 100 mm in 55 to 61 seconds. A human operator in control of the dig used judgment as to when to stop or to allow the digging to continue based on visual and audible clues. Again, this highlights the stochasticity of a digging event as failure is difficult to predict.

Table 1: Trial End States labeled with **R**PM, **A**mps, **D**epth, **P**rogress, and **T**ime, in terms of **L**ow, **M**edium, and **H**igh as defined in Figure 3

Successful Attempts		**Failed Attempts**	
End State	Quantity	End State	Quantity
$R_L A_M D_H P_L T_H$	5	$R_L A_M D_M P_L T_H$	3
$R_L A_M D_H P_M T_H$	2	$R_L A_H D_L P_L T_H$	1
$R_L A_M D_H P_H T_M$	2	$R_L A_H D_M P_L T_H$	11
$R_L A_H D_H P_L T_H$	4	$R_L A_H D_M P_M T_H$	1
$R_L A_H D_H P_M T_H$	2	$R_M A_L D_M P_L T_H$	2
$R_M A_L D_H P_H T_M$	1	$R_M A_M D_L P_L T_H$	1
$R_M A_M D_H P_L T_H$	3	$R_M A_M D_M P_L T_H$	1
$R_M A_M D_H P_M T_H$	1	$R_M A_H D_M P_L T_H$	4
$R_M A_M D_H P_H T_L$	1	$R_H A_L D_L P_L T_M$	2
$R_M A_M D_H P_H T_M$	4	$R_H A_L D_L P_L T_H$	7
$R_M A_H D_H P_M T_H$	4	$R_H A_L D_L P_M T_M$	1
$R_H A_L D_H P_L T_H$	2	$R_H A_L D_M P_L T_H$	12
$R_H A_L D_H P_M T_H$	4	$R_H A_L D_M P_M T_M$	1
$R_H A_L D_H P_H T_L$	3	$R_H A_L D_M P_H T_L$	1
$R_H A_L D_H P_H T_M$	12	$R_H A_L D_M P_H T_M$	1
$R_H A_M D_H P_L T_H$	7	$R_H A_M D_L P_L T_H$	2
$R_H A_M D_H P_M T_H$	2	$R_H A_M D_M P_L T_H$	20
$R_H A_M D_H P_H T_L$	2	$R_H A_M D_M P_M T_H$	1
$R_H A_M D_H P_H T_M$	19	$R_H A_H D_M P_L T_H$	1

4.1 MDP

We tested our MDP-derived policy against the 137 remaining data sets not used for training by iterating over the timed sequence of data and using the policy to decide whether to continue to the next time step or stop. Our MDP policy had an overall success rate of 82.5% in predicting the outcome of a dig. While overall success is an important metric, we are more concerned with instances when our MDP predicts success when success is not possible (a false positive) and how quickly it correctly decides to abort a dig when success is not achievable (a true negative). Recall that our worst case is when energy is wasted by a system that continues to dig when it will not succeed as this energy could be used to fly to a new location for another attempt at sensor emplacement. Figure 9 shows the true/false positive/negative rates of our four decision making methods. All "stop" decisions based on the true negatives recognized by our MDP occurred within 20 seconds.

There were nine instances where the MDP policy correctly made predictions counter to the human operator. We highlight the following two cases as examples.

In our 40th trial, the human operator observed what appeared to be a stalled state where the auger had stopped advancing at about midway to the target depth and had remained there for approximately 13 seconds. In the two seconds prior to the operator making a prediction, the auger began a slightly perceptible advance downward, but the human operator determined this was not enough to allow it to achieve the desired depth and recorded a prediction of "failure." However, that slight advance in depth was enough to put the state of progress into the medium category (P_M) which the MDP policy recognized as a "continue" condition, and the auger eventually reached the desired depth.

Conversely, in our 85th trial, the human operator noted a very reasonable advance in depth that had reached three quarters of the way to the target depth by the time a prediction was required. The human operator recorded a prediction of "success" but did not realize the commensurate increase in motor current draw and decrease in motor RPM were indicating a situation where the auger bit might stop turning completely. However, as the system was in a state with high amperage (A_H) and low RPM (R_L), the MDP policy recognized a "stop" condition and correctly predicted failure.

MDP Predicts	Condition: Successful Dig	Condition: Unsuccessful Dig
Success	TPR: 81%	FPR: 08%
Failure	FNR: 19%	TNR: 92%

Dec Tree Predicts	Condition: Successful Dig	Condition: Unsuccessful Dig
Success	TPR: 85%	FPR: 66%
Failure	FNR: 15%	TNR: 34%

SVM Predicts	Condition: Successful Dig	Condition: Unsuccessful Dig
Success	TPR: 76%	FPR: 100%
Failure	FNR: 24%	TNR: 0%

Human Predicts	Condition: Successful Dig	Condition: Unsuccessful Dig
Success	TPR: 44%	FPR: 17%
Failure	FNR: 56%	TNR: 83%

Fig. 9: Decision Method True/False Positive/Negative Rates out of 137 trials

4.2 Decision Tree

The decision tree model was similarly tested against the 137 remaining data sets not used for training by iterating through each time-step in the trial data and applying the model to determine if digging should continue or stop. Our decision tree-based model had an overall success rate of 61.3%. Its false positive rate was 65.6%.

4.3 Support Vector Machine

The SVM model was also tested against the 137 trials not used for its training using the same method as described above. It had an overall success rate of 23.3% with a false positive rate of 100% and a true negative rate of 0%, which means that the SVM did not successfully predict a failure in any of the 137 trials it was simulated against.

4.4 Human Operator

The goal of human operator was to make a prediction at approximately the 20-second mark of each digging evolution. As 61 of the 153 digging trials reached their target depth within 20 seconds, the human operator made 91 predictions with an overall success rate of 69.5% and a false positive rate of 16.7%.

5 Main Experimental Insights

Initially we believed the human operator was the benchmark for correct prediction due to the significant experience with the system and additional knowledge of the environmental conditions (e.g. the type of soil the system was digging in). Our expert human operator was only able to successfully predict the outcome of a digging evolution 69.5% of the time, but also with a 16.7% false positive rate. The decision tree and SVM both underperform the human operator (although the decision tree only by a few percentage points). While both methods can handle noise within their training data, the amount of noise within the particular training sets may have impacted their overall effectiveness. The MDP-generated policy performs better than the other three methods with an overall success rate of 82.5% and only an 8% false positive rate, and this low false positive rate is critical in our application as premature stopping is preferred to continued digging when success is not possible.

A key contribution of this paper is the insight gained due to the nature of the problem. It was initially perceived that a simple timer to determine failure would be sufficient. However, in practice, many surprising cases emerged in which little auger progress was observed for over a minute just prior to rapid success. This showcases the variability in digging operations and how algorithmic decision makers can enhance performance.

6 Conclusion and Future Work

A digging UAS provides an excellent means of reaching and placing sensors in a variety of locations. This mobility comes with costs associated with the digging mechanism including the need to predict success/failure. Since failure of a single digging activity is likely, predicting the event and relocating for another attempt as quickly as possible is critical for mission success.

Our work shows that a MDP can be developed, trained, and successfully used to predict the outcome of a digging evolution with 82.5% accuracy which exceeds those of other decision making methods, including expert human operators. Our next goal is to deploy our algorithm on an improved version of our UAS digging system for real-time testing. Furthermore, we would like to extend the capability of algorithm to not only predict success or failure, but to also select appropriate courses of action based on the specific type of failure detected.

Acknowledgments. Thanks to Dr. Sebastian Elbaum, Dr. Brittany Duncan, Andrew Rasmussen, Ajay Shankur, Jacob Hogberg, and Aaron Clare and the staff at Horning State Farm for their assistance.

References

1. Anthony, D., Basha, E., Ostdiek, J., Ore, J.P., Detweiler, C.: Surface classification for sensor deployment from uav landings. In: 2015 IEEE International Conference on Robotics and Automation (ICRA). pp. 3464–3470 (May 2015)

2. Bethke, B., Bertuccelli, L., How, J.: Experimental demonstration of adaptive mdp-based planning with model uncertainty. In: AIAA Guidance, Navigation and Control Conference and Exhibit. p. 6322 (2008)
3. Chadès, I., Chapron, G., Cros, M.J., Garcia, F., Sabbadin, R.: Mdptoolbox: a multi-platform toolbox to solve stochastic dynamic programming problems. Ecography 37(9), 916–920 (2014)
4. Chanel, C.P.C., Teichteil-Königsbuch, F., Lesire, C.: Multi-target detection and recognition by uavs using online pomdps. In: AAAI. pp. 1381–1387 (2013)
5. Corke, P., Hrabar, S., Peterson, R., Rus, D., Saripalli, S., Sukhatme, G.: Autonomous deployment and repair of a sensor network using an unmanned aerial vehicle. In: Robotics and Automation, 2004. Proceedings. ICRA'04. 2004 IEEE International Conference on. vol. 4, pp. 3602–3608. IEEE (2004)
6. Dadkhah, N., Mettler, B.: Survey of motion planning literature in the presence of uncertainty: Considerations for uav guidance. Journal of Intelligent & Robotic Systems 65(1-4), 233–246 (2012)
7. The Mathworks, Inc., Natick, Massachusetts: MATLAB version 9.3.0.713579 (R2017b) (2017)
8. Pister, K.S.: Tracking vehicles with a uav-delivered sensor network (2001)
9. Rokach, L., Maimon, O.: Data Mining with Decision Trees: Theory and Application, vol. 2. World Scientific (2014)
10. Russell, S.J., Norvig, P.: Artificial intelligence: a modern approach (3rd edition), vol. 3. Prentice hall Upper Saddle River (2009)
11. Scott, K.K.: Occlusion-Aware Sensing and Coverage in Unmanned Aerial Vehicle (UAV) Networks. Ph.D. thesis, University of Cincinnati (2016)
12. Shafri, H., Ramle, F.: A comparison of support vector machine and decision tree classifications using satellite data of langkawi island. Information Technology Journal 8(1), 64–70 (2009)
13. Sun, Y., Plowcha, A., Nail, M., Elbaum, S., Terry, B., Detweiler, C.: Unmanned aerial auger for underground sensor installation. In: Accepted to appear 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2018). (2018)
14. Waharte, S., Trigoni, N.: Supporting search and rescue operations with uavs. In: Emerging Security Technologies (EST), 2010 International Conference on. pp. 142–147. IEEE (2010)