

UAV Coverage Path Planning under Varying Power Constraints using Deep Reinforcement Learning

Mirco Theile¹, Harald Bayerlein², Richard Nai¹, David Gesbert², and Marco Caccamo¹

Abstract—Coverage path planning (CPP) is the task of designing a trajectory that enables a mobile agent to travel over every point of an area of interest. We propose a new method to control an unmanned aerial vehicle (UAV) carrying a camera on a CPP mission with random start positions and multiple options for landing positions in an environment containing no-fly zones. While numerous approaches have been proposed to solve similar CPP problems, we leverage end-to-end reinforcement learning (RL) to learn a control policy that generalizes over varying power constraints for the UAV. Despite recent improvements in battery technology, the maximum flying range of small UAVs is still a severe constraint, which is exacerbated by variations in the UAV’s power consumption that are hard to predict. By using map-like input channels to feed spatial information through convolutional network layers to the agent, we are able to train a double deep Q-network (DDQN) to make control decisions for the UAV, balancing limited power budget and coverage goal. The proposed method can be applied to a wide variety of environments and harmonizes complex goal structures with system constraints.

I. INTRODUCTION

Whereas the CPP problem for ground-based robotics has already found its way into our everyday life in the form of vacuum cleaning robots [1], autonomous coverage with UAVs, while not yet having attained the same level of prominence, is being considered for a wide range of applications, such as photogrammetry, smart farming and especially disaster management [2]. UAVs can be deployed rapidly to gather initial or continuous survey data of areas hit by natural disasters, or mitigate their consequences. In the aftermath of the 2019-20 Australian bushfire season, wildlife officers inventively used quadcopter drones with infrared sensors to conduct a search-and-rescue operation for koalas affected by the blaze [3].

As its name suggests, covering all points inside an area of interest with CPP is related to conventional path planning where the goal is to find a path between start and goal positions. In general, CPP aims to cover as much of the target area as possible within given energy or path-length constraints while avoiding obstacles or no-fly zones. Due to the limitations in battery energy density, available power limits mission duration for quadcopter UAVs severely. Finding a CPP control policy that generalizes over varying power constraints and setting a specific movement budget can be

seen as a way to model the variations in actual power consumption during the mission, e.g. caused by environmental factors that are hard to predict. Similar to conventional path planning, CPP can usually be reduced to some form of the travelling salesman problem, which is NP-hard [1]. Lawn-mowing and milling [4] are other examples of closely related problems.

The most recent survey of UAV coverage path planning is given by Cabreira *et al.* [2]. Galceran and Carreras [1] provide a survey of general (ground robotics) approaches to CPP. Autonomous UAVs for applications in wireless communications have also sparked a lot of interest recently. Some scenarios, e.g. deep RL trajectory planning for UAVs providing wireless connectivity under battery power constraints, are related to CPP. An overview of UAV applications in wireless communications can be found in [5].

To guarantee complete coverage, most existing CPP approaches split the target area and surrounding free space into cells, by means of exact or approximate cellular decomposition. Choset and Pignon [6] proposed the classical boustrophedon (“the way of the ox”, back and forth motion) cellular decomposition, an exact decomposition method that guarantees full coverage but offers no bounds on path-length. This algorithm was extended by Mannadiar and Rekleitis [7] through encoding the cells of the boustrophedon decomposition as a Reeb graph and then constructing the Euler tour that covers every edge in the graph exactly once. Cases where the mobile agent does not have enough power to cover the whole area are not considered. The authors in [8] adapted this method for use in a non-holonomic, fixed-wing UAV and conducted extensive experimental validation. Two other approaches combining CPP and the travelling salesman problem to find near-optimal solutions for coverage of target regions enclosed by non-target areas are proposed by the authors in [9]: grid-based and dynamic programming-based, respectively. Both approaches suffer from exponential increase in time complexity with the number of target regions and do not consider obstacles or UAV power limitations.

Non-standard approaches have made use of neural networks (NNs) before. The authors in [10] design a network of neurons with only lateral connections that each represent one grid cell in a cleaning robot’s non-stationary 2D environment. The path planning is directly based on the neural network’s activity landscape, which is computationally simple and can support changing environments, but does not take path-length or power constraints into account.

Reinforcement learning with deep neural networks has only recently started to be considered for UAV path planning.

¹Mirco Theile, Richard Nai, and Marco Caccamo are with the TUM Department of Mechanical Engineering, Technical University of Munich, Germany {mirco.theile, richard.nai, mcaccamo}@tum.de

²Harald Bayerlein and David Gesbert are with the Communication Systems Department, EURECOM, Sophia Antipolis, France {harald.bayerlein, david.gesbert}@eurecom.fr

Maciel-Pearson *et al.* [11] proposed a method using an extended double deep Q-network (EDDQN) to explore and navigate from a start to a goal position in outdoor environments by combining map and camera information from the drone. Their approach is focused on obstacle avoidance under changing weather conditions. The authors in [12] investigate the CPP-related drone patrolling problem where a UAV patrols an area optimizing the relevance of its observations through the use of a single-channel relevance map fed into a convolutional layer of a DDQN agent. However, there is no consideration for power constraints and the relevance map is preprocessed showing only local information. To the best of our knowledge, deep RL has not been considered for UAV control in coverage path planning under power constraints before.

The main contributions of this paper are the following:

- Introduction of a novel UAV control method for coverage path planning based on double deep Q-learning;
- The usage of map-like channels to feed spatial information into convolutional network layers of the agent;
- Learning a control policy that generalizes over random start positions and varying power constraints and decides between multiple landing positions.

The remainder of this paper is organized as follows: Section II introduces the CPP problem formulation, Section III describes our DDQN learning approach and in Section IV follow simulation results and their discussion. We conclude the paper with a summary and outlook onto future work in Section V.

II. PROBLEM FORMULATION

A. Setup

The sensors of the UAV forming the input of the reinforcement learning agent are depicted in Figure 1: camera and GPS receiver. The camera gives a periodic frame of the current coverage view and the GPS yields the drone’s position. Power constraints determined by external factors are modelled as a movement budget for the drone that is fixed at mission start. Two additional software components are running on the UAV. The first is the mission algorithm which is responsible for the analysis of the camera data. We assume that any mission algorithm can give feedback on the area that was already covered. The second component is a safety controller that evaluates the proposed action of the agent and accepts or rejects it based on the safety constraints (entering into no-fly zones or landing in unsuitable areas). Note that the safety controller does not assist the agent in finding the landing area. The last component is a map which is provided by the operator on the ground. While this map could be dynamic throughout the mission, we focus on static maps for the duration of one mission in this paper.

B. 3-Channel Map

The coverage problem to be solved can be represented by a two dimensional grid map with three channels. Each cell in the grid represents a square area of the coverage region. The three channels describe starting and landing zones, target

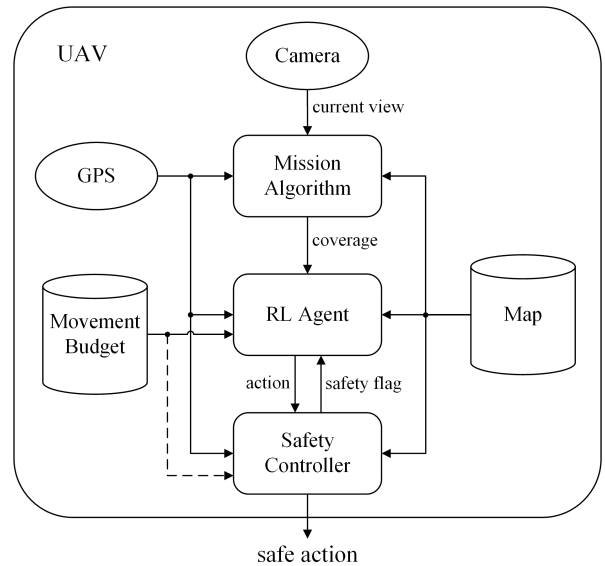


Fig. 1: System-level diagram depicting sensor and software components on the UAV during a coverage mission.

zones, and no-fly zones. The start and landing zones are areas the agent can start from and land on after finishing a coverage path. Target zones have to be covered at least once by the field of view (FoV) of the UAV’s camera. No-fly zones represent areas which the drone is prohibited from entering. Note that it is possible that a cell is declared as none, or more than one of these zones, with the exception that starting and landing zones can not be no-fly zones at the same time.

C. Markov Decision Process

In order to solve the described coverage path planning problem with reinforcement learning, it is converted into a Markov decision process (MDP). An MDP is described by the tuple $(\mathcal{S}, \mathcal{A}, R, P)$, with the set of possible states \mathcal{S} , the set of possible actions \mathcal{A} , the reward function R , and the deterministic state transition function $P : \mathcal{S} \times \mathcal{A} \mapsto \mathcal{S}$.

In a $N \times N$ grid, the state space \mathcal{S} has the following dimensions:

$$\mathcal{S} = \underbrace{\mathbb{B}^{N \times N \times 3}}_{\text{Map}} \times \underbrace{\mathbb{B}^{N \times N}}_{\text{Coverage}} \times \underbrace{\mathbb{R}^2}_{\text{Position}} \times \underbrace{\mathbb{N}}_{\text{Movement Budget}} \times \underbrace{\mathbb{B}}_{\text{Safety Flag}}$$

where \mathbb{B} is the Boolean domain $\{0, 1\}$. The action space \mathcal{A} contains the following five actions:

$$\mathcal{A} = \{\text{north, east, south, west, land}\}$$

The reward function $R : \mathcal{S} \times \mathcal{A} \mapsto \mathbb{R}$, mapping the current state $s \in \mathcal{S}$ and current action $a \in \mathcal{A}$ to a real-valued reward, consists of multiple components:

- r_{cov} (positive) coverage reward for each target cell that is covered by the UAV’s field of view for the first time;
- r_{sc} (negative) safety penalty in case the safety controller (SC) rejects the agent’s proposed action;

- r_{mov} (*negative*) constant movement penalty that is applied for every unit of the movement budget the UAV uses
- r_{crash} (*negative*) penalty in case the UAV runs out of movement budget without having safely landed in a landing zone.

III. METHODOLOGY

A. Q-Learning

Reinforcement learning, in general, proceeds in a cycle of interactions between an agent and its environment. At time t , the agent observes a state $s_t \in \mathcal{S}$, performs an action $a_t \in \mathcal{A}$ and subsequently receives a reward $r(s_t, a_t) \in \mathbb{R}$. The time index is then incremented and the environment propagates the agent to a new state s_{t+1} , from where the cycle restarts. The goal of the agent is to maximize the discounted cumulative return R_t from the current state up to a terminal state at time T . It is given as

$$R_t = \sum_{k=t}^T \gamma^{k-t} r(s_k, a_k). \quad (1)$$

with $\gamma \in [0, 1]$ being the discount factor, balancing the importance of immediate and future rewards. The return is maximized by adapting the agent's behavioral policy π . The policy can be deterministic with $\pi(s)$ such that $\pi : \mathcal{S} \mapsto \mathcal{A}$, or probabilistic with $\pi(a|s)$ such that $\pi : \mathcal{S} \times \mathcal{A} \mapsto \mathbb{R}$, yielding a probability distribution over the action space for each $s \in \mathcal{S}$.

To find a policy which maximizes the return, we utilize Q-learning, a model-free reinforcement learning approach. It is based on learning the state-action-value function, or *Q-function* $Q : \mathcal{S} \times \mathcal{A} \mapsto \mathbb{R}$, defined as

$$Q^\pi(s, a) = \mathbb{E}_\pi [R_t | s_t = s, a_t = a]. \quad (2)$$

Q-learning relies on iteratively updating the current knowledge of the Q-function. When the optimal Q-function is known, it is easy to construct an optimal policy by taking actions that maximize the Q-function. For convenience, s_t and a_t are abbreviated to s and a and s_{t+1} and a_{t+1} to s' and a' in the following.

B. Deep Q-Learning

The Q-function from (2) can be represented through a table of Q-values with the dimension $\mathcal{S} \times \mathcal{A}$. This is not feasible for large state or action spaces, but it is possible to approximate the Q-function by a neural network in those cases. A deep Q-network (DQN) parameterizing the Q-function with the parameter vector θ is trained to minimize the expected temporal difference (TD) error given by

$$L(\theta) = \mathbb{E}_\pi [(Q_\theta(s, a) - Y(s, a, s'))^2] \quad (3)$$

with the target value

$$Y(s, a, s') = r(s, a) + \gamma \max_{a'} Q_\theta(s', a'). \quad (4)$$

While a DQN is significantly more data efficient compared to a Q-table due to its ability to generalize, the

deadly triad [13] of function approximation, bootstrapping and off-policy training can make its training unstable and cause divergence. In 2015, Mnih *et al.* [14] presented a methodology to stabilize the DQN learning process. Their training approach makes use of an experience replay memory \mathcal{D} which stores experience tuples (s, a, r, s') collected by the agent during each interaction with the environment. Training the agent on uniformly sampled batches from the replay memory decorrelates the individual samples and rephrases the TD-error as

$$L^{\text{DQN}}(\theta) = \mathbb{E}_{s, a, s' \sim \mathcal{D}} [(Q_\theta(s, a) - Y^{\text{DQN}}(s, a, s'))^2]. \quad (5)$$

Additionally, Mnih *et al.* used a separate target network for the estimation of the next maximum Q-value changing the target value to

$$Y^{\text{DQN}}(s, a, s') = r(s, a) + \gamma \max_{a'} Q_{\bar{\theta}}(s', a') \quad (6)$$

with $\bar{\theta}$ representing the parameters of the target network. The parameters of the target network $\bar{\theta}$ can either be updated as a periodic hard copy of θ or as a soft update with

$$\bar{\theta} \leftarrow (1 - \tau)\bar{\theta} + \tau\theta \quad (7)$$

after each update of θ . $\tau \in [0, 1]$ is the update factor determining the adaptation pace. The combination of replay memory and target network separation to stabilize the training process laid the groundwork for the rise in popularity of DQN methods.

An additional improvement was proposed by Van Hasselt *et al.* [15], who showed that under certain conditions, action values in (6) get overestimated. To solve this issue, the double deep Q-network (DDQN) was introduced. The target value is then given by

$$Y^{\text{DDQN}}(s, a, s') = r(s, a) + \gamma Q_{\bar{\theta}}(s', \arg\max_{a'} Q_\theta(s', a')) \quad (8)$$

and the corresponding loss function

$$L^{\text{DDQN}}(\theta) = \mathbb{E}_{s, a, s' \sim \mathcal{D}} [(Q_\theta(s, a) - Y^{\text{DDQN}}(s, a, s'))^2], \quad (9)$$

in which the overestimation of action values is reduced by selecting the best action using θ but estimating the value of that action using $\bar{\theta}$. When calculating $\nabla_\theta L^{\text{DDQN}}(\theta)$ the target value is taken as is, hence, the back-propagating gradient is stopped before $Y^{\text{DDQN}}(s, a, s')$.

C. Neural Network Model and Training Procedure

The DQN solving the MDP from Section II consists of convolutional and fully-connected layers. It is visualized in Figure 2. The UAV's own position is converted to a 2D one-hot representation, i.e. the encoding of the occupied cell inside the whole grid. With the position encoded in this way, it can be stacked with the three-channel map and the coverage grid to form the five-channel input of the network's convolutional layers. The kernels of the convolutional layers are then able to form direct spatial connections between the current position and nearby cells. The remaining movement budget is fed into the network after the convolutional layers.




Fig. 2: Neural network structure for the reinforcement learning agent.

The convolutional layers are padded so that their output shape remains the same as their input shape. All layers are zero-padded for all channels, with the exception of the first layer’s no-fly zone channel, which is one-padded. This is an explicit representation of the no-fly zone surrounding the mission grid. The rectified linear unit (ReLU) is chosen as activation function for the convolutional layers. The last layer of the convolutional network is flattened and concatenated with the movement budget input. Fully-connected layers with ReLU activation are attached to this flatten layer.

The last fully-connected layer is of size $|\mathcal{A}|$ and has no activation function. It directly represents the Q-values for each action given the input state. Choosing the argmax of the Q-values is called the greedy policy and exploits already learned knowledge. The greedy policy given by

$$\pi(s) = \operatorname{argmax}_{a \in \mathcal{A}} Q_{\theta}(s, a) \quad (10)$$

is deterministic and used when evaluating the agent’s learning progress. During training, the sampled soft-max policy for exploration of the state and action space is used instead. It is given by

$$\pi(a_i | s) = \frac{e^{Q_{\theta}(s, a_i) / \beta}}{\sum_{a_j \in \mathcal{A}} e^{Q_{\theta}(s, a_j) / \beta}} \quad (11)$$

with the temperature parameter $\beta \in \mathbb{R}$ scaling the balance of exploration versus exploitation. When β is increased so does exploration. The limit $\beta \rightarrow 0$ of the soft-max policy (11) is the greedy policy (10). The soft-max policy was chosen over the ϵ -greedy policy because it offers variable exploration based on the relative difference of Q-values and does not depend on the number of training steps or episodes. This appeared to be beneficial for this particular problem.

Algorithm 1 describes the training procedure for the double deep Q-network in more detail. After replay memory and network parameters are initialized, a new training episode begins with resetting the state, choosing a random UAV starting position and random movement budget $b_0 \in \mathcal{B}$. The episode continues as long as the movement budget is greater than zero and the UAV has not landed. A new action $a \in \mathcal{A}$

Algorithm 1 DDQN training for coverage path planning

```

Initialize  $\mathcal{D}$ , initialize  $\theta$  randomly,  $\bar{\theta} \leftarrow \theta$ 
1: for  $n = 0$  to  $N_{max}$  do
2:   Initialize state  $s_0$  with random starting position and
   sample initial movement budget  $b_0$  uniformly from  $\mathcal{B}$ 
3:   while  $b > 0$  and not landed do
4:     Sample  $a$  according to (11)
5:     Observe  $r, s'$ 
6:     Store  $(s, a, r, s')$  in  $\mathcal{D}$ 
7:     for  $i = 1$  to  $m$  do
8:       Sample  $(s_i, a_i, r_i, s'_i)$  uniformly from  $\mathcal{D}$ 
9:        $Y_i = \begin{cases} r_i, & \text{if } s'_i \text{ terminal} \\ \text{according to (8),} & \text{otherwise} \end{cases}$ 
10:      Compute loss  $L_i(\theta)$  according to (9)
11:    end for
12:    Update  $\theta$  with gradient loss  $\frac{1}{m} \sum_{i=1}^m L_i(\theta)$ 
13:    Soft update of  $\bar{\theta}$  according to (7)
14:     $b = b - 1$ 
15:  end while
16: end for

```

is chosen according to (11) and the subsequent experience stored in the replay memory buffer \mathcal{D} .

Sampling a minibatch of size m from the replay memory, the primary network parameters θ are updated by performing a gradient step using the Adam optimizer. Subsequently, the target network parameters $\bar{\theta}$ are updated using the soft update (7) and the movement budget is decremented. The episode ends when either the drone lands or the movement budget decreases to zero. Then, a new episode starts unless the maximum number of episodes N_{max} is reached. The hyperparameters that were used during training are listed in Table I.

IV. EXPERIMENTS

A. Simulation Setup

The agent can move in a two dimensional grid through action commands in \mathcal{A} if accepted by the safety controller.

Parameter	Value	Description
$ \theta $	1,190,389	number of trainable parameters
$ \mathcal{D} $	50,000	replay memory buffer size
N_{max}	10,000	maximum number of training episodes
β	0.1	temperature parameter (11)
m	128	minibatch size
γ	0.95	discount factor for target value in (8)
τ	0.005	target network update factor (7)

TABLE I: Hyperparameters for DDQN training.

Each action, no matter if accepted or rejected, consumes one unit of movement budget since energy is spent during hovering as well. The agent’s initial state $s_0 \in \mathcal{S}$ consists of a fixed map, a zero-initialized coverage grid and a position, which is uniformly sampled from the starting and landing zone of the map. Additionally, the initial movement budget is uniformly sampled from a movement budget range \mathcal{B} , which is set to 25-75 for the purpose of this evaluation. The value of the safety flag in s_0 is initialized to zero and the UAV’s camera field of view (FoV) is set to a fixed 3-by-3-cell area centered underneath the agent. After each step the mission algorithm marks the FoV as seen in the coverage grid map.

Three evaluation scenarios were chosen, each with a unique problem for the agent to solve. Map A depicted in Figure 3 (a)-(c) has a large starting and landing zone, which yields high variation during training. Additionally, the shape of the target area is challenging to cover. The difficulty of map B in Figure 3 (d)-(f) lies in the yellow area that is marked as target zone, but also marked as a no-fly zone, and therefore must be covered by flying adjacent to it. Map C, in Figure 3 (g)-(i) with a narrow passage between no-fly zones, while easy to cover is very difficult for training as discussed later.

B. Evaluation

After being trained on their respective scenario with varying movement budgets and varying starting positions under the exploration policy $\pi(a|s)$ from (11), the agents are evaluated under their exploitation policy $\pi(s)$ from (10). Their performance during coverage for the full movement budget range and all possible starting positions is evaluated. The performance is described through Figures 3 and 4 and Table II.

The agents’ ability to plan a trajectory that ends with a safe landing inside the landing zone over the full movement budget range and starting at each possible position is evaluated through Table II, showing the ratio of landing for all scenario variations. Despite the agent’s good landing performance, the safety controller’s capabilities on a real-world UAV would likely be extended to force navigation to the closest landing zone in the rare cases when the RL agent misses the right moment to return.

To evaluate the impact of movement budget on the achieved coverage ratio, the starting position was fixed. For the selected starting positions the agents successfully landed after completing a trajectory for each movement budget. Figure 4 shows the coverage ratio of each agent with respect

to initial movement budget. Selected trajectories for each map under three different movement budgets are depicted in Figure 3. Whereas the movement budget is increasing from left to right, the agent does not necessarily utilize the whole allocated budget if it determines that there is a risk of not returning to the landing area in time or the coverage goal is already fulfilled. It can be seen that the agent finds a trajectory balancing the goals of safe landing and maximum coverage ratio.

Figure 5 shows the training process of an agent on map C. The curve describes the cumulative reward of the exploitation strategy when evaluated during the training process. Three major phases appear during the training process, which are highlighted in the graph. In phase one the agent learns to land safely, but does not venture far enough from the landing zone to find the target zone. When transitioning to phase two, the agent discovers the target zone, yielding high immediate reward. Due to the focus on mid-term reward through the choice of discount factor γ , this strategy represents a local optimum. In phase three the agent discovers the path back to the landing zone, avoiding the crashing penalty r_{crash} . After refining the trajectory, the agent finds the optimal path at the end of phase three. The phase transitions are highly dependent on the exploration strategy. Soft-max exploration appeared to be more effective than the ϵ -greedy policy to guide these transitions. The basic pattern of this incremental learning process is also visible when applied to other maps, albeit with less pronounced transitions due to bigger variations in coverage ratios.

V. CONCLUSION

We have introduced a new deep reinforcement learning approach for the longstanding problem of coverage path planning. While existing approaches might offer guarantees on the (near)-optimality of their solutions, the case where available power constrains the path planning is usually not considered. By feeding spatial information through map-like input channels to the agent, we train a double deep Q-network to learn a UAV control policy that generalizes over varying starting positions and varying power constraints modelled as movement budgets. Using this method, we observed an incremental learning process that successfully balances safe landing and coverage of the target area on three different environments, each with unique challenges.

In the future we will investigate the possibilities of transfer learning for this set of problems. At first we will train the agents on a curriculum of problems based on individual map channels to further accelerate the training process described in this work. From there we will examine approaches to transfer the two dimensional grid agent to higher dimensions and dynamics, e.g. adding altitude and heading. To this effect, it might be beneficial to investigate other RL techniques such as actor-critic methods and policy optimization. The proposed approach can also be seen as an initial step for handling variable power consumption in a real-world scenario. Through these steps an application on physical hardware will likely be within reach.




Fig. 3: Coverage plots for three different maps (map A: (a)-(c), map B: (d)-(f), map C: (g)-(i)) with three different movement budgets each; red, blue, and green are no-fly zones, starting/landing zones, and target zones, respectively; the red arrows describe the trajectory and the yellow and white cell describe start and landing position, respectively; lighter cells were covered by the agent’s FoV.

ACKNOWLEDGMENTS

Marco Caccamo was supported by an Alexander von Humboldt Professorship endowed by the German Federal Ministry of Education and Research. Harald Bayerlein and David Gesbert are supported by the PERFUME project funded by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program (grant agreement no. 670896).

REFERENCES

- [1] E. Galceran and M. Carreras, “A survey on coverage path planning for robotics,” *Robotics and Autonomous Systems*, vol. 61, no. 12, pp. 1258–1276, 2013.
- [2] T. Cabreira, L. Brisolara, and P. R. Ferreira, “Survey on coverage path planning with unmanned aerial vehicles,” *Drones*, vol. 3, no. 1, 2019.
- [3] D. Gimsey, “Drones and thermal imaging: saving koalas injured in the bushfires - [news],” *The Guardian*, 10 Feb 2020.
- [4] E. M. Arkin, S. P. Fekete, and J. S. Mitchell, “Approximation algorithms for lawn mowing and milling,” *Computational Geometry*, vol. 17, no. 1-2, pp. 25–50, 2000.
- [5] Y. Zeng, Q. Wu, and R. Zhang, “Accessing from the sky: A tutorial on UAV communications for 5G and beyond,” *Proceedings of the IEEE*, vol. 107, no. 12, pp. 2327–2375, 2019.
- [6] H. Choset and P. Pignon, “Coverage path planning: The boustrophedon cellular decomposition,” in *Field and service robotics*, pp. 203–209, Springer, 1998.

	Map A	Map B	Map C
Landing ratio	99.37%	99.78%	98.26%

TABLE II: Landing ratio for each map evaluated on the full range of movement budgets and possible starting positions.




Fig. 4: Coverage ratio with varying movement budget for the three maps.




Fig. 5: Training process of an agent trained on map C with dashed lines indicating training phase transitions.

- [7] R. Mannadiar and I. Rekleitis, “Optimal coverage of a known arbitrary environment,” in *2010 IEEE International conference on robotics and automation*, pp. 5525–5530, 2010.
- [8] A. Xu, C. Viriyasuthee, and I. Rekleitis, “Optimal complete terrain coverage using an unmanned aerial vehicle,” in *2011 IEEE International conference on robotics and automation*, pp. 2513–2519, 2011.
- [9] J. Xie, L. R. G. Carrillo, and L. Jin, “An integrated traveling salesman and coverage path planning problem for unmanned aircraft systems,” *IEEE control systems letters*, vol. 3, no. 1, pp. 67–72, 2018.
- [10] S. X. Yang and C. Luo, “A neural network approach to complete coverage path planning,” *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, vol. 34, no. 1, pp. 718–724, 2004.
- [11] B. G. Maciel-Pearson, L. Marchegiani, S. Akcay, A. Atapour-Abarghouei, J. Garforth, and T. P. Breckon, “Online deep reinforcement learning for autonomous UAV navigation and exploration of outdoor environments,” *arXiv preprint arXiv:1912.05684*, 2019.
- [12] C. Piciarelli and G. L. Foresti, “Drone patrolling with reinforcement learning,” in *Proceedings of the 13th International Conference on Distributed Smart Cameras*, ACM, 2019.
- [13] R. S. Sutton and A. G. Barto, *Reinforcement Learning: an introduction*. MIT Press, second ed., 2018.
- [14] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, et al., “Human-level control through deep reinforcement learning,” *Nature*, vol. 518, no. 7540, pp. 529–533, 2015.
- [15] H. Van Hasselt, A. Guez, and D. Silver, “Deep reinforcement learning with double Q-learning,” in *Thirtieth AAAI conference on artificial intelligence*, pp. 2094–2100, 2016.