

# Heterogeneous Robot Teams for Informative Sampling

Travis Manderson, Sandeep Manjanna and Gregory Dudek  
Mobile Robotics Lab, McGill University, Montreal, Canada

(travism, msandeeep, dudek)@cim.mcgill.ca

**Abstract**—In this paper we present a cooperative multi-robot strategy to adaptively explore and sample environments that are unfavorable for humans. We propose a methodology for a team of heterogeneous robots to collaborate on information based planning for applications like sampling thermal imagery in a wildfire affected site to assist with detecting spot fires and areas of residual fires, fire mapping and monitoring fire progression or applications in marine domain for coral reef monitoring and survey. We use Gabor filter based texture classifier on aerial images from an Unmanned Aerial Vehicle (UAV) to segment the region of interest into classes. A policy gradient based path planner is used on the texture classified aerial image to plan a path for the Unmanned Ground Vehicle (UGV). The UGV then uses a local planner to reach the goals set by the global planner by avoiding obstacles. The UGV also learns the labels for the segmented classes as drivable and non-drivable using the feedback from the performance while reaching the planned waypoints. We evaluated the building blocks of our approach and present the results with application of these strategies to different domains.

- 2) proximity to human-life and property
- 3) fuel type (i.e. trees or grass)
- 4) terrain topography (flatness, rivers, mountain slope or valley)
- 5) weather conditions
- 6) fire size (height and area)

## I. INTRODUCTION

We propose a strategy for a team of autonomous vehicles to collaborate on information based planning for sampling dynamic environments that are dangerous for the humans. Some of the example applications include – collecting thermal imagery in a wildfire-affected site to assist with detecting spot fires, fire mapping and monitoring fire progression (Fig.1b), visual inspection and monitoring of coral reefs by planning on aerial image from an UAV and learning to navigate (Fig.1f) tough underwater environments with an Autonomous Underwater Vehicle (AUV) (Fig.1e).

Each year over 5 million hectares are burned due to forest fires in Canada and the United States. A profound number of people, ground-vehicles and aerial-vehicles are used to detect, suppress and extinguish forest fires. Wildfire suppression is a complex task that involves the dispatch of limited resources to the most at-risk fires where detection is often done through thermal imaging (Fig.1a) and human observations (Fig.1b). Manned helicopters are used to survey areas for *hot-spots* both before and after a wildfire which can spread beyond control. Several techniques are used for fire suppression, but the most notable is Initial Attack which can contain/control over 90% of wildfires<sup>1</sup>. Several factors determine the initial attack strategy including:

- 1) location

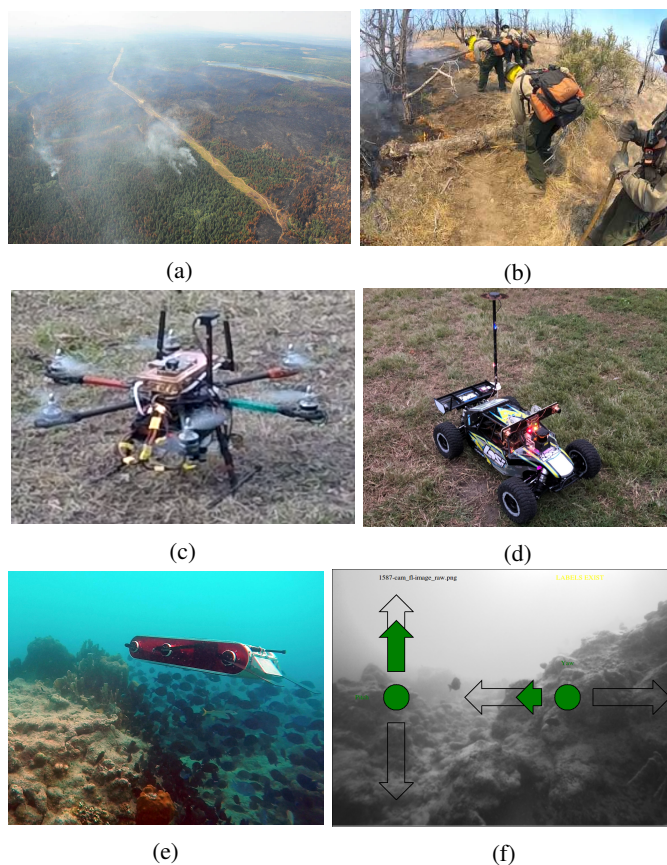


Fig. 1: a) Aerial view of hot-spots where fires may spread if not extinguished, b) initial attack crew work to extinguish hot-spots, c) hexacopter UAV used to take aerial photos, d) off-road UGV used for experiments, e) an Aqua class robot autonomously swimming in dense coral, f) a forward-facing image taken from the robot showing the predicted yaw and pitch steering angles.

Initial Attack techniques include extinguishing fires directly by firefighters using hand tools for small fires such as those under one meter or smouldering in the ground. Slightly

<sup>1</sup><https://www.nifc.gov/PUBLICATIONS/redbook/2000/Chapter9.pdf>

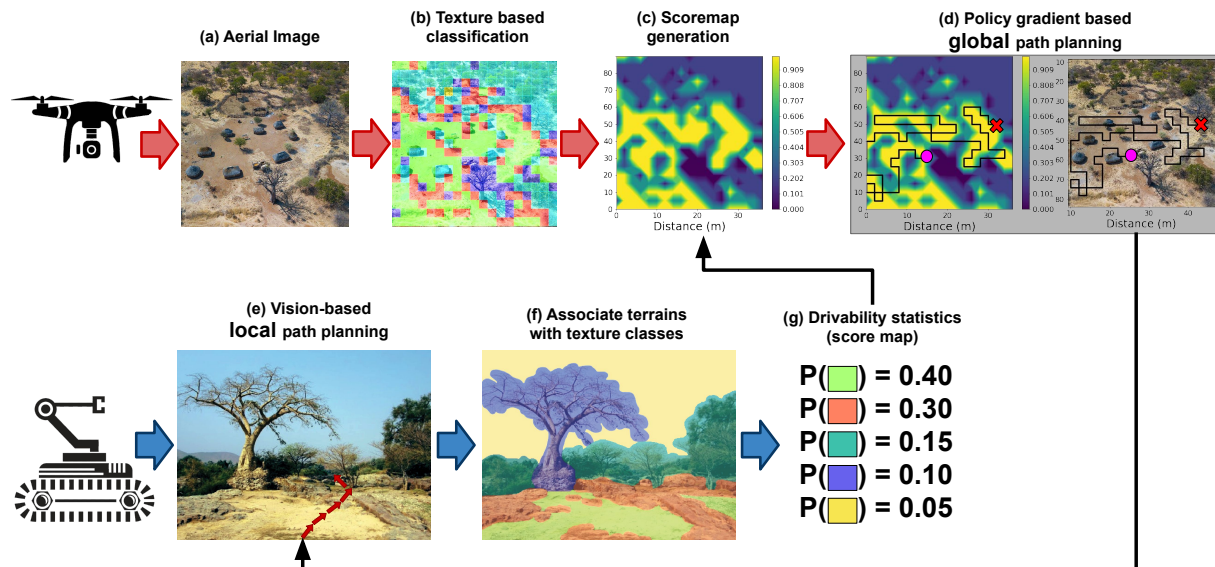


Fig. 2: Overview of our approach

larger fires which are too intense for firefighters are directly extinguished using bulldozers and fire engines or by dropping retardant. For large fires that cannot be attacked directly, a perimeter of natural (such as roads) or artificial barriers (created by limiting fuel sources) are used to control the fire and is allowed to burn out.

Recent advances in aerial and ground robotics could one day be used to assist in the detection and suppression of fires [2, 15]. Low-cost UAVs can replace much of the current work being done by manned helicopters by carrying thermal cameras to detect *hot-spots* at a high altitude as well as take local temperature measurements on the ground where accessible. Further, UGVs can be used to perform in situ measurements, digging up hot-spots, extinguishing by spraying retardant, bulldozing fuel sources to create perimeters and even supplying resources to manned crews. In literature, collaborative teams of UAVs and UGVs have been used for localizing each other [16], or for shared coverage, or in search and rescue scenarios or for target tracking [17] and surveillance [6]. We focus this work on using an heterogeneous robotic team for coordinated information gathering in outdoor environments.

In this work, we propose a methodology for global and local path planning of an UGV using imagery from an aerial vehicle. We seek to generate efficient paths for the UGV by segmenting the aerial image based on textures and labeling each class with a score associated to its *drivability*. As the ground-vehicle navigates, the performance in terms of progress towards the destination and incidence of obstacles measures the drivability of the terrain class and the drivability score is updated. Robots can play an important role in the process through aerial thermal imaging. In this work we focus on the global and local planning for an UGV for sampling applications. The presented methodology of collaborative sampling can find use in various real world applications, such as search and rescue operations,

sampling large-scale environmental scientific data, and large-scale coral reef monitoring using an UAV and an AUV.

## II. OUR APPROACH

The flowchart in Fig.2 presents an overview of our strategy for collaborative coverage and sampling on an outdoor scene. We split our approach into two planning phases: *Global planning*, where the high-level planner chooses locations that need to be visited to achieve good sampling and plans an efficient path (encoded as a discrete set of waypoints) to visit these locations; and *Local planner*, which handles short-term navigation between the waypoints (preferring good drivable terrains) and small deviations caused by obstacles.

A collaborative phase is responsible for labeling the terrain classes (used by the global planner) with the drivability feedback from the UGV. This update will result in an improved global plan and re-planning is necessary when the local planner on the UGV is unable to circumnavigate an obstacle. We discuss the building blocks of our system in the following subsections and later present some preliminary results from the evaluation of these building blocks.

### A. Texture classification and learning the scoremap

Our texture classification (as shown in Fig.3a) classifies patches of the image into a number of distinct texture classes. The texture classifier uses several Gabor filters to describe the texture [4]. The filters are convolved with the image to produce robust energy statistics, and several filters allow the classifier to be rotation and scale invariant.

The distribution of hue content is parameterized as a histogram. The hue histogram parameters and Gabor order statistics are combined into a feature vector, which are fed into a K-means classifier to discriminate between a preset number of texture classes. Fig.3c illustrates the texture based segmentation on an aerial image in Fig.3b using our classification pipeline.




Fig. 3: a) Segmentation pipeline, (b) Aerial image from an UAV, (c) Segmented aerial image based on a textural classifier.

Independent of the texture classifier, we label each texture class with a drivability score which is learned and updated on-line using feedback from the performance of the UGV (Fig. 2e, f, and g), such as rate of progress towards goal, incidence of obstacles, and terrain topology and rugosity.

### B. Policy Gradient based global path planning

We use policy search on aggregated state space [10] to generate paths for the UGV such that the drivable regions are covered. In this approach [11] a continuous two-dimensional sampling region is discretized into uniform grid-cells, such that the robot’s position  $\mathbf{x}$  can be represented by a pair of integers  $\mathbf{x} \in \mathbb{Z}^2$ . Each grid-cell  $(i, j)$  is assigned a score  $q(i, j)$  indicating the expected drivability in that cell. The goal is to maximize the total accumulated score  $J$  over a trajectory  $\tau$  within a fixed amount of time  $T$ . To specify the robot’s behavior we use a parametrized policy  $\pi_\theta(\mathbf{s}, \mathbf{a}) = p(\mathbf{a}|\mathbf{s}; \theta)$  that maps the current state  $\mathbf{s}$  of sampling to a distribution over possible actions  $\mathbf{a}$ . The aim here is to automatically find good parameters  $\theta$ , after which the policy can be deployed without additional training on new problems. We use a multi-resolution feature representation centered around the robot as explained in [10]. In this representation, the feature cells grow in size along with the distance from the robot. This results in

high resolution features close to the robot and lower resolution features further from the robot’s current position.

We use the GPOMDP and REINFORCE algorithms [1, 14, 3, 7] for computing the policy gradient as they have fewer hyper-parameters, making it easier to deploy,

$$\nabla_\theta J_\theta = \frac{1}{m} \sum_{i=1}^m \sum_{t=0}^{H-1} \nabla_\theta \log \pi_\theta(a_t^{(i)} | s_t^{(i)}) \left( \sum_{j=t}^{H-1} r(s_j^{(i)}, a_j^{(i)}) - b(s_t^{(i)}) \right).$$

In this equation, the gradient is based on  $m$  sampled trajectories with horizon  $H$  and state  $\mathbf{s}_j^{(i)}$  at the  $j^{\text{th}}$  time-step of the  $i^{\text{th}}$  sampled roll-outs. Furthermore,  $b$  is a variance-reducing baseline. In our experiments, we set the baseline to the observed average reward.

The path planned on the generated scoremap is presented in Fig. 2(d). Fig. 2(d) also illustrates the planned path overlaid on the initial aerial image.

### C. Vision based local path planning for UGV

The UGV uses vision-based navigation which steers the vehicle towards the next way-point while avoiding obstacles and preferring paths which are good for driving and provide good texture for visual state estimation.

A Convolutional Neural Network (CNN) is used to predict steering angles that avoid obstacles and leads to driving on good terrain. To train the CNN, we perform behavioural cloning not unlike the DAGGER algorithm[12]. The UGV is remotely controlled in both good and poor configurations such as very close to and far from obstacles. A Resnet-18 based CNN is trained to predict steering angle  $\theta$  from a set  $C = \{-M, \dots, 0, \dots, M\}$ , where  $M$  is a discrete number of steering angles in each the left and right direction. From the data collected in the remotely controlled step, a set of images and one-hot encoded steering angles make up a dataset  $D$ . The cost function for training the network is as follows:

$$\mathcal{L}(D, \mathbf{w}) = \mathcal{L}_{\text{pred}}(D, \mathbf{w}) + \lambda_1 \mathcal{L}_{\text{reg}}(D, \mathbf{w})$$

$$\mathcal{L}_{\text{pred}} = \sum_{(\mathbf{x}_i, \theta_i) \in D} l(\mathbf{f}_i^{(\theta)}, \theta_i, \mathbf{w})$$

where the regularization term corresponds to the KL-divergence term in [5]. The prediction loss for yaw steering actions was given by

$$l(\mathbf{f}_i^{(\theta)}, \theta_i, \mathbf{w}) = - \sum_j \theta_{ij} \log f_{ij}^{(\theta)} - \lambda_2 \sum_j f_{ij}^{(\theta)} \log f_{ij}^{(\theta)}$$

where first term corresponds to the cross-entropy between the network predictions and the smoothed labels, and the second term is the penalty for overconfident predictions, which aims to maximize the entropy of the predictive distributions. The hyper-parameters  $\lambda_1$  and  $\lambda_2$  determine the weights of the KL-divergence regularization and the entropy penalty and are selected manually.

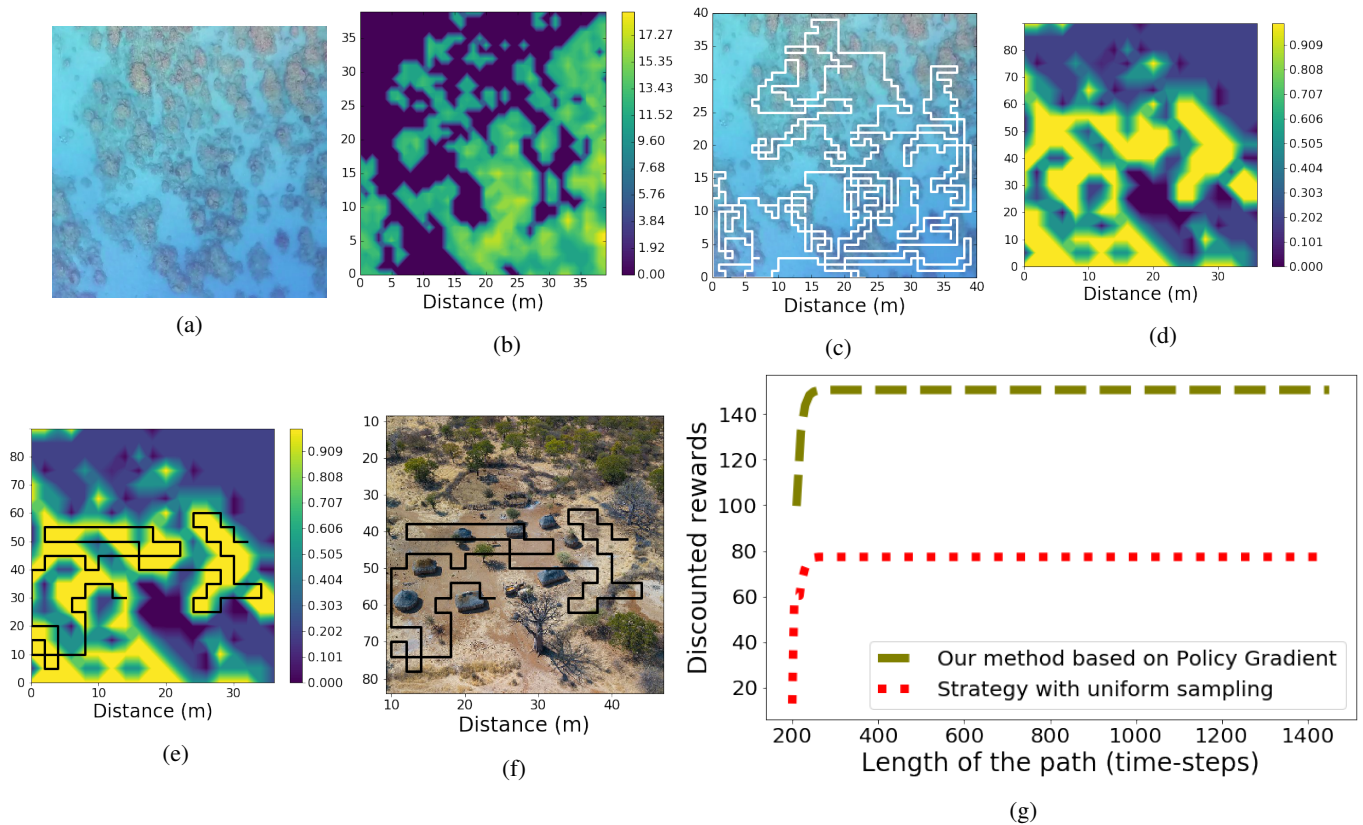


Fig. 4: a) Aerial view of a coral reef, b) Scoremap generated using texture information in the aerial image, c) Trajectory using the policy gradient search based planner on the scoremap, d) Scoremap generated using the segmentations from Fig.3c, e) The path planned using the policy gradient based planner, f) The path generated is overlaid on top of the initial aerial image, g) Plot comparing the total discounted rewards collected by our strategy and the uniform sampling strategy.

### III. PRELIMINARY RESULTS

We evaluated the building blocks of our system on both marine and terrestrial setups. In the marine example, the application is to cooperatively collect better visual samples of the coral reefs in a limited time budget using the UAV for global planning and an AUV for local planning and execution. The presented results are from experiments conducted in field at Folkestone Marine Reserve in Barbados, over a region known to have several coral outcrops. The AUV used in these experiments is a Aqua class underwater robot [13]. The AUV has a vision based local planner to swim between reefs collecting high-resolution visual samples. We used the aerial images from the UAV to plan trajectories for the AUV (as presented in Fig.4a, Fig.4b, and Fig.4c), such that the information gain about the coral reefs is maximized with a limited time budget. In the terrestrial example, we used the aerial images from the UAV presented in Fig.3b to generate a scoremap indicating the drivability of the terrain as illustrated in Fig.4d. The policy gradient based path planner then plans a trajectory for the UGV (as presented in Fig.4e and Fig.4f), such that the coverage of drivable terrain is maximized with a limited time.

In one of our recent work [10], we evaluated our policy gradient based path planner in the marine domain to monitor coral reefs. The results presented in Fig.4g illustrate that we are able to achieve higher discounted rewards in comparison to a uniform sampling technique. The goal here is to maximize the discounted rewards as the time budget is limited for the robotic vehicle.

We have performed preliminary validation of our local planner in the same marine environment as we validated the global planner. In [8], [9] we performed long-distance autonomous underwater navigation in close-proximity to coral reefs (on average 43 cm from coral) without collision. We manually labelled approximately 13,000 images taken in varying configurations and several lighting configurations with expected yaw and pitch and used them to train a CNN. Our experimental validation resulted in over one km of collision-free navigation in the open sea, while maintaining good coverage of coral which was our target observation. Fig. 5 shows the percentage of the forward-facing images that were covered with coral, showing that our robot maintained consistent coverage of the target terrain (coral). On average, 33% of the image showed the presence of coral indicating the effectiveness of navigating




Fig. 5: The percentage of image containing coral. It's expected that at most 50% of the image can be coral since the top half of the image is *normally* above the horizon.

coral region as opposed to barren regions. Since the horizon is located in the center of image, we expect the coral in view to be much lower than 50%. In an upcoming field trial, we plan to evaluate all the building blocks and the whole pipeline.



Fig. 6: The final paths swam by the Aqua robot covering over 1 km.

#### IV. CONCLUSION

In this paper we have outlined a system for cooperative path planning using information acquired from an UAV to plan paths for an UGV. Drivability performance on the UGV is used to update a scoremap to plan the path iteratively. We demonstrate the effectiveness of the subsystems in our preliminary experiments in a marine domain with collaboration between an aerial and an underwater vehicles. In our current and future work, we are tightening the interaction between the UAV and UGV along with a unified sampling goal. We expect that this work will lead to a robust and generalized framework for deploying cooperative aerial and ground robots in complex, unstructured and dynamic environments. The active-learning

framework allows the system to work from previously seen environments, but also incorporate new unseen information into the model for improved planning and execution.

#### REFERENCES

- [1] Jonathan Baxter and Peter L Bartlett. Infinite-horizon policy-gradient estimation. *Journal of Artificial Intelligence Research*, 15:319–350, 2001.
- [2] Evan Beachly, Carrick Detweiler, Sebastian Elbaum, Brittany Duncan, Carl Hildebrandt, Dirac Twidwell, and Craig Allen. Fire-aware planning of aerial trajectories and ignitions. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 685–692. IEEE, 2018.
- [3] Marc Peter Deisenroth, Gerhard Neumann, and Jan Peters. A survey on policy search for robotics. *Foundations and Trends® in Robotics*, 2(1–2):1–142, 2013.
- [4] I. Fogel and D. Sagi. Gabor filters as texture discriminator. *Biological Cybernetics*, 61(2):103–113, 1989. ISSN 0340-1200. doi: 10.1007/BF00204594. URL <http://dx.doi.org/10.1007/BF00204594>.
- [5] Yarin Gal, Jiri Hron, and Alex Kendall. Concrete Dropout. In *Advances in Neural Information Processing Systems 30 (NIPS)*, 2017.
- [6] Ben Grocholsky, James Keller, Vijay Kumar, and George Pappas. Cooperative air and ground surveillance. *IEEE Robotics & Automation Magazine*, 13(3):16–25, 2006.
- [7] Jens Kober, J Andrew Bagnell, and Jan Peters. Reinforcement learning in robotics: A survey. *The International Journal of Robotics Research*, 32(11):1238–1274, 2013.
- [8] Travis Manderson, Ran Cheng, Dave Meger, and Gregory Dudek. Navigation in the Service of Enhanced Pose Estimation. In *International Symposium on Experimental Robotics (ISER)*, 2018.
- [9] Travis Manderson, Juan Camilo Gamboa Higuera, Ran Cheng, and Gregory Dudek. Vision-based autonomous underwater swimming in dense coral for combined collision avoidance and target selection. In *2018 IEEE/RSJ*

- International Conference on Intelligent Robots and Systems (IROS)*, pages 1885–1891. IEEE, 2018.
- [10] Sandeep Manjanna, Herke van Hoof, and Gregory Dudek. Policy search on aggregated state space for active sampling. In *2018 International Symposium on Experimental Robotics (ISER)*, pages 1–7, 2018.
  - [11] Sandeep Manjanna, Herke van Hoof, and Gregory Dudek. Reinforcement learning with non-uniform state representations for adaptive search. In *2018 IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR)*, pages 1–7. IEEE, 2018.
  - [12] Stéphane Ross, Geoffrey Gordon, and Drew Bagnell. A reduction of imitation learning and structured prediction to no-regret online learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pages 627–635, 2011.
  - [13] Junaed Sattar, Gregory Dudek, Olivia Chiu, Ioannis Rekleitis, Philippe Giguere, Alec Mills, Nicolas Plamondon, Chris Prahacs, Yogesh Girdhar, Meyer Nahon, et al. Enabling autonomous capabilities in underwater robotics. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 3628–3634, 2008.
  - [14] Richard S Sutton, David A McAllester, Satinder P Singh, and Yishay Mansour. Policy gradient methods for reinforcement learning with function approximation. In *Advances in neural information processing systems*, pages 1057–1063, 2000.
  - [15] Chiara Torresan, Andrea Berton, Federico Carotenuto, Salvatore Filippo Di Gennaro, Beniamino Gioli, Alessandro Matese, Franco Miglietta, Carolina Vagnoli, Alessandro Zaldei, and Luke Wallace. Forestry applications of uavs in europe: A review. *International Journal of Remote Sensing*, 38(8-10):2427–2447, 2017.
  - [16] Richard T Vaughan, Gaurav S Sukhatme, Francisco J Mesa-Martinez, and James F Montgomery. Fly spy: Lightweight localization and target tracking for cooperating air and ground robots. In *Distributed autonomous robotic systems 4*, pages 315–324. Springer, 2000.
  - [17] Huili Yu, Randal W Beard, Matthew Argyle, and Caleb Chamberlain. Probabilistic path planning for cooperative target tracking using aerial and ground vehicles. In *Proceedings of the 2011 American Control Conference*, pages 4673–4678. IEEE, 2011.