

# Informative Path Planning for Active Learning in Aerial Semantic Mapping

Julius Rückin, Liren Jin, Federico Magistri, Cyrill Stachniss, Marija Popović

**Abstract**—Semantic segmentation of aerial imagery is an important tool for mapping and earth observation. However, supervised deep learning models for segmentation rely on large amounts of high-quality labelled data, which is labour-intensive and time-consuming to generate. To address this, we propose a new approach for using unmanned aerial vehicles (UAVs) to autonomously collect useful data for model training. We exploit a Bayesian approach to estimate model uncertainty in semantic segmentation. During a mission, the semantic predictions and model uncertainty are used as input for terrain mapping. A key aspect of our pipeline is to link the mapped model uncertainty to a robotic planning objective based on active learning. This enables us to adaptively guide a UAV to gather the most informative terrain images to be labelled by a human for model training. Our experimental evaluation on real-world data shows the benefit of using our informative planning approach in comparison to static coverage paths in terms of maximising model performance and reducing labelling efforts.

## I. INTRODUCTION

Semantic segmentation of remote sensing data provides valuable information for monitoring and understanding the environment. Unmanned aerial vehicles (UAVs) are popular for aerial image acquisition in many applications, including urban planning [1–4], precision agriculture [5–7], and wildlife conservation [8], due to their relatively low cost and high operational capabilities [8, 9]. In parallel, the advent of deep learning for semantic segmentation with fully convolutional neural networks (FCNs) has opened new frontiers [10, 11]. These developments have expanded the possibilities for the automated interpretation of aerial imagery towards more complex tasks and larger environments. Classical semantic segmentation approaches rely on supervised learning procedures. They require large amounts of training data that need to be labelled in a pixel-wise fashion, which is labour- and time-consuming.

This paper examines the problem of active learning (AL) for efficient training data acquisition in aerial semantic mapping scenarios. Our goal is to guide a UAV to collect highly informative images for training a FCN for semantic segmentation in a targeted way, and thus minimise the total amount of labelled data necessary. This is achieved by replanning the UAV paths online as new data is gathered to maximise an existing FCN’s segmentation performance

This work has been funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy, EXC-2070 – 390732324 (PhenoRob). Authors are with the Cluster of Excellence PhenoRob, Institute of Geodesy and Geoinformation, University of Bonn. Cyrill Stachniss is also with the Lamarr Institute for Machine Learning and Artificial Intelligence, Germany. Corresponding: j.rueckin@uni-bonn.de.

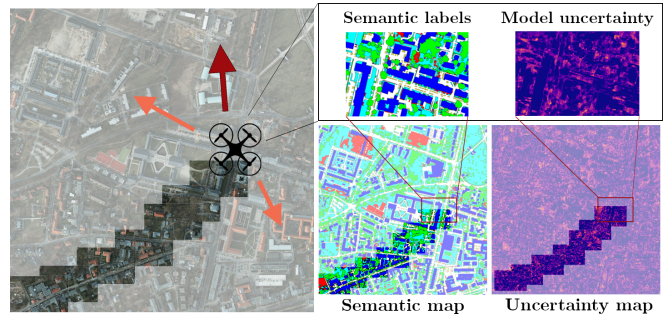


Fig. 1: Our planning strategy for active learning in UAV-based terrain mapping. During a mission, we estimate model uncertainty in semantic segmentation (top-right) and fuse it in a global terrain map (bottom-right). Based on the map, our approach guides a UAV to collect most useful (most uncertain) training images for labelling (left). Orange arrows indicate candidate paths and red shows the chosen path. In this way, our pipeline reduces the number of images that must be labelled by a human.

when it is retrained on images collected during the mission and labelled by a human operator.

To alleviate the labelling effort required to train FCNs, AL methods have been widely used to optimise the collection of training data. In AL, the aim is to selectively label data in order to maximise model performance. Recently, several studies have proposed AL frameworks for deep learning models with image data [12, 13], which effectively reduce labelling requirements. However, applying such methods in the context of autonomous robotic monitoring missions has been largely unaddressed. Recent works examining AL with UAV-acquired imagery [2, 8] only consider selecting images from a static pool obtained from previous exhaustive aerial surveys. Therefore, linking the AL objective to active robotic decision-making and planning remains an open challenge.

To tackle this issue, we propose a new informative path planning framework for AL in aerial semantic mapping tasks, bridging the gap between recent advances in AL and robotic applications. Our main contribution is a novel approach to integrate AL into a planning strategy for efficient training data acquisition. Our method exploits Bayesian deep learning [12] to estimate the model uncertainty of a pretrained FCN based on the light-weight ERFNet architecture [11]. The FCN provides pixel-wise semantic labels and model uncertainties which are fused online into a global terrain map. Based on the model uncertainty, we plan paths for a UAV to adaptively collect the most uncertain, i.e. the most informative, images for labelling and retraining the FCN after the mission. This pipeline enables reducing labelling efforts

without relying on large static pools of unlabelled data. Fig. 1 shows an overview of our approach.

In sum, we make the following three claims. First, our Bayesian ERFNet achieves higher segmentation performance than its non-Bayesian counterpart and provides consistent model uncertainty estimates for AL. Second, our online planning framework reduces the number of labelled images needed to maximise segmentation performance compared to static coverage paths [14], which are commonly used in aerial surveys. Third, we show that map-based planning with a lookahead in our integrated framework yields better AL performance compared to greedy one-step planning strategies. We open-source our code for usage by the community.<sup>1</sup>

## II. RELATED WORK

Our work combines recent developments in AL and Bayesian deep learning with informative path planning. The aim is to efficiently collect image data for training a deep learning model using autonomous robots.

### A. Active Learning with Image Data

AL has been intensively researched in the last decades. It aims to maximise a model’s performance gain while manually labelling the fewest amount of samples possible from a pool of unlabelled data. Early techniques for AL with images focus on kernel-based methods such as support vector machines [4] and Gaussian processes [15]. Recently, FCNs have become the state of the art in computer vision due to their superior performance [9–11]. AL has significant potential to reduce human labelling effort for FCNs, whose performance depends on large high-quality labelled datasets [12, 16]. To select informative samples, AL methods use acquisition functions to estimate either model uncertainty [12] or train data diversity [17]. We focus on model uncertainty measures since we aim to detect misclassifications in heterogenous aerial imagery, which may cover large areas. Building up on ideas by Gal et al. [12] to extract model uncertainty, we combine concepts from Bayesian deep learning into an AL framework for semantic segmentation.

AL techniques hold particular promise in remote sensing applications where unlabelled data is abundant but labelled benchmark datasets are scarce. As a result, they are experiencing rapid uptake in various scenarios, e.g. wildlife detection [8], land cover analysis [2, 6, 18], and agricultural harvesting [13]. Kellenberger et al. [8] combine AL with UAV imagery to improve the performance of an object-detection model subject to domain shifts. However, their AL objective is not suitable for quantifying model uncertainty, and assumes large unlabelled data pools from manually executed data campaigns. Lenczner et al. [2] propose an interactive AL framework to obtain accurate aerial segmentation maps. Although they exploit model uncertainty to target misclassified regions, their framework also requires a static annotated database. In contrast to these works, we integrate the AL objective into a path planning framework.

Our approach enables guiding a robot to collect the most informative new training images during a mission, without relying on previously acquired data.

### B. Uncertainty in Deep Learning

Uncertainty quantification is a long-standing problem in machine learning and a key component of AL. In our approach, quantifying model uncertainty is crucial as it can be reduced by adding more training samples. In deep learning, efficient uncertainty estimation using analytic methods is challenging since the model posterior distribution is often intractable [19]. One approach to estimate model uncertainty is to measure the disagreement between predictions of different networks in an ensemble [20]. To quantify uncertainty of a single model with reduced computational costs, Monte-Carlo (MC) dropout is commonly applied [19, 21]. At test time, this approach averages multiple model predictions with independently sampled dropout masks. Kendall et al. [22] leverage MC dropout in a Bayesian FCNs for semantic segmentation. They show that modelling uncertainty improves segmentation performance and produces a reliable measure for decision-making. Similar techniques also have been used to select informative samples in AL [6, 12, 13]. Our Bayesian ERFNet follows this line of work to obtain computationally efficient model uncertainty calculation for online path replanning.

### C. Informative Path Planning for Active Sensing

In active sensing, informative path planning aims to maximise the information of terrain measurements collected subject to platform constraints, such as mission time. Recent advances have unlocked the potential of efficient data-gathering using autonomous robots in various domains [5, 18, 20, 23]. Our work considers an adaptive planning strategy which allows a robot to focus on informative training images by replanning online as new data are accumulated [23].

Adaptive approaches seek to maximise the expected information gain of new measurements with respect to terrain maps. Such methods are used to, e.g., find hotspots [24] or perform inspection [23] in a targeted way. Several studies consider image processing using FCNs as a basis for mapping [5, 7, 25]. However, they focus on planning with respect to the semantic predictions. As such, they leverage fully trained models and do not consider model uncertainty.

The combination of path planning with AL is a relatively unexplored research area, but presents a natural way to exploit current robotic capabilities for reducing labelling efforts. Georgakis et al. [20] propose a framework for active semantic goal navigation which uses an ensemble to estimate model uncertainty as the planning objective. In contrast, our Bayesian ERFNet exploits MC dropout and a light-weight architecture for fast online replanning in UAV-based scenarios. Most resembling our work is the planning approach of Blum et al. [18] for AL in semantic mapping. A key difference to our approach is that their algorithm relies on novelty, which requires reasoning about the entire growing training set. In contrast, our MC dropout procedure permits fast model uncertainty estimation suitable for replanning.

<sup>1</sup>[github.com/dmar-bonn/ipp-al](https://github.com/dmar-bonn/ipp-al)

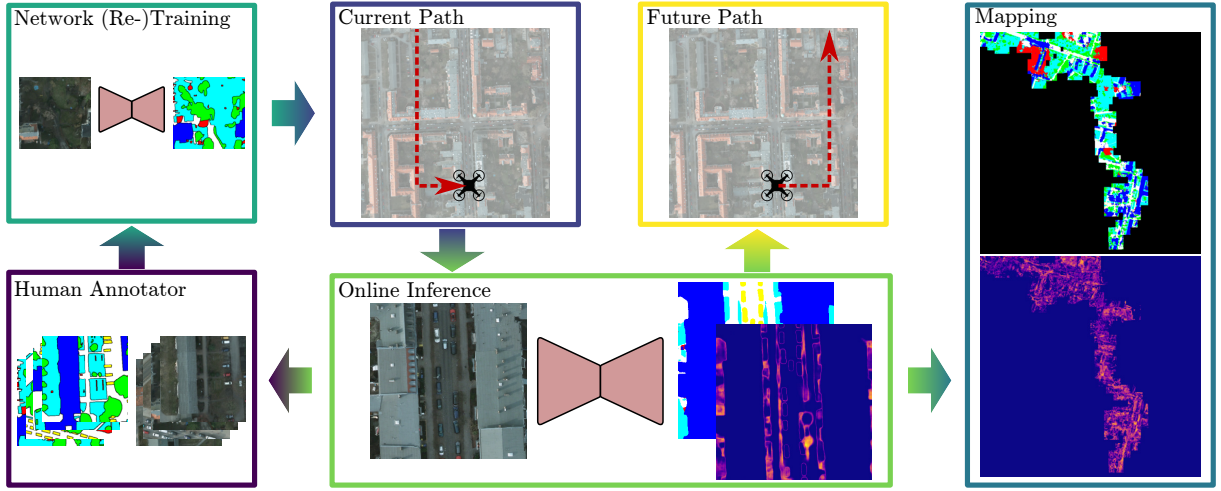


Fig. 2: Overview of our proposed approach. We start with a pretrained network for probabilistic semantic segmentation, deployed on a UAV. During a mission, the network processes RGB images to predict pixel-wise semantic labels and model uncertainties (Sec. III-A), which are projected onto the ground to build global maps capturing these variables (Sec. III-B). Based on the estimated model uncertainty, the current UAV position, and the current map state, our algorithm plans paths for the UAV to collect the most uncertain (most informative) training data for improving the network performance (Sec. III-C). After the mission, the collected images are labelled by an annotator and used for network retraining. By guiding the UAV to collect informative training data, our pipeline reduces the labelling effort.

Moreover, we leverage model uncertainties to maintain a consistent terrain map as a basis for global planning.

### III. OUR APPROACH

We propose a new informative path planning framework for AL in UAV-based semantic mapping scenarios. Our goal is to deploy a UAV to adaptively collect the most informative training images to be labelled by a human, and thereby reduce labelling efforts. Fig. 2 depicts an overview of our framework. During a mission, we use a pretrained FCN to perform semantic segmentation on aerial images of terrain and estimate the associated model uncertainty. The model prediction uncertainty are then used as input for semantic terrain mapping. In this setup, we exploit model uncertainty estimates to plan paths online for gathering new data that maximises an AL information objective. The following subsections describe the key elements in our proposed approach. Note that these modules are generic and can be adapted for any given monitoring scenario. Thus, we see our main contribution as a general robotic planning framework for AL that enables efficient training data acquisition.

#### A. Bayesian Segmentation Network

Our goal is to perform the semantic segmentation on RGB images and predict the associated pixel-wise model uncertainty as a basis for AL. To achieve this, we extend the ERFNet architecture proposed by Romera et al. [11] to account for this probabilistic model interpretation. Fig. 3 depicts our new network. Our main motivation for using ERFNet is its effectiveness for real-time inference, as required for online replanning, which is essential in our setup.

ERFNet is an encoder-decoder FCN structure with 16 blocks in the encoder and 5 blocks in the decoder. A key feature is the use of non-bottleneck-1D blocks, separating 2D convolutional layers into two subsequent 1D convolutional

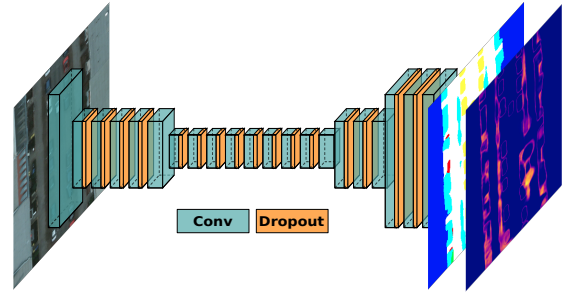


Fig. 3: Our Bayesian ERFNet architecture for probabilistic semantic segmentation. We extend the network of Romera et al. [11] with Monte-Carlo dropout (orange layers) to predict model uncertainty. Our network takes as input RGB (left) and outputs semantic labels (second from right) and pixel-wise uncertainty (first from right).

layers to reduce the learnable parameter count without sacrificing model capacity. The model  $f^W(z)$  is parameterised by weights  $W$  and its output is the likelihood  $p(y | f^W(z)) = \text{softmax}(f^W(z))$ , where  $z$  is the input RGB image, and  $y$  is the predicted pixel-wise class probability simplex.

We leverage advances in Bayesian deep learning to create a probabilistic version of ERFNet quantifying model uncertainty for AL [19, 21]. Assume a train set with images  $Z = \{z_1, z_2, \dots\}$  and labels  $Y = \{y_1, y_2, \dots\}$ . To quantify the model uncertainty, we estimate the posterior  $p(W | Z, Y)$  with MC dropout [21] performing multiple forward passes at test time with independently sampled dropout masks. Thus, we add dropout layers after all non-bottleneck-1D blocks in the encoder and decoder to create a fully Bayesian ERFNet. Our network is trained to minimise cross entropy [19]:

$$\mathcal{L}(\theta) = -\frac{1}{|Z|} \sum_{i=1}^{|Z|} \log p(y_i | f^{\hat{W}_i}(z_i)) + \lambda \|W\|_2^2, \quad (1)$$



where  $\lambda$  is the weight decay and weights  $\hat{\mathbf{W}}_i$  are sampled by applying dropout to the full set of parameters  $\mathbf{W}$ .

At deploy time, the prediction  $\mathbf{y}$  for an image  $\mathbf{z}$  is computed by marginalising over the approximated model posterior. To do so in a tractable fashion, we use MC integration [19]:

$$\hat{p}(\mathbf{y} | \mathbf{z}, \mathbf{Z}, \mathbf{Y}) = \frac{1}{T} \sum_{t=1}^T \text{softmax}(\mathbf{f}^{\hat{\mathbf{W}}_t}(\mathbf{z})), \quad (2)$$

where  $\hat{\mathbf{W}}_t$  are independently sampled by applying dropout to the full set of parameters  $\mathbf{W}$  in each forward pass.

Next, we outline our approach for linking the model uncertainty to an AL acquisition function. This function encodes the expected value of adding a new image to the training dataset and is used to select most informative samples for annotation in Fig. 2. The main idea is to pick images with the highest model uncertainty, i.e. highest information value, for training. To measure information, we use the Bayesian active learning by disagreement acquisition function [26], where images from an unlabelled pool are chosen to maximise the mutual information  $\mathbb{I}(\mathbf{y}, \mathbf{W} | \mathbf{z}, \mathbf{Z}, \mathbf{Y}) \in [0, 1]$  between model predictions and the posterior. Intuitively, this quantity is high when the predictions induced by the sampled network weights differ substantially, while the model is certain in each single prediction. However, for our Bayesian ERFNet, the model posterior prediction  $\hat{p}(\mathbf{y} | \mathbf{z}, \mathbf{Z}, \mathbf{Y})$  is intractable. This prevents us from computing model uncertainty, and hence information gain, analytically. Instead, following Gal et al. [12], we approximate mutual information by using Eq. (2):

$$\begin{aligned} \mathbb{I}(\mathbf{y}, \mathbf{W} | \mathbf{z}, \mathbf{Z}, \mathbf{Y}) &\approx -\hat{p}(\mathbf{y} | \mathbf{z}, \mathbf{Z}, \mathbf{Y})^T \log(\hat{p}(\mathbf{y} | \mathbf{z}, \mathbf{Z}, \mathbf{Y})) \\ &+ \frac{1}{T} \sum_{t=1}^T p(\mathbf{y} | \mathbf{z}, \hat{\mathbf{W}}_t)^T \log(p(\mathbf{y} | \mathbf{z}, \hat{\mathbf{W}}_t)), \end{aligned} \quad (3)$$

where  $\log(\cdot)$  is applied element-wise. For the derivation, we refer to the work of Gal et al. [12].

### B. Terrain Mapping

A key aspect of our approach is a probabilistic map, which captures the semantic and uncertainty information about the 2D terrain at a given time. This map is updated online during a mission as new measurements arrive, and can be exploited for global uncertainty-based planning. Our method leverages sequential Bayesian fusion and probabilistic occupancy grid mapping [27] to map semantic terrain information and model uncertainty respectively. The terrain  $\xi \subset \mathbb{R}^2$  is discretised using two grid maps: a semantic map  $\mathcal{X}_s$  and an model uncertainty map  $\mathcal{X}_u$ . The semantic information is represented by  $\zeta : \mathcal{X}_s \rightarrow [0, 1]^C$  and we assign a  $C$ -dimensional probability simplex to each discrete grid cell on the terrain  $\xi$ , where  $C$  is the number of possible class labels. Similarly, the model uncertainty in each cell is given by  $\eta : \mathcal{X}_u \rightarrow [0, 1]$ .

To map multiple semantic classes, we maintain one independent grid layer in  $\mathcal{X}_s$  for each of the  $C$  labels. In each layer  $i \in [C]$ , the prior semantic map distribution  $p(\zeta_i | \xi) \sim \mathcal{N}(\boldsymbol{\mu}^-, \mathbf{P}^-)$ , is defined by a mean vector  $\boldsymbol{\mu}^- = \frac{1}{C} \mathbf{1}$  and

a diagonal covariance matrix  $\mathbf{P}^- = \varepsilon \mathbf{I}$  with variances  $\varepsilon \in (0, 1]$ . During a mission, when an image is taken at time step  $t$ , the resulting output from our network is projected to the flat terrain given the camera intrinsics and UAV position. To update the map, a Kalman filter is used to recursively fuse the pixel-wise probabilistic segmentation predictions  $\mathbf{s}_t = \hat{p}(\mathbf{y} | \mathbf{z}_t, \mathbf{Z}, \mathbf{Y})$  (Eq. (2)) and estimated model uncertainties  $\mathbf{u}_t = \mathbb{I}(\mathbf{y}, \mathbf{W} | \mathbf{z}_t, \mathbf{Z}, \mathbf{Y})$  (Eq. (3)) based on input image  $\mathbf{z}_t$  taken at measurement position  $\mathbf{x}_t \in \mathbb{R}^3$  above the terrain  $\xi$  into a posterior mean  $\boldsymbol{\mu}^+$  and a posterior diagonal covariance matrix  $\mathbf{P}^+$  of model uncertainties defining the posterior map state  $\mathcal{X}_{s,t} \sim \mathcal{N}(\boldsymbol{\mu}^+, \mathbf{P}^+)$  at time step  $t$ . Note that we apply standard Bayesian fusion to update the map state  $\mathcal{X}_{s,t}$ .

To handle model uncertainties in  $\mathcal{X}_u$ , our goal is to fuse new uncertainties at time step  $t$  into the map as well as to keep track of exploration frontiers for planning. To achieve this, we update: (1) a maximum likelihood map state  $\mathcal{X}_{u,t}$  of model uncertainty estimates  $\mathbf{u}_t$  at measurement positions  $\mathbf{x}_t$  and (2) a hit map  $\mathcal{H}_{u,t}$  counting total updates of grid cells in  $\mathcal{X}_{u,t}$ , i.e. hits. By combining standard robotic mapping techniques with Bayesian deep learning, we maintain a terrain map suitable for global planning with an AL objective. We note the differences between our approach and previous works [18, 20], which assess the expected information gain of new data for AL only on a local per-image basis, without considering its global distribution in the environment.

### C. Informative Path Planning

We develop informative path planning algorithms to guide a UAV to adaptively collect useful training data for our FCN. Our goal is to plan paths to gather new images that maximise the model uncertainty, i.e. maximise the AL acquisition function, as described in Sec. III-A.

In general, the informative path planning problem aims to maximise the information gain  $\mathbb{I}(\cdot)$  of measurements collected within an initially unknown environment  $\xi$ :

$$\psi^* = \underset{\psi \in \Psi}{\operatorname{argmax}} \mathbb{I}(\psi), \text{ s.t. } C(\psi) \leq B, \quad (4)$$

where  $C : \Psi \rightarrow \mathbb{R}_{\geq 0}$  maps a path  $\psi$  to its execution costs,  $B \in \mathbb{R}_{\geq 0}$  is the robot's budget limit, e.g. time or energy, and  $\mathbb{I} : \Psi \rightarrow \mathbb{R}_{\geq 0}$  is the information criterion, computed from the new measurements taken along  $\psi$ . In our work, the costs  $C(\psi)$  of a path  $\psi = (\mathbf{x}_1, \dots, \mathbf{x}_N)$  of length  $N$  are defined by the total flight time:

$$C(\psi) = \sum_{i=1}^{N-1} c(\mathbf{x}_i, \mathbf{x}_{i+1}), \quad (5)$$

where  $\mathbf{x}_i \in \mathbb{R}^3$  is a 3D position above the terrain  $\xi$  an image is registered from. The function  $c : \mathbb{R}^3 \times \mathbb{R}^3 \rightarrow \mathbb{R}_{\geq 0}$  computes the flight time between measurement positions by a constant acceleration-deceleration  $\pm u_a$  with maximum speed  $u_v$ . A key feature of our approach is to relate the information criterion  $\mathbb{I}(\psi)$  to the AL acquisition function, i.e. we plan paths to collect new images along  $\psi$  that maximise the model's performance gain when added to the training set.



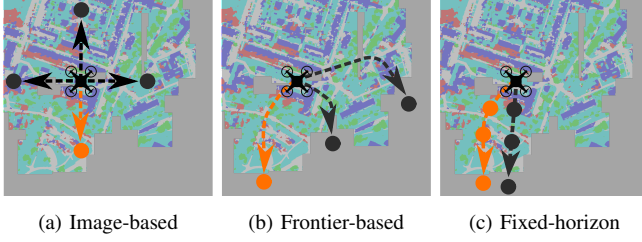


Fig. 4: Our planning strategies for training data acquisition. Black dots and arrows indicate candidate measurements and paths evaluated based on their information gain adding new images to the training set. Orange dots and arrows indicate the best chosen measurement and paths, respectively. Grey depicts unexplored areas.

Next, we propose three different informative planning strategies to optimise Eq. (4) for replanning during a mission. The approaches are illustrated in Fig. 4. In our experiments in Sec. IV-B, we compare each planner in terms of segmentation performance over the total labelling cost.

**Image-based planner.** Our image-based planner follows the direction of the highest estimated model uncertainty in the image recorded at the current UAV position (Fig. 4a). To this end, we greedily plan a next-best measurement position  $\mathbf{x}_t^*$  at time step  $t$  based on the previously estimated model uncertainty  $\mathbf{u}_{t-1} \in [0, 1]^{W \times H}$  over segmentation  $\mathbf{s}_{t-1} \in [0, 1]^{C \times W \times H}$  of RGB image  $\mathbf{z}_{t-1}$  of size  $W \times H$ , and the current hit map state  $\mathcal{H}_{u,t}$ . The main idea is to choose the next-best measurement position  $\mathbf{x}_t^*$  towards the image edge  $\mathbf{e}_t^*$  maximising the model uncertainty over hit counts in  $\mathcal{H}_{u,t}$ . To do so, we set a width  $\leq \min(\frac{W}{2}, \frac{H}{2})$  pixels defining the edge area, and sum pixel-wise model uncertainty  $\mathbf{u}_{t-1}$  along each edge normalised over the respective sum of grid cell counts in  $\mathcal{H}_{u,t}$ . Finally,  $\mathbf{x}_t^*$  is computed by moving from the previous position  $\mathbf{x}_{t-1}$  in direction of  $\mathbf{e}_t^*$  for a user-defined step size at a fixed altitude. Note that this strategy relies only on model uncertainty in the current local image and does not exploit our global uncertainty map. This resembles the planner proposed by Blum et al. [18].

**Frontier-based planner.** Our frontier-based planner navigates to the most uncertain boundaries of explored terrain, given the current state of the global uncertainty map (Fig. 4b). We greedily select a next-best measurement point  $\mathbf{x}_t^*$  at the boundary of known and unknown space of the terrain  $\xi$  at time step  $t$ . Candidate frontiers are generated by sampling positions equidistantly along the frontiers at a fixed altitude. Then,  $\mathbf{x}_t^*$  is chosen to be the frontier position maximising the model uncertainty in the global map  $\mathcal{X}_{u,t}$  normalised over hit counts in  $\mathcal{H}_{u,t}$ , corresponding to the camera field of view (FoV) at the respective frontier position.

**Fixed-horizon planner.** In contrast to the previous two strategies, our fixed-horizon planner reasons about images taken over multiple time-steps, given the current global uncertainty map state (Fig. 4c). To this end, we propose a fixed-horizon strategy inspired by the informative path planning framework for terrain monitoring introduced by Popović et al. [5]. To ensure efficient online replanning, we use a two-step approach. First, we greedily optimise a path

$\psi_t^g = (\mathbf{x}_t^g, \dots, \mathbf{x}_{t+N-1}^g)$  of length  $N$  at time step  $t$  using grid search at a fixed altitude. Second, an optimisation routine is applied to refine  $\psi_t^g$  in the continuous UAV workspace and obtain the next-best path  $\psi_t^* = (\mathbf{x}_t^*, \dots, \mathbf{x}_{t+N-1}^*)$ .

In the first step,  $\psi_t^g$  is computed sequentially for  $N$  steps by choosing the  $n$ -th measurement point  $\mathbf{x}_{t+n}^g$ ,  $n \leq N-1$ , over a sparse grid of candidate positions  $\mathbf{x}_c$ :

$$\mathbf{x}_{t+n}^g = \underset{\mathbf{x}_c}{\operatorname{argmax}} \frac{\|\mathcal{X}_{u,t}(\mathbf{x}_c)\|_1}{\|\mathcal{H}_{u,t+n}(\mathbf{x}_c)\|_1 \cdot c(\mathbf{x}_{t+n-1}^g, \mathbf{x}_c)}, \quad (6)$$

where  $\|\cdot\|_1$  is the matrix norm summing all its elements, and  $\mathcal{X}_{u,t}(\mathbf{x}_c)$  and  $\mathcal{H}_{u,t+n}(\mathbf{x}_c)$  are the model uncertainty map's subset at time step  $t$  and forward-simulated hit map's subset at future time step  $t+n$  spanned by the camera FoV at position  $\mathbf{x}_c$ , respectively.

In the second step, we apply an optimisation routine initialised with  $\psi_t^g$  using an objective function similar to Eq. (6), but generalised to candidate paths  $\psi_t^c = (\mathbf{x}_t^c, \dots, \mathbf{x}_{t+N-1}^c)$ :

$$\psi_t^c = \frac{\sum_{n=0}^{N-1} \|\mathcal{X}_{u,t}(\mathbf{x}_{t+n}^c)\|_1}{\sum_{n=0}^{N-1} \|\mathcal{H}_{u,t+n}(\mathbf{x}_{t+n}^c)\|_1 \cdot c(\mathbf{x}_{t+n-1}^c, \mathbf{x}_{t+n}^c)}, \quad (7)$$

where  $\mathbf{x}_{t-1}^c = \mathbf{x}_{t-1}$  is the previous measurement position. After convergence, the best-performing candidate path  $\psi_t^* = (\mathbf{x}_t^*, \dots, \mathbf{x}_{t+N-1}^*)$  with respect to Eq. (7) is returned, and  $\psi_{t,0}^* = \mathbf{x}_t^*$  is executed as the next-best measurement position.

Note that we normalise by the flight time to efficiently use the budget  $B$ , forward-simulate  $\mathcal{H}_{u,t}$  to penalise inspecting the same terrain regions again, and assume a uniform prior in unknown regions of  $\mathcal{X}_{u,t}$  to foster exploration.

## IV. EXPERIMENTAL RESULTS

Our experimental evaluation aims to analyse the performance of our approach and support the claims that our Bayesian FCN achieves higher segmentation performance than its non-Bayesian counterpart and provides consistent model uncertainty estimates for AL (Sec. IV-A). Further, the experiments will show that our online planning framework reduces the number of labelled images needed to maximise segmentation performance compared to static coverage paths (Sec. IV-B), and our proposed planning with lookahead strategy is more beneficial for AL performance than greedily optimising for the next-best measurement (Sec. IV-B).

### A. Bayesian ERFNet Ablation Study

The experimental results presented in this section aim to support the claim that our Bayesian ERFNet proposed in Sec. III-A achieves higher segmentation performance than the non-Bayesian ERFNet and provides consistent model uncertainty estimates for AL. To this end, we perform an ablation study with varying Bayesian units of the ERFNet base architecture to find the best-performing trade-off between segmentation performance and model uncertainty estimation. We assess four different probabilistic variants of ERFNet with varying dropout probabilities  $p = \{10\%, 30\%, 50\%\}$ :

- **Standard:** Dropout layers after all non-bt-1D layers in the encoder as in the normal ERFNet implementation.

Variant (dropout probabilities)	Accuracy [%] $\uparrow$	mIoU [%] $\uparrow$	ECE [%] $\downarrow$
Non-Bayesian (10% / 30% / 50%)	82.28 / 81.99 / 82.71	64.47 / 64.35 / 66.24	59.98 / 59.09 / 59.43
Standard (10% / 30% / 50%)	82.47 / 82.21 / 83.94	64.81 / 64.70 / 68.00	57.32 / 55.46 / 55.41
Center (10% / 30% / 50%)	81.58 / 82.26 / 82.91	62.34 / 64.47 / 65.78	60.05 / 57.61 / 58.55
Classifier (10% / 30% / 50%)	80.80 / 82.14 / 81.04	62.02 / 63.94 / 62.16	60.62 / 60.40 / 61.02
All (10% / 30% / 50%, <b>Sec. III-A</b> )	84.00 / 82.20 / <b>84.24</b>	67.75 / 63.93 / <b>68.74</b>	55.76 / <b>53.26</b> / 54.87

TABLE I: Ablation study of our Bayesian ERFNet trained with varying dropout layers and probabilities  $p = \{10\%, 30\%, 50\%\}$  on the ISPRS Potsdam dataset [1] with  $T = 50$  MC dropout samples. Our best Bayesian ERFNet (all,  $p = 50\%$ ) outperforms non-Bayesian ERFNet (standard,  $p = 50\%$ ) in terms of segmentation performance (by 3.8% mIoU) and model uncertainty calibration (by 8.3% ECE).



Fig. 5: The urban ISPRS Potsdam orthomosaic [1]. We simulate images and labels with a square footprint, downwards-facing camera, and GSD of 15 cm/px to generate training (orange), validation (blue), and testing (red) data from disjoint areas.

- **Center:** Dropout layers after the last four and first two encoder and decoder non-bt-1D layers respectively.
- **Classifier:** Single dropout layer after the last decoder non-bt-1D layer before the classification head.
- **All:** Dropout layers after all non-bt-1D layers in the encoder and decoder as proposed in **Sec. III-A**.

All models are trained and evaluated on the 7-class urban aerial ISPRS Potsdam orthomosaic dataset [1]. We simulate images and labels at random uniformly chosen positions from 30 m altitude with a square footprint, downwards-facing camera, and ground sample distance (GSD) of 15 cm/px. In total, we create 4000, 1000, and 3500 training, validation, and testing datapoints, respectively, from disjoint areas as depicted in **Fig. 5**. We assess segmentation performance by pixel-wise global accuracy and mean Intersection-over-Union (mIoU) [28], and model uncertainty estimate quality by the expected calibration error (ECE). At test time, we use a reasonably large number of MC dropout samples  $T = 50$ .

**Table I** summarises our results. With highest accuracy and mIoU, Bayesian ERFNet-All, proposed in **Sec. III-A**, trained with  $p = 50\%$  performs best. Noticeably, our Bayesian ERFNet-All ( $p = 50\%$ ) outperforms its strongest non-Bayesian counterpart (standard,  $p = 50\%$ ) by 3.8% mIoU while resulting in 8.3% improved ECE. Qualitatively, **Fig. 6** confirms high uncertainty in misclassified, cluttered regions, hence providing a reliable acquisition function for AL (**Eq. (3)**). These results validate that our probabilistic

model interpretation via MC dropout improves performance and provides reliable uncertainty estimates.

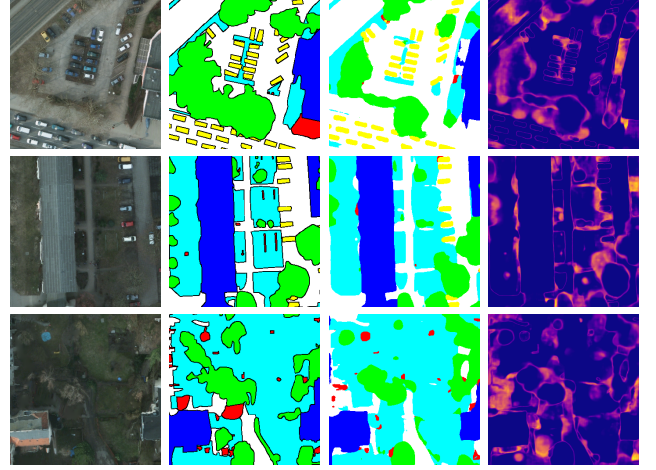


Fig. 6: Qualitative results with Bayesian ERFNet (all,  $p = 50\%$ ) trained on the ISPRS Potsdam dataset [1]. High model uncertainty in misclassified regions validates that our Bayesian ERFNet provides consistent uncertainty estimates.

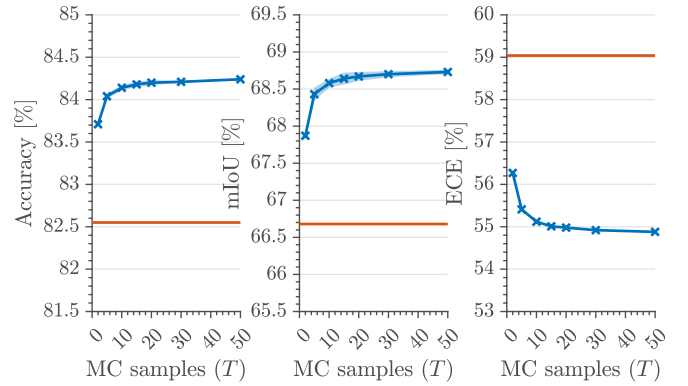


Fig. 7: Our Bayesian ERFNet (all,  $p = 50\%$ , blue) with varying MC dropout samples compared against non-Bayesian ERFNet (all,  $p = 50\%$ ) (orange) on the ISPRS Potsdam dataset [1]. Metrics are averaged over three trials with shaded regions indicating standard deviations. For  $T = 50$ , Bayesian ERFNet improves mIoU by 3.1% (left, middle) and reduces ECE by 7.6% (right).

To assess computational requirements and support the applicability of our method for online replanning, we study the performance of our Bayesian ERFNet with varying numbers of MC dropout samples  $T = \{2, 5, 10, 15, 20, 30, 50\}$  in

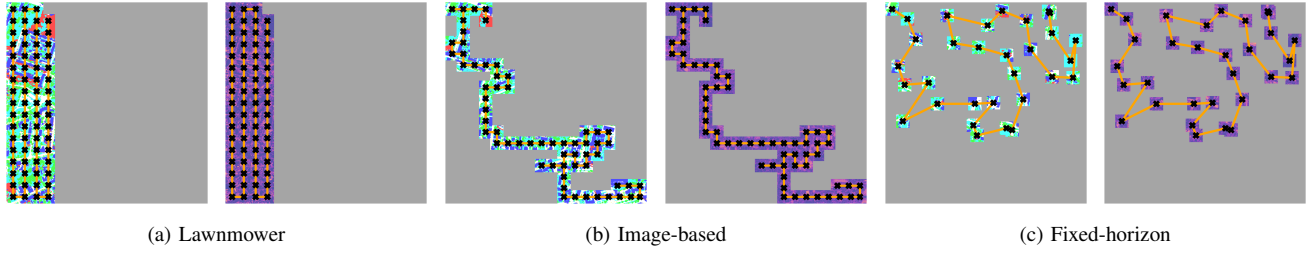


Fig. 8: Qualitative mapping and planning results for one mission with (a) the coverage baseline, (b) image-based, and (c) fixed-horizon strategy on the ISPRS Potsdam orthomosaic [1]. Orange lines show planned paths with black crosses indicating collected training images. Semantic and model uncertainty maps are shown, where grey areas depict unexplored terrain. Our planners in (b)-(c) adaptively plan to maximise the information value of collected images leading to superior AL performance over the coverage pattern in (a).

Fig. 7. As the number of MC dropout samples increases, segmentation performance and ECE both improve. Favourably for active robotic decision-making,  $T \approx 20$  samples are already sufficient for converging performance gains.

### B. Planning for Active Learning Evaluation

The following results on planning for AL suggest that our online planning framework reduces the number of labelled images needed to maximise segmentation performance compared to static coverage paths. Moreover, we show that fixed-horizon planning with a lookahead is better for AL performance than a greedy, one-step approach.

These claims are evaluated on the ISPRS Potsdam orthomosaic [1]. As in Sec. IV-A, we simulate images and labels from 30 m altitude to create 1000 validation and 3500 test datapoints. Data collection missions are executed in the disjoint  $900\text{m} \times 900\text{m}$  training area shown in Fig. 5. Our proposed Bayesian ERFNet (all,  $p = 50\%$ ) is pretrained on the Cityscapes dataset [29], so that each training starts from the same checkpoint. Each experiment considers 10 subsequent UAV data collection missions with a flight budget of  $B = 1800\text{ s}$  per mission, starting at  $(30, 30, 30)\text{ m}$ . After each mission, we retrain the model until convergence with batch size 8 and weight decay  $\lambda = (1 - p)/2N$ , where  $p = 0.5$  and  $N$  is the number of collected images up to the current mission. After each mission, the model is reset to its pretrained checkpoint to avoid catastrophic forgetting and other effects from simply accumulating train time.

To verify our claims, we compare AL performance of the planning methods introduced in Sec. III-C to commonly used static coverage paths. We assess segmentation performance by pixel-wise accuracy and mIoU over the number of training samples after each retraining. For a fair comparison, we implement a coverage baseline by changing path orientations and distances between measurement positions in the missions to maximise both spatial coverage and training data diversity. In our fixed-horizon planner, we use a lookahead of  $N = 5$ . Following Popović et al. [5], we apply the Covariance Matrix Adaptation Evolution Strategy [30] optimisation routine and tune its hyperparameters to trade-off between runtime and path quality. The frontier-based planner uses a candidate sampling distance of 15 m, and the image-based planner uses a step size of 50 m and an edge width of 10 px.

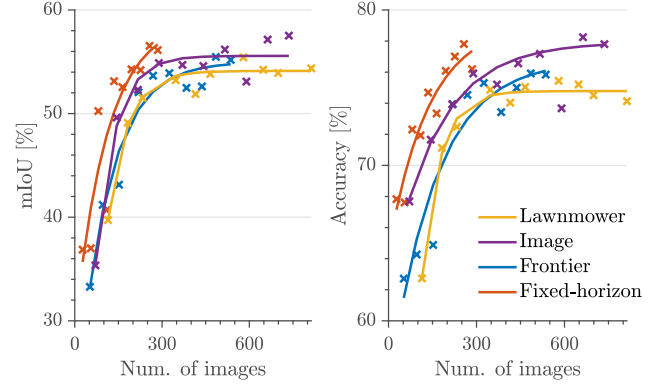


Fig. 9: Comparison of AL performance of our planners on the ISPRS Potsdam orthomosaic [1]. Steeper curves indicate better performance. Our informative planners either outperform or perform on par with the strong coverage baseline (yellow). Our fixed-horizon strategy (orange) performs best as it can plan multiple steps ahead based on the uncertainty map.

Fig. 9 summarises the planning results. Better AL performance is reflected by maximising accuracy and mIoU with less acquired images that must be labelled by a human operator. As indicated by the steeper-rising curves, our proposed fixed-horizon (orange) and local image-based (purple) strategies clearly outperform the strong coverage baseline (yellow), since they allow active decision-making for data collection. Frontier-based planning (blue) performs on par with the coverage baseline. Qualitatively, Fig. 8 confirms the adaptive behaviour of our approaches, which plan paths to target regions of high model uncertainty and gather images about different environment features. These results verify that our online planning framework effectively reduces human labelling effort while maximising segmentation performance.

Particularly, Fig. 9 shows that our fixed-horizon with lookahead strategy performs better for AL than the two greedy one-step methods. The fixed-horizon planner already reaches  $\sim 57\%$  mIoU and  $\sim 78\%$  accuracy with  $\sim 250$  labelled images while the second-best local image-based strategy requires  $\sim 520$  images, i.e. around twice the number of training samples, to achieve these levels. The superior performance of the fixed-horizon method demonstrates the benefits of planning multiple steps ahead in our integrated framework to achieve efficient training data acquisition.



## V. CONCLUSIONS AND FUTURE WORK

This paper introduced a general informative path planning framework for active learning in aerial semantic mapping scenarios. Our framework exploits a Bayesian FCN to estimate model uncertainty in semantic segmentation. During a mission, the semantic predictions and model uncertainty are used as an input for terrain mapping. A key aspect of our work is that we link the mapped model uncertainty to a planning objective guiding UAV to collect the most informative images for training the FCN. This pipeline reduces the total number of images that must be labelled to maximise segmentation performance, thus conserving human effort, time, and cost.

We validated our Bayesian FCN in an ablation study, showing that it both improves segmentation performance compared to a non-Bayesian baseline and provides consistent model uncertainty estimates as a basis for AL. The integrated system for informative planning with the FCN was evaluated in an urban mapping scenario using real-world data. Results show that our framework effectively reduces labelled training data requirements compared to static coverage paths. Moreover, we demonstrate that planning with a lookahead strategy in our approach improves data-gathering efficiency compared to simpler greedy one-step planners. Future work will investigate path planning at different altitudes and exploiting semantic map information to target specific informative classes, e.g. buildings or cars.

## ACKNOWLEDGEMENT

We would like to thank Matteo Sodano and Tiziano Guadagnino for help with our experiments and proofreading. We would like to thank Jan Weyler for providing a PyTorch Lightning implementation of ERFNet.

## REFERENCES

- [1] ISPRS. (2018) 2D Semantic Labeling Contest. [Online]. Available: <https://www.isprs.org/education/benchmarks/UrbanSemLab/semantic-labeling.aspx>
- [2] G. Lenczner, A. Chan-Hon-Tong, B. L. Saux, N. Luminari, and G. L. Besnerais, "DIAL: Deep Interactive and Active Learning for Semantic Segmentation in Remote Sensing," *arXiv preprint arXiv:2201.01047*, 2022.
- [3] R. Kemker, C. Salvaggio, and C. Kanan, "Algorithms for semantic segmentation of multispectral remote sensing imagery using deep learning," *ISPRS Journal of Photogrammetry and Remote Sensing (JPRS)*, vol. 145, pp. 60–77, 2018.
- [4] D. Tuia, F. Ratle, F. Pacifici, M. F. Kanevski, and W. J. Emery, "Active Learning Methods for Remote Sensing Image Classification," vol. 47, no. 7, pp. 2218–2232, 2009.
- [5] M. Popović, T. Vidal-Calleja, G. Hitz, J. J. Chung, I. Sa, R. Siegwart, and J. Nieto, "An informative path planning framework for UAV-based terrain monitoring," *Autonomous Robots*, vol. 44, no. 6, pp. 889–911, 2020.
- [6] A. C. Rodríguez, S. D'Aronco, K. Schindler, and J. D. Wegner, "Mapping oil palm density at country scale: An active learning approach," *Remote Sensing of Environment*, vol. 261, 2021.
- [7] F. Stache, J. Westheider, F. Magistri, M. Popović, and C. Stachniss, "Adaptive Path Planning for UAV-based Multi-Resolution Semantic Segmentation," in *Proc. of the Europ. Conf. on Mobile Robotics (ECMR)*. IEEE, 2021.
- [8] B. Kellenberger, D. Marcos, S. Lobry, and D. Tuia, "Half a Percent of Labels is Enough: Efficient Animal Detection in UAV Imagery Using Deep CNNs and Active Learning," *IEEE Trans. on Geoscience and Remote Sensing*, vol. 57, no. 12, pp. 9524–9533, 2019.
- [9] L. P. Osco, J. Marcato Junior, A. P. Marques Ramos, L. A. de Castro Jorge, S. N. Fatholahi, J. de Andrade Silva, E. T. Matsubara, H. Pistori, W. N. Gonçalves, and J. Li, "A review on deep learning in uav remote sensing," *International Journal of Applied Earth Observation and Geoinformation*, vol. 102, 2021.
- [10] A. Garcia-Garcia, S. Orts-Escolano, S. Oprea, V. Villena-Martinez, and J. G. Rodríguez, "A Review on Deep Learning Techniques Applied to Semantic Segmentation," *arXiv preprint arXiv:1704.06857*, 2017.
- [11] E. Romera, J. M. Álvarez, L. M. Bergasa, and R. Arroyo, "ERFNet: Efficient Residual Factorized ConvNet for Real-Time Semantic Segmentation," *IEEE Trans. on Intelligent Transportation Systems (ITS)*, vol. 19, no. 1, pp. 263–272, 2018.
- [12] Y. Gal, R. Islam, and Z. Ghahramani, "Deep Bayesian Active Learning with Image Data," in *Proc. of the Int. Conf. on Machine Learning (ICML)*, 2017, p. 1183–1192.
- [13] P. M. Blok, G. Kootstra, H. E. Elghor, B. Diallo, F. K. van Evert, and E. J. van Henten, "Active learning with MaskAL reduces annotation effort for training Mask R-CNN," *arXiv preprint arXiv:2112.06586*, 2021.
- [14] E. Galceran and M. Carreras, "A survey on coverage path planning for robotics," *Journal on Robotics and Autonomous Systems (RAS)*, vol. 61, no. 12, pp. 1258–1276, 2013.
- [15] X. Li and Y. Guo, "Adaptive Active Learning for Image Classification," in *Proc. of the IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2013, pp. 859–866.
- [16] P. Ren, Y. Xiao, X. Chang, P.-Y. Huang, Z. Li, B. B. Gupta, X. Chen, and X. Wang, "A Survey of Deep Active Learning," *ACM Computing Surveys*, vol. 54, no. 9, 2021.
- [17] O. Sener and S. Savarese, "Active learning for convolutional neural networks: A core-set approach," in *International Conference on Learning Representations*, 2018.
- [18] H. Blum, S. Rohrbach, M. Popović, L. Bartolomei, and R. Siegwart, "Active Learning for UAV-based Semantic Mapping," in *Robotics: Science and Systems 2nd Workshop on Informative Path Planning and Adaptive Sampling*, 2019.
- [19] A. Kendall and Y. Gal, "What uncertainties do we need in Bayesian deep learning for computer vision?" in *Proc. of the Advances in Neural Information Processing Systems (NIPS)*. Curran Associates, 2017, pp. 5575–5585.
- [20] G. Georgakis, B. Bucher, K. Schmeckpeper, S. Singh, and K. Daniilidis, "Learning to Map for Active Semantic Goal Navigation," *arXiv preprint arXiv:2106.15648*, 2021.
- [21] Y. Gal and Z. Ghahramani, "Dropout as a bayesian approximation: Representing model uncertainty in deep learning," in *Proc. of the Int. Conf. on Machine Learning (ICML)*. JMLR, 2016, p. 1050–1059.
- [22] A. Kendall, V. Badrinarayanan, and R. Cipolla, "Bayesian SegNet: Model Uncertainty in Deep Convolutional Encoder-Decoder Architectures for Scene Understanding," in *Proc. of British Machine Vision Conference (BMVC)*. BMVA Press, 2017, pp. 57.1–57.12.
- [23] G. A. Hollinger, B. Englot, F. S. Hover, U. Mitra, and G. S. Sukhatme, "Active planning for underwater inspection and the benefit of adaptivity," *Intl. Journal of Robotics Research (IJRR)*, vol. 32, no. 1, pp. 3–18, 2013.
- [24] A. Singh, A. Krause, and W. J. Kaiser, "Nonmyopic Adaptive Informative Path Planning for Multiple Robots," in *International Joint Conference on Artificial Intelligence (IJCAI)*, 2009, pp. 1843–1850.
- [25] T. Dang, C. Papachristos, and K. Alexis, "Autonomous exploration and simultaneous object search using aerial robots," in *IEEE Aerospace Conference*, 2018.
- [26] N. Houlsby, F. Huszár, Z. Ghahramani, and M. Lengyel, "Bayesian active learning for classification and preference learning," *arXiv preprint arXiv:1112.5745*, 2011.
- [27] A. Elfes, "Using occupancy grids for mobile robot perception and navigation," *Computer*, vol. 22, no. 6, pp. 46–57, 1989.
- [28] M. Everingham, L. Van Gool, C. K. Williams, J. Winn, and A. Zisserman, "The pascal visual object classes (voc) challenge," *Intl. Journal of Computer Vision (IJCV)*, vol. 88, no. 2, pp. 303–338, 2010.
- [29] M. Cordts, M. Omran, S. Ramos, T. Rehfeld, M. Enzweiler, R. Benenson, U. Franke, S. Roth, and B. Schiele, "The Cityscapes Dataset for Semantic Urban Scene Understanding," in *Proc. of the IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [30] N. Hansen and A. Ostermeier, "Completely derandomized self-adaptation in evolution strategies," *Evolutionary Computation*, vol. 9, no. 2, pp. 159–195, 2001.