

An Informative Path Planning Framework for Active Learning in UAV-based Semantic Mapping

Julius Rückin, Federico Magistri, Cyrill Stachniss, Marija Popović

Abstract—Unmanned aerial vehicles (UAVs) are frequently used for aerial mapping and general monitoring tasks. Recent progress in deep learning enabled automated semantic segmentation of imagery to facilitate the interpretation of large-scale complex environments. Commonly used supervised deep learning for segmentation relies on large amounts of pixel-wise labelled data, which is tedious and costly to annotate. The domain-specific visual appearance of aerial environments often prevents the usage of models pre-trained on publicly available datasets. To address this, we propose a novel general planning framework for UAVs to autonomously acquire informative training images for model re-training. We leverage multiple acquisition functions and fuse them into probabilistic terrain maps. Our framework combines the mapped acquisition function information into the UAV’s planning objectives. In this way, the UAV adaptively acquires informative aerial images to be manually labelled for model re-training. Experimental results on real-world data and in a photorealistic simulation show that our framework maximises model performance and drastically reduces labelling efforts. Our map-based planners outperform state-of-the-art local planning.

Index Terms—Informative Path Planning, Active Learning, Bayesian Deep Learning, Semantic Segmentation and Mapping

I. INTRODUCTION

Unmanned aerial vehicles (UAVs) enable highly agile, low-cost operations in various aerial imaging applications [1, 2], such as precision agriculture [3, 4], wildlife conservation [2], and urban planning [5–8]. Combined with advances in deep learning for semantic segmentation through fully convolutional neural networks (FCNs) [9, 10], deploying UAVs accelerates automated scene understanding in large-scale and complex aerial environments [11]. Classical deep learning-based semantic segmentation models often used in this context are usually trained on a static curated dataset in a supervised fashion only once before deployment. This leads to two major drawbacks. First, training a semantic segmentation model requires enormous amounts of pixel-wise labelled images, which is a repetitive and time-consuming process often executed by costly domain experts. Second, visual appearance can differ significantly between environments or change over time. Thus, a critical requirement for robot autonomy is the ability to learn about an environment by continuously improving the robot’s semantic perception with minimal expert guidance.

In this work, we examine the problem of active learning (AL) in UAV-based semantic mapping. Our goal is to

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

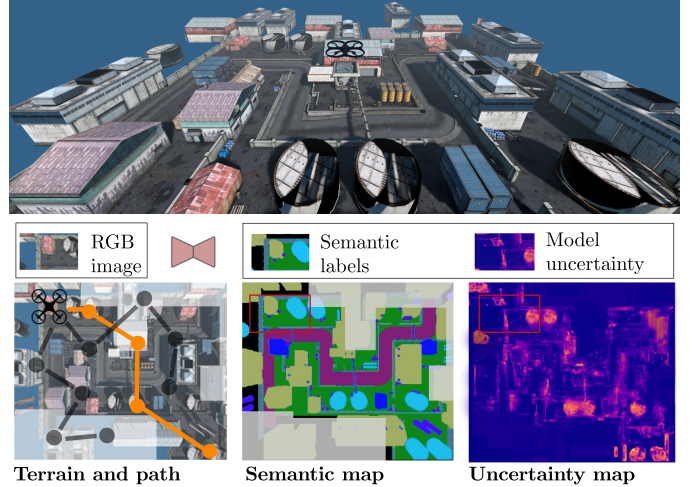


Fig. 1. Our general planning framework for active learning in UAV-based semantic mapping deployed in a photo-realistic simulator [12] (top). We compute an acquisition function, e.g. model uncertainty, and predict semantic segmentation online (centre-right) and fuse both in terrain maps (bottom-right). Our map-based planners replan a UAV’s path (orange, bottom-left) to collect the most informative, e.g. most uncertain (yellow), images for network re-training. Our approach reduces the number of images that must be manually labelled to maximise semantic segmentation performance.

improve the robot’s vision capabilities in initially unknown environments while minimising the total amount of human-labelled data. To this end, our approach exploits ideas from AL research and incorporates them into a new informative path planning (IPP) framework. The framework replans the UAV’s path online as new observations are collected to actively target regions of informative training data. The newly gathered images are labelled by a human annotator and used to re-train an FCN, maximising its semantic segmentation performance.

Various AL methods for machine learning effectively reduce the requirements for human-labelled training data [13–19]. Recently, AL approaches for deep learning models are gaining attention [20–25]. These works develop acquisition functions for selecting to-be-labelled training data to maximise model performance. However, they cannot be directly applied to robotic missions as they assume access to large pre-recorded unlabelled in-domain data pools. An open problem is how to leverage AL to improve robot perception with minimal expert guidance when operating in initially unknown environments. More recent AL works for aerial imagery consider the UAV to be a passive data collection device to record static data pools [2, 6]. In contrast, we aim to utilise the UAV’s decision-making capabilities to improve its perception and, thus, its reasoning about the environment for downstream tasks.

The main contribution of this work is a novel and generally applicable active planning framework linking ideas from AL to robotic planning objectives, as illustrated in Fig. 1. The key benefit of our approach is that it reduces human labelling effort for continuous robotic perception improvement. We exploit various model uncertainty and training data novelty estimation techniques for deep learning models [26–28] and apply them to semantic segmentation with a pre-trained FCN [29]. The inferred pixel-wise semantic labels and estimated model uncertainty and novelty scores are fused sequentially into a probabilistic terrain map as new observations are acquired. As a key feature, our new IPP framework iteratively replans the UAV’s path to collect the most informative, i.e. the most uncertain or novel, images for labelling and model re-training in a targeted fashion.

This article builds upon our previous conference paper [30]. In our previous work [30], we proposed an IPP approach linking globally mapped model uncertainties to a robotic planning objective for AL in UAV-based terrain monitoring. This submission extends our previous method and generalises experimental findings in the following ways. First, we introduce a general IPP framework for AL in UAV-based semantic mapping by linking various uncertainty- and representation-driven acquisition functions to planning objectives, as opposed to just utilising model uncertainties computed via Monte-Carlo (MC) dropout. Second, we propose new terrain mapping features to improve map-based planning compared to our conference version. Third, we systematically evaluate mapping, planners, and AL planning objectives on new datasets from different domains and in a photorealistic simulator. We present a thorough empirical analysis of combining AL acquisition functions and IPP approaches, giving new insights into how to connect AL and autonomous robotic decision-making.

In sum, we make the following four claims. First, our active planning framework for AL in UAV-based semantic mapping reduces the number of labelled images needed to maximise segmentation performance compared to both traditionally used coverage and random walk data collection. Second, probabilistic global mapping of gathered information enhances map-based planning performance for AL. Third, our map-based planners outperform state-of-the-art local planning for AL [31]. Fourth, we demonstrate the generality of our approach, showing that it significantly reduces labelling effort in largely varying domains irrespective of the used uncertainty estimation methods, planning strategies, and AL-based objectives. We open-source our code for usage by the community at: <https://github.com/dmar-bonn/ipp-al-framework>.

II. RELATED WORK

Our goal is to collect the most informative images to train a semantic segmentation model with a minimal amount of labelled data using UAVs in aerial mapping missions. Our approach combines advances in AL with IPP. This section overviews how our work is placed within these research areas.

A. Active Learning

Active learning aims to maximise model performance while minimising labelled training data. It assumes the existence

of a large unlabelled data pool, then iteratively selects a data point from the pool by maximising an acquisition function until a labelling budget is exceeded [13–16]. Settles et al. [32] provide a comprehensive overview of AL approaches for low-dimensional machine learning problems. Recent AL approaches focus on training deep learning models from high-dimensional inputs, e.g. images, where a single data point has a negligible effect on model performance. AL methods for deep learning collect a batch of data from the pool instead of single data points, called batch-mode AL [20, 21, 33, 34]. However, these strategies are not applicable in robotic settings since they reason about which images from an existing large data pool should be labelled. In contrast, we propose an IPP framework for AL collecting new batches of to-be-labelled data directly during a mission in initially unknown environments. We link the AL acquisition function to an IPP objective, adaptively guiding the UAV towards regions of informative training data. Further, we answer the following two open research questions. First, how to incorporate recently proposed acquisition functions [20, 31, 35] into our IPP framework and second, in which ways planners, planning objectives, and terrain mapping influence AL performance.

Uncertainty-based AL methods select data with the highest model uncertainty [17, 20, 34, 35]. Early methods use Gaussian processes [17] or support-vector machines [34] to quantify model uncertainty in tasks with low-dimensional inputs. Measuring model uncertainty in deep neural networks is computationally challenging due to their parameter space dimensionality. One approach aims at estimating the model uncertainty deterministically in a single forward pass. Although computationally efficient, these methods are often not well-calibrated in real-world vision tasks [28]. Alternatively, Gal et al. [26] propose using dropout at test time, called MC dropout, to efficiently approximate the Bayesian posterior over the network parameters. They utilise MC dropout in acquisition functions applied to image classification maximising model uncertainty [20]. Other works use neural network ensembles for uncertainty estimation [27, 36]. Each network is independently initialised and trained. Ensembles achieve higher prediction performance and better calibration than MC dropout [35, 37]. Further, recent advances make ensemble training computationally more efficient [37, 38]. In this work, we study the applicability of different uncertainty-based AL objectives in a robotic planning context.

Representation-based AL methods maximise training data diversity by selecting data points with novel representations in feature space [21–23]. Generative adversarial network-inspired approaches use a generator learning the joint data representation, while the discriminator distinguishes labelled and unlabelled data [22, 23]. Sener et al. [21] select a number of data points, called a core-set, geometrically covering a data pool in the model’s latent space with a minimal number of data points. However, both approaches require large in-domain data pools to learn rich representations of the data-generating distribution. These methods are impractical in our scenario as autonomous robots operate in unknown and visually varying environments. In contrast, Blum et al. [31] propose a method for quantifying data novelty in semantic segmentation tasks

without access to large in-domain data pools by using kernel-density estimation of unlabelled images in the network’s latent space. They use this novelty estimation in a local planning objective and apply it for AL in aerial semantic mapping. We integrate their novelty estimation into new global map-based planning objectives. We rigorously analyse its AL performance using various global planning schemes and datasets, outperforming their local planning strategy.

B. Informative Path Planning

Informative path planning enables autonomous robots to efficiently and actively explore initially unknown environments subject to platform constraints, such as battery capacity [39]. IPP methods have been applied to various environmental monitoring scenarios, including lake monitoring [40], underwater inspection [41], infrastructure surface inspection [42], and agricultural monitoring [43]. We distinguish between non-adaptive approaches, which precompute paths before a mission starts, e.g. coverage planning [44], and adaptive approaches, which replan paths online as new data is collected [3, 40, 41]. We focus on adaptive methods as our goal is to collect informative training data on-the-fly.

Combinatorial approaches solve IPP problems in a near-optimal fashion [45–48]. However, they exhaustively query the search space scaling exponentially in problem size, which makes most of them impractical for online replanning. In contrast, sampling approaches break the curse of dimensionality to increase the computational efficiency of online IPP [49, 50]. Hollinger et al. [49] propose receding-horizon rapidly exploring information gathering algorithms to sample motion plans. Choudhury et al. [50] combine MC planning with cost-benefit rollouts to increase sampling efficiency.

Similarly, optimisation approaches directly optimise IPP objectives [3, 40, 51]. Vivaldini et al. [51] utilise Bayesian optimisation to choose a sequence of informative measurement positions for UAV-based tree disease monitoring. In a similar problem setup, Hitz et al. [40] leverage the covariance matrix adaptation evolution strategy (CMA-ES) to optimise a sequence of measurement positions. Popović et al. [3] extend this approach by introducing a greedily optimised initial sequence of measurement positions, then using the CMA-ES to fine-tune the initial solution resulting in more informative paths.

Geometric approaches collect candidate measurement positions for efficient exploration. The candidate maximising an objective function is chosen as the most informative one [52–54]. Gonzalez et al. [54] choose the position maximising the potentially visible unexplored space. Ghaffari et al. [52] generate candidate positions along probabilistic frontiers of explored space, greedily selecting the one which maximises the expected information gain. Similarly, Cheng et al. [53] train an agent choosing frontiers to minimise localisation uncertainty and maximise information gain.

The above-mentioned works consider adaptive IPP for mapping environmental phenomena. In contrast, our framework applies planning algorithms to the problem setting of improving robot vision with minimal human labelling effort. We design new IPP objective functions to replan paths towards

informative training data as new observations are collected and demonstrate their integration into map-based planners.

C. Informative Path Planning for Active Learning

Using autonomous robots to reduce manual labelling effort for training deep learning models is a relatively unexplored research area. Georgakis et al. [55] propose a framework for active semantic goal navigation which uses ensembles to estimate model uncertainty in their planning objective. Other methods introduce self-supervised approaches to improve or adapt the robot’s perceptions to new environments without the need for manual labelling. Frey et al. [56] introduce a self-improving continual learning framework for semantic segmentation in indoor scenes without manual labelling by generating pseudo-labels from 3D maps. Zurbrugg et al. [57] extend this approach to an embodied agent autonomously navigating towards high training data novelty viewpoints. Chaplot et al. [58] suggest a similar self-supervised approach for semantic segmentation in indoor scenes training an exploration policy with reinforcement learning (RL) to target uncertain 3D map parts. The policy training depends on the simulation environment and the currently trained network at the same time. As RL performance degrades with simulation to real-world gaps, this method requires the availability of realistic domain-specific simulators and introduces policy re-training costs after each network re-training. Further, as discussed by Chaplot et al. [58], the approaches [56–58] rely on large labelled indoor datasets for pre-training a semantic segmentation model to produce high-quality pseudo-labels in new indoor scenes. If the pre-trained model misclassifies objects, these errors not only prevent learning semantics but could even be reinforced in the case of over-confident predictions. Zurbrugg et al. [57] experimentally show that the expected model improvement strongly depends on the chosen pre-training dataset and environment the robot is deployed in. Aerial mapping missions, as in our problem setting, present much more visual variability, with very little and often small pre-training datasets being available, further exacerbating these issues. As the environment and domain are initially unknown, these purely self-supervised methods require enormous engineering work to relax the above-mentioned assumptions and are not directly applicable to our use case.

Most similar to our work is the local planning approach of Blum et al. [31] for AL in semantic mapping. Their planning objective aims to promote training data novelty in semantic prediction tasks. We combine their ideas on novelty estimation for AL with our previous work on IPP for AL [30]. In contrast to Blum et al. [31], we propose a general and unified IPP framework supporting probabilistic semantic mapping, various acquisition functions, planning objectives, and map-based planning algorithms. Further, we provide in-depth empirical analyses and show that our map-based planners outperform existing methods [30, 31].

III. OUR APPROACH

We present our general IPP framework for AL in UAV-based semantic mapping. Our setup considers a UAV collecting images of a flat terrain using a downwards-facing

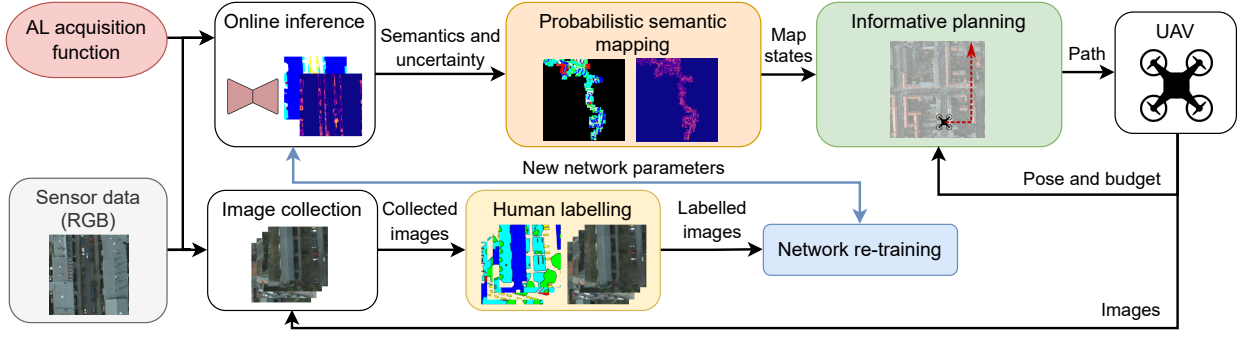


Fig. 2. Overview of our approach. We start with a pre-trained semantic segmentation network deployed on a UAV. During a mission, the network processes RGB images to predict pixel-wise semantic labels, model uncertainties (Sec. III-A1), and novelty scores (Sec. III-A2), which are projected onto the terrain to build global maps capturing these variables (Sec. III-B). Based on the current UAV position, budget, and posterior map state, our algorithm plans paths for the UAV to collect 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 re-training. By guiding the UAV to collect informative training data, our pipeline reduces the human labelling effort.

RGB camera. Assuming no further prior knowledge about the terrain, the goal is to autonomously collect informative training data to improve the robot’s perception with minimal human labelling effort. As shown in Fig. 2, our framework links AL with planning objectives guiding the UAV to regions of informative training data. As new data is collected, we utilise a lightweight FCN to predict pixel-wise semantics. Further, we estimate the pixel-wise model uncertainty associated with the prediction and training data novelty of the collected image and then fuse them into probabilistic terrain maps. The UAV position, its remaining budget, and the current map state are combined into new AL-based information objectives used to replan the future path towards informative training data. A key feature of our framework is its general applicability, as it is agnostic to the chosen network and supports different uncertainty estimation techniques, mapping methods, and map-based planners. The following subsections detail the framework’s individual modules and the specific methods we investigate in this work.

A. Active Learning Acquisition Functions

We first derive measures for an image’s information value when a network is re-trained on this data. To this end, AL works propose two main paradigms, uncertainty-based and representation-based acquisition functions. We demonstrate the generality of our approach using either paradigm.

We adapt the ERFNet encoder-decoder architecture proposed by Romera et al. [29] depicted in Fig. 3 to our AL use case. Although our framework is agnostic to the chosen network architecture, the lightweight ERFNet is particularly suitable for online robot deployment with limited computational resources. In the following, the model $f^{\mathbf{W}}(\cdot)$ is parameterised by weights \mathbf{W} and outputs a probability tensor $p(\mathbf{y} | f^{\mathbf{W}}(\mathbf{z})) = \text{softmax}(f^{\mathbf{W}}(\mathbf{z})) \in [0, 1]^{K \times w \times h}$, where \mathbf{z} is the input RGB image with width w and height h , and \mathbf{y} is the pixel-wise semantic label over the K classes. The training set contains N images $\mathbf{Z} = \{\mathbf{z}_1, \dots, \mathbf{z}_N\}$ and semantic labels $\mathbf{Y} = \{\mathbf{y}_1, \dots, \mathbf{y}_N\}$. Our network is trained to minimise cross-




Fig. 3. ERFNet architecture proposed by Romera et al. [29]. The network takes an RGB image (left) as input and outputs semantic labels (right). We utilise the network in our ensemble method to predict model uncertainty.

entropy with weight decay regularisation factor λ :

$$\mathcal{L}(\theta) = -\frac{1}{N} \sum_{i=1}^N \log p(\mathbf{y}_i | f^{\mathbf{W}}(\mathbf{z}_i)) + \lambda \|\mathbf{W}\|_2^2. \quad (1)$$

The following subsections describe different methods to estimate the information value of a candidate image for AL.

1) *Bayesian Uncertainty-based Methods*: We estimate pixel-wise model uncertainty over the prediction $p(\mathbf{y} | f^{\mathbf{W}}(\mathbf{z}))$ as a measure for the informativeness of image \mathbf{z} for re-training [20, 26, 35, 59, 60]. We leverage advances in Bayesian deep learning, transforming the deterministic ERFNet into a probabilistic version [30]. We consider using two alternative methods: Monte-Carlo (MC) dropout [26] and ensembles [35]. To measure model uncertainty, we utilise Bayesian active learning by disagreement [60], which computes the mutual information between the unknown labels \mathbf{y} and the posterior distribution over weights $p(\mathbf{W} | \mathbf{Z}, \mathbf{Y})$. However, the weights’ posterior is intractable for FCNs [26]. Thus, we approximate the true posterior prediction [59]:

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

where we independently sample T weights $\hat{\mathbf{W}}_i$ from a prior weight distribution $q(\mathbf{W})$ performing MC integration.




Fig. 4. Representation-based image novelty score [31]. An RGB image z is passed through our ERFNet encoder and its latent vectors $r_{i,j}^z \in \mathbf{r}^z$ are extracted along the channel dimension. We compute the cosine distance between each $r_{i,j}^z$ and its k -nearest neighbors from the training representation vector database. The resulting novelty image is upsampled to the spatial dimensions of z . Last, we add all $r_{i,j}^z$ to the representation vector database. Lighter colours indicate higher novelty, i.e. higher informativeness for AL.

MC dropout and ensemble methods provide two alternative approaches to construct the prior $q(\mathbf{W})$. In MC dropout, dropout is applied independently to the weights \mathbf{W} before each of the T forward passes at test time. In the ensemble method, we train T independently randomised ERFNet models with stochastic mini-batch gradient descent. For further details on MC dropout and ensembles, we refer to [26, 30] and [35], respectively. Following Gal et al. [20], we approximate the mutual information using Eq. (2):

$$\mathbb{I}(\mathbf{y}, \mathbf{W} | \mathbf{z}, \mathbf{Z}, \mathbf{Y}) \approx -\hat{p}(\mathbf{y} | \mathbf{z}, \mathbf{Z}, \mathbf{Y})^T \log(\hat{p}(\mathbf{y} | \mathbf{z}, \mathbf{Z}, \mathbf{Y})) + \frac{1}{T} \sum_{i=1}^T p(\mathbf{y} | \mathbf{z}, \hat{\mathbf{W}}_i)^T \log(p(\mathbf{y} | \mathbf{z}, \hat{\mathbf{W}}_i)), \quad (3)$$

where $\log(\cdot)$ is applied element-wise. Intuitively, model uncertainty is high whenever the posterior prediction entropy is high, while single prediction entropy is low, but disagreeing with each other. We exploit this measure to guide the UAV towards more informative areas, i.e. regions of high model uncertainty. Note that our framework is agnostic to both, the model uncertainty estimation method and the chosen network.

2) *Representation-based Method*: Inspired by recent AL works [21–23], we study a representation-based planning objective as an alternative to uncertainty-based objectives. We deterministically quantify the network’s confidence in its prediction by estimating the image’s novelty to the network $f^{\mathbf{W}}$ given training images and labels \mathbf{Z}, \mathbf{Y} [28, 31, 61, 62]. Intuitively, the image’s novelty is high whenever the network’s latent representation of a new image z and training images \mathbf{Z} is dissimilar. Although confidence measures for classification are well-known [61, 62], they are not directly applicable for semantic segmentation as they do not provide pixel-wise scores and are not invariant to object locations. Thus, we utilise the novelty measure for semantic segmentation proposed by Blum et al. [31].

We perform kernel-density estimation in the network latent space by computing the average cosine distance between the latent representations of image z and its k -nearest latent

representations of training images \mathbf{Z} . We exploit the FCN’s architecture, where the network $f^{\mathbf{W}}(\cdot) = d^{\mathbf{W}_d}(e^{\mathbf{W}_e}(\cdot))$ consists of an encoder $e^{\mathbf{W}_e}$ parameterised by \mathbf{W}_e , a decoder $d^{\mathbf{W}_d}$ parameterised by \mathbf{W}_d , and $\mathbf{W} = \{\mathbf{W}_e, \mathbf{W}_d\}$. Specifically, we extract representations $e^{\mathbf{W}_e}(z) = \mathbf{r}^z \in \mathbb{R}^{\frac{w}{8} \times \frac{h}{8} \times C}$ after the encoder’s last convolutional layer with spatial dimensions downsampled by a factor of 8 compared to the image, and C channel dimensions as induced by the ERFNet architecture. Hence, $\mathbf{r}_{i,j}^z$ is a C -dimensional latent vector of the (i, j) -th 8×8 pixels patch of image z . After model training, we generate a database $R = \{\mathbf{r}_{1,1}^z, \dots, \mathbf{r}_{\frac{w}{8} \cdot \frac{h}{8}, N}^z\}$ of $\frac{w}{8} \cdot \frac{h}{8} \cdot N$ patch-wise representations of the training images \mathbf{Z} . Given an image z at inference time, its (i, j) -th novelty score is:

$$r(z)_{i,j} = \frac{1}{k} \sum_{\mathbf{r} \in NN(\mathbf{r}_{i,j}^z)} 1 - \frac{|\mathbf{r}^\top \mathbf{r}_{i,j}^z|}{\|\mathbf{r}_{i,j}^z\|_2 \|\mathbf{r}\|_2}, \quad (4)$$

where $NN(\mathbf{r}_{i,j}^z)$ is the set of k -nearest neighbors of $\mathbf{r}_{i,j}^z$ in R with respect to the cosine distance. Intuitively, higher novelty indicates higher informativeness of image z for re-training. Fig. 4 provides a schematic of an image’s novelty score computation. For more details, we refer to Blum et al. [31].

A key feature of our framework is that it can easily be adapted to other acquisition functions and FCNs. This work shows its generality using the uncertainty- and representation-based objectives with ERFNet, as described above.

B. Probabilistic Semantic Mapping

An important basis for our new planning objective functions is our 2D multi-layer terrain map. This map captures global semantics, model uncertainties, representation novelties, and training data statistics to provide different sources of information for informative planning. We propose a probabilistic mapping module updating this information online as the UAV collects new images of the terrain. To achieve this, we utilise sequential probabilistic occupancy grid mapping [63] to update each map layer when a new measurement arrives. We discretise the terrain into three 2D maps $\mathcal{G}_S : G \rightarrow \{0, 1\}^{K \times W \times L}$, $\mathcal{G}_U : G \rightarrow [0, 1]^{W \times L}$, $\mathcal{G}_R : G \rightarrow [0, 1]^{W \times L}$ defined over a grid lattice G with $W \times L$ spatially independent cells capturing the discrete semantic classes, continuous model uncertainties, and continuous novelty scores.

The semantic map \mathcal{G}_S consists of K independent layers $\mathcal{G}_{S_i} : G \rightarrow \{0, 1\}^{W \times L}$ to map $i \in [K] > 2$ classes. Each grid cell’s $G^{m,n}$ random state follows a uniform prior distribution $\mathcal{G}_{S_i}^{m,n} \sim p(\mathcal{G}_{S_i}^{m,n} = 1) = \frac{1}{K}$. When a new image z_t arrives at time step t , the semantic predictions $\hat{p}(\mathbf{y} | f^{\mathbf{W}}(z_t))$, see Eq. (2), are projected to the flat terrain given the UAV position $\mathbf{p}_t \in \mathbb{R}^3$ and camera intrinsics. We utilise standard occupancy grid mapping for each layer i and cell $G^{m,n}$ computing the posterior belief $\mathcal{G}_{S_i}^{m,n} \sim p(\cdot | z_{1:t}, \mathbf{p}_{1:t})$:

$$l(\mathcal{G}_{S_i}^{m,n} | z_{1:t}, \mathbf{p}_{1:t}) = l(\mathcal{G}_{S_i}^{m,n} | z_t, \mathbf{p}_t) + l(\mathcal{G}_{S_i}^{m,n} | z_{1:t-1}, \mathbf{p}_{1:t-1}) - l(\mathcal{G}_{S_i}^{m,n}), \quad (5)$$

where $l(\cdot)$ are the log odds of the binary random variable, $p(\mathcal{G}_{S_i}^{m,n} | z_t, \mathbf{p}_t)$ is given by the projected semantic predic-

tions, $p(\mathcal{G}_{S_i}^{m,n} | \mathbf{z}_{1:t-1}, \mathbf{p}_{1:t-1})$ is the recursive map belief, and $p(\mathcal{G}_{S_i}^{m,n})$ is the map prior.

The model uncertainties and novelties are stored in the maps \mathcal{G}_U and \mathcal{G}_R with prior means $\mu_{U,0}$ and $\mu_{R,0}$ respectively. We fuse projected uncertainties \mathbf{u}_t given by Eq. (3) and novelty scores $r(\mathbf{z}_t)$ given by Eq. (4) using maximum likelihood estimation assuming normally distributed \mathcal{G}_U and \mathcal{G}_R . We maintain a hit map $H : G \rightarrow \mathbb{N}^{W \times L}$ counting the total number of times a grid cell was updated during a mission. Then, we update the means $\mu_{U,t}^{m,n}$ and $\mu_{R,t}^{m,n}$ for a grid cell $G^{m,n}$ by:

$$\begin{aligned} \mu_{U,t}^{m,n} &= \mu_{U,t-1}^{m,n} + \frac{1}{H(G^{m,n})}(\mathbf{u}_t^{m,n} - \mu_{U,t-1}^{m,n}), \\ \mu_{R,t}^{m,n} &= \mu_{R,t-1}^{m,n} + \frac{1}{H(G^{m,n})}(r(\mathbf{z}_t)^{m,n} - \mu_{R,t-1}^{m,n}). \end{aligned} \quad (6)$$

Last, we store a map $T : G \rightarrow \mathbb{N}^{W \times L}$ to count how often grid cells occur in the training data set to foster data diversity in our proposed planning objectives. Note that the maps $H(\cdot)$ and $T(\cdot)$ are different as the camera could provide a high-frequency image stream for mapping while images only at the planned measurement position are collected for training.

A key feature of our mapping approach is that we accumulate and update the information between missions by updating the map prior. After each UAV mission, the network is re-trained on the collected training data. Re-training changes the semantic predictions, model uncertainty, and representation novelty estimates. Thus, we store all previously collected data and corresponding UAV positions. After re-training, we predict semantics, model uncertainties, and representation novelties of the stored data and sequentially fuse them according to Eq. (5) and Eq. (6). Our informed map prior strategy enhances map-based planning by avoiding exploring from scratch or replanning with outdated terrain knowledge.

C. Informative Path Planning

We develop IPP algorithms to guide a UAV to adaptively collect useful training data for our FCN. Our key idea is to link acquisition functions introduced in Sec. III-A to planning objective functions. Our planning strategies use the probabilistic terrain maps presented in Sec. III-B to guide the UAV online towards informative training data in an unknown terrain.

In general, IPP algorithms optimise an information criterion $I : \Psi \rightarrow \mathbb{R}_{\geq 0}$ over paths $\psi = (\mathbf{p}_1, \dots, \mathbf{p}_P) \in \Psi$ defined by P measurement positions $\mathbf{p}_i \in \mathbb{R}^3$:

$$\psi^* = \operatorname{argmax}_{\psi \in \Psi} I(\psi), \text{ s.t. } C(\psi) \leq B, \quad (7)$$

where $B \geq 0$ is the mission budget, e.g. flight time, and Ψ is the set of all possible paths of length P . The function $C : \Psi \rightarrow \mathbb{R}_{\geq 0}$ defines the cost of executing a path ψ :

$$C(\psi) = \sum_{i=1}^{P-1} c(\mathbf{p}_i, \mathbf{p}_{i+1}), \quad (8)$$

where $c : \mathbb{R}^3 \times \mathbb{R}^3 \rightarrow \mathbb{R}_{\geq 0}$ computes the flight time between two measurement positions assuming constant acceleration and deceleration $\pm a$, and maximum velocity v . The key insight of our work is to couple the AL acquisition functions with IPP




Fig. 5. Our planning strategies for training data collection. Dark gray dots and lines indicate candidate measurements and paths evaluated based on their estimated information adding these images to the training set. Orange dots and lines indicate the most informative chosen measurements and paths. Light gray depicts unexplored terrain. In (c), black indicates the greedy initialisation, and ellipses indicate the optimisation of candidate paths in continuous space.

information criteria $I(\cdot)$. This allows us to maximise model performance and minimise the labelling effort resulting from collecting images along a planned path ψ .

We propose four different replanning strategies in our framework, one local image-based and three global frontier, sampling and optimisation schemes with information criteria $I(\cdot)$ optimising Eq. (7) given the current terrain map states. The planners are illustrated in Fig. 5. In our experimental evaluation, we compare each planner's performance in terms of segmentation performance over the total labelling cost.

In the following, we exemplarily present our planning objectives with respect to the globally mapped model uncertainties $\mathcal{G}_{U,t}$ at a time step t (Sec. III-A1). In case of the representation-based objective (Sec. III-A2), we substitute the uncertainties $\mathcal{G}_{U,t}$ with novelties $\mathcal{G}_{R,t}$, see Eq. (6). This variable can also be changed to capture other AL acquisition functions.

Local planner. Our local image-based planner follows the direction of the highest estimated training data information in the image recorded at the current UAV position. Specifically, we choose the direction of the image edge $e_{\mathbf{z}_t}^*$ with the highest average AL value normalised by the current training data counts $T_t(\mathbf{p}_t)$ in the grid's subset spanned by the camera field of view from position \mathbf{p}_t projected to the flat terrain. This way, we select neighboring informative images while locally fostering training data diversity. Then, \mathbf{p}_{t+1}^* is reached by taking a predefined step size towards the direction of edge $e_{\mathbf{z}_t}^*$ at a fixed altitude. This resembles the planner proposed by Blum et al. [31] and generalises it to any AL objective.

Frontier-based planner. Our global geometric planner guides the UAV towards frontiers of the explored terrain with the highest AL objective in the terrain map. We use the hit map $H(\cdot)$ to identify exploration frontiers. Particularly, we greedily choose the next-best measurement position \mathbf{p}_{t+1}^* from a set of candidate positions \mathbf{p}_{t+1}^c equidistantly sampled along the frontiers at a fixed altitude. As the planner acts greedily, optimising Eq. (7) reduces to selecting the path $\psi^* = (\mathbf{p}_{t+1}^*)$:

$$\mathbf{p}_{t+1}^* = \operatorname{argmax}_{\mathbf{p}_{t+1}^c} I((\mathbf{p}_{t+1}^c)) = \operatorname{argmax}_{\mathbf{p}_{t+1}^c} \frac{\|\mathcal{G}_{U,t}(\mathbf{p}_{t+1}^c)\|_1}{\|T_t(\mathbf{p}_{t+1}^c)\|_1}, \quad (9)$$

where $\mathcal{G}_{U,t}(\mathbf{p}_{t+1}^c)$ and $T_t(\mathbf{p}_{t+1}^c)$ are the globally mapped model uncertainties and training data counts within the camera field of view from position \mathbf{p}_{t+1}^c , and $\|\cdot\|_1$ is the norm summing all elements in these subsets. This way, our frontier

planner trades off both exploration of unknown space for data diversity and focusing on regions potentially valuable for AL. **Optimisation-based planner.** Our optimisation-based planner selects a path ψ_{t+1}^* over a fixed horizon of multiple time steps. We utilise a two-step approach for efficient online replanning inspired by Popović et al. [3]. First, we greedily select a path ψ_{t+1}^g of length P over a grid above the terrain. Second, we use an optimisation procedure to fine-tune ψ_{t+1}^g in the continuous UAV workspace and return the next-best path ψ_{t+1}^* .

First, we iteratively select a path $\psi_{t+1}^g = (\mathbf{p}_{t+1}^g, \dots, \mathbf{p}_{t+P}^g)$, where each measurement position \mathbf{p}_{t+i}^g , $i \in \{1, \dots, P\}$, is greedily chosen over a sparse lattice F of discrete candidate positions \mathbf{p}^c at a fixed altitude:

$$\mathbf{p}_{t+i}^g = \operatorname{argmax}_{\mathbf{p}^c \in F} \frac{\|\mathcal{G}_{U,t}(\mathbf{p}^c)\|_1}{c(\mathbf{p}_{t+i-1}, \mathbf{p}^c) \|T_{t+i-1}(\mathbf{p}^c)\|_1}, \quad (10)$$

where $T_{t+i-1}(\mathbf{p}^c)$ is the subset of the forward-simulated training data count map given by the camera field of view at position \mathbf{p}^c . The forward simulation of the current map T_t based on the previously selected positions $(\mathbf{p}_{t+1}^g, \dots, \mathbf{p}_{t+i-1}^g)$ is crucial as one cannot forward-simulate model uncertainties. Forward-simulating T_t linearly decreases uncertainty with the number of a grid cell's training set occurrences. This fosters data diversity and terrain exploration.

Second, we refine the greedy positions of ψ_{t+1}^g in parallel in the continuous UAV workspace. To this end, we initialise an optimisation procedure with the greedy solution ψ_{t+1}^g and extend Eq. (10) to an information criterion $I(\cdot)$ evaluating candidate paths $\psi_{t+1}^o = (\mathbf{p}_{t+1}^o, \dots, \mathbf{p}_{t+P}^o)$:

$$I(\psi_{t+1}^o) = \frac{\sum_{i=1}^P \|\mathcal{G}_{U,t}(\mathbf{p}_{t+i}^o)\|_1}{c(\mathbf{p}_{t+i-1}^o, \mathbf{p}_{t+i}^o) \|T_{t+i-1}(\mathbf{p}_{t+i}^o)\|_1}. \quad (11)$$

The candidate path $\psi_{t+1}^* = (\mathbf{p}_{t+1}^*, \dots, \mathbf{p}_{t+P}^*)$ maximising Eq. (11) is chosen and measurement position \mathbf{p}_{t+1}^* is executed. We found that normalising AL information of a path by its execution costs leads to more efficient budget allocation. This planning strategy supports any optimisation algorithm, which can optimise objective function Eq. (11).

Sampling-based planner. Our sampling-based planner utilises Monte-Carlo tree search (MCTS) [64] to optimise a next-best measurement position \mathbf{p}_{t+1}^* in a non-myopic fashion. We simulate a number of future paths $\psi_{t+1} = (\mathbf{p}_{t+1}^{n_1}, \dots, \mathbf{p}_{t+P}^{n_P})$ of length P at a fixed altitude. Each tree node n_i at depth $i \in \{0, \dots, P\}$ is uniquely defined by its state $S^{n_i} = \{\mathbf{p}_{t+i}^{n_i}, T_{t+i}^{n_i}, \bar{B}^{n_i}\}$ consisting of a measurement position $\mathbf{p}_{t+i}^{n_i}$, forward-simulated training data count map $T_{t+i}^{n_i}$ along the tree's traversed path to node n_i , and remaining mission budget \bar{B}^{n_i} . The tree's root node n_0 is defined by $S^{n_0} = \{\mathbf{p}_t, T_t, B\}$, where \mathbf{p}_t , T_t and B are the current UAV position, training data count map, and mission budget. At each node, the planner selects the next position from a discrete set of actions with different step sizes and orientations. While traversing the search tree, we use the upper confidence bound bandit algorithm [64] to choose a child node. When reaching a leaf node, we roll out the path by sampling actions uniformly at random in each subsequent node until the remaining budget is exceeded or path length P is reached. A simulated path's




Fig. 6. Examples of paths planned by the baseline strategies on ISPRS Potsdam [5]. Orange lines show paths planned in one mission, black crosses indicate collected training images, and gray areas depict unexplored terrain.

information $I(\psi_{t+1})$ is computed by summing rewards along subsequent parent and child nodes n_i, n_{i+1} given by:

$$R(n_i, n_{i+1}) = \frac{\|\mathcal{G}_{U,t}(\mathbf{p}_{t+i+1}^{n_{i+1}})\|_1}{c(\mathbf{p}_{t+i}^{n_i}, \mathbf{p}_{t+i+1}^{n_{i+1}}) \|T_{t+i}^{n_i}(\mathbf{p}_{t+i+1}^{n_{i+1}})\|_1}. \quad (12)$$

Note that $T_{t+i}^{n_i}(\mathbf{p}_{t+i+1}^{n_{i+1}})$ are the training data occurrences at the child node's position assuming the training data count after collecting a measurement at the parent's node position. This way, the reward estimates the next position's information value given the map state at replanning time $t+i$. After simulating a number of paths ψ_{t+1} , we select the root's child node n_1 with the highest average information value and the UAV moves to its associated measurement position $\mathbf{p}_{t+1}^{n_1}$.

To show that our approach supports various planning algorithms, we proposed the four diverse planners above and showcase their integration into our modular framework. Further, we highlight that our planning strategies are agnostic to the acquisition functions introduced in Sec. III-A.

IV. EXPERIMENTAL RESULTS

Our experiments evaluate our proposed method and show the benefits of our mapping module (Sec. IV-B), Bayesian ensemble (Sec. IV-C), and AL planning objective functions (Sec. IV-D). We verify our planning framework's generality and analyse its AL performance on vastly differing real-world aerial datasets and in a photo-realistic simulator against classical coverage and random walk exploration data collection (Sec. IV-E). We conduct a sensitivity analysis validating the framework's robustness to the choice of model architecture, pre-training schemes, and UAV starting positions (Sec. IV-F). Our framework consistently maximises semantic segmentation performance while minimising human labelling effort. Notably, our experiments show that our map-based planners outperform the local planning strategy for AL proposed by Blum et al. [31], which, to the best of our knowledge, is the only directly comparable approach to date.

A. Experimental Setup

Baselines. We compare our planning framework against three baselines: a traditionally-used coverage-based collection strategy [44], and two random walk-based exploration planners.

The coverage strategy precomputes a static path maximising the area covered by the UAV to foster spatial coverage of

TABLE I
ENVIRONMENT, SENSOR, AND UAV MISSION PARAMETERS FOR THE THREE DATASETS IN OUR EXPERIMENTS

Dataset	Type	Classes	Task	Area [m×m]	FoV [px×px]	GSD [cm/px]	Altitude [m]	Budget [s]	Test Images
Potsdam [5]	Orthomosaic	7	Urban	900×900	400×400	15.0	30	1800	3500
RIT-18 [7]	Orthomosaic	6	Land Cover	261×568	400×400	8.0	15	400	3500
Flightmare [12]	Unity/Gazebo	10	Industrial	150×130	480×720	8.3	20	150	1000

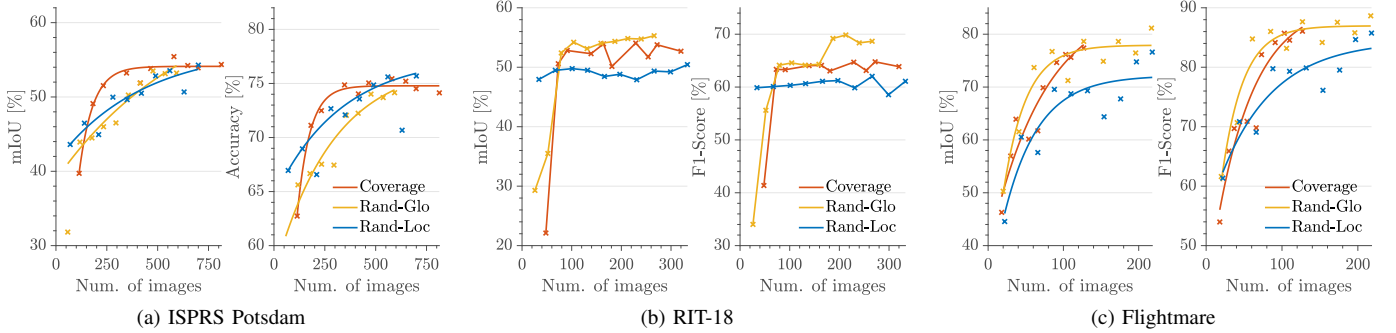


Fig. 7. Comparison of AL performance of our three baseline planning approaches with MC dropout inference (a) on the ISPRS Potsdam dataset [5], (b) on the RIT-18 dataset [7], and (c) in the photo-realistic Flightmare UAV simulator [12]. Steeper curves indicate better AL performance. We compare our planning strategies to the best-performing baseline in each setting: coverage on ISPRS Potsdam and the global random walk on RIT-18 and in Flightmare.

training data. We precompute lawnmower-like patterns before each mission starts, alternate the pattern’s orientations, and vary the step size between measurement positions.

We consider two random walk exploration planners, local and global planning. Similar to the local planner, the local exploration planner chooses for a given UAV position one of the four image edges at random and follows the edge direction with predefined step size. The global exploration planner randomly selects a UAV position in the continuous space above the terrain, similar to our map-based planners. For better budget management, we sample a step size uniformly at random between a minimum and maximum radius around the UAV, then also set its heading uniformly at random.

This way, both exploration planners aim to foster data diversity while handling the budget properly. As they resemble the action spaces of the planners introduced in Sec. III-C, we can study the influence of our action space design and verify that our active planners maximise AL performance beyond random effects. Fig. 6 exemplifies the paths planned by all three baselines on the ISPRS Potsdam dataset [5].

Datasets. We evaluate our planning framework on two real-world orthomosaic datasets and in a photorealistic physics-based UAV simulator resembling real-world deployment conditions. Detailed environment, sensor, and UAV mission settings are shown in Table I. Below, we highlight the key differences between the three scenarios.

First, we use the large 7-class urban aerial ISPRS Potsdam orthomosaic dataset [5]. This dataset is characterised by a dense spatial distribution of classes, such that the coverage and exploration baselines can collect visually and semantically different features easily. We sample 4000 train, 1000 validation, and 3500 test images uniformly at random from non-overlapping regions in the orthomosaic. We use the ISPRS Potsdam dataset for the main experiments evaluating our

mapping module (Sec. IV-B), Bayesian ensemble (Sec. IV-C), and planning objectives (Sec. IV-D).

Second, we evaluate our approach on the land cover RIT-18 orthomosaic dataset [7] consisting of semantics covering large connected areas, e.g. asphalt, vegetation, and lake, and local regions, e.g. building, with six classes in total. As the RIT-18 dataset does not provide different orthomosaics for training and testing, we evaluate the UAV’s vision capabilities by sampling the test set from the same area. In contrast to the ISPRS Potsdam dataset, this does not allow us to draw conclusions about the model’s generalisability, but about its performance in the deployed environment only. This is still a crucial skill for autonomous robot deployment. Our evaluation protocol on RIT-18 resembles that of Blum et al. [31].

Last, we test our framework in Flightmare, a photorealistic simulator with a physics engine for emulating UAV dynamics simulation [12]. We deploy a UAV in the provided ‘Industrial’ environment introducing 10 semantic classes of different spatial distributions, e.g. hangar, container, road, fence, and pipe. The scene covers a dense area leading to compactly distributed semantics easily explorable by the baseline approaches. As the ‘Industrial’ terrain is small, we evaluate the UAV’s semantic segmentation performance in the deployed environment only.

We perform a study comparing the AL performance of the baseline strategies in Fig. 7. On the ISPRS Potsdam dataset, the coverage pattern is the superior baseline, while the global random walk exploration performs best on the RIT-18 dataset and in the Flightmare simulator. Note that, while MC dropout is used in Fig. 7 to predict semantic segmentation, we found that similar results hold true for deterministic network and ensemble inference. For visual clarity, we only compare our framework to the baselines with the strongest AL performance.

Evaluation Metrics. Our AL planning pipeline aims to maximise semantic segmentation performance with minimal

human labelling effort, i.e. minimal training data. In line with the standard in AL literature [8, 16, 19–21, 24, 31, 33–35, 60], our key evaluation metrics assess semantic segmentation performance (dependent variable) over the number of collected training images (independent variable). Higher semantic segmentation performance thanks to newly added images indicate better AL, and thus, planning performance. We choose mean Intersection-over-Union (mIoU), per-pixel accuracy, and per-pixel F1-score to assess semantic segmentation performance. mIoU is used in popular semantic segmentation benchmarks [65, 66]. It is defined as $mIoU = \frac{TP}{TP+FP+FN}$, where TP , FP , TN , and FN are the true and false positives, and true and false negatives. Per-pixel accuracy acc and F1-score $f1$ are typically used in classification benchmarks [67]. They are defined as $acc = \frac{TP}{TP+FP+TN+FN}$ and $f1 = \frac{2TP}{2TP+FP+FN}$. RIT-18 and Flightmare have strongly imbalanced class distributions. Thus, we use the F1-score instead of accuracy for these scenarios. Note that, as training datasets are incrementally collected while exploring an initially unknown environment, the training image distribution changes during deployment as new visual features or semantics are discovered. Hence, the training image distribution could differ from the true image distribution, which could lead to non-monotonic model improvement. To make model performance trends easier to follow, we additionally fit trend lines for the experiments conducted on the ISPRS Potsdam and Flightmare datasets. As performance trends are less regular on the RIT-18 dataset due to the more challenging exploration of semantics, we show piecewise linear line plots for these experiments.

Training Procedure. We utilise a lightweight Bayesian ERFNet for semantic segmentation as described in Sec. III-A. The model is pre-trained on the Cityscapes dataset [66] to start experiments and training after each of the 10 subsequent data collection missions from the same checkpoint. This also avoids catastrophic forgetting and accumulating train time. We re-train the model until convergence with batch size 8 and weight decay $\lambda = (1 - p)/2N$ in Eq. (1), where $p = 0.5$ is the dropout probability, and N is the number of training images [20]. All other model hyperparameters follow the standard ERFNet [29], not tuned for maximal performance in our setting, and kept fixed with changing datasets and planners. **Planning Hyperparameters.** Our optimisation-based planner leverages the CMA-ES procedure as it has been shown to yield competitive performance in terrain monitoring tasks [3, 40]. We fix a set of hyperparameters for all planners with reasonable length scales on the ISPRS Potsdam dataset, i.e. UAV step sizes, minimum and maximum action space radii, grid discretisation, and initial CMA-ES covariance. Only these hyperparameters dependent on the aerial dimensions are scaled accordingly with changing environment sizes. The scale-independent hyperparameters, e.g. number of MCTS simulations, are set in line with prior works [3, 64]. We fix the UAV's starting position to the top-left corners of each terrain. **Planning Strategies.** We outline our planning strategies in detail in Sec. III-C. In our experiments, we refer to the planners in the legends as follows: the local planner is named *Local*, the frontier-based planner is named *Frontier*, the optimisation-based planner is named *Optimisation*, and the sampling-based

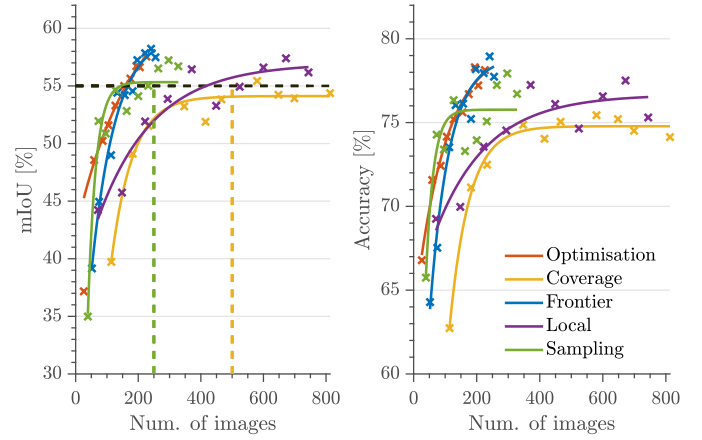


Fig. 8. Comparison of AL performance with a Bayesian model uncertainty-based planning objective estimated by MC dropout and computing informative prior maps before each mission starts. All active planners exceed the coverage baseline's performance (yellow) with less training data as shown by the dashed lines. Our map-based planners outperform the local planner (purple).

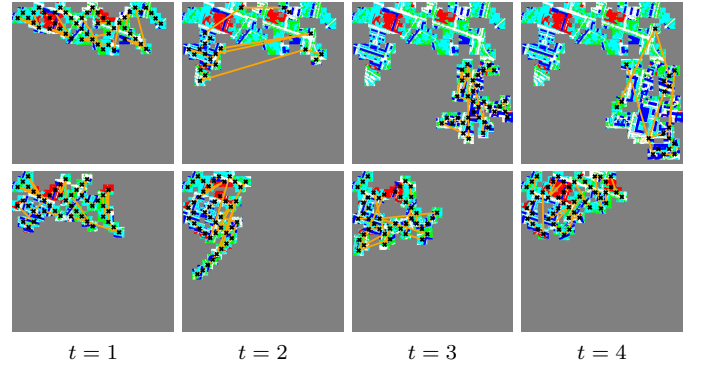


Fig. 9. Examples of paths planned on ISPRS Potsdam using the frontier-based planning strategy with (top) and without (bottom) our approach for precomputing informative prior maps. The priors are computed before each of four subsequent missions, with the UAV starting in the top-left corner. As shown by the planned paths (orange lines) and measurement positions (black crosses), using informative priors facilitates spatial exploration across missions and leads to more targeted training data collection within missions.

planner is named *Sampling*. The baseline approaches are referred to as follows: the coverage pattern is called *Coverage*, and the local and global random walk exploration strategies are abbreviated with *Rand-Glo* and *Rand-Loc*.

B. Informative Mapping

The first set of experiments analyses the performance of our approach. It (i) verifies the superior AL performance of our planning framework over the baselines; (ii) shows that our global map-based planners outperform state-of-the-art local planning; and (iii) quantifies the benefit of our mapping module for the global map-based planners. The experiments are evaluated on the ISPRS Potsdam dataset [5]. To focus on evaluating the effect of our mapping module (Sec. III-B) on the map-based planners, we fix Bayesian model uncertainty as our planning objective estimated with MC dropout (Eq. (3)). Map priors are recomputed before each mission starts to allow for maximally informed global planning.




Fig. 10. Comparison of effects of our mapping module using our framework. The planners consistently benefit from recomputing informative map priors before a mission starts (purple). The performance gain of mapping a continuous RGB sensor stream (blue) instead of only mapping images at planned measurement positions (orange) is less significant. Combining both continuous sensor streams and informative map priors (yellow) also leads to consistent performance improvements. In particular, our informative mapping approach drastically improves the greedy frontier-based strategy.

Fig. 8 summarises the AL performance with the informed mapping strategy for each planner. All planners reach higher final prediction performance than the coverage baseline (yellow). This supports the claim that our framework is generally applicable to different planning algorithms. Further, it suggests that active replanning is key to efficiently improving robot vision. Notably, our global map-based planners (orange, blue, green) exceed the coverage baseline’s maximum prediction performance ($\approx 55\%$ mIoU, black dashed line) after ≤ 250 labelled images (dashed green line), while the baseline requires ≈ 500 labelled images (yellow dashed line) to reach this performance. Particularly, for the uncertainty-based objective, our map-based planners show stronger AL performance than the local planner (purple) proposed by Blum et al. [31] as the map-based planners’ performances upper-bound the local planner’s performance for any fixed number of labelled images. In contrast to local planning, map-based planners drastically reduce training data requirements and tend to achieve higher final prediction performance.

To better understand the benefits of our active planning framework, Fig. 11 exemplarily compares the per-class AL performance of the map-based frontier planner (dashed lines) to the coverage baseline (solid lines) in the ISPRS Potsdam scenario. Our active frontier-based planning strategy shows higher AL performance in almost all classes, irrespective of their training data support. Interestingly, the ‘car’ class (blue) has lower training data support than the ‘tree’ (green) and ‘vegetation’ (red) classes but shows stronger IoU performance, even with non-targeted coverage planning. However, active planning improves the ‘car’ prediction performance even faster than the non-targeted baseline showing the benefit of our framework for classes with little training data support. Further, although the ‘tree’, ‘background’ (orange), and ‘vegetation’ classes have high training data support, they are difficult to distinguish as their visual appearance from a top-down view depends on the image resolution, altitude, and season. This leads to challenging predictions, which may be partially attributed to data instead of model uncertainty, which cannot be explained away with more training data [59]. Thus, not all classes with high training data support benefit to the same extent from active planning. At the same time, although

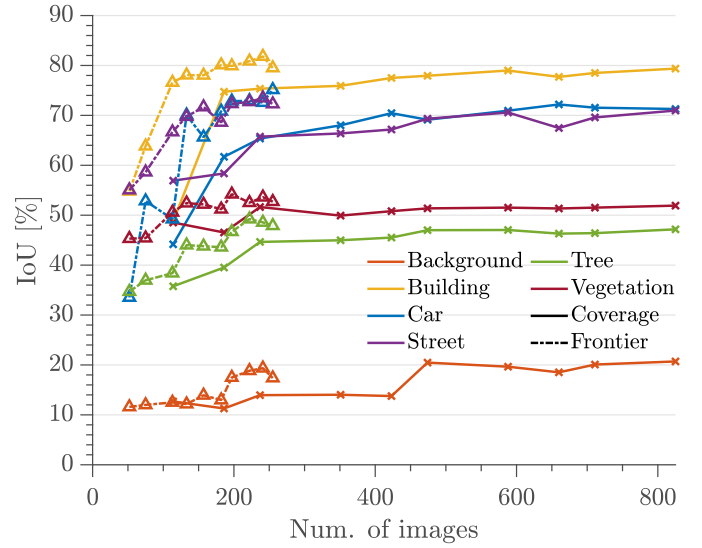


Fig. 11. Comparison of per-class AL performance of map-based frontier vs. coverage planning with a Bayesian model uncertainty-based planning objective estimated by MC dropout and computing informative prior maps before each mission starts. The frontier planner outperforms the coverage baseline (yellow) in almost all classes as our framework can capture complex task-dependent inter-class and intra-class model uncertainties.

the ‘building’ class (yellow) has high training data support and is reliably detected by both planners, the frontier-based planner still shows faster performance improvement as our framework can account for the differing visual appearance and geometry of office buildings, historical buildings, and townhouses. Overall, the results suggest that our framework can capture complex task-dependent inter-class and intra-class model uncertainties, which are too complex to capture with a single training data support heuristic, leading to superior AL performance over non-targeted baselines.

To support the claim that our new mapping module is important for the planning framework’s performance, we perform an ablation study to measure its effect on our map-based planners. We consider two mapping setups where the UAV either maps training images at planned measurement positions only (pointwise sensor) or maps the images continuously as it moves (continuous sensor stream). Fig. 10 displays the

AL performance of our map-based planners (i) recomputing informative prior maps before each mission starts based on previously collected data and the re-trained network (purple), (ii) mapping a continuous RGB image stream (blue) instead of mapping training images at planned measurement positions only (orange), and (iii) combining both informative prior maps and mapping continuous sensor streams (yellow).

All map-based planners show better AL performance with recomputed map priors as they exploit already mapped heterogeneous terrain information. This suggests that mapping and updating knowledge collected across missions with re-trained networks, i.e. changing vision capabilities, is key to strong planning performance. In contrast, mapping more information during a single mission with a fixed network is less crucial. Mapping a continuous image stream instead of mapping training data information at planned measurement positions only leads to performance improvements for the greedy frontier-based planner, while both non-greedy planners do not benefit from mapping more information during a mission. Accordingly, combining both mapping continuous sensor streams and recomputing map priors leads to higher AL performance of the frontier-based and optimisation-based planner. The sampling-based planner does not show a performance gain when combining mapping of continuous sensor streams and recomputing map priors. Particularly, our greedy frontier-based planning strategy shows significant improvements by leveraging the informative mapping procedure. Our non-greedy optimisation- and sampling-based planners are more robust to less informed map priors, perhaps because they utilise non-myopic planning along multiple waypoints, while the frontier-based planner only reasons about the next waypoint. Qualitatively, Fig. 9 verifies that informative prior maps for frontier-based planning leads to more efficient terrain exploration across missions and targeted data collection within missions resulting in higher model performance with fewer training images.

C. Bayesian ERFNet Ensemble Study

The second experiment shows that our Bayesian ensemble provides reliable uncertainty estimates for AL planning objectives. Moreover, the ensemble achieves higher prediction performance than non-Bayesian and Bayesian ERFNet with MC dropout, presented in our prior work [30].

To confirm that our Bayesian ensemble of ERFNets delivers informative model uncertainties for planning and yields superior prediction performance, we train an ensemble on the ISPRS Potsdam dataset [5] with 4000 training images and compare it to the Bayesian ERFNet with MC dropout developed in our previous work [30] and a deterministic ERFNet. Qualitatively, Fig. 12 verifies high model uncertainty of our ensemble in misclassified or hard-to-predict regions. Thus, the ensemble's model uncertainties provide reliable information for planning objectives.

To assess our Bayesian ensemble's prediction capabilities and computational efficiency for online inference on UAVs, we study its performance with varying numbers of ERFNet models $T = \{2, \dots, 8\}$ in Fig. 13. We compare the ensemble's performance to the deterministic ERFNet [29] using a

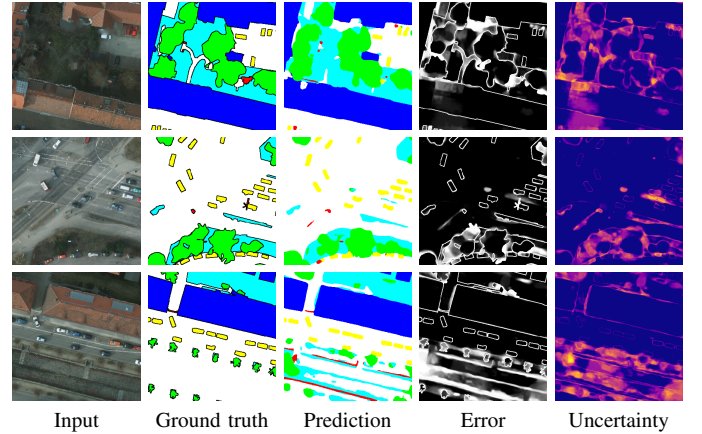


Fig. 12. Qualitative results with our ensemble of $T = 8$ ERFNets trained on ISPRS Potsdam. Columns from left to right: RGB input, ground truth, prediction, error image (negative prediction log-likelihood), model uncertainty (Eq. (3)). High model uncertainty (yellow) in misclassified regions (whiter) validates that our ensemble provides consistent uncertainty estimates as a basis for an AL planning objective in our framework.

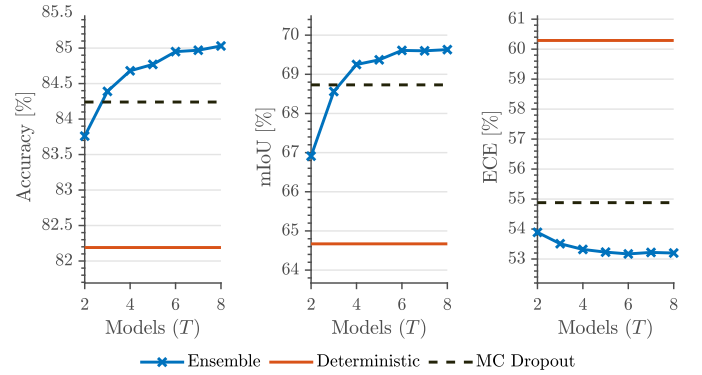


Fig. 13. Performance of our Bayesian ensemble with varying number T of ERFNets (blue), deterministic ERFNet (orange), and Bayesian ERFNet with $T = 50$ MC dropout samples [30] (black, dashed) on ISPRS Potsdam. For $T = 8$, the ensemble improves mIoU by 4.96% (middle) and reduces ECE by 7.09% (right) over a deterministic ERFNet, and improves mIoU by 0.90% and reduces ECE by 1.68% over the Bayesian ERFNet with MC Dropout.

single forward pass and to our Bayesian ERFNet utilising $T = 20$ MC dropout samples for converging to maximal performance [30]. To quantify the reliability of estimated uncertainties, we measure model calibration using the expected calibration error (ECE) metric [68]. Intuitively, model calibration is high, i.e. ECE low, when the model's probabilistic predictions match its accuracy on a test set.

For $T = 8$ models, our ensemble (blue) improves segmentation performance by 4.96% mIoU and ECE by 7.09% over the deterministic ERFNet (orange). Additionally, for $T = 8$ models, our ensemble improves segmentation performance by 0.90% mIoU and ECE by 1.68% compared to the Bayesian ERFNet with MC dropout. Overall, as the number of models increases, segmentation performance and calibration both improve. Favourably for online inference, with $T \approx 6$ models, performance gains already converge. Further, the Bayesian ensemble performs on par with the Bayesian ERFNet ($T = 20$ MC dropout samples) already




Fig. 14. Comparison of AL performance with a Bayesian model uncertainty-based planning objective estimated by an ensemble of $T = 4$ ERFNets and computing informative prior maps before each mission starts. All active planners exceed the coverage baseline’s performance with less training data. Our global map-based planners outperform the local planning scheme.

with $T = 3$ ERFNet models. Thus, our ensemble requires substantially fewer forward passes ($\approx 6\times$), i.e. compute resources, at deployment to achieve the same performance. At train time, the ensemble’s compute requirements scale linearly with the number of models T , while the MC dropout Bayesian ERFNet has constant compute requirements. However, training is performed offline, hence it is not time-critical. For details about efficient ensemble training, we refer to Huan et al. [38].

D. Comparison of Planning Objectives

Our third set of experiments shows that Bayesian model uncertainty-based objectives guarantee strong AL performance irrespective of the uncertainty estimation technique. Further, it verifies that our general framework supports various AL acquisition function paradigms, including representation-based and uncertainty-based objectives.

As we show in this experiment, Bayesian model uncertainty-based planning objectives outperform baselines with different uncertainty estimation techniques. We investigate our Bayesian ensemble’s AL performance on the ISPRS Potsdam dataset. For a fair assessment, we evaluate the coverage baseline with ensemble inference. Fig. 14 summarises the results using our Bayesian ensemble of $T = 4$ ERFNets for all planning approaches. All planners show better performance than the coverage baseline (yellow), which confirms the intuition that active planning for AL benefits from Bayesian model uncertainty-based objective functions. Similar to our MC dropout-based uncertainty estimation in Fig. 8, map-based planners (orange, blue, green) achieve higher prediction performance with fewer training images compared to the local planner (purple), further illustrating the advantage of map-based planning in our framework.

To further support our framework’s generality under various uncertainty-based objective functions, we investigate its performance using a classical non-Bayesian entropy-based acquisition function [20, 28]. Given an image z and a model with

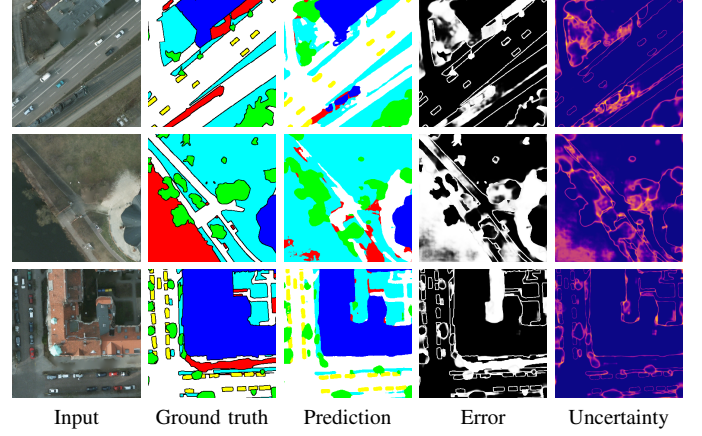


Fig. 15. Qualitative results of a deterministic ERFNet trained on ISPRS Potsdam. Columns from left to right: RGB input, ground truth, prediction, error image (negative prediction log-likelihood), prediction entropy. High entropy (yellow) is only weakly correlated with high errors (whiter).




Fig. 16. Comparison of AL performance with a non-Bayesian entropy-based planning objective over a deterministic forward pass and computing informative prior maps before each mission starts. The frontier-based, optimisation-based, and local planner outperform the coverage baseline’s AL performance.

deterministic parameters \mathbf{W} , $p(\mathbf{y}|\mathbf{z}, \mathbf{W})$ is the maximum likelihood estimate over labels \mathbf{y} . Then, the prediction entropy

$$\mathbb{H}(\mathbf{y}) = -p(\mathbf{y}|\mathbf{z}, \mathbf{W})^\top \log(p(\mathbf{y}|\mathbf{z}, \mathbf{W})) \quad (13)$$

is highest when the prediction is uniform, i.e. most uncertain. Qualitatively, Fig. 15 shows that non-Bayesian entropy is weakly correlated with prediction errors as it fails to estimate globally calibrated uncertainties.

We replace the Bayesian model uncertainty, see Eq. (3), with the entropy of a deterministic forward pass. For a fair comparison, the coverage baseline uses a deterministic forward pass as well. As shown in Fig. 16, the optimisation-based, frontier-based and local planners outperform the baseline, while the sampling-based planner performs similarly to the baseline. In line with results for Bayesian model uncertainty-based objectives, the optimisation-based and frontier-based planners show high prediction performance with substantially fewer training images compared to the local planner.

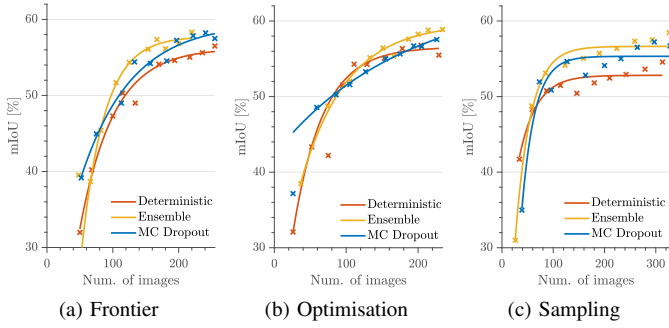


Fig. 17. Comparison of AL performance of uncertainty-based planning objectives on ISPRS Potsdam. Our Bayesian uncertainty-based objectives (blue, yellow) tend to perform better than the non-Bayesian entropy-based objective (orange). The Bayesian ensemble (yellow) achieves the highest AL performance across the planning strategies.

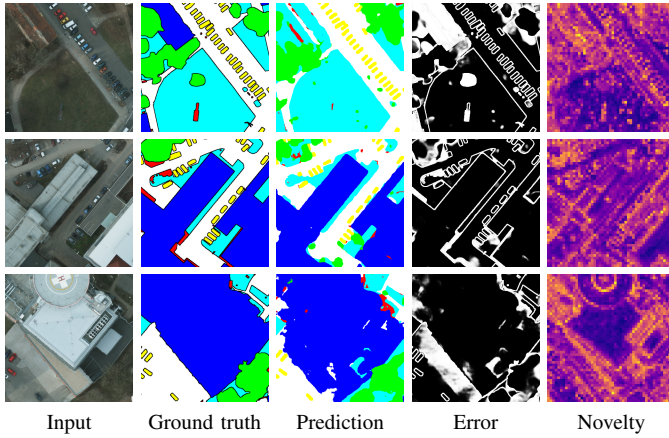


Fig. 18. Qualitative results of a deterministic ERFNet trained on the ISPRS Potsdam dataset [5]. Columns from left to right: RGB input, ground truth, prediction, error image (negative prediction log-likelihood), representation novelty (Eq. (4)). High novelty scores (yellow) in case of rare visual cues, such as the helipad (bottom row), suggest that our representation-based objective provides useful information for AL planning scenarios.

Fig. 17 shows the effect of non-Bayesian entropy-based (orange) and Bayesian model uncertainty-based planning objectives estimated by either MC dropout (blue) or an ensemble (yellow) on the planners' performances. Particularly, the map-based planners achieve higher AL performance using Bayesian model uncertainty-based objectives irrespective of the uncertainty estimation technique. Although the non-Bayesian objective yields competitive performance with multiple planners in early missions, generally, the Bayesian ensemble method leads to the best AL results. This could be due to two reasons. First, the ensemble shows higher prediction power (Fig. 13). Second, as suggested by our qualitative results (Fig. 15), non-Bayesian uncertainty is weakly calibrated, which results in a less informative planning objective.

To confirm that our framework is applicable to representation-based acquisition functions, we utilise the novelty score shown in Eq. (4) computed over the latent space of a deterministic ERFNet in our planning objective. For a fair assessment, we also utilise a deterministic ERFNet for the coverage baseline. Qualitatively, Fig. 18 visualises the

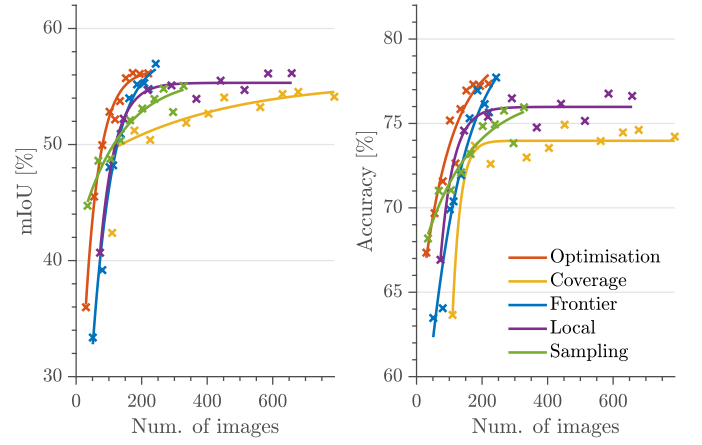


Fig. 19. Comparison of AL performance with representation-based novelty objective and computing informative prior maps before missions start. All planners outperform the baseline (yellow) with fewer training images.

representation novelties of a network trained and tested on disjoint areas of ISPRS Potsdam. Although the novelties do not correlate strongly with prediction errors (whiter), high novelty (yellow) is assigned to rare visual cues, such as the helipad (bottom row), which could be an informative objective to collect diverse training images.

Fig. 19 depicts the AL results using representation novelties in the planning objective. All adaptive planners achieve higher segmentation performance than the coverage baseline (yellow). Further, our map-based optimisation (orange) and frontier (blue) planners require fewer training images than the local planner (purple) to reach high prediction performance. This validates that our framework generally supports various acquisition function paradigms and ensures higher AL performance than the baseline approaches, irrespective of the planning objective. Our experiments suggest that the map-based planners outperform the local planner more significantly using Bayesian uncertainty objectives. This could be due to the better-calibrated Bayesian uncertainty estimates (Fig. 13) leading to more informative planning objectives.

E. Other Scenarios

The fourth set of experiments suggests that (i) our planning framework reduces the number of labelled images required to maximise segmentation performance across substantially different environments, and (ii) our global map-based planning strategies outperform state-of-the-art local planning in most cases, irrespective of the chosen planning objective.

We support these claims with an evaluation of our framework on the RIT-18 dataset [7] and in the Flightmare simulator [12]. The RIT-18 semantics cover large areas leading to challenging exploration. The Flightmare simulator resembles real-world UAV control over an easy-to-explore photorealistic industrial terrain with strong random walk baseline performance. We access the framework's performance using the Bayesian model uncertainty estimated with MC dropout, see Eq. (3), and the representation novelty score given by Eq. (4).

Fig. 20 summarises our planning results on the RIT-18 dataset [7]. Note that non-monotonic model performance

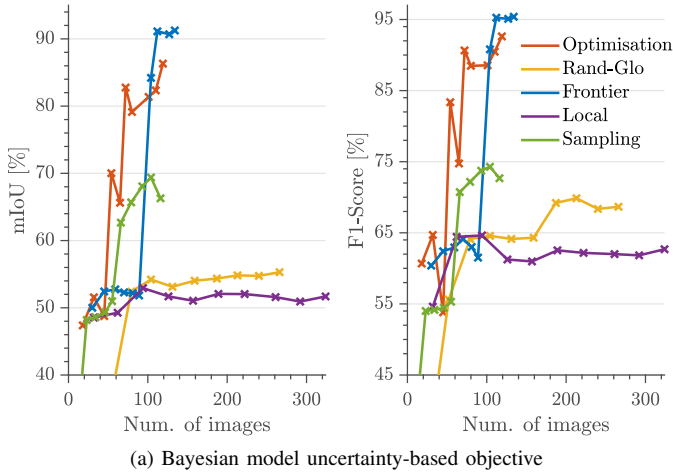


Fig. 20. AL results on the RIT-18 dataset [7] using informative prior maps with (a) the Bayesian model uncertainty objective estimated by MC dropout (Eq. (3)), and (b) the representation novelty objective (Eq. (4)). All map-based planners significantly outperform the random walk baseline (yellow). In most cases, our map-based planners lead to substantially higher AL performance than the local planning strategy (purple).

improvements on RIT-18 are expected as semantics cover large areas leading to challenging exploration influencing the training class distribution. All map-based planning strategies show significantly higher final segmentation performance than the random walk baseline (yellow), irrespective of the chosen planning objective. This confirms that our framework reduces human labelling effort while maximising segmentation performance over vastly differing terrains. Particularly, in most cases, our map-based planners require fewer training images to achieve segmentation performance on par or higher than the local planner (purple). Notably, the local planner performs worse than the baseline using Bayesian model uncertainty showing that our map-based planners are more generally applicable than the local planner.

Fig. 21 illustrates our planning results in the Flightmare simulator [12]. All planners using the Bayesian model uncertainty objective show higher AL performance than the random walk baseline (yellow). Using the representation-based objective, only our two map-based optimisation (orange) and sampling (green) planners result in higher final prediction per-

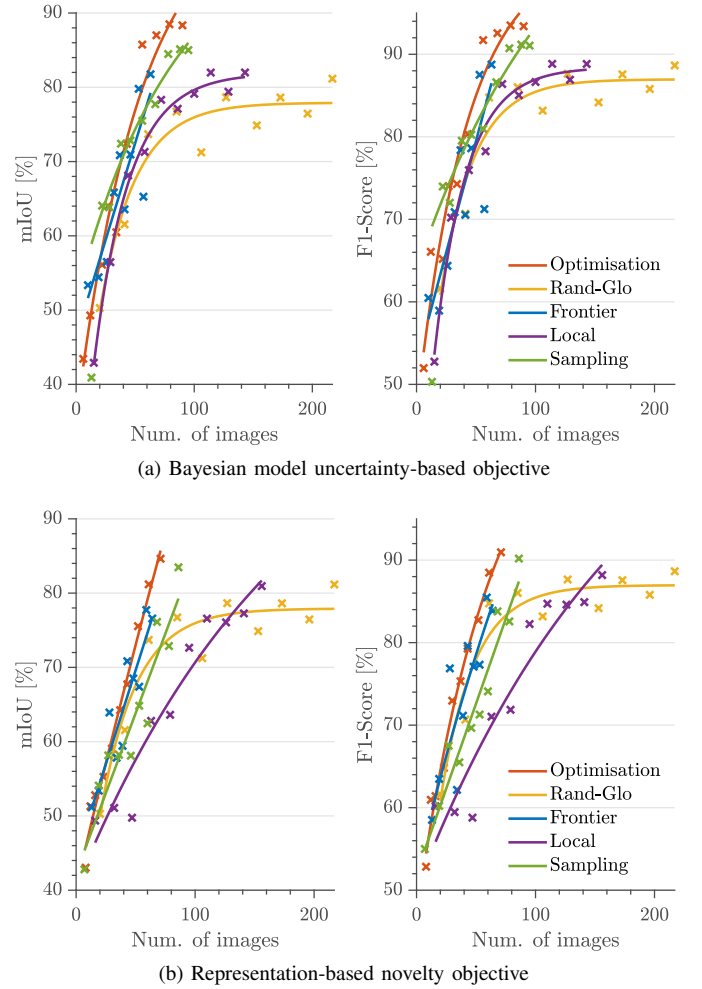


Fig. 21. AL results in the Flightmare simulator [12] using informative prior maps with (a) the Bayesian model uncertainty objective (Eq. (3)), and (b) the representation novelty objective (Eq. (4)). All planners outperform the random walk baseline (yellow) using the Bayesian model uncertainty objective. Using the representation novelty objective, only our two map-based optimisation and sampling planners show higher final prediction performance than the baseline.

formance than the baseline. Combined with the RIT-18 results (Fig. 20), this suggests that our Bayesian model uncertainty-based objectives are more robustly applicable across varying terrains compared to the representation novelty score proposed by Blum et al. [31]. One possible explanation could be that Bayesian model uncertainty is more strongly correlated with the prediction errors, as indicated by our qualitative results in Fig. 12 and Fig. 18. In most cases, our map-based planners show higher AL performance in both terrains than local planning. This verifies that our map-based planners are crucial for informative data collection, while local planning is not robustly applicable to varying terrains and planning objectives.

F. Sensitivity Analysis

The fifth set of experiments analyses our framework under various task-dependent design choices. It (i) verifies our framework's AL performance with varying UAV starting positions; (ii) validates our framework's robustness to different pre-training schemes; and (iii) showcases our framework's applicability and superior performance over baselines with different

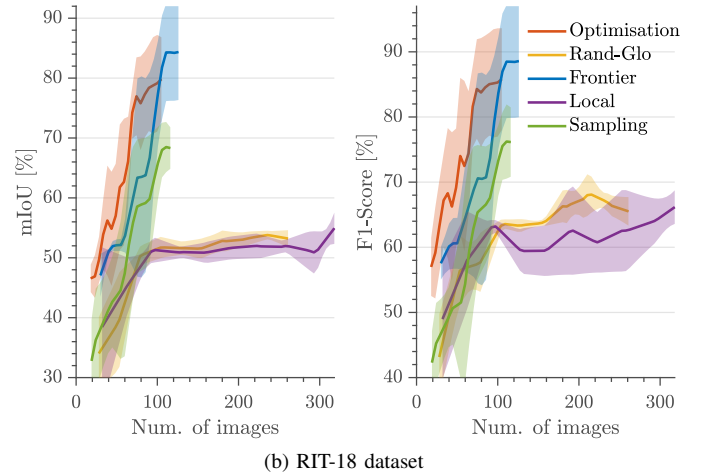
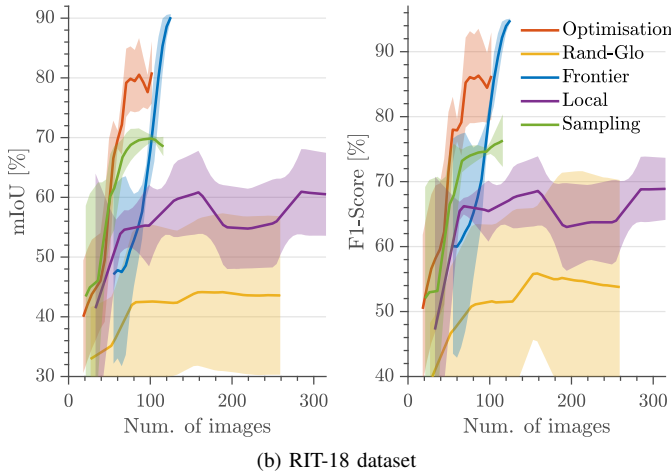
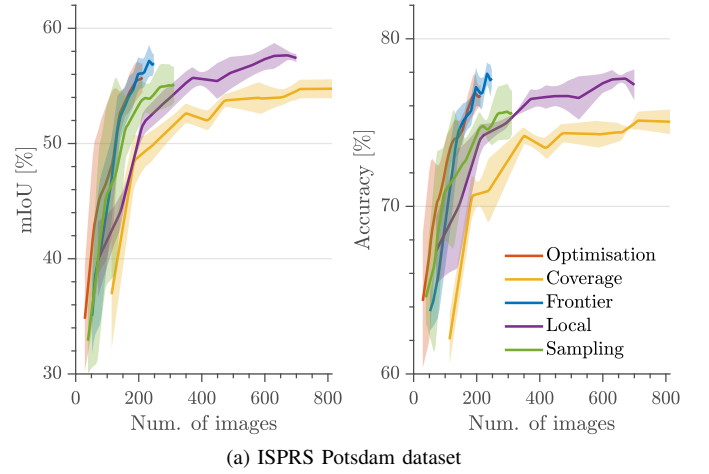



Fig. 22. Comparison of AL performance on the (a) ISPRS Potsdam dataset [5] and (b) RIT-18 dataset [7] with the Bayesian model uncertainty-based planning objective estimated by MC dropout and computing informative prior maps before each mission starts. Results are averaged over three different UAV starting positions. Shaded regions indicate one standard deviation. Our map-based planners consistently outperform the baseline approaches (yellow) and local planning (purple) on both datasets with less training data while showing less sensitivity to the UAV starting position.

Fig. 23. Comparison of AL performance on the (a) ISPRS Potsdam dataset [5] and (b) RIT-18 dataset [7] with the Bayesian model uncertainty-based planning objective estimated by MC dropout and computing informative prior maps before each mission starts. Results are averaged over three differently pre-trained Bayesian ERFNets. Shaded regions indicate one standard deviation. Our map-based planners, on average, outperform the baseline approaches (yellow) and local planning (purple) on both datasets with less training data, irrespective of the pre-training scheme.

model architectures. The experiments are evaluated on the ISPRS Potsdam [5] and RIT-18 [7] datasets using the Bayesian model uncertainty-based planning objective estimated by MC dropout. If not stated otherwise, we utilise the Bayesian ERFNet (Sec. III-A) pre-trained on Cityscapes [66].

Fig. 22 summarises the AL performance for each planner averaged over three different starting positions at the top-left, top-right, and bottom-right corners of the ISPRS Potsdam and RIT-18 datasets. All our map-based planners, on average, reach higher AL performance than the coverage baseline (yellow) and local planner (purple) on both datasets. In contrast, the local planner, on average, does not perform better than the coverage baseline on the ISPRS Potsdam dataset, as indicated by their largely overlapping means and standard deviations. Further, as indicated by the large standard deviations of the local planner and random walk baseline (yellow) on the RIT-18 dataset, the local planning and random walk AL performances heavily depend on the UAV starting position in challenging to explore terrains. This verifies that our map-based planners are

robust to varying UAV starting positions, while local planning and the baselines are sensitive to the UAV starting position.

Fig. 23 summarises the AL performance for planner averaged over three differently pre-trained Bayesian ERFNets. Each mission starts from the top-left corner of the ISPRS Potsdam and RIT-18 datasets with Bayesian ERFNet being randomly initialised, pre-trained on the Cityscapes dataset [66], or pre-trained on the Flightmare dataset [12]. Note that the standard deviations are mainly a result of the randomly initialised models having, as expected, weaker prediction performance than the pre-trained models irrespective of the planning approach. All our map-based planners, on average, show stronger AL performance than the baseline approaches (yellow) and the local planner (purple) on both datasets. Particularly, on the RIT-18 dataset, the local planner fails to outperform the random walk (yellow) irrespective of the pre-training scheme. These findings validate our map-based planners' robustness to varying model pre-training schemes.

Fig. 24 summarises the AL performance of our planning




Fig. 24. Comparison of AL performance on the ISPRS Potsdam dataset [5] using a Bayesian version of U-Net [10] pre-trained on the Flightmare dataset [12]. The Bayesian model uncertainty-based planning objective is estimated by MC dropout, and informative prior maps are computed before each mission starts. Our map-based optimisation (orange) and frontier (blue) planners outperform coverage (yellow) and local planning (purple), while the sampling planner (green) performs on par with local planning.

framework utilising a Bayesian variant of U-Net [10]. We extend the U-Net architecture by adding dropout layers after each convolutional block with a dropout probability of 10% to perform MC dropout for computing the Bayesian uncertainty-based planning objective. We conduct experiments with the Bayesian U-Net pre-trained on the Flightmare dataset [12] using the ISPRS Potsdam dataset starting each mission from the top-left corner. All active planners exceed the maximum semantic segmentation performance of the coverage base-line (yellow) with less than half of the training images. This confirms the effectiveness of active planning for AL irrespective of the chosen model architecture. Further, our map-based frontier (blue) and optimisation (orange) planners outperform local planning (purple), while the sampling planner (green) performs on par with local planning. This showcases strong AL performance of our map-based planners and the applicability of our framework to different model architectures.

V. CONCLUSION AND FUTURE WORK

This paper proposed a novel and unified planning framework for active learning in aerial semantic mapping to improve a robot's semantic perception with minimal expert guidance. A key aspect of our work is to link our planning objectives to active learning acquisition functions, enabling us to adaptively replan the robot's paths towards regions of informative training data. To ensure maximally informed online decision-making, our global planning algorithms leverage a sequentially updated probabilistic terrain map capturing semantics and acquisition function information. The framework is generally applicable to aerial robotic missions as it provides diverse acquisition functions, proposes various planning algorithms, is agnostic to the model architecture, and can be easily extended to other acquisition functions and planners.

Our experimental results show that our framework reduces the human labelling effort and maximises segmentation performance across varying terrains compared to traditionally used

coverage and random walk data collection. Further, our map-based planners outperform state-of-the-art local planners used in active learning. The results also verify the benefit of our mapping module for the active learning performance. Overall, our findings demonstrate how active learning combined with online planning enables efficient training data collection to improve robotic perception in initially unknown environments.

Future work concerns integrating varying altitudes into the planning algorithms and estimating the resulting data uncertainty to select multi-view consistent informative training data. To further reduce human labelling effort, combining the supervised AL paradigm with self-supervised training and continual learning across different terrains could be a promising avenue for future research.

REFERENCES

- [1] 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," *Int. Journal of Applied Earth Observation and Geoinformation*, vol. 102, 2021.
- [2] 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.
- [3] 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.
- [4] 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, p. 112479, 2021.
- [5] ISPRS. (2018) 2D Semantic Labeling Contest. [Online]. Available: <https://www.isprs.org/education/benchmarks/UrbanSemLab/semantic-labeling.aspx>
- [6] G. Lenczner, A. Chan-Hon-Tong, B. Le Saux, N. Luminari, and G. Le Besnerais, "DIAL: Deep Interactive and Active Learning for Semantic Segmentation in Remote Sensing," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 15, pp. 3376–3389, 2022.
- [7] 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.
- [8] D. Tuia, F. Ratle, F. Pacifici, M. F. Kanevski, and W. J. Emery, "Active Learning Methods for Remote Sensing Image Classification," *IEEE Trans. on Geoscience and Remote Sensing*, vol. 47, no. 7, pp. 2218–2232, 2009.
- [9] J. Long, E. Shelhamer, and T. Darrell, "Fully convolutional networks for semantic segmentation," in *Proc. of the IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2015, pp. 3431–3440.
- [10] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional Networks for Biomedical Image Segmentation," in *Proc. of the Int. Conf. on Medical Image Computing and Computer-Assisted Intervention*, 2015, pp. 234–241.
- [11] A. Garcia-Garcia, S. Orts-Escolano, S. Oprea, V. Villena-Martinez, and J. G. Rodriguez, "A Review on Deep Learning Techniques Applied to Semantic Segmentation," *arXiv preprint arXiv:1704.06857*, 2017.
- [12] Y. Song, S. Naji, E. Kaufmann, A. Loquercio, and D. Scaramuzza, "Flightmare: A Flexible Quadrotor Simulator," in *Proc. of the Conf. on Robot Learning (CoRL)*, 2021, pp. 1147–1157.
- [13] Y. Freund, H. S. Seung, E. Shamir, and N. Tishby, "Selective Sampling Using the Query by Committee Algorithm," *Machine Learning*, vol. 28, no. 2, pp. 133–168, 1997.
- [14] D. D. Lewis and W. A. Gale, "A Sequential Algorithm for Training Text Classifiers," in *Proc. of the Int. ACM-SIGIR Conf. on Research and Development in Information Retrieval*, 1994, pp. 3–12.
- [15] C. Campbell, N. Cristianini, A. Smola *et al.*, "Query Learning with Large Margin Classifiers," in *Proc. of the Int. Conf. on Machine Learning (ICML)*, 2000.

- [16] S. Tong and D. Koller, "Support Vector Machine Active Learning with Applications to Text Classification," *Journal on Machine Learning Research (JMLR)*, vol. 2, no. Nov, pp. 45–66, 2001.
- [17] A. Kapoor, K. Grauman, R. Urtasun, and T. Darrell, "Active Learning with Gaussian Processes for Object Categorization," in *Proc. of the IEEE/CVF Int. Conf. on Computer Vision (ICCV)*, 2007, pp. 1–8.
- [18] 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.
- [19] Z. Wang and J. Ye, "Querying Discriminative and Representative Samples for Batch Mode Active Learning," *ACM Transactions on Knowledge Discovery from Data (TKDD)*, vol. 9, no. 3, pp. 1–23, 2015.
- [20] Y. Gal, R. Islam, and Z. Ghahramani, "Deep Bayesian Active Learning with Image Data," in *Proc. of the Int. Conf. on Machine Learning (ICML)*. PMLR, 2017, pp. 1183–1192.
- [21] O. Sener and S. Savarese, "Active Learning for Convolutional Neural Networks: A Core-Set Approach," in *iclr*, 2018.
- [22] S. Sinha, S. Ebrahimi, and T. Darrell, "Variational Adversarial Active Learning," in *Proc. of the IEEE/CVF Int. Conf. on Computer Vision (ICCV)*, 2019, pp. 5972–5981.
- [23] S. Ebrahimi, W. Gan, D. Chen, G. Biamby, K. Salahi, M. Laielli, S. Zhu, and T. Darrell, "Minimax Active Learning," *arXiv preprint arXiv:2012.10467*, 2020.
- [24] L. Yang, Y. Zhang, J. Chen, S. Zhang, and D. Z. Chen, "Suggestive Annotation: A Deep Active Learning Framework for Biomedical Image Segmentation," in *Proc. of the Int. Conf. on Medical Image Computing and Computer-Assisted Intervention*, 2017, pp. 399–407.
- [25] F. Zhdanov, "Diverse Mini-batch Active Learning," *arXiv preprint arXiv:1901.05954*, 2019.
- [26] 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)*, 2016, pp. 1050–1059.
- [27] B. Lakshminarayanan, A. Pritzel, and C. Blundell, "Simple and Scalable Predictive Uncertainty Estimation Using Deep Ensembles," *Proc. of the Conf. on Neural Information Processing Systems (NIPS)*, 2017.
- [28] J. Postels, M. Segu, T. Sun, L. D. Sieber, L. Van Gool, F. Yu, and F. Tombari, "On the practicality of deterministic epistemic uncertainty," in *Proc. of the Int. Conf. on Machine Learning (ICML)*, 2022, pp. 17 870–17 909.
- [29] 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.
- [30] J. Rückin, L. Jin, F. Magistri, C. Stachniss, and M. Popović, "Informative Path Planning for Active Learning in Aerial Semantic Mapping," in *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2022.
- [31] 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.
- [32] B. Settles, "Active Learning Literature Survey," 2009. [Online]. Available: <https://minds.wisconsin.edu/bitstream/handle/1793/60660/TR1648.pdf>
- [33] Y. Guo and D. Schuurmans, "Discriminative Batch Mode Active Learning," *Proc. of the Conf. on Neural Information Processing Systems (NIPS)*, 2007.
- [34] A. J. Joshi, F. Porikli, and N. Papanikolopoulos, "Multi-class Active Learning for Image Classification," in *Proc. of the IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2009, pp. 2372–2379.
- [35] W. H. Beluch, T. Genewein, A. Nürnberger, and J. M. Köhler, "The Power of Ensembles for Active Learning in Image Classification," in *Proc. of the IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2018, pp. 9368–9377.
- [36] N. Pawłowski, M. Jaques, and B. Glocker, "Efficient Variational Bayesian Neural Network Ensembles for Outlier Detection," *arXiv preprint arXiv:1703.06749*, 2017.
- [37] N. Durasov, T. Bagautdinov, P. Baque, and P. Fua, "Masksembles for Uncertainty Estimation," in *Proc. of the IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2021, pp. 13 539–13 548.
- [38] G. Huang, Y. Li, G. Pleiss, Z. Liu, J. E. Hopcroft, and K. Q. Weinberger, "Snapshot Ensembles: Train 1, Get M for Free," in *International Conference on Learning Representations*, 2017.
- [39] M. Dunbabin and L. Marques, "Robots for Environmental Monitoring: Significant Advancements and Applications," *IEEE Robotics and Automation Magazine (RAM)*, vol. 19, no. 1, pp. 24–39, 2012.
- [40] G. Hitz, E. Galceran, M.-È. Garneau, F. Pomerleau, and R. Siegwart, "Adaptive Continuous-Space Informative Path Planning for Online Environmental Monitoring," *Journal of Field Robotics (JFR)*, vol. 34, no. 8, pp. 1427–1449, 2017.
- [41] G. A. Hollinger, B. Englot, F. S. Hover, U. Mitra, and G. S. Sukhatme, "Active planning for Underwater Inspection and the Benefit of Adaptivity," *Int. Journal of Robotics Research (IJRR)*, vol. 32, no. 1, pp. 3–18, 2013.
- [42] A. Bircher, M. Kamel, K. Alexis, H. Oleynikova, and R. Siegwart, "Receding Horizon Path Planning for 3D Exploration and Surface Inspection," *Autonomous Robots*, vol. 42, no. 2, pp. 291–306, 2018.
- [43] M. Popović, T. Vidal-Calleja, G. Hitz, I. Sa, R. Siegwart, and J. Nieto, "Multiresolution Mapping and Informative Path Planning for UAV-based Terrain Monitoring," in *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2017, pp. 1382–1388.
- [44] 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.
- [45] C. Chekuri and M. Pal, "A Recursive Greedy Algorithm for Walks in Directed Graphs," in *IEEE Symp. on Foundations of Computer Science (FOCS)*, 2005, pp. 245–253.
- [46] C.-W. Ko, J. Lee, and M. Queyranne, "An Exact Algorithm for Maximum Entropy Sampling," *Operations Research*, vol. 43, no. 4, pp. 684–691, 1995.
- [47] J. Binney and G. S. Sukhatme, "Branch and Bound for Informative Path Planning," in *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*. IEEE, 2012, pp. 2147–2154.
- [48] A. Singh, A. Krause, C. Guestrin, and W. J. Kaiser, "Efficient Informative Sensing Using Multiple Robots," *Journal of Artificial Intelligence Research (JAIR)*, vol. 34, pp. 707–755, 2009.
- [49] G. A. Hollinger and G. S. Sukhatme, "Sampling-based Robotic Information Gathering Algorithms," *Int. Journal of Robotics Research (IJRR)*, vol. 33, no. 9, pp. 1271–1287, 2014.
- [50] S. Choudhury, N. Gruver, and M. J. Kochenderfer, "Adaptive Informative Path Planning with Multimodal Sensing," in *Proc. of the Int. Conf. on Automated Planning and Scheduling (ICAPS)*, vol. 30, 2020, pp. 57–65.
- [51] K. C. Vivaldini, T. H. Martinelli, V. C. Guizilini, J. R. Souza, M. D. Oliveira, F. T. Ramos, and D. F. Wolf, "UAV Route Planning for Active Disease Classification," *Autonomous Robots*, vol. 43, no. 5, pp. 1137–1153, 2019.
- [52] M. Ghaffari Jadidi, J. Valls Miro, and G. Dissanayake, "Gaussian Processes Autonomous Mapping and Exploration for Range-Sensing Mobile Robots," *Autonomous Robots*, vol. 42, no. 2, pp. 273–290, 2018.
- [53] F. Chen, J. D. Martin, Y. Huang, J. Wang, and B. Englot, "Autonomous Exploration Under Uncertainty via Deep Reinforcement Learning on Graphs," in *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2020, pp. 6140–6147.
- [54] H. H. González-Banos and J.-C. Latombe, "Navigation Strategies for Exploring Indoor Environments," *Int. Journal of Robotics Research (IJRR)*, vol. 21, no. 10-11, pp. 829–848, 2002.
- [55] G. Georgakis, B. Bucher, K. Schmeckpeper, S. Singh, and K. Daniilidis, "Learning to Map for Active Semantic Goal Navigation," in *Proc. of the Int. Conf. on Learning Representations (ICLR)*, 2022.
- [56] J. Frey, H. Blum, F. Milano, R. Siegwart, and C. Cadena, "Continual Adaptation of Semantic Segmentation using Complementary 2D-3D Data Representations," *IEEE Robotics and Automation Letters (RA-L)*, vol. 7, no. 4, pp. 11 665–11 672, 2022.
- [57] R. Zurbrügge, H. Blum, C. Cadena, R. Siegwart, and L. Schmid, "Embodied active domain adaptation for semantic segmentation via informative path planning," *IEEE Robotics and Automation Letters (RA-L)*, vol. 7, no. 4, pp. 8691–8698, 2022.
- [58] D. S. Chaplot, M. Dalal, S. Gupta, J. Malik, and R. R. Salakhutdinov, "SEAL: Self-supervised embodied active learning using exploration and 3d consistency," *Proc. of the Conf. on Neural Information Processing Systems (NeurIPS)*, pp. 13 086–13 098, 2021.
- [59] A. Kendall and Y. Gal, "What Uncertainties Do We Need in Bayesian Deep Learning for Computer Vision?" *Proc. of the Conf. on Neural Information Processing Systems (NIPS)*, 2017.
- [60] N. Houlsby, F. Huszar, Z. Ghahramani, and M. Lengyel, "Bayesian Active Learning for Classification and Preference Learning," *arXiv preprint arXiv:1112.5745*, 2011.
- [61] N. Papernot and P. McDaniel, "Deep k-Nearest Neighbors: Towards Confident, Interpretable and Robust Deep Learning," *arXiv preprint arXiv:1803.04765*, 2018.
- [62] A. Mandelbaum and D. Weinshall, "Distance-based Confidence Score for Neural Network Classifiers," *arXiv preprint arXiv:1709.09844*, 2017.
- [63] H. Moravec and A. Elfes, "High resolution maps from wide angle sonar,"

in *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, 1985, pp. 116–121.

- [64] C. B. Browne, E. Powley, D. Whitehouse, S. M. Lucas, P. I. Cowling, P. Rohlfshagen, S. Tavener, D. Perez, S. Samothrakis, and S. Colton, “A Survey of Monte Carlo Tree Search Methods,” *IEEE Trans. on Computational Intelligence and AI in Games*, vol. 4, no. 1, pp. 1–43, 2012.
- [65] M. Everingham, L. Van Gool, C. K. Williams, J. Winn, and A. Zisserman, “The Pascal Visual Object Classes (VOC) Challenge,” *Int. Journal of Computer Vision (IJCV)*, vol. 88, no. 2, pp. 303–338, 2010.
- [66] 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, pp. 3213–3223.
- [67] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, “ImageNet: A Large-Scale Hierarchical Image Database,” in *Proc. of the IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2009, pp. 248–255.
- [68] C. Guo, G. Pleiss, Y. Sun, and K. Q. Weinberger, “On Calibration of Modern Neural Networks,” in *Proc. of the Int. Conf. on Machine Learning (ICML)*, 2017, pp. 1321–1330.



Julius Rückin received a MSc degree in “Mathematics in Data Science” from the Technical University of Munich in 2020. Previously, he received a B.Sc. degree in IT Systems Engineering from the University of Potsdam in 2018. He is currently a doctoral student in the Decision-Making for Autonomous Mobile Robots Group at the University of Bonn. His research focuses on integrated planning and learning for autonomous robots at the intersection of computer vision, active learning and reinforcement learning.



Federico Magistri is a Ph.D. student at the Photogrammetry Lab at the Rheinische Friedrich-Wilhelms-Universität Bonn since November 2019. He received his M.Sc. in Artificial Intelligence and Robotics from “La Sapienza” University of Rome with a thesis on Swarm Robotics for Precision Agriculture in collaboration with the National Research Council of Italy and the Wageningen University and Research. During his master, he spent one semester at the Albert-Ludwigs Universität Freiburg as an Erasmus student.



Cyrill Stachniss is a full professor at the University of Bonn, a Visiting Professor in Engineering at the University of Oxford, and is with the Lamarr Institute for Machine Learning and AI. He is the Spokesperson of the DFG Cluster of Excellence PhenoRob at the University of Bonn. Before his appointment in Bonn, he was with the University of Freiburg and ETH Zurich. His research focuses on probabilistic techniques and learning approaches for mobile robotics, perception, and navigation. The main application areas of his research are agricultural

robotics, self-driving cars, and service robots. He has co-authored over 300 peer-reviewed publications and has coordinated multiple large-scale research projects on the national and European levels.



Marija Popović is a Junior Research Group Leader at the University of Bonn and Cluster of Excellence “PhenoRob”. Before starting in Bonn, she was a postdoctoral research associate at the Smart Robotics Lab at Imperial College London, UK. She received her Ph.D. from the Autonomous Systems Lab at ETH Zurich, Switzerland (2019) and did her Master in Engineering in Integrated Mechanical & Electrical Engineering at the University of Bath, UK (2015). Her research centres around developing algorithms that enable intelligent robotic decision-

making, including in planning/coordination, environmental mapping, and computer vision.