

Adaptive-Resolution Field Mapping Using Gaussian Process Fusion with Integral Kernels

Liren Jin¹, Julius Rückin¹, Stefan H. Kiss², Teresa Vidal-Calleja², Marija Popović¹

Abstract—Unmanned aerial vehicles are rapidly gaining popularity in many environmental monitoring tasks. A prerequisite for their autonomous operation is the ability to perform efficient and accurate mapping online, given limited on-board resources constraining operation time and computational capacity. To address this, we present an online adaptive-resolution approach for field mapping based on Gaussian Process fusion, a strategy in which Bayesian fusion is applied to update a Gaussian Process prior map. A key aspect of our approach is an integral kernel encoding spatial correlation over the areas of grid cells. This enables efficient information compression in uninteresting areas to achieve a compact map representation while maintaining spatial correlations in a theoretically sound fashion. We evaluate the performance of our approach on both synthetic and real-world data. Results show that our method is more efficient in terms of mapping time and memory consumption without compromising on map quality. Further, we integrate our mapping strategy into an adaptive path planning framework to show that it facilitates information gathering efficiency in online settings.

I. INTRODUCTION

Environmental monitoring plays a central role in helping us better understand the Earth and its natural processes. However, many commonly observed natural phenomena, e.g., temperature, humidity, etc., exhibit complex spatial variations that are difficult to capture using traditional sensing methods, such as manual sampling or static sensor networks [1, 2]. Recently, unmanned aerial vehicles (UAVs) have emerged as a more flexible, cost-efficient alternative for data acquisition in a wide range of applications, including biomass measurement [1, 3], signal strength monitoring [4], weed detection [5] and thermal mapping [1]. To fully exploit these platforms, a key challenge is developing map representations that can accurately capture heterogeneous natural variables, while being compact and computationally efficient for online interpretability and decision-making on resource-constrained systems.

This paper focuses on mapping methods for terrain monitoring scenarios, where the task is to reconstruct a continuous, non-uniform 2D scalar field, e.g., of temperature, biomass cover, etc., using measurement information from on-board sensors. In this setup, our goal is to develop an online mapping strategy that can accommodate *both* high-fidelity field reconstruction in targeted areas of interest, e.g., hotspots or anomalies, as well as mapping with low computational and memory requirements. By catering for these two aspects simultaneously, our work bridges the gap between environmental




Fig 1: Our adaptive-resolution Gaussian Process fusion approach for online field mapping. Left: Synthetic ground truth distribution. Yellow lower shades indicate higher values we would like to map in greater detail. Right: Mapping result with uncertainty. Our approach maps areas of interest at higher resolutions while compressing information in less interesting regions to increase computational and memory efficiency. The checkerboard is added for visual interpretation of the map uncertainty (high opacity means low uncertainty).

monitoring problems and autonomous robotic applications, e.g., adaptive path planning based on the current map state.

There are several methods for field mapping in environmental monitoring contexts. In the remote sensing community, most existing approaches exploit aerial data to create high-resolution reconstructions, e.g., terrain orthomosaics [1, 2]. Although they produce very detailed models, such procedures often involve heavy postprocessing and are thus not suitable for online applications. A common strategy to tackle this problem is to discretize the environment in a grid map and fuse new measurements into it during a monitoring mission. However, traditional grid-based methods [6–8] assume independence between cells, neglecting important spatial correlations which characterize environmental phenomena, and thereby limiting the map quality.

We propose a new method for online field mapping. Our approach is based on *Gaussian Process fusion* (GPF) [9, 10]. In GPF, a GP model is exploited to initialize a spatially correlated grid map, which serves as a prior for recursive Bayesian fusion. The key goal of our work is to *adaptively* adjust the GPF map resolution online based on the information value of associated measurements, such that only areas of interest are mapped at high resolutions. Different from GP regression, which pools the entire measurement history to predict the posterior map state at any resolution at once, the usage of GPF, although more efficient, poses a major challenge: in order to account for resolution changes, we need to modify online not only the map mean, but also the map covariance. In other words, adapting the map resolution leads to varying mapping locations in the environment; however, correlations at these new locations cannot be easily obtained

¹Cluster of Excellence PhenoRob, Institute of Geodesy and Geoinformation, University of Bonn. ²UTS Robotics Institute, Faculty of Engineering and IT, University of Technology Sydney. This work was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy - EXC 2070 - 390732324. Corresponding: ljin@uni-bonn.de.

from the previous measurements or the current map state [11]. The posterior after map resolution change is thus difficult to retrieve in a theoretically sound and efficient manner. This hinders the recursive update step and constitutes an open research question.

To address this, we propose a novel approach adopting an *integral kernel* to describe the spatial correlation over the areas of grid cells instead of points, e.g., grid cell center point. Combined with an ND-tree structure, we can adapt the map resolution online while preserving its spatial correlations. This enables us to retain high-resolution details in targeted areas of the field, while using coarser resolutions otherwise, as shown in Fig. 1. This way, we achieve memory and computationally efficient mapping without sacrificing map quality, as necessary for online application on platforms, e.g., UAVs, with limited computing power. In sum, our contributions are:

- 1) A new method for incrementally mapping scalar fields online. Our approach combines GPF with the ND-tree data structure to allow for resolution adaptation of spatially correlated maps.
- 2) An integral kernel function to encode the spatial correlations over the areas of 2D grid cells. This enables efficient merging operation of grid cells to compress information at any scale in uninteresting regions, while preserving spatial correlations in the map.

Our mapping approach is evaluated and benchmarked against state-of-the-art approaches using synthetic and real-world data. Experimental results show that our method reduces memory consumption and improves computational efficiency when compared against mapping baselines. Further, we demonstrate its applicability for online adaptive path planning.

II. RELATED WORK

A large body of literature has studied mapping methods for monitoring continuous phenomena in different application domains [3, 9, 12–15]. Our work focuses on online mapping methods suitable for robotic monitoring scenarios. Our new approach introduces an integral kernel for adaptive-resolution mapping which brings together two key concepts: (1) GP models and (2) ND-tree structure. The following subsections review previous studies related to these topics.

A. Gaussian Processes Mapping

Grid maps are the most commonly used representation for robotic mapping [16]. Despite their successful application, traditional occupancy grid models assume the stochastic independence of grid cells to enhance computational efficiency [17]. However, this representation often poorly captures the spatial correlations found in natural phenomena, e.g., distributions of temperature, humidity, etc. To address this, GP models are applied in environmental monitoring. For instance, GPs are used to incorporate uncertainty and represent spatially correlated data in 2.5D pipe thickness mapping [9]. Vasudevan et al. [15] apply GP regression to predict elevation on a field where sensory information is incomplete. Other applications include gas distribution mapping [12], occupancy mapping [17], and

aquatic monitoring [13, 14]. Our work follows these lines by using GPs to model the latent scalar field.

The main limitation of applying standard GP regression for online robotic mapping is its cubically growing computational complexity as measurements accumulate over time [18]. Previous work has tackled this problem by storing measurements in a KD-tree structure and using local models to approximate GPs [15, 17, 19]. To predict the mean and variance of query points, only nearby measurements are considered. However, local GPs require performing regression for each query point individually. To alleviate this problem, the concept of extended blocks was introduced [20], which applies GPs to the query points in individual blocks of the map only using the measurements in neighboring blocks. This approach decomposes a large GP into sub-models and applies regression to infer the posterior of each block. The multiple regression results are then fused using a Bayesian Committee Machine (BCM) [21], whose computational complexity scales cubically with the number of query points. Based on that, Wang and Englot [22] introduce test-data octrees, which prune nodes of the same state to condense the number of query points in regression.

In GP-based occupancy grid mapping with range sensors, O’Callaghan and Ramos [23] propose an integral kernel to handle beam line observations directly rather than discretizing them into point observations, thereby reducing the number of measurements used for GP regression. Most similar to our approach, Reid et al. [24] use an integral kernel to capture spatial correlations between image areas and infer a high-resolution estimate from a low-resolution observation in a UAV-based setup. However, inference over the map is still performed using standard GP regression, which suffers from poor scalability, especially with dense image data.

In contrast to regression-based methods, our method leverages GPF [9, 10] to reduce the computational burden for online mapping. This procedure removes the need to preserve the measurement history and infer the map posterior from scratch each time new data arrive [11]. A key difference in our approach with respect to previous fusion-based works [9, 10] is the proposed integral kernel, which bridges the gap between GPF and efficient online adaptive-resolution mapping.

B. Multi-Resolution Mapping

In practice, many monitoring scenarios exhibit a non-uniform distribution of information in the environment, i.e., some regions are considered more interesting or informative for mapping than others. Therefore, maintaining a map with constant resolution over the whole environment is redundant and costly. A common method to generate compact map representations is by using tree structures. A well-known algorithm in this category is OctoMap [6], which prunes child nodes with the same state, e.g., occupied, to achieve both memory-savings and highly precise maps. Funk et al. [7] use the octree structure in an online mapping system that adjusts map resolution based on occupancy state. Similarly, Chen et al. [25] apply quadrees to build multi-resolution 2D maps. The ND-tree generalizes these approaches by subdividing any d -dimensional volume recursively with N^d children [8]. Rather

than compressing a map only in a postprocessing step, as in OctoMap, we adapt the map resolution online based on incoming measurements similarly to Einhorn et al. [8] and Funk et al. [7]. Our approach shares the same motivation, as we tailor the map structure to reduce memory consumption and computation time in applications requiring online mapping, such as adaptive path planning [4, 10, 14].

Previous works in adaptive-resolution mapping assume cell independence [6–8], such that no correlation information needs to be maintained. This substantially simplifies mapping at the cost of map quality. However, in our online mapping setup, the covariance must be modified to account for resolution changes, which is challenging in the GPF framework. Popović et al. [3] introduce an approach for incrementally fusing variable-resolution measurements into a spatially correlated map. However, their method still considers a fixed-resolution map. In contrast, our strategy supports adaptive-resolution mapping while preserving spatial correlations.

III. ADAPTIVE-RESOLUTION GAUSSIAN PROCESS FUSION

This section introduces our online field mapping approach. We initialize the map using a GP model and store it in an ND-tree structure. This map is then recursively updated with new measurements using Bayesian fusion. We first present the theory behind GPs with the integral kernel function and define an ‘average measurement’ model, in which the state of a grid cell represents the average value of a latent scalar function in the area it covers. Then, we explain our Bayesian fusion update and the merging operation for incrementally building multi-resolution field maps. Bringing together these elements, our key contribution is the ability to efficiently merge grid cells in GPF without losing spatial correlations. Note that our setting in this work considers a UAV-based terrain mapping scenario. However, our approach is also applicable for general 2.5D mapping problems.

A. Gaussian Processes and Integral Kernel

A GP is the generalization of a Gaussian distribution over a finite vector space to an infinite-dimensional function space. It is fully described by its mean $\mu(\mathbf{x})$ and kernel function $k(\mathbf{x}, \mathbf{x}')$, where \mathbf{x} is an arbitrary point in input space. In practice, a GP regression model is used to encode spatial correlations in a probabilistic non-parametric manner and infer function values at a finite set of query points given observed data [18]. Different from GP regression, previous studies in GPF [9, 10] exploit the GP mean and kernel function to calculate the prior in predefined mapping positions, e.g., grid cell center points. The posteriors at these points are then recursively updated with grid cell measurements using Bayesian fusion. This GPF setting largely enhances the mapping efficiency compared to standard GP regression. However, as the map posterior is only maintained in fixed mapping positions, adaptive-resolution is hard to achieve.

To address this problem, we propose a new GPF approach leveraging an integral kernel. The mapping target in our problem is assumed to be a stationary continuous function described by a GP: $f(\mathbf{x}) \sim \mathcal{GP}(\mu, k) : \mathcal{E} \rightarrow \mathbb{R}$, where

$\mathcal{E} \subset \mathbb{R}^2$ is the 2D rectangular input space and $\mathbf{x} \in \mathcal{E}$. Similar to Reid et al. [24], we modify the given kernel function $k(\mathbf{x}, \mathbf{x}')$ to correctly encode the spatial correlations between areas of grid cells in our mapping framework. We define $\zeta(R) = \frac{1}{A} \int_R f(\mathbf{x}) d\mathbf{x}$ to represent the average of the latent function f over a rectangular domain $R \subset \mathbb{R}^2$ with area $A \in \mathbb{R}$. Since applying a linear operator to a GP leads to another GP [26], we obtain the new GP: $\zeta(R) \sim \mathcal{GP}(\mu_1, k_1)$, whose mean and kernel function are described as follows:

$$\mu_1(R_i) = \frac{1}{A_i} \int_{R_i} \mu(\mathbf{x}) d\mathbf{x}, \quad (1)$$

$$k_1(R_i, R_j) = \frac{1}{A_i A_j} \iint_{R_i \times R_j} k(\mathbf{x}, \mathbf{x}') d\mathbf{x} d\mathbf{x}', \quad (2)$$

where \mathbf{x} and \mathbf{x}' are the point positions contained within the rectangular domains R_i with area A_i and R_j with area A_j respectively. The area-related terms in Eqs. (1) and (2) simply transform the integral into average, which makes the physical meaning of mean and covariance in accordance with our measurement model introduced in Sec. III-B.

We initialize our grid map using this new GP model. For rectangular cells and squared exponential (SE) kernel, we can find a closed-form solution to Eq. (2). In general, numerical integration is required to determine the kernel integration [23]. Note that, the integral calculation is only conducted in initialization step and does not burden online mapping. The fact that our model is a GP allows us to initialize the prior map at any resolution. We sequentially discretize the input space into rectangular grid cells using an ND-tree until maximal depth t is reached. Only leaf grid cells are shown and updated in the map $C = \{c_1, \dots, c_n\}$, where $n = (N^d)^t$ with $d = 2$ as we focus on 2D field mapping; $c_i = [x_i^{\min}, x_i^{\max}] \times [y_i^{\min}, y_i^{\max}]$ is the parametrization of grid cell $c_i \subset \mathcal{E}$ and $\mathbf{C} = [c_1, \dots, c_n]^T$ is the vectorization of C . This prior map, with prior mean vector μ^- and covariance matrix \mathbf{K}^- calculated by Eqs. (1) and (2), can be seen as a multivariate Gaussian distribution from the perspective of recursive update introduced in Sec. III-C:

$$\mu^- = \begin{bmatrix} \mu_1(c_1) \\ \vdots \\ \mu_1(c_n) \end{bmatrix}, \mathbf{K}^- = \begin{bmatrix} k_1(c_1, c_1) & \dots & k_1(c_1, c_n) \\ \vdots & \ddots & \vdots \\ k_1(c_n, c_1) & \dots & k_1(c_n, c_n) \end{bmatrix}. \quad (3)$$

In our online mapping framework, we initialize the map to the highest resolution and adaptively merge uninteresting grid cells during mapping. Note that, since our underlying model is a GP, we can still infer the function values at arbitrary resolutions using mean function, kernel function, and map posteriors as described by Reece and Roberts [11]. As this procedure is computationally heavy, we only consider it as a post-processing step to recover a high-resolution map after an online mission is complete.

B. Sensor Model

In our GPF setting, we consider a Gaussian sensor model to account for noisy measurement data. For each observed




Fig 2: The ‘average measurement’ sensor model provides the measurements of averaged function value over a grid cell. For instance, the measurement z_2 observed from c_2 is the average of 4 single measurement values $\{m_2, m_3, m_5, m_6\}$. For calculating $\sigma_{c,i}^2$, A_{cover} is the green area on the terrain and A_c is the area of c_2 itself.

grid cell $c_i \in C$, the sensor provides a measurement z_i capturing the average value of function f over the area of this cell as $z_i \sim \mathcal{N}(\mu_{s,i}, \sigma_{s,i}^2)$, where $\mu_{s,i}$ is the mean and $\sigma_{s,i}^2$ is the variance expressing uncertainty in z_i . The variance can be further decomposed into two parts. First, we assume measurements taken from higher altitudes are more susceptible to environmental noise such as light conditions. To this end, we describe the degraded accuracy of sensor information at higher altitude by $\sigma_{a,i}^2 = \alpha h$, where $\alpha \in \mathbb{R}^+$ and h is the sensor’s altitude. Second, we consider uncertainty caused by observing incomplete grid cells. In our mapping framework, some grid cells are only partially covered by the current sensor footprint, especially when the grid cells occupy larger area after they are merged. Directly assigning the average measurements as the observation of these grid cells would be an over-confident assumption, as the unobserved part of these grid cells may contradict the current measurements, e.g., when grid cells span over the domain of heterogeneous function values. To tackle this problem, we propose the coverage-ratio-dependent variance $\sigma_{c,i}^2 = \beta \left(1 - \frac{A_{cover}}{A_c}\right)$ in our sensor model, where $\beta \in \mathbb{R}^+$ is a weighting and A_c, A_{cover} are the area of the grid cell and the part covered by the footprint. Note that this variance term decreases linearly with respect to coverage ratio and disappears under complete observation. This captures the intuition that partial observations contribute less to the update.

For each new measurement, its data are generated as follows: the sensor footprint is determined given the known intrinsic and extrinsic parameters. First, we query the grid cells having overlap with the footprint using depth-first tree search with pruning. Second, for observed grid cell c_i , we calculate the corresponding averaged measurement value z_i as illustrated in Fig. 2. Third, we sum $\sigma_{a,i}^2$ and $\sigma_{c,i}^2$ as the total variance of each measurement z_i .

C. Sequential Map Update

A major difference between GP regression and our GPF approach lies in the map update rule. During the online mapping mission, the map state is fully described by the mean vector $\mu(C)$ and covariance matrix $K(C, C)$, which is initialized by a GP model as shown in Eq. (3) and recursively updated by Bayesian fusion with new measurements.

For each update step, z is defined as a vector consisting of m new average function value measurements observed from m corresponding grid cells as introduced in Sec. III-B. The posterior density $p(\zeta|z, C) \propto p(z|\zeta, C) \cdot p(\zeta|C)$ can be directly computed using the Kalman Filter update equations [11]:

$$\mu^+ = \mu^- + \Gamma v, \quad (4)$$

$$K^+ = K^- - \Gamma H K^-, \quad (5)$$

where $\Gamma = K^- H^\top S^{-1}$ is the Kalman gain; $v = z - H\mu^-$ and $S = H K^- H^\top + R$ are the measurement and covariance innovations; R is a diagonal $m \times m$ matrix composed of variance term $\sigma_{a,i}^2 + \sigma_{c,i}^2$ associated with each measurement z_i and H is a $m \times n$ observation matrix denoting the part of the map observed by z , where n and m are the number of grid cells in the current map and observed grid cells respectively. Note that the current map only contains leaf grid cells and a small matrix $S \in \mathbb{R}^{m \times m}$ is inverted at each update.

D. Merging

Given a non-uniform target field for mapping, our goal is to use coarser (larger) grid cells to map uninteresting regions and denser (smaller) grid cells to retain details in interesting parts. Previous works in GPF [9, 10] do not support efficient resolution changes. By using our new GPF method with the integral kernel, however, we naturally encode the states of parent nodes in their children, which enables efficiently retrieving a parent’s posterior from its children on-the-fly.




Fig 3: Illustration of the merging operation in our map. Top and bottom rows show the grid cell map and its corresponding ND-tree (with $N = d = 2$) structure. Only leaf nodes (blue and orange) are considered in the map update. After merging (right), the states of child grid cells are summarized into their parent in the new map.

The online merging operation allows us to summarize information in larger areas and thus monotonically reduce the total number of grid cells in the map, which facilitates mapping efficiency and memory usage. For this, we subdivide our map into uninteresting regions (UR) and hotspots (HS):

$$C_{UR} = \{c_i \in C \mid \mu_i^+ + \gamma K_{i,i}^+ \leq f_{th}\}, C_{HS} = C \setminus C_{UR}, \quad (6)$$

where μ_i^+ and $K_{i,i}^+$ are the posterior mean and variance of grid cell c_i ; the design parameter γ is chosen to specify the margin to the threshold f_{th} [13]. f_{th} can be defined by expert knowledge in a certain application, e.g., in agricultural scenarios, high temperature may indicate crop health issues, see Sec. IV-B. This setting avoids merging grid cells with possibly high mean values, which would cause detail loss in interesting regions. Thus, grid cells are only merged if there is a strong belief that they fall within uninteresting bounds.

For a parent grid cell, if all of its $P = N^d$ child grid cells are uninteresting leaves (grid cells in C_{UR}), these child grid cells can be replaced by their parent grid cell. When we merge the information of P children into their parent, based on the definition of grid cell variable and the correlation encoded by the integral kernel, we have the parent grid cell defined as:

$$\zeta_{\text{parent}} = \frac{1}{P} \sum_{i=1}^P \zeta_{\text{child}_i}. \quad (7)$$

The parent grid cell now represents the average function value of the entire region covered by its children. For the grid map, the merging operation can be described as a linear transformation of a multivariate Gaussian as follows:

$$\mu^+(C_{\text{merged}}) = M\mu^+(C_{\text{original}}), \quad (8)$$

$$K^+(C_{\text{merged}}, C_{\text{merged}}) = MK^+(C_{\text{original}}, C_{\text{original}})M^\top, \quad (9)$$

where C_{original} represents the map including n grid cells before merging and C_{merged} is the newly-merged map. In the simplest case, where only one parent's children cells are merged, C_{merged} , C_{original} , and M can be expressed as:

$$C_{\text{original}} = \begin{bmatrix} c_1 \\ \vdots \\ c_{n-P} \\ c_{n-P+1} \\ \vdots \\ c_n \end{bmatrix}, \quad C_{\text{merged}} = \begin{bmatrix} c_1 \\ \vdots \\ c_{n-P} \\ c_{n-P+1} \end{bmatrix}, \quad (10)$$

$$M = \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ (n-P) \times (n-P) & (n-P) \times P \\ \mathbf{0} & \mathbf{Q} \\ 1 \times (n-P) & 1 \times P \end{bmatrix}, \quad (11)$$

assuming that c_{n-P+1} in C_{merged} now represents the parent of grid cells $\{c_{n-P+1}, \dots, c_n\}$ in C_{original} and \mathbf{Q} is $[\frac{1}{P}, \dots, \frac{1}{P}]$. A simple illustration is given in Fig. 3. The merging operation is performed for eligible grid cells after every map update. As the multivariate Gaussian distribution is closed under linear transformations, we now treat the map after the merging operation as a prior map for the next Bayesian fusion update.

IV. EXPERIMENTAL RESULTS

This section presents our experimental results. We first evaluate our proposed mapping strategy by comparing it against different benchmarks in terrain mapping scenarios. Then, we validate our approach using real-world surface temperature data and integrate it into an adaptive path planning framework to demonstrate its benefits for online robotic applications.

A. Mapping Evaluation

We evaluate mapping performance with total mapping time, mapping quality in terms of root mean square error (RMSE), intersection over union (IoU) of hotspots, memory consumption ratio, and number of grid cells in the final maps. The total mapping time is obtained by aggregating the individual update times over the mapping task; RMSE and IoU are calculated by comparing resulting maps and ground truth at the ground truth

map resolution; memory usage is reported as a ratio relative to the approach with the highest memory consumption. We compare six different mapping approaches:

- *FR-IDP*: vanilla fixed-resolution mapping under independence assumption [16];
- *AR-IDP*: adaptive-resolution mapping under independence assumption. Uninteresting grid cells are pruned during mapping as proposed in [8];
- *AR-BCM*: adaptive-resolution mapping using BCM and test-data tree, as adapted from [22]. Uninteresting grid cells are pruned to reduce the number of query points in BCM. Note that we do not follow nested BCM approach as our whole map can be seen as a block in their case;
- *AR-GPR-IK*: adaptive-resolution GP regression with integral kernel based on the original approach proposed in [24]. We take one step further to recursively merge grid cells if they are uninteresting after each regression update;
- *FR-GPF*: fixed-resolution GPF proposed in [3];
- *Ours*: our new adaptive-resolution mapping strategy based on GPF with integral kernel as described in Sec. III.

We simulate 20 $20m \times 20m$ Gaussian random fields with 400×400 resolution as ground truth environments representing a spatially correlated field on a terrain. For simplification, the ground truth field values are normalized to $[0, 1]$ and we define regions with values greater than 0.7 as hotspots of interest. To assess mapping performance at different scales, we conduct experiments at 3 different maximal resolutions: 16×16 , 32×32 , and 64×64 grid cell maps corresponding to adaptive-resolution approaches with maximum tree depths of 4, 5, and 6, respectively, in a quadtree configuration. Note that different maximal resolution settings are possible by adopting the general ND-tree decomposition.

The terrains are mapped using a lawnmower pattern to focus on comparing the methods in terms of mapping performance only, excluding the influence of path variations. To simulate a UAV mission, we take 16 non-overlapping measurements as shown in Fig. 4a to fully cover the terrain, assuming a flight altitude of $2.5m$ and $5m \times 5m$ sensor footprint on the ground. All GP-based mapping approaches (*AR-BCM*, *AR-GPR-IK*, *FR-GPF*, *Ours*) use the SE kernel function with hyperparameters $\theta = \{\sigma^2, l\} = \{1, 2.36\}$ and a constant prior mean of 0.5. This design choice is based on our assumption on smooth latent function and access to the ground truth. In general, the domain knowledge should be exploited in kernel function selection and the hyperparameters can be optimized using prior information, e.g., data sets from an earlier sampling campaign or similar fields. For approaches using an integral kernel (*AR-GPR-IK*, *Ours*), we follow Eqs. (1) and (2) to calculate the prior maps. For mapping under independence assumption (*FR-IDP*, *AR-IDP*), we use the same prior mean and variance, while ignoring all cross-correlation terms to isolate each grid cell. We consider the ‘average measurement’ model in all mapping approaches. For merging operation in adaptive-resolution approaches, we choose $\{\gamma, f_{\text{th}}\} = \{2, 0.7\}$ in Eq. (6). Note that all these hyperparameters are manually tuned but used consistently in all experiments.

The results are summarized in Tab. I and Fig. 4. In all




Fig 4: Qualitative comparison of our approach (g) against benchmarks (b)-(f). The terrain is mapped using a lawnmower strategy, as shown in (a). The red line and black dots indicate the travelled path and measurement locations. All approaches use a map size of 32×32 grid cells. By mapping adaptively, our method compresses information in areas with low information value (blue) while preserving details in higher-value areas of interest (yellow) to achieve an efficient, compact map representation for online applications. (h) shows the offline higher-resolution (50×50) reconstruction from our online mapping result (g), illustrating how the map can be decompressed after the mission.

Map size	Method	RMSE \downarrow	RMSE (hotspots) \downarrow	IoU (hotspots) \uparrow	Mapping time [ms] \downarrow	Memory usage ratio [%] \downarrow	Number of map cells \downarrow
16×16	FR-IDP	0.045 ± 0.002	0.045 ± 0.002	0.813 ± 0.024	5.575 ± 0.466	4.187 ± 0	256 ± 0
	AR-IDP	0.071 ± 0.003	0.046 ± 0.002	0.812 ± 0.024	7.299 ± 0.901	2.155 ± 1.127	125.2 ± 18.258
	AR-BMC	0.071 ± 0.003	0.038 ± 0.003	0.856 ± 0.023	273.546 ± 69.549	49.892 ± 11.433	115.22 ± 16.167
	AR-GPR-IK	0.065 ± 0.003	0.037 ± 0.002	0.856 ± 0.024	59.379 ± 20.680	40.692 ± 8.282	118.62 ± 16.907
	FR-GPF	0.037 ± 0.002	0.037 ± 0.002	0.857 ± 0.023	11.375 ± 0.965	100 ± 0	256 ± 0
	Ours	0.065 ± 0.003	0.037 ± 0.002	0.857 ± 0.026	10.168 ± 1.119	38.729 ± 9.416	114.4 ± 19.754
32×32	FR-IDP	0.065 ± 0.004	0.065 ± 0.004	0.723 ± 0.021	28.968 ± 0.834	1.067 ± 0	1024 ± 0
	AR-IDP	0.079 ± 0.005	0.067 ± 0.003	0.725 ± 0.020	63.796 ± 6.880	0.529 ± 0.068	508 ± 65.121
	AR-BMC	0.071 ± 0.005	0.025 ± 0.003	0.864 ± 0.023	13100.225 ± 2009.709	26.433 ± 5.359	356.25 ± 74.012
	AR-GPR-IK	0.066 ± 0.004	0.026 ± 0.003	0.866 ± 0.024	1747.631 ± 677.670	17.538 ± 5.640	371.3 ± 74.279
	FR-GPF	0.027 ± 0.002	0.026 ± 0.002	0.867 ± 0.024	430.843 ± 7.447	100 ± 0	1024 ± 0
	Ours	0.065 ± 0.004	0.026 ± 0.003	0.867 ± 0.025	261.763 ± 24.443	16.632 ± 5.112	360.4 ± 71.003
64×64	FR-IDP	0.123 ± 0.003	0.123 ± 0.008	0.625 ± 0.056	123.562 ± 4.271	0.268 ± 0	4096 ± 0
	AR-IDP	0.127 ± 0.004	0.124 ± 0.008	0.623 ± 0.058	586.427 ± 43.362	0.218 ± 0.023	2687.14 ± 383.364
	AR-BMC	0.104 ± 0.004	0.025 ± 0.003	0.869 ± 0.028	68645.365 ± 4606.239	15.782 ± 4.125	1258.5 ± 473.78
	AR-GPR-IK	0.073 ± 0.004	0.025 ± 0.003	0.872 ± 0.021	18620.558 ± 6448.663	11.552 ± 4.221	1387 ± 443.744
	FR-GPF	0.024 ± 0.002	0.024 ± 0.002	0.875 ± 0.023	9977.749 ± 251.003	100 ± 0	4096 ± 0
	Ours	0.073 ± 0.004	0.024 ± 0.002	0.876 ± 0.023	4098.029 ± 548.881	8.877 ± 4.824	1271.25 ± 484.191

TABLE I: Comparison of our approach against benchmarks for varying map sizes. By combining GPF and adaptive-resolution mapping using an integral kernel, our strategy reduces runtime and memory consumption while delivering highly accurate maps.

cases, approaches relying on the cell independence assumption yield least accurate maps with highest RMSE and lowest IoU, since they are most vulnerable to sensor noise or sparse measurements. This is because they neglect correlations for mapping, which are key for capturing continuous variables. In contrast, the four GP-based approaches reflect the smooth structure of the Gaussian random field, as they incorporate covariance information into the map update. As expected, the averaging effect caused by merging cells in adaptive-resolution approaches leads to higher total RMSE compared to *FR-GPF*. However, all GP-based approaches show the same accuracy in mapping hotspots as well as IoU scores, as required in our problem setup. In terms of mapping efficiency, *AR-BCM*

performs the worst as it executes large matrix inversion and BCM fusion at every update step, leading to prohibitively slow mapping. Note that BCM benefits from parallelizing several GP regressions. However, in online mapping scenarios, where measurements are accumulated incrementally, BCM loses this strength. *AR-GPR-IK* is slower than two GPF approaches (*FR-GPF* and *Ours*), due to regression using accumulated measurements. We point out that by using the integral kernel together with the ‘average measurement’ model, *AR-GPR-IK* already achieves significant speed-up compared to vanilla GP regression. In all cases, *AR-IDP* is slower than *FR-IDP* due to overhead caused by tree search. The same overhead is expected in our approach, however, as the major bottleneck in fusion

is the matrix inversion and multiplication in Eqs. (4) and (5), this can be compensated by faster Bayesian fusion update with less grid cells in our approach. In terms of memory usage, *FR-GPF* consumes the most memory space as it maintains a large constant number of grid cells and a large covariance matrix. Among the adaptive-resolution approaches, *AR-IDP* shows the worst merging ability, as indicated by the number of grid cells in the final map. This can be explained by heterogeneous states in children nodes caused by inaccurate mapping, which potentially reduces the chances of merging operation. Among the GP-based methods, our new approach achieves the fastest mapping updates and best memory compression ratios with competitive map quality. The benefit of online merging operation can be seen by comparing *Ours* and *FR-GPF*. In all cases, our approach outperforms *AR-IDP* and *FR-IDP* in terms of map quality. In Fig. 4h, we show how our mapping result can be decompressed to recover a high-resolution reconstruction in an offline post-processing step.

B. Validation on Real-World Data

We demonstrate our mapping approach using real-world surface temperature data. The data was collected in a $150m \times 150m$ crop field (50.86° lat., 6.45° lon.) near Jülich, Germany on June 25, 2021 using a DJI Matrice 600 UAV platform equipped with a Vue Pro R 640 thermal sensor. During data acquisition, the UAV followed a lawnmower path at $100m$ altitude to collect images at $15cm$ spatial ground resolution. The images were processed using Pix4D software to create an orthomosaic used as a proxy for ground truth in our experiment. We use a maximal map resolution of 64×64 . The entire mapping takes $28.31s$ considering 81 measurements with 50% overlap. The aim is to validate our method for adaptively mapping temperature hotspots ($> 28^\circ C$) at finer resolutions using this real data. The mapping result in Fig. 5 confirms that our approach can adapt the map resolution in a targeted way.

C. Application for Adaptive Path Planning

Finally, we integrate our mapping approach into an adaptive path planning framework for UAV-based terrain monitoring [10] to demonstrate its benefits for online robotic scenarios. The planning task aims at efficient detection of regions of interest in an initially unknown environment under time constraints. For this, the UAV must adaptively plan its path based on the current map to trade off between exploration and exploitation. This experiment considers the same setup as described in Sec. IV-A except setting our prior mean to 0.7 to initially encourage exploration. We compare the *FR-IDP*, *AR-IDP*, *FR-GPF* methods to our approach, as regression-based mapping approaches are prohibitively slow for online planning. We use a 3D lattice consisting of 300 total waypoints at altitudes of $2m$ and $5m$ to represent the discrete action space. The planner applies greedy search among these candidates to find the next best measurement position by forward-simulating the map update and calculating the reward. The information-theoretic reward is defined by the posterior variance reduction in regions of interest divided by the flight

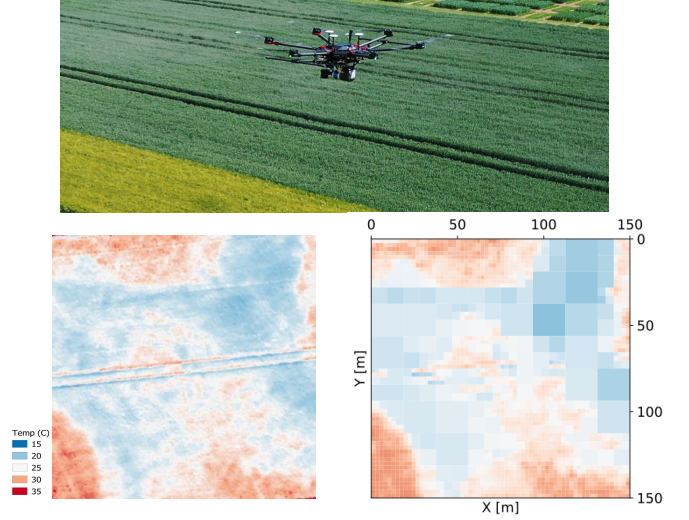


Fig 5: Validation of our approach for surface temperature mapping. Top: Experimental setup showing our UAV over a crop field. Bottom-left: Temperature data of the crop field. Bottom-right: Map generated by our method. High-temperature areas (red) are mapped at higher resolutions to preserve detail in these regions of interest.

time to the waypoint candidate. For more technical details, please refer to the planning framework of Popović et al. [10].

We conduct experiments on 10 simulated Gaussian random fields and plot the evolution of RMSE (hotspots) and IoU over mission time in Fig. 6. The mission time is the sum of planning time, mapping time, and flight time. The results show that planning using our mapping approach achieves the best IoU and RMSE (hotspots) scores with the shortest mission time, which is favorable for autonomous monitoring tasks using resource-constrained UAVs. Planning using our approach outperforms *FR-GPF* due to more efficient map updates, which significantly accelerates forward-simulation during predictive planning. The poorer planning performance using *FR-IDP* and *AR-IDP* is a direct consequence of inaccurate mapping results. As observed in Sec. IV-A, mapping approaches using independence assumption neglect important spatial correlation and are thus more susceptible to sensor noise. Due to inaccurate mapping, the false positive interesting areas mislead the UAV, leading to the close inspection of actually uninteresting regions. This inaccuracy deprives *FR-IDP* and *AR-IDP* of their advantage in fast planning.

V. CONCLUSIONS AND FUTURE WORK

This paper proposes a new approach for online field mapping. We introduce a novel use of an integral kernel in GPF framework and use ND-tree to store our map. Combining these two elements enables us to adapt the map resolution on-the-fly while neatly maintaining spatial correlations. Results show that our approach achieves competitive performance in terms of mapping efficiency, memory usage, and map quality. Our mapping approach is validated using real-world surface temperature data. Moreover, we demonstrate that faster and more accurate map updates facilitate adaptive path planning for efficient information gathering in robotic applications.




Fig 6: Comparison of mapping approaches for adaptive path planning in terrain monitoring. Our strategy (red) performs best to efficiently reconstruct hotspot areas in an unknown environment with highest mapping accuracy (top) and map quality (bottom). Solid lines and shaded regions represent means and standard deviations over 10 trials.

Our approach has several limitations. First, partitioning the map to higher resolutions is not considered in our online mapping framework, which could limit its applicability in dynamic environments. Second, the tree structure in our approach prevents merging neighboring grid cells if they have different parents, which constrains its compression capabilities. Third, we assume deterministic sensor localization. Several works have tackled this problem by incorporating localization noise in GP mapping [17, 27], which improves the robustness. Finally, the map scale is still limited by maintaining a global covariance matrix. For mapping at larger scales, our approach could be used to generate local sub-maps which are fused by a BCM or the approach proposed by Sun et al. [28]. Future work will address these issues and apply our approach in domains with higher spatial variability.

ACKNOWLEDGEMENT

We would like to thank Mr. Jordan Bates from Forschungszentrum Jülich for providing the temperature data.

REFERENCES

- [1] S. Manfreda, M. F. McCabe, P. E. Miller, R. Lucas, V. Pajuelo Madrigal, G. Mallinis, E. Ben Dor, D. Helman, L. Estes, G. Ciraolo *et al.*, “On the Use of Unmanned Aerial Systems for Environmental Monitoring,” *Remote Sensing*, vol. 10, no. 4, 2018.
- [2] G. Tmušić, S. Manfreda, H. Aasen, M. R. James, G. Gonçalves, E. Ben-Dor, A. Brook, M. Polinova, J. J. Arranz, J. Mészáros, R. Zhuang, K. Johansen, Y. Malbeteau, I. P. de Lima, C. Davids, S. Herban, and M. F. McCabe, “Current Practices in UAS-based Environmental Monitoring,” *Remote Sensing*, vol. 12, no. 6, 2020.
- [3] 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 *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2017, pp. 1382–1388.
- [4] G. A. Hollinger and G. S. Sukhatme, “Sampling-based robotic information gathering algorithms,” *The International Journal of Robotics Research*, vol. 33, pp. 1271 – 1287, 2014.

- [5] F. Stache, J. Westheider, F. Magistri, M. Popović, and C. Stachniss, “Adaptive path planning for uav-based multi-resolution semantic segmentation,” in *European Conference on Mobile Robots*, 2021.
- [6] A. Hornung, K. M. Wurm, M. Bennewitz, C. Stachniss, and W. Burgard, “OctoMap: An efficient probabilistic 3D mapping framework based on octrees,” *Autonomous Robots*, vol. 34, no. 3, pp. 189–206, 2013.
- [7] N. Funk, J. Tarrio, S. Papatheodorou, M. Popović, P. F. Alcantarilla, and S. Leutenegger, “Multi-Resolution 3D Mapping with Explicit Free Space Representation for Fast and Accurate Mobile Robot Motion Planning,” *Robotics and Automation Letters*, vol. 6, pp. 3553–3560, 2021.
- [8] E. Einhorn, C. Schröter, and H. M. Gross, “Finding the adequate resolution for grid mapping - Cell sizes locally adapting on-the-fly,” in *IEEE International Conference on Robotics and Automation*. IEEE, 2011, pp. 1843–1848.
- [9] T. Vidal-Calleja, D. Su, F. De Bruijn, and J. V. Miro, “Learning spatial correlations for Bayesian fusion in pipe thickness mapping,” in *IEEE International Conference on Robotics and Automation*. IEEE, 2014, pp. 683–690.
- [10] 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, pp. 889–911, 2020.
- [11] S. Reece and S. Roberts, “An introduction to Gaussian processes for the Kalman filter expert,” in *International Conference on Information Fusion*. IEEE, 2010.
- [12] C. Stachniss, C. Plagemann, and A. Lilienthal, “Learning Gas Distribution Models using Sparse Gaussian Process Mixtures,” *Autonomous Robots*, vol. 26, pp. 187–202, 2009.
- [13] G. Hitz, A. Gotovos, F. Pomerleau, M.-E. Garneau, C. Pradalier, A. Krause, and R. Siegwart, “Fully Autonomous Focused Exploration for Robotic Environmental Monitoring,” in *IEEE International Conference on Robotics and Automation*. IEEE, 2014, pp. 2658–2664.
- [14] 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*, vol. 34, no. 8, pp. 1427–1449, 2017.
- [15] S. Vasudevan, F. Ramos, E. Nettleton, H. Durrant-Whyte, and A. Blair, “Gaussian Process modeling of large scale terrain,” in *IEEE International Conference on Robotics and Automation*. IEEE, 2009, pp. 1047–1053.
- [16] A. Elfes, “Using occupancy grids for mobile robot perception and navigation,” *Computer*, vol. 22, no. 6, pp. 46–57, 1989.
- [17] S. T. O’Callaghan and F. T. Ramos, “Gaussian process occupancy maps,” *The International Journal of Robotics Research*, vol. 31, no. 1, pp. 42–62, 2012.
- [18] C. Rasmussen and C. Williams, *Gaussian Processes for Machine Learning*. MIT Press, 2006.
- [19] Y. Shen, A. Ng, and M. Seeger, “Fast Gaussian Process Regression using KD-Trees,” *Neural Information Processing Systems*, vol. 18, 2005.
- [20] S. Kim and J. Kim, “Recursive bayesian updates for occupancy mapping and surface reconstruction,” in *IEEE International Conference on Robotics and Automation*, 2014.
- [21] V. Tresp, “A Bayesian Committee Machine,” *Neural computation*, vol. 12, pp. 2719–41, 12 2000.
- [22] J. Wang and B. Englot, “Fast, accurate Gaussian process occupancy maps via test-data octrees and nested Bayesian fusion,” in *IEEE International Conference on Robotics and Automation*. IEEE, 2016, pp. 1003–1010.
- [23] S. O’Callaghan and F. Ramos, “Continuous Occupancy Mapping with Integral Kernels,” in *AAAI Conference on Artificial Intelligence*. AAAI, 2011.
- [24] A. Reid, F. Ramos, and S. Sukkarieh, “Bayesian fusion for multi-modal aerial images,” in *Robotics: Science and Systems*, 2013, pp. 1–8.
- [25] Y. Chen, W. Shuai, and X. Chen, “A probabilistic, variable-resolution and effective quadtree representation for mapping of large environments,” in *International Conference on Advanced Robotics*. IEEE, 2015, pp. 605–610.
- [26] S. Särkkä, “Linear operators and stochastic partial differential equations in Gaussian process regression,” *Lecture Notes in Computer Science*, vol. 6792, no. 2, pp. 151–158, 2011.
- [27] M. G. Jadidi, J. V. Miro, and G. Dissanayake, “Sampling-based incremental information gathering with applications to robotic exploration and environmental monitoring,” *The International Journal of Robotics Research*, vol. 38, no. 6, pp. 658–685, 2019.
- [28] L. Sun, T. Vidal-Calleja, and J. V. Miro, “Bayesian fusion using conditionally independent submaps for high resolution 2.5D mapping,” in *IEEE International Conference on Robotics and Automation*, vol. 2015-June, no. June, 2015, pp. 3394–3400.