

Bayesian Active Learning for Sample Efficient 5G Radio Map Reconstruction

Konstantinos D. Polyzos, Alireza Sadeghi, Wei Ye, Steven Sleder, Kodjo Houssou, Jeff Calder, Zhi-Li Zhang, and Georgios B. Giannakis

Abstract—The advent of diverse frequency bands in 5G networks has promoted measurement studies focused on 5G signal propagation, aiming to understand its pathloss, coverage, and channel quality characteristics. Nonetheless, conducting a thorough 5G measurement campaign is markedly laborious given the large number of samples that must be collected. To alleviate this burden, the present contribution leverages principled active learning (AL) methods to prudently select only a *few*, yet most *informative* locations to collect samples. The core idea is to rely on a Gaussian Process (GP) model to efficiently extrapolate measurements throughout the coverage area. Specifically, an ensemble (E) of GP models is adopted that not only provides a rich learning function space, but also quantifies uncertainty, and can offer accurate predictions. Building on this EGP model, a suite of acquisition functions (AFs) are advocated to query new locations on-the-fly. To account for realistic scenarios, the proposed AFs are augmented with a novel distance-based AL rule that selects informative samples, while penalizing queries at long distances. Numerical tests on 5G data generated by the Sionna simulator and on real urban and suburban datasets, showcase the merits of the novel EGP-AL approaches.

Index Terms—Active Learning, 5G measurement, Radio Map Reconstruction

I. INTRODUCTION

The fifth-generation cellular network, commonly known as 5G, expands its spectrum resources across both the low/mid-frequency bands (<6GHz) and the high-frequency bands (>24GHz) [1]. This brings a remarkable improvement in service quality and outperforms 4G (LTE¹), which typically operates at frequencies below 2 GHz. In some regions, however, higher bands are also used, including 2.1 GHz, 2.5 GHz, 2.6 GHz, and even up to 6 GHz [2]. In particular, major 5G operators like ATT and Verizon in the US have been actively deploying C-band (3.7GHz) and mmWave (26GHz) technologies, achieving impressive gigabits-per-second-level throughput and millisecond-level latency [1], [3], [4].

Nevertheless, applying these frequency bands introduces new characteristics to mobile networks. For instance, mmWave

has a much shorter wavelength (about 1~12mm) [4], making its propagation highly directional, easily blocked and reflected by obstacles. This can lead to complex signal strength behaviors in the real world beyond just distance-based dependency [5], [6].

A plethora of wireless network-related tasks, such as network/radio resource/base station planning, handover optimization, quality of Service (QoS) management, etc., critically rely on a better understanding of radio propagation characteristics in the real world and require quantifying the signal coverage quality of specific base stations. To gain such insight, a large number of in-field measurements need to be conducted.

Obtaining these measurements is a laborious process, since the collection of valid measurements entails a significant amount of effort and the collected data may not even cover all of the locations within the coverage area. The lack of immediate feedback on data quality during in-field measurements can further lead to many redundant and/or useless measuring efforts. To cope with these challenges, existing works [6]–[18] have attempted to reconstruct radio maps only using *limited* measurements (see Sec. II). Albeit interesting, these approaches still fail to fully address the aforementioned practical issues.

The goal of this work is to improve the efficiency of obtaining measurements for reconstructing 5G radio maps by leveraging Bayesian-based active learning (AL) methods. Building on a statistical model to capture the mapping of any location within the coverage area to its corresponding received signal strength (RSS) value, AL leverages a collection of acquisition criteria (i.e., *acquisition functions* (AFs)) to select a *few* highly *informative* locations to obtain RSS measurements. We utilize the Gaussian Process (GP) for the underlying statistical model, which is a well-motivated Bayesian model that is able to effectively extrapolate and interpolate RSS measurements in regions with limited or even no observed data. Along with accurate RSS predictions at unobserved locations, the GP model provides valuable information about the confidence and reliability of predictions by offering a well-quantifiable notion of uncertainty around each predicted RSS value [19]. GPs are non-parametric models that do not require explicit assumptions about the functional form or distribution of the data. While GP models define a Gaussian *prior* over the target function, they effectively handle various data distributions by employing appropriate *likelihood* functions [19]. With the right choice of (i) likelihood function; and, (ii) the underlying kernel, GPs excel in settings where the predictive distribution may be intricate, or not easily captured by simple parametric models,

Konstantinos D. Polyzos, Alireza Sadeghi and Georgios B. Giannakis are with the Department of Electrical and Computer Engineering at University of Minnesota, Twin Cities (emails: {polyz003, sadeg012, georgios}@umn.edu).

Wei Ye, Steven Sleder and Zhi-Li Zhang are with the Department of Computer Science and Engineering at University of Minnesota, Twin Cities (emails: {ye000094, slede001}@umn.edu and zh Zhang@cs.umn.edu).

Kodjo Houssou and Jeff Calder are with the School of Mathematics at University of Minnesota, Twin Cities (emails: {houss001, jwcalder}@umn.edu).

This work was supported by NSF grants 2212318, 2220292, 2126052, 2102312, 2103256, 2312547. The work of Konstantinos D. Polyzos was also supported by the Onassis Foundation Scholarship.

¹Long Term Evolution.

making them versatile for numerous applications [19], [20].

Nonetheless, their performance hinges on a pre-selected kernel function to evaluate the pairwise similarity of different locations, which subsequently affects the covariance of the corresponding RSS values. Apriori selection of the kernel function is a nontrivial task, that may require domain-knowledge as well. The use of a single pre-selected kernel in GP modeling confines the expressiveness of the sought learning function for estimating RSS values. In addition, the complexity of these methods significantly grows as the number of available measurements increases. To cope with expressiveness challenges of single kernel-based GP learning, and to allow for *scalable* and *online* model updates, we will rely on an ensemble of parametric approximants of GP models, similarly as in [21] (see also Sec.II).

Contributions. Our contributions in this work can be summarized in the following aspects:

- To the best of our knowledge, this is the first work that formulates and implements Bayesian-AL for the 5G radio map reconstruction problem using only a small number of measurements.
- We advocate intuitive AL methods that leverage an ensemble of GPs (EGPs) to enhance the model’s expressiveness when dealing with multi-modal and complex data distributions. We then leverage parametric function approximants of the GP models in the ensemble using random features (RFs) to offer online and scalable model updates. These are particularly appealing in AL settings where new measurements are processed online.
- To allow for path-efficient acquisition of RSS measurements in new locations, we combine a number of well-motivated AFs with a novel *path efficient* rule that penalizes querying far-distant locations at each iteration of AL.
- We conduct thorough experiments on a total of six simulated datasets with three scenarios and two radio frequencies, and on two real urban and suburban datasets, to corroborate the efficiency and effectiveness of the proposed EGP-based AL in reconstructing 5G radio maps. Our numerical tests show that our EGP-based AL approaches achieve up to 54.8% improvement in terms of normalized mean square error (NMSE) over the EGP model with random sampling. In addition, compared to the EGP-AL methods that do not impose any constraints on the acquisition of new locations to measure, the EGP-based AL approaches that consider the novel path efficient rule, achieve up to 57.1% reduction in terms of required cumulative traveling distance to obtain measurements.

The remainder of the paper is organized as follows. In Sec. II related works are discussed. Active learning formulation for 5G radio map reconstruction is introduced in Sec. III, while the advocated EGP-based AL approach with path-efficient acquisition functions is presented in IV. Upon discussing details about data generation in Sec. V, numerical tests on simulated and real 5G datasets are provided in Sec. VI. Finally, Sec. VII concludes the paper.

Notations. Scalars are denoted by lowercase, column vectors by bold lowercase, matrices by bold uppercase, and sets by

TABLE I: Symbol Definitions

Symbol	Definition
\mathcal{X}	Spatial coverage area
\mathbf{x}	2-dimensional Cartesian coordinate
$y(\mathbf{x})$	Received Signal Strength (RSS) at location $\mathbf{x} \in \mathcal{X}$
\mathbf{x}_i	i -th probed location in the coverage area \mathcal{X}
y_i	RSS value measured at the i -th probed location \mathbf{x}_i
$f(\mathbf{x})$	Function mapping location \mathbf{x} to associated RSS $y(\mathbf{x})$
$\hat{f}(\mathbf{x})$	Nonparametric function estimate at location \mathbf{x}
$\tilde{f}(\mathbf{x})$	Linear and parametric function approximant of $f(\mathbf{x})$
ν	Gaussian observation noise $\nu \sim \mathcal{N}(\nu, 0, \sigma_{\text{noise}}^2)$
\mathcal{D}_n	Dataset of n pairs of data samples $\{(\mathbf{x}_i, y_i)\}_{i=1}^n$
$\alpha(\mathbf{x}; \mathcal{D}_n)$	Acquisition function
$\kappa(\mathbf{x}_i, \mathbf{x}_j)$	Kernel function to evaluate similarity of \mathbf{x}_i and \mathbf{x}_j
$\mathbf{X}_n := [\mathbf{x}_1, \dots, \mathbf{x}_n]^\top$	Data matrix containing all n probed locations
\mathbf{K}_n	Covariance (kernel) matrix of the GP prior over \mathbf{X}_n
$\mathbf{k}_n(\mathbf{x})$	Kernel vector containing similarity between location \mathbf{x} and all n locations in \mathbf{X}_n
\mathbf{y}_n	Vector of observed RSS values at locations \mathbf{X}_n
$\bar{y}_{n+1 \mathbf{n}}(\mathbf{x})$	Predictive mean of RSS at location \mathbf{x} given \mathcal{D}_n
$\sigma_{n+1 \mathbf{n}}^2(\mathbf{x})$	Predictive variance of RSS at location \mathbf{x} given \mathcal{D}_n
\mathcal{K}	Set of M kernels $\{\kappa^1, \dots, \kappa^M\}$ in mixture of GPs
w_n^m	Weight of the m -th GP model in the ensemble
$\phi_\zeta(\mathbf{x})$	Random feature vector
$\boldsymbol{\theta}$	Parameter vector to linearly approximate GP
$\mathcal{N}(\cdot; \boldsymbol{\mu}, \mathbf{K})$	Gaussian distribution with mean $\boldsymbol{\mu}$ and covariance \mathbf{K}
$n+1 \mathbf{n}$	Processing location $n+1$ uses all previous $\mathbf{n} := [1, \dots, n]^\top$ measurements
$n+1 n$	Processing location $n+1$ uses only previous measurement n
$\hat{\boldsymbol{\theta}}_{n+1}^m$	Posterior mean of parametric model m when processing location $n+1$
$\hat{\boldsymbol{\Sigma}}_{n+1}^m$	Posterior covariance of parametric model m when processing location $n+1$
\mathcal{H}	Reproducing Kernel Hilbert Space
$H(\hat{f}(\mathbf{x}) \mathcal{D}_n)$	Entropy of the learned function $\hat{f}(\mathbf{x}) \forall \mathbf{x} \in \mathcal{X}$

calligraphic uppercase fonts. Superscripts $^\top$ and $^{-1}$ denote transpose and inverse respectively; while $\mathbf{1}_N$ stands for the $N \times 1$ all-ones vector; and $\mathcal{N}(\mathbf{x}; \boldsymbol{\mu}, \mathbf{K})$ for the PDF of a Gaussian random vector \mathbf{x} with mean $\boldsymbol{\mu}$, and covariance matrix \mathbf{K} . The identity matrix will be represented by \mathbf{I} , and the all-zeros matrix by $\mathbf{0}$. Table I lists the most common symbols for reference purposes.

II. RELATED WORKS

Radio Map Reconstruction. Numerous approaches have been proposed to reconstruct 5G radio maps, comprising both non-parametric methods [10]–[12], [22]–[24] and classical parametric ones [13], [14], [25]. More recently, researchers have explored deep learning-based approaches [6], [15]–[18], [26]–[28] to enhance the accuracy of radio propagation modeling and better capture its characteristics. In addition, physics-informed conditional generative models have demonstrated enhanced ability to enrich geometric information, and reconstruct dense radio maps from sparse observations [29]. However, these methods primarily concentrate on interpolating or extrapolating unmeasured locations based on existing measured points, lacking the ability to guide future sampling—a critical practical challenge. In contrast, the novel active learning method adopted in this paper effectively addresses this challenge by providing a principled method to search for the most informative location to query the next measurement.

Active Learning. AL methods can be broadly categorized into parametric AL [30]–[32] and non-parametric ones [33]–[35]. Although most existing parametric approaches in the literature excel at learning complex representations from data, they are *deterministic* and hence fall short in inherently providing uncertainty quantification. In addition, they are tailored solely for classification tasks, and do not account for regression tasks, such as the RSS prediction problem in this paper; see e.g [36]. On the other hand, Bayesian non-parametric approaches such as GPs, inherently offer uncertainty quantification, and demonstrate superior sample efficiency [19], [34]. Existing GP-based AL approaches typically rely on a single pre-selected kernel, which limits their expressiveness when dealing with data exhibiting multi-modal behavior. To cope with this challenge, existing deep-learning based approaches aim to learn the kernel with neural networks, but they typically require a sufficient number of labels which are not available in several AL settings; see e.g [37]. The proposed AL method in this paper uses a more expressive statistical model that relies on an ensemble of GP models with scalable online model updates as in [38]. Finally, the proposed acquisition strategies in the existing literature do not impose constraints on the data to label, which discourages their application in the 5G map reconstruction problem where the travelling cost of collecting 5G measurements should be minimized. A recent work [39] has proposed AL methods that consider similarities of query data samples with previously queried ones, though the method is not applicable to 5G signal reconstruction. Specifically, the proposed method operates in batch form, is tailored for classification, and the selection of query samples is based on enforcing at least a certain pairwise distance between samples for the sake of accelerating data annotation through parallel processing. In contrast, this work aims to *avoid* querying locations that are far from each other. To cope with this challenge, this paper incorporates a novel path-efficient acquisition rule into the AL process, where locations at long distances are penalized.

III. PROBLEM STATEMENT

We now formulate the 5G active measurement problem (see Fig.1) and then briefly discuss the limitations of AL when it is directly applied in practice.

Active Learning for 5G Measurements. For a 5G coverage area $\mathcal{X} \subset \mathbb{R}^2$, we are interested in estimating the real-valued received signal strength (RSS) $y(\mathbf{x}) \in \mathbb{R}$ at any desired location (or feature vector) $\mathbf{x} := [\text{Cartesian Coordinate-x, Cartesian Coordinate-y}] \in \mathbb{R}^2$. Let \mathbf{x}_i denote the i -th probed location with associated scalar RSS value $y_i \in \mathbb{R}$. We postulate that there exists a ground truth function $f(\mathbf{x}) : \mathbb{R}^2 \rightarrow \mathbb{R}$ which maps a desired location \mathbf{x} to its associated RSS through

$$y = f(\mathbf{x}) + \nu, \quad (1)$$

where $\nu \sim \mathcal{N}(\nu; 0, \sigma_{\text{noise}}^2)$ represents the observation noise.

To *efficiently* learn an estimate $\hat{f}(\mathbf{x})$ at any desired \mathbf{x} with only few, but representative, measured RSS data, one can rely on the *active learning* (AL) paradigm; see Fig. 1 for an illustration. AL begins with a relatively small dataset of

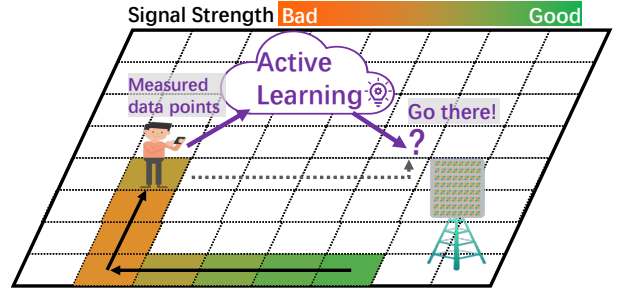


Fig. 1: Illustration of active learning for efficient RSS radiomap reconstruction. The area of interest is partitioned into small grids with measurements shown in colors from green (high RSS) to orange (low RSS). The objective is to identify the next most informative measurement location, enabling accurate and sample-efficient RSS reconstruction over the entire area.

labeled (measured) data $\mathcal{D}_n := \{(\mathbf{x}_i, y_i)\}_{i=1}^n$ (where n is small). To judiciously update \mathcal{D}_n and build an informative training set, AL leverages an *acquisition function* (AF) $\alpha(\cdot)$ to search for the *most informative* query locations on-the-fly while updating the underlying model $\hat{f}(\cdot)$. The next query location \mathbf{x}_{n+1} is obtained by solving

$$\mathbf{x}_{n+1} := \arg \min_{\mathbf{x} \in \mathcal{X}} \alpha(\mathbf{x}; \mathcal{D}_n). \quad (2)$$

To obtain the RSS value $y_{n+1} := y(\mathbf{x}_{n+1})$ corresponding to the *next* queried location \mathbf{x}_{n+1} , a field-test² can be carried out. Upon obtaining y_{n+1} , AL augments the training dataset $\mathcal{D}_{n+1} := \mathcal{D}_n \cup \{(\mathbf{x}_{n+1}, y_{n+1})\}$ and updates the estimate $\hat{f}(\cdot)$ on-the-fly. Having an updated $\hat{f}(\cdot)$, the active learner can predict the RSS at any desired (unmeasured) location $\mathbf{x} \in \mathcal{X}$.

Therefore, to learn $\hat{f}(\cdot)$ in a *sample-efficient manner* with AL, we aim to properly select a learning model and design AFs relying on the chosen model, to effectively guide the acquisition of new query locations to measure in an online manner.

Challenges Conventional AL methods rely on the ‘independent data acquisition’ assumption, meaning that the active learner can query measurements from any location within the coverage area independently from the previous measurement locations. This limits the applicability of conventional AL in wireless communications, since there exist many practical considerations that one needs to consider when querying a new location, such as traveling cost and measurement runtime complexity. For example, simply selecting a location that may be informative but requires a much longer travel distance is not cost-optimized, as the user may have to travel a prohibitively long distance to take a sufficient number of measurements and may repeatedly traverse the same regions within the measurement area (see Sec. VI). In the ensuing sections, we provide principled methods to facilitate a *path-* and *sample-*efficient data acquisition process.

²One can obtain RSS measurements at different locations within \mathcal{X} , using some off-the-shelf tools, such as Android-based API [40] running on smartphones or more professional scanners [41].

IV. ACTIVE LEARNING WITH ENSEMBLES OF GPs

To effectively estimate the sought function f with *well-quantifiable uncertainty* in a *sample-efficient fashion*, this work leverages GPs. As nonparametric Bayesian models, GPs will estimate unknown RSS values and guide the acquisition of new measurement locations, as outlined next.

A. Gaussian Processes

We leverage GPs to probabilistically model $f(\mathbf{x})$ for all $\mathbf{x} \in \mathcal{X}$, and subsequently obtain a model over RSS values $\{y(\mathbf{x}), \forall \mathbf{x} \in \mathcal{X}\}$. To this aim, let us define a *single GP-based prior* distribution over $f(\mathbf{x}) \sim \mathcal{GP}(\mu(\mathbf{x}), \kappa(\mathbf{x}, \mathbf{x}_j))$, with mean $\mu(\mathbf{x})$ and a pre-specified kernel function $\kappa(\mathbf{x}, \mathbf{x}_j)$. Usually, for notational simplicity it is assumed that $\mu(\mathbf{x}) = 0, \forall \mathbf{x} \in \mathcal{X}$. This prior on $f(\mathbf{x})$ implies that for any finite number $n \in \mathbb{N}^+$ of samples, the random vector of function values $\mathbf{f}_n := [f(\mathbf{x}_1), \dots, f(\mathbf{x}_n)]^\top$ at locations $\mathbf{X}_n := [\mathbf{x}_1, \dots, \mathbf{x}_n]^\top$ has a *joint Gaussian* distribution; that is $p(\mathbf{f}_n | \mathbf{X}_n) = \mathcal{N}(\mathbf{f}_n; \mathbf{0}_n, \mathbf{K}_n)$, where \mathbf{K}_n represents the $n \times n$ covariance matrix with (i, j) -th entry $[\mathbf{K}_n]_{i,j} := \text{cov}(f(\mathbf{x}_i), f(\mathbf{x}_j)) = \kappa(\mathbf{x}_i, \mathbf{x}_j)$ [19].

The random vector \mathbf{f}_n is related to the (possibly) noisy observations $\mathbf{y}_n := [y_1, \dots, y_n]^\top$ with $y_i := y(\mathbf{x}_i)$, for $i = 1, \dots, n$, through the batch likelihood $p(\mathbf{y}_n | \mathbf{f}_n; \mathbf{X}_n)$ that is assumed to be factored as $p(\mathbf{y}_n | \mathbf{f}_n; \mathbf{X}_n) = \prod_{i=1}^n p(y_i | f(\mathbf{x}_i))$. Capitalizing on the GP prior and the batch likelihood, one can express the joint probability density function (PDF) of observation vector \mathbf{y}_n and $y(\mathbf{x})$ at any unmeasured location as

$$\begin{bmatrix} \mathbf{y}_n \\ y(\mathbf{x}) \end{bmatrix} \sim \mathcal{N} \left(\mathbf{0}, \begin{bmatrix} \mathbf{K}_n + \sigma_{\text{noise}}^2 \mathbf{I}_n & \mathbf{k}_n^\top(\mathbf{x}) \\ \mathbf{k}_n(\mathbf{x}) & \kappa(\mathbf{x}, \mathbf{x}) + \sigma_{\text{noise}}^2 \end{bmatrix} \right), \quad (3)$$

where $\mathbf{k}_n(\mathbf{x}) := [\kappa(\mathbf{x}_1, \mathbf{x}), \dots, \kappa(\mathbf{x}_n, \mathbf{x})]$ is the kernel vector containing similarity between location \mathbf{x} and all n observed locations $\{\mathbf{x}_i\}_{i=1}^n$. Leveraging this joint Gaussian distribution in (3), and *marginalizing* it yields the predictive PDF of $y(\mathbf{x})$ as [19]

$$p(y(\mathbf{x}) | \mathcal{D}_n, \mathbf{x}) = \mathcal{N}(y(\mathbf{x}); \bar{y}_{n+1|\mathbf{n}}(\mathbf{x}), \sigma_{n+1|\mathbf{n}}^2(\mathbf{x})), \quad (4)$$

where

$$\bar{y}_{n+1|\mathbf{n}}(\mathbf{x}) = \mathbf{k}_n^\top(\mathbf{x}) (\mathbf{K}_n + \sigma_{\text{noise}}^2 \mathbf{I}_n)^{-1} \mathbf{y}_n \quad (5a)$$

$$\begin{aligned} \sigma_{n+1|\mathbf{n}}^2(\mathbf{x}) &= \kappa(\mathbf{x}, \mathbf{x}) - \mathbf{k}_n^\top(\mathbf{x}) (\mathbf{K}_n + \sigma_{\text{noise}}^2 \mathbf{I}_n)^{-1} \mathbf{k}_n(\mathbf{x}) \\ &\quad + \sigma_{\text{noise}}^2, \end{aligned} \quad (5b)$$

with the notation $n+1|\mathbf{n}$ signifying that all n data samples are employed to obtain the predictive PDF at the next $(n+1)$ -th location. The mean in (5a) provides a point estimate of the RSS value of location \mathbf{x} and the variance in (5b) quantifies the associated uncertainty.

Although intuitive, the predictive PDF in (5) is subject to certain limitations. Primarily, the predictive mean and variance in equation (5) depend on a *single* pre-selected kernel function which restricts the expressive capacity of the learned function. In addition, the mean and variance in (5a) and (5b) entail $\mathcal{O}(n^2)$ storage requirements and $\mathcal{O}(n^3)$ computational complexity due to the matrix inversion operation involved, which

may become computationally intractable when n is large. In the subsequent sections, we will outline how to circumvent these challenges.

B. Ensemble of Gaussian Processes

Instead of relying on a single kernel to estimate $f(\mathbf{x})$, we target a more expressive function space by utilizing a set of M kernels, denoted by $\mathcal{K} := \{\kappa^1, \dots, \kappa^M\}$ to form a *mixture* of GP models. The set \mathcal{K} constitutes a diverse collection of kernels with different hyper-parameters. Each $\{\kappa^m\}_{m=1}^M \in \mathcal{K}$ induces a unique GP prior m over the function $f(\mathbf{x}) \sim \mathcal{GP}(0, \kappa^m(\mathbf{x}_i, \mathbf{x}_j)), \forall \mathbf{x}_i, \mathbf{x}_j \in \mathcal{X}$. Constructing a weighted combination of these GP models yields a mixture of GP priors over $f(\mathbf{x})$, as $f(\mathbf{x}) \sim \sum_{m=1}^M w_n^m \mathcal{GP}(0, \kappa(\mathbf{x}, \mathbf{x}_j)), \forall j$ subject to $\sum_{m=1}^M w_n^m = 1$. The per-model weight w_n^m signifies the contribution of m -th kernel in the GP prior over $f(\mathbf{x})$. Using this mixture of GP priors, the posterior distribution of $f(\mathbf{x})$ can be obtained by

$$p(f(\mathbf{x}) | \mathcal{D}_n) = \sum_{m=1}^M w_n^m p(f(\mathbf{x}) | \mathcal{D}_n, \kappa = \kappa^m), \quad (6)$$

with $w_n^m \propto \Pr(\kappa = \kappa^m) \Pr(\mathcal{D}_n | \kappa = \kappa^m)$, where $\Pr(\mathcal{D}_n | \kappa = \kappa^m)$ represents the marginal likelihood of data \mathcal{D}_n for the m -th GP. For any desired location $\mathbf{x} \in \mathcal{X}$, each GP induces a Gaussian predictive PDF $p(y(\mathbf{x}) | \mathcal{D}_n, m, \mathbf{x}) = \mathcal{N}(y(\mathbf{x}); \bar{y}_{n+1|\mathbf{n}}^m(\mathbf{x}), (\sigma_{n+1|\mathbf{n}}^m(\mathbf{x}))^2)$ to estimate the corresponding RSS value. By appropriately updating the mean and variance for each GP (c.f., (5a) and (5b)) and adjusting the weights $\{w_n^m\}_{m=1}^M$ for all GP models, one can obtain [21] the overall predictive PDF of $y(\mathbf{x})$

$$p(y(\mathbf{x}) | \mathcal{D}_n, \mathbf{x}) = \sum_{m=1}^M w_n^m \mathcal{N}(y(\mathbf{x}); \bar{y}_{n+1|\mathbf{n}}^m(\mathbf{x}), (\sigma_{n+1|\mathbf{n}}^m(\mathbf{x}))^2). \quad (7)$$

Although more expressive compared to (4), the predictive PDF in (7) requires large computational complexity $\mathcal{O}(Mn^3)$.

C. Random Feature-based EGPs

To bypass this cubic complexity, one can employ parametric approximants to the non-parametric models using the so-called random features (RF) [20]. The main idea of the RF-based approximation is to rely on a shift-invariant kernel $\bar{\kappa}(\mathbf{x}_i, \mathbf{x}_j) = \bar{\kappa}(\mathbf{x}_i - \mathbf{x}_j)$, normalized as $\bar{\kappa} = \kappa / \sigma_\theta^2$, and express it via the inverse Fourier transformation of its power spectral density $\pi_\kappa(\boldsymbol{\zeta})$ [20]; that is

$$\begin{aligned} \bar{\kappa}(\mathbf{x}_i - \mathbf{x}_j) &:= \mathbb{E}_{\pi_\kappa} \left[e^{j\boldsymbol{\zeta}^\top (\mathbf{x}_i - \mathbf{x}_j)} d\boldsymbol{\zeta} \right] \\ &= \int \pi_\kappa(\boldsymbol{\zeta}) e^{j\boldsymbol{\zeta}^\top (\mathbf{x}_i - \mathbf{x}_j)} d\boldsymbol{\zeta} \end{aligned} \quad (8)$$

where $\int \pi_\kappa(\boldsymbol{\zeta}) d\boldsymbol{\zeta} = 1$, so that it can be thought of as a PDF. Since the values of κ are always real, it holds that $\bar{\kappa}(\mathbf{x}_i - \mathbf{x}_j) := \mathbb{E}_{\pi_\kappa} \left[\cos(\boldsymbol{\zeta}^\top (\mathbf{x}_i - \mathbf{x}_j)) \right]$. Leveraging this new representation, one can approximate the kernel by drawing a

sufficiently large number of i.i.d samples $\{\zeta_d\}_{d=1}^D$ from $\pi_\kappa(\zeta)$ yielding

$$\hat{\kappa}(\mathbf{x}_i - \mathbf{x}_j) := \frac{1}{D} \sum_{d=1}^D \cos(\zeta_d^\top (\mathbf{x}_i - \mathbf{x}_j)) = \phi_\zeta(\mathbf{x}_i)^\top \phi_\zeta(\mathbf{x}_j), \quad (9)$$

where the random feature vector $\phi_\zeta(\cdot) \in \mathbb{R}^{2D \times 1}$ is defined as

$$\phi_\zeta(\mathbf{x}) := \frac{1}{\sqrt{D}} \left[\sin(\zeta_1^\top \mathbf{x}), \cos(\zeta_1^\top \mathbf{x}), \dots, \cos(\zeta_D^\top \mathbf{x}) \right]. \quad (10)$$

Relying on the RF vector $\phi_\zeta(\cdot)$, a *linear* and *parametric* approximant of the sought function can be obtained as [20]

$$\check{f}(\mathbf{x}) = \boldsymbol{\theta}^\top \phi_\zeta(\mathbf{x}), \quad \boldsymbol{\theta} \sim \mathcal{N}(\boldsymbol{\theta}; \mathbf{0}_{2D}, \sigma_\theta^2 \mathbf{I}_{2D}). \quad (11)$$

To elaborate more on (11), and the parameter vector $\boldsymbol{\theta}$, assuming a GP prior over f induced by a kernel function κ implies that the function f belongs to a reproducing kernel Hilbert space (RKHS) [19]; that is $\mathcal{H} := \{f | f(\mathbf{x}) = \sum_{i=1}^\infty \alpha_i \kappa(\mathbf{x}, \mathbf{x}_i)\}$ for some constants $\{\alpha_i\}$ to be determined from data. From the so-termed Representer theorem, the optimal function approximant for f can be written [42], [43]

$$\check{f}(\mathbf{x}) = \sum_{i=1}^n \alpha_i \kappa(\mathbf{x}, \mathbf{x}_i) \quad (12)$$

where $\{\alpha_i\}_{i=1}^n$ are to be learned, and κ is the pre-selected kernel function. Leveraging the RF vector one can have a kernel approximant expressed as $\kappa(\mathbf{x}, \mathbf{x}_i) \approx \phi_\zeta(\mathbf{x})^\top \phi_\zeta(\mathbf{x}_i)$, and upon replacing it in equation (12) yields

$$\check{f}(\mathbf{x}) = \sum_{i=1}^n \alpha_i \kappa(\mathbf{x}, \mathbf{x}_i) \approx \sum_{i=1}^n \alpha_i \phi_\zeta(\mathbf{x}_i)^\top \phi_\zeta(\mathbf{x}) = \underbrace{\left(\sum_{i=1}^n \alpha_i \phi_\zeta(\mathbf{x}_i) \right)^\top}_{\boldsymbol{\theta} :=} \phi_\zeta(\mathbf{x}) \quad (13)$$

The parameter vector $\boldsymbol{\theta}$ represents the weights of a *linear function* that approximates the desired f in RKHS. From a Bayesian perspective, all $\{\alpha_i\}_{i=1}^n$ are assumed to be random variables which are Gaussian distributed, and it follows that $\boldsymbol{\theta}$ is also Gaussian distributed as a sum of Gaussians, yielding $\boldsymbol{\theta} \sim \mathcal{N}(\boldsymbol{\theta}; \mathbf{0}_{2D}, \sigma_\theta^2 \mathbf{I}_{2D})$ (c.f., (11)).

This parametric model over $\check{f}(\mathbf{x})$ enables the propagation of the posterior $p(\boldsymbol{\theta} | \mathcal{D}_n) = \mathcal{N}(\boldsymbol{\theta}; \boldsymbol{\theta}_n, \boldsymbol{\Sigma}_n)$ using a recursive Bayesian iteration, which can be updated as new data arrive on-the-fly, as will be shown next.

D. Ensemble of Parametric GPs

To allow for reduced complexity and online model updates that are particularly appealing in the AL setting, we consider an ensemble of parametric GPs. Each GP model m , relies on a shift-invariant and normalized kernel $\bar{\kappa}^m = \kappa^m / \sigma_{\theta^m}^2$ to construct its RF vector $\phi_\zeta^m(\cdot)$ by drawing i.i.d random vectors $\{\zeta_d^m\}_{d=1}^D$ from its corresponding power spectral density

$\pi_{\bar{\kappa}^m}(\zeta)$. Similar to (11), each learner m forms a generative parametric model as

$$p(\boldsymbol{\theta}^m) = \mathcal{N}(\boldsymbol{\theta}^m; \mathbf{0}_{2D}, \sigma_{\theta^m}^2 \mathbf{I}_{2D}) \quad (14)$$

$$p(f(\mathbf{x}) | \kappa = \bar{\kappa}^m, \boldsymbol{\theta}^m) = \delta(f(\mathbf{x}) - \phi_\zeta^m(\mathbf{x})^\top \boldsymbol{\theta}^m) \quad (15)$$

$$p(y(\mathbf{x}) | \boldsymbol{\theta}^m, \mathbf{x}) = \mathcal{N}(y(\mathbf{x}); \phi_\zeta^m(\mathbf{x})^\top \boldsymbol{\theta}^m, \sigma_{\text{noise}}^2). \quad (16)$$

This generative model allows model m to form the posterior PDF $p(\boldsymbol{\theta}^m | \mathcal{D}_n) = \mathcal{N}(\boldsymbol{\theta}^m; \hat{\boldsymbol{\theta}}_n^m, \boldsymbol{\Sigma}_n^m)$ using available data \mathcal{D}_n , and the advocated parametric EGP model combines them with the weights $\{w_n^m\}_{m=1}^M$ to form the ensemble predictive PDF over the target $y(\mathbf{x})$, for all $\mathbf{x} \in \mathcal{X}$. This predictive PDF not only offers an estimate of the RSS value at location \mathbf{x} , but will further guide the design of a set of acquisition functions to find the next locations to obtain measurements, as will be discussed in the next subsection.

We now show how the EGP model parameters will be updated as a newly acquired measurement at a new location indexed by $n + 1$ is obtained. Each model m leverages the learned posterior $p(\boldsymbol{\theta}^m | \mathcal{D}_n)$ to find the predictive PDF of $y(\mathbf{x})$ at any desired target location \mathbf{x} as

$$\begin{aligned} p(y(\mathbf{x}) | \kappa = \bar{\kappa}^m, \mathcal{D}_n, \mathbf{x}) &= \int p(y(\mathbf{x}) | \boldsymbol{\theta}^m, \mathbf{x}) p(\boldsymbol{\theta}^m | \mathcal{D}_n) d\boldsymbol{\theta}^m \\ &= \mathcal{N}(y(\mathbf{x}); \hat{y}_{n+1|n}^m(\mathbf{x}), (\sigma_{n+1|n}^m(\mathbf{x}))^2), \end{aligned} \quad (17)$$

with mean and variance given by

$$\hat{y}_{n+1|n}^m(\mathbf{x}) = \phi_\zeta^m(\mathbf{x})^\top \hat{\boldsymbol{\theta}}_n^m \quad (18a)$$

$$(\sigma_{n+1|n}^m(\mathbf{x}))^2 = \phi_\zeta^m(\mathbf{x})^\top \boldsymbol{\Sigma}_n^m \phi_\zeta^m(\mathbf{x}) + \sigma_{\text{noise}}^2. \quad (18b)$$

To estimate the RSS value $y(\mathbf{x})$, the EGP model forms a Gaussian mixture (GM) model by combining all M predictive PDFs to form the ensemble predictive PDF

$$p(y(\mathbf{x}) | \mathcal{D}_n, \mathbf{x}) = \sum_{m=1}^M w_n^m \mathcal{N}(y(\mathbf{x}); \hat{y}_{n+1|n}^m(\mathbf{x}), (\sigma_{n+1|n}^m(\mathbf{x}))^2). \quad (19)$$

Having this Gaussian mixture model over $y(\mathbf{x})$, one can obtain an estimate of the RSS value along with the corresponding uncertainty, by considering the minimum mean-square error (MMSE) estimator of $y(\mathbf{x})$ along with the variance of the estimator; that is

$$\hat{y}_{n+1|n}(\mathbf{x}) = \sum_{m=1}^M w_n^m \hat{y}_{n+1|n}^m(\mathbf{x}) \quad (20a)$$

$$\begin{aligned} \sigma_{n+1|n}^2(\mathbf{x}) &= \sum_{m=1}^M w_n^m \left[(\sigma_{n+1|n}^m(\mathbf{x}))^2 \right. \\ &\quad \left. + (\hat{y}_{n+1|n}(\mathbf{x}) - \hat{y}_{n+1|n}^m(\mathbf{x}))^2 \right]. \end{aligned} \quad (20b)$$

where “ $n + 1|n$ ” denotes that only the measurement from location n is needed to predict y at the next location $n + 1$.

Leveraging the learned PDF in (19) along with the mean and variance in (20), one can find the next query location \mathbf{x}_{n+1} to obtain the RSS measurement y_{n+1} by minimizing the AFs as will be elaborated in the next subsection IV-E. Upon obtaining

\mathbf{x}_{n+1} and y_{n+1} , one can update the EGP model parameters. Specifically, each model m uses Bayes' rule along with (18) to update its weight w_n^m as

$$w_{n+1}^m = \frac{w_n^m \mathcal{N}(y_{n+1}; \hat{y}_{n+1|n}^m, (\sigma_{n+1|n}^m)^2)}{\sum_{m'=1}^M w_n^{m'} \mathcal{N}(y_{n+1}; \hat{y}_{n+1|n}^{m'}, (\sigma_{n+1|n}^{m'})^2)}, \quad (21)$$

and its posterior PDF as

$$\begin{aligned} p(\boldsymbol{\theta}^m | \mathcal{D}_{n+1}) &= \frac{p(\boldsymbol{\theta}^m | \mathcal{D}_n) p(y_{n+1} | \boldsymbol{\theta}^m, \mathbf{x}_{n+1})}{p(y_{n+1} | \boldsymbol{\theta}^m, \mathcal{D}_n, \mathbf{x}_{n+1})} \\ &= \mathcal{N}(\boldsymbol{\theta}^m; \hat{\boldsymbol{\theta}}_{n+1}^m, \boldsymbol{\Sigma}_{n+1}^m), \end{aligned} \quad (22)$$

where

$$\hat{\boldsymbol{\theta}}_{n+1}^m = \hat{\boldsymbol{\theta}}_n^m + \frac{\boldsymbol{\Sigma}_n^m \phi_{\boldsymbol{\zeta}}^m(\mathbf{x}_{n+1})(y_{n+1} - \hat{y}_{n+1|n}^m)}{(\sigma_{n+1|n}^m)^2} \quad (23a)$$

$$\boldsymbol{\Sigma}_{n+1}^m = \boldsymbol{\Sigma}_n^m - \frac{\boldsymbol{\Sigma}_n^m \phi_{\boldsymbol{\zeta}}^m(\mathbf{x}_{n+1}) \phi_{\boldsymbol{\zeta}}^m(\mathbf{x}_{n+1})^\top \boldsymbol{\Sigma}_n^m}{(\sigma_{n+1|n}^m)^2}. \quad (23b)$$

The required complexity of updating the model parameters in (23) is $\mathcal{O}(M((2D)^2 + 2Dn))$, which is smaller than the $\mathcal{O}(Mn^3)$ complexity of the original EGP-based model.

E. Acquisition functions for EGP-based active learning.

This section introduces a number of EGP-based AFs, which depend on the learned predictive PDF $p(y(\mathbf{x}) | \mathcal{D}_n)$ in (19), to guide the acquisition process; that is to find the next most informative location within the coverage area, i.e., $\mathbf{x}_{n+1} \in \mathcal{X}$, and obtain its associated RSS measurement y_{n+1} .

1) *Weighted variance*: To exploit the uncertainty of each GP model, we leverage the variance of all GP models' posterior PDFs and combine them to form the weighted variance AF as

$$\alpha^{\text{wVar}}(\mathbf{x}; \mathcal{D}_n) := \sum_{m=1}^M w_n^m (\sigma_{n+1|n}^m(\mathbf{x}))^2. \quad (24)$$

The intuition is that locations with high uncertainty can be informative for the training dataset. Although $\alpha^{\text{wVar}}(\cdot)$ well captures the uncertainty of each GP model, it ignores the information offered by the predicted means (c.f. (18a)), motivating the next AF.

2) *Query-by-committee*: Considering M GP models as members of a committee, the Query-by-committee (QBC) AF chooses the next location to be measured where the committee members exhibit the most disagreement; that is

$$\alpha^{\text{QBC}}(\mathbf{x}; \mathcal{D}_n) := \sum_{m=1}^M w_n^m \left(\hat{y}_{n+1|n}^m(\mathbf{x}) - \bar{y}_n(\mathbf{x}) \right), \quad (25)$$

where $\hat{y}_{n+1|n}^m(\mathbf{x})$ represents the estimated mean of RSS values at location \mathbf{x} by learner m (c.f., (18a)), and $\bar{y}_n(\mathbf{x}) := \sum_{m=1}^M w_n^m \hat{y}_{n+1|n}^m(\mathbf{x})$ denotes the weighted mean of the committee. Note that only the predictive means are used in $\alpha^{\text{QBC}}(\cdot)$ and the predictive variances are not considered.

3) *Variance of GP mixture*: To simultaneously account for the *means* and *variances* associated with each GP model's predictive PDF, the variance of GP mixture AF combines the the per-learner uncertainty in (24) and committee disagreement in ((25) (c.f., (20b)) as

$$\alpha^{\text{GPM-var}}(\mathbf{x}; \mathcal{D}_n) := \sigma_{n+1|n}^2(\mathbf{x}). \quad (26)$$

4) *Weighted entropy*: Relying on the entropy as an alternative measure of the associated uncertainty, one can form an AF using a weighted combination of the entropies associated with the predictive PDFs of all learners as

$$\alpha^{\text{wEnt}}(\mathbf{x}; \mathcal{D}_n) := \sum_{m=1}^M w_n^m \log((\sigma_{n+1|n}^m(\mathbf{x}))^2). \quad (27)$$

5) *Entropy of GP mixtures*: Instead of employing the entropies of individual GP models independently as in (27), we can leverage the entropy of the learned GP-mixture. While an analytical closed-form expression for this entropy does not exist, we leverage a closed-form lower bound; see e.g [38]. To this aim, let us first define the entropy $H(\hat{f}(\mathbf{x}) | \mathcal{D}_n)$ of the learned function $\hat{f}(\mathbf{x})$ at any location \mathbf{x}

$$\begin{aligned} H(\hat{f}(\mathbf{x}) | \mathcal{D}_n) &:= - \int_{-\infty}^{\infty} \sum_{m=1}^M w_n^m \mathcal{N}(\hat{f}(\mathbf{x}) | \hat{y}_{n+1|n}^m(\mathbf{x}), (\sigma_{n+1|n}^m(\mathbf{x}))^2) \\ &\quad \times \log(p(\hat{f}(\mathbf{x}) | \mathcal{D}_n)) d\hat{f}(\mathbf{x}) \end{aligned} \quad (28)$$

where the integration here is taken over all the *random function values* that $\hat{f}(\mathbf{x})$ can take at any desired, but fixed $\mathbf{x} \in \mathcal{X}$. Using Jensen's inequality yields a lower bound for the entropy H at any location $\mathbf{x} \in \mathcal{X}$ as follows

$$\begin{aligned} H(\hat{f}(\mathbf{x}) | \mathcal{D}_n) &\geq \\ &- \sum_{m=1}^M w_n^m \log \left(\int_{-\infty}^{\infty} \mathcal{N}(\hat{f}(\mathbf{x}) | \hat{y}_{n+1|n}^m(\mathbf{x}), (\sigma_{n+1|n}^m(\mathbf{x}))^2) \right. \\ &\quad \left. \times p(\hat{f}(\mathbf{x}) | \mathcal{D}_n) d\hat{f}(\mathbf{x}) \right). \end{aligned} \quad (29)$$

The term inside the logarithm in (29) admits a closed-form expression and hence the lower-bound can be expressed analytically as

$$\alpha^{\text{GPM-Ent}}(\mathbf{x}; \mathcal{D}_n) := \sum_{m=1}^M \log \left(\sum_{m'=1}^M w_n^{m'} \psi_n^{m, m'} \right), \quad (30)$$

where $\psi_n^{m, m'}$ captures the intra-GP interactions and is given by

$$\begin{aligned} \psi_n^{m, m'} &:= \int_{-\infty}^{\infty} \mathcal{N}(\hat{f}(\mathbf{x}) | \hat{y}_{n+1|n}^m(\mathbf{x}), (\sigma_{n+1|n}^m(\mathbf{x}))^2) \\ &\quad \times \mathcal{N}(\hat{f}(\mathbf{x}) | \hat{y}_{n+1|n}^{m'}(\mathbf{x}), (\hat{\sigma}_{n+1|n}^{m'}(\mathbf{x}))^2) d\hat{f}(\mathbf{x}) \\ &= \mathcal{N}(\hat{y}_{n+1|n}^m; \hat{y}_{n+1|n}^{m'}, (\sigma_{n+1|n}^m)^2 + (\sigma_{n+1|n}^{m'})^2). \end{aligned} \quad (31)$$

While these AFs cover a diverse range of important ideas within AL, they impose no constraints on the spatial locations of the next query point, which can lead to prohibitively long travelling distances required to collect the necessary measurements dictated by the AFs. To address this limitation, we will couple these EGP-based AFs with a minimum distance-based criterion, described in the next subsection.




Fig. 2: Radio coverage maps of diverse scenarios rendered with Blender (top) and Sionna ray-tracing simulator at 3.7GHz (bottom left) and 26GHz (bottom right).

TABLE II: Simulation configurations.

	Parameters	Descriptions		
Data	Scenario	OpenSpace	Reflection	UrbanCanyon
	3.7GHz Samples	780.5K	841.7K	47.9K
	26GHz Samples	779.8K	744.5K	41.3K
Antenna Configs	Rx Pattern	Isotropic		
	Tx Pattern	TR 38.901		
	Tx Power	24.0 dB		
	Ray Interactions	9		
	Clipping Range	(0.0, 160.0] dB		
	Cell Size	5m x 5m		

F. Path efficient EGP-based AFs

To allow for efficient and cost-effective data acquisition strategies, thereby enhancing the practical feasibility and economic viability of acquiring measurements and reconstructing 5G radio maps, we consider AFs with the following form

$$\alpha(\mathbf{x}, \mathcal{D}_n) := \begin{cases} \alpha^{\text{EGP}}(\mathbf{x}, \mathcal{D}_n) & \text{w.p. } \varepsilon \\ \alpha^{\text{EGP}}(\mathbf{x}, \mathcal{D}_n) c(\mathbf{x}, \mathbf{x}_{n-1}) & \text{w.p. } 1 - \varepsilon \end{cases} \quad (32)$$

where $c(\mathbf{x}, \mathbf{x}_{n-1}) := \frac{1}{\lambda + \text{dist}(\mathbf{x}, \mathbf{x}_{n-1})}$ is a distance-penalizing function with λ being a hyperparameter to be tuned, $\alpha^{\text{EGP}}(\cdot)$ is one of the above EGP-based AFs, and $\varepsilon > 0$ is a small hyperparameter.³ With a small probability $0 < \varepsilon < 1$ our novel AF in (32) allows exploring \mathcal{X} while adhering to the EGP-based AF $\alpha^{\text{EGP}}(\mathbf{x}, \mathcal{D}_n)$. With a much larger probability $1 - \varepsilon$ our AF is strongly encouraged to select points *nearby* the most recently obtained measurement \mathbf{x}_n by penalizing far distant points via the path-cost function $c(\mathbf{x}, \mathbf{x}_{n-1})$. This allows the AF to explicitly acknowledge that obtaining measurements from diverse locations in 5G scenarios can be costly and resource-intensive, while still allowing a user-controlled degree of exploration through the parameter ε . This is crucial in several AL applications including the 5G radiomap reconstruction that this paper focuses on.

V. DATASET GENERATION

We now introduce the dataset utilized for evaluation, which comprises simulated and real measurements. The simulated




Fig. 3: Radio coverage maps for real RSS measurements. Green color represents strong RSS, while red poor RSS.

dataset replicates three distinct scenes: open space, reflection & blockage, and urban canyon. Meanwhile, the real measurements consist of two scenes: suburban and urban canyons.

Simulated Dataset. Following [44], we employ Blender [45] and Sionna [46] to produce radio maps and use them as ground truth for the evaluation. Blender is an open-source software that can create different 3D environments. We utilize its Open Street Map (OSM) and Mitsuba plugins to load the real-world map and render the digital scenes. Subsequently, these scenes are imported into Sionna, an open-source link-level ray-tracing simulator that is compatible with the 3GPP channel models, to generate radio coverage maps that we sample from.

The upper panel of Fig. 2 depicts the loaded digital scenes, illustrating three distinct scenarios: (1) a large open space with the line-of-sight propagation of radio signals; (2) a more complex wall scenario featuring multiple walls to demonstrate reflection, diffraction, and signal blockage; and (3) a highly complex downtown urban canyon, where the signals are reflected and obstructed by the surrounding buildings. The lower panel of Fig. 2 shows the simulated radio coverage maps. We also investigate the propagation characteristics of different frequency radio waves across 3.7GHz (mid-band) and 26GHz (high-band) in light of commercial 5G networks. For each scenario-frequency pair, we place the transmitter (Tx) at the center location with a directional antenna oriented at the negative y -axis. The map is divided into sets of cells, each containing a receiver (Rx) with an omnidirectional antenna, generating a total of N unique samples of the form [Cartesian Coordinate- x , Cartesian Coordinate- y , RSS]. To determine the RSS value, rays are fired from a Tx at the origin

³The abbreviation w.p. stands for *with probability*.




Fig. 4: NMSE vs iteration for 3.7GHz mid-band data under (a) open space; (b) reflect scenarios; and (c) urban canyon.




Fig. 5: NMSE vs iteration for 26GHz high-band data under (a) open space; (b) reflect scenarios; and (c) urban canyon.




Fig. 6: NMSE vs iteration for the real (a) urban and (b) suburban datasets.

and are allowed to interact with the environment, i.e., reflect, diffract, etc., up to a bounded number of times. The RSS is aggregated from per intersecting rays [46], clipped to a range. For specific configuration parameters of our dataset, we refer the reader to Table II ⁴.

Real Measurement Dataset. In addition to simulated datasets, we have gathered real RSS radio measurements from urban and suburban environments. Specifically, we used a smartphone as an Rx, which was connected to a commercial 5G network by running the “iperf3” application to maintain active connection. The phone is also connected to a laptop equipped with the professional software XCAL [41] to read RSS measurements from the phone’s chipset-level diagnostics. This enabled us to collect a total of 27k data samples from an urban area of size 0.8 km^2 , and 18k data samples from a suburban area of size 1.0 km^2 . Similar to the simulated data, each real data sample is rep-

resented by [Cartesian Coordinate-x, Cartesian Coordinate-y, RSS]. The visualization of the collected real radio RSS map is depicted in Fig. 3.

VI. EVALUATION

In this section, we evaluate the performance of our EGP-based AL methods on the simulated 5G mid-band, high-band and real urban and suburban datasets described above. We will first quantify the advantages of the advocated EGP-based AL approaches over random sampling with EGPs and then highlight the benefits of our novel path-aware method. We will also conduct an in-depth study with visualized results to demonstrate the efficacy of those AL algorithms in the 5G measurement problem.

A. Experiment Settings

We denote the advocated EGP-based AL methods that employ the AFs in (24)-(30) as ‘EGP-WVar’, ‘EGP-QBC’, ‘EGP-GM-Var’, ‘EGP-Went’, and ‘EGP-GM-Ent’ respectively. We

⁴The dataset is publicly available at: <https://github.com/StrongWeiUMN/Active-Learning-for-5G-Measurement>.

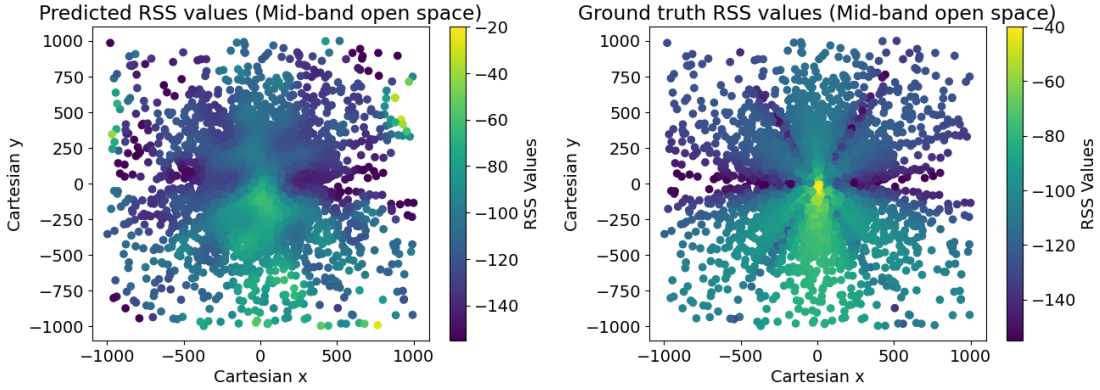


Fig. 7: Left: Predicted RSS values of the EGP-QBC method for 5000 test instances on the Mid-band open space dataset; Right: Ground truth RSS values.

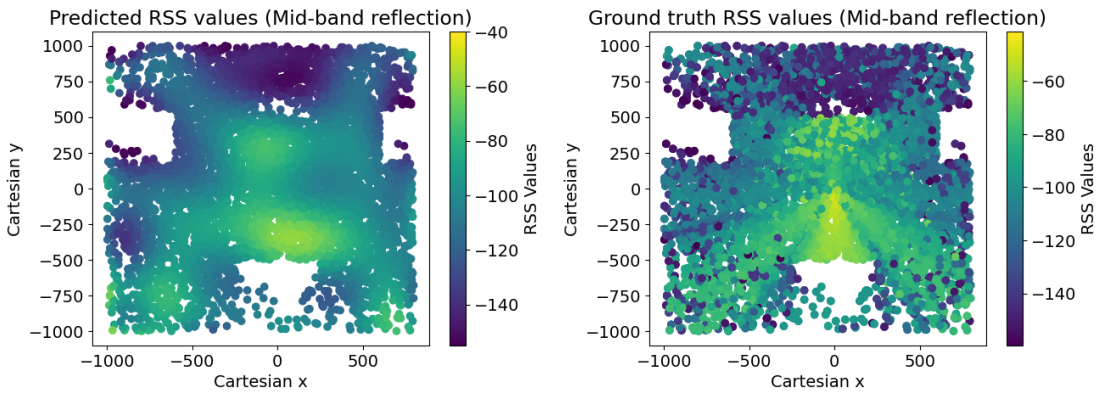


Fig. 8: Left: Predicted RSS values of the EGP-GM-Var method for 5000 test instances on the Mid-band reflection dataset; Right: Ground truth RSS values.

TABLE III: NMSE performance for Mid-band (a) open space; (b) reflection; and (c) urban datasets.

Method	Mid-band open space	Mid-band reflection	Mid-band urban
EGP-GM-Var	0.4904 ± 0.0117	0.5041 ± 0.0052	0.3417 ± 0.0055
EGP-Wvar	0.4811 ± 0.0183	0.5087 ± 0.0010	0.3459 ± 0.0076
EGP-Went	0.4775 ± 0.0208	0.5115 ± 0.0032	0.3457 ± 0.0130
EGP-GM-Ent	0.4901 ± 0.0091	0.5083 ± 0.0033	0.3466 ± 0.0064
EGP-QBC	0.4079 ± 0.0271	0.5123 ± 0.0142	0.3442 ± 0.0107
Ordinary Kriging	1.0052 ± 0.0035	0.4558 ± 0.0037	1.0075 ± 0.0042
Universal Kriging	0.8255 ± 0.0087	0.4614 ± 0.0052	1.0248 ± 0.0089
Residual Kriging	0.7695 ± 0.1161	0.7287 ± 0.0616	1.0246 ± 0.0081

also implement the ‘EGP-Random’ baseline for comparison, which relies on the EGP model as well but randomly selects new locations to query at each time slot. For all competing approaches, 100 initial labeled (measured) data samples are considered for training, and at each iteration of the AL process, one sample is queried from an unlabeled set consisting of 700 available unmeasured locations. Their performance is evaluated on a test set $\mathcal{T} := \{(\mathbf{x}_{n'}^{\text{test}}, y_{n'}^{\text{test}})\}_{n'=1}^{N^{\text{test}}}$ consisting of $N^{\text{test}} = 5000$ test locations. As a figure of merit the normalized mean squared error (NMSE) is used, which is expressed as $\text{NMSE}_n := \frac{1}{N^{\text{test}}} \sum_{n'=1}^{N^{\text{test}}} (\hat{y}_{n'|n}^{\text{test}} - y_{n'}^{\text{test}})^2 / \sigma_y^2$, where $\hat{y}_{n'|n}^{\text{test}}$ represents the point estimate of the RSS value of test location n' upon having processed location n , and $\sigma_y^2 := \mathbb{E}\|\mathbf{y}_{N^{\text{test}}}^{\text{test}} - \mathbb{E}\{\mathbf{y}_{N^{\text{test}}}^{\text{test}}\}\|^2$, where $\mathbf{y}_{N^{\text{test}}}^{\text{test}} := [y_1^{\text{test}}, \dots, y_{N^{\text{test}}}^{\text{test}}]^T$.

For each method, the EGP model consists of 11 GP learners,

each capitalizing on a distinct radial basis function (RBF) kernel with characteristic length scale chosen from the set $\{10^c\}_{c=-4}^6$. The kernel hyperparameters of each GP learner are obtained by maximizing the marginal log-likelihood utilizing the *sklearn* Python package. To allow for scalability and online model updates, $D = 50$ RFs are employed to yield the RF-based parametric and approximate GP models in the ensemble.

B. Performance of AL in the 5G Measurement Problem

Effectiveness of EGP-based AFs. To demonstrate the significance of the EGP-based AFs, we first compare them with the ‘EGP-Random’ baseline. To this end, the path-cost function $c(\mathbf{x}, \mathbf{x}_{n-1})$ is not taken into account; that is $\varepsilon = 1$ in (32).

Figs. 4, 5 and 6 illustrate the average NMSE performance of all competing approaches over 10 independent runs along

TABLE IV: NMSE performance for the High-band (a) open space; (b) reflection; and (c) urban datasets.

Method	High-band open space	High-band reflection	High-band urban
EGP-GM-Var	0.3810 ± 0.0062	0.6638 ± 0.0112	0.3831 ± 0.0037
EGP-Wvar	0.3803 ± 0.0079	0.6663 ± 0.0114	0.3849 ± 0.0050
EGP-Went	0.3870 ± 0.0074	0.6677 ± 0.0140	0.3915 ± 0.0136
EGP-GM-Ent	0.3841 ± 0.0033	0.6617 ± 0.0092	0.3802 ± 0.0048
EGP-QBC	0.3833 ± 0.0070	0.6593 ± 0.0169	0.4023 ± 0.0123
Ordinary Kriging	0.4356 ± 0.0054	0.6626 ± 0.0061	1.0068 ± 0.0066
Universal Kriging	0.4352 ± 0.0051	0.6641 ± 0.0062	0.9949 ± 0.0027
Residual Kriging	0.5019 ± 0.0032	0.7114 ± 0.0062	0.8952 ± 0.1392




Fig. 9: NMSE vs cumulative distance for 3.7GHz mid-band data under (a) open space; (b) urban canyon; and (c) reflect datasets.




Fig. 10: NMSE vs cumulative distance 26GHz high-band data under (a) high-band open space; (b) high-band urban canyon; and (c) high-band reflect datasets.

TABLE V: NMSE performance for the Urban; and Suburban datasets.

Method	Urban	Suburban
EGP-GM-Var	0.3831 ± 0.0037	0.1701 ± 0.0074
EGP-Wvar	0.3849 ± 0.0050	0.1691 ± 0.0144
EGP-Went	0.3915 ± 0.0136	0.1623 ± 0.0127
EGP-GM-Ent	0.3802 ± 0.0048	0.1646 ± 0.0081
EGP-QBC	0.4023 ± 0.012	0.1931 ± 0.0100
Ordinary Kriging	0.7136 ± 0.0746	0.3351 ± 0.0232
Universal Kriging	0.7131 ± 0.0767	0.3323 ± 0.0222
Residual Kriging	0.7259 ± 0.0572	0.5316 ± 0.0523

with the corresponding highlighted standard deviation, on the mid-band, high-band and real urban and suburban datasets, respectively, for 70 iterations of the AL process; that is 70 queried locations that are measured. It is evident that in all

mid-band, in almost all high-band and in all real datasets, the advocated EGP-based AL methods markedly outperform the ‘EGP-Random’ baseline that does not rely on AL, achieving up to 54.8% improvement. This showcases the significance of judiciously selecting the few, but most informative, locations to measure so as to accurately predict the RSS values in unknown locations. For visualization purposes, Figs. 7 and 8 illustrate the predicted RSS values for 5000 test instances compared to the ground truth ones for the Mid-band open space and Mid-band reflection datasets respectively.

In addition, we compare our EGP-based AL methods with three additional baselines, namely, (a) Ordinary Kriging [47]; (b) Universal Kriging [48]; and, (c) Residual Kriging [49]. In Tables III, IV, V we report the prediction (test) performance of all competing methods over 5000 testing data samples. Training has been carried out using in total 170 data samples




Fig. 11: NMSE vs cumulative distance for the real (a) urban and (b) suburban datasets.




Fig. 12: Trajectory of the queried locations from the EGP-GM-Ent method for (a) $\varepsilon = 1$; (b) $\varepsilon = 0.3$; and (c) $\varepsilon = 0$ in the mid-band open space dataset. The background corresponds to the radio coverage information in an open space environment as shown in Fig. 2a.

(the same 100 initial samples plus 70 newly acquired samples). Table III reports the results of our methods compared with these three baselines in the three Mid-band cases namely (i) open space; (ii) reflection; and, (iii) urban setting. Clearly, our proposed methods enjoy a substantially lower prediction error compared to the newly considered baselines in open and urban spaces, while they have similar performance on the reflection scenario. Similar results are reported in Table IV, where we use a high-frequency setting, where again our advocated methods outperform the considered baselines in the open-space and urban settings, while having similar performance in the reflection scenario. Finally, Table V reports the prediction performance on the real urban and suburban datasets where it is evident that the EGP-based approaches markedly outperform the kriging-based baselines.

Although effective, these AFs by themselves do not impose constraints on the traveling distance needed to collect the necessary measurements, which may limit their application in settings with restricted travel budgets.

Path-aware AF. We now show the importance of adopting the novel path-cost related AF in (32) that allows for accurate RSS estimates with less required traveling distance. Regarding the exploration-exploitation parameter ε in (32), we set $\varepsilon = 0.3$, meaning that our method selects a location to query not far away with probability 0.7 or selects any location without any path constraint with probability 0.3, allowing for further exploration.

In Figs. 9, 10 and 11, the average NMSE performance along with the corresponding standard deviation are reported, with

respect to the cumulative traveling distance for 70 iterations of the AL process. It can be clearly seen that in both simulated mid-band and high-band, and real urban and suburban datasets, all EGP-AL methods that use (32) with $\varepsilon = 0.3$ have comparable or even lower NMSE compared to the EGP-AL counterparts that do not consider any path constraint ($\varepsilon = 1$), while at the same time requiring significantly less traveling distance (up to 57.1% reduction in terms of required cumulative traveling distance).

C. In-depth Results Study

Sampling Trajectory. To further demonstrate the efficiency of the advocated method, Figs. 12 and 13 depict the trajectories of the queried locations obtained by the ‘EGP-GM-Ent’ and ‘EGP-GM-Var’ AFs for the mid-band and high-band open space datasets, respectively, for $\varepsilon \in \{0, 0.3, 1\}$ ⁵. As expected, in both datasets the corresponding required traveling distance significantly reduces as the value of ε decreases. Interestingly, in the mid-band open space dataset, the EGP-GM-Ent approaches with $\varepsilon = 1$ and $\varepsilon = 0.3$ query very similar measurement locations; however, the latter entails much less traveling cost, as shown in Fig. 12.

Hyper-Parameter Sensitivity. Fig. 14a illustrates the NMSE performance of the ‘EGP-GM-Ent’ methods with $\varepsilon = 1, 0.3, 0$ respectively for a fixed cumulative distance budget, where it can be seen that the latter two can query more locations and

⁵Trajectories of other AL methods and datasets are omitted due to space limitations.

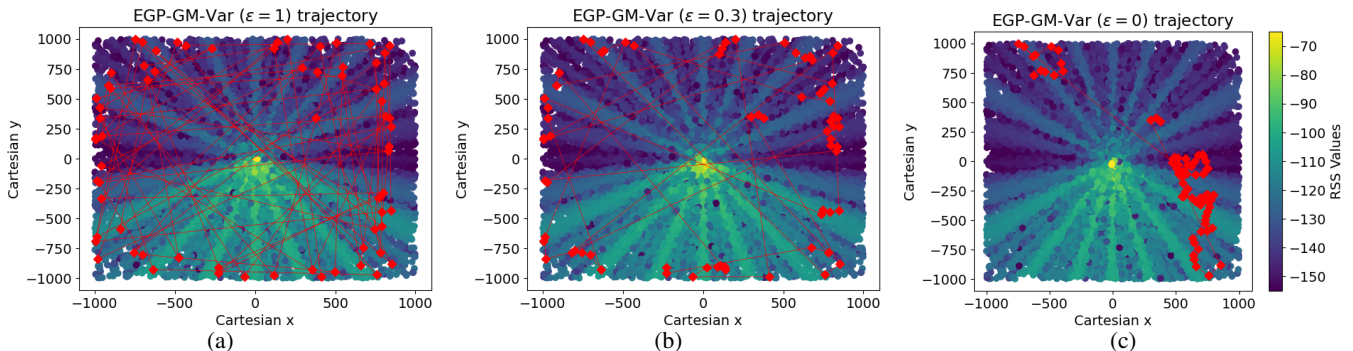


Fig. 13: Trajectory of the queried locations from the EGP-GM-Var method for (a) $\varepsilon = 1$; (b) $\varepsilon = 0.3$; and (c) $\varepsilon = 0$ in the high-band open space dataset. The background corresponds to the radio coverage information in an open space environment as shown in Fig. 2a.




Fig. 14: NMSE vs (a) cumulative distance of the EGP-GM-Ent (b) method for different values of ε for (a) mid-band open space dataset; (b) high-band open space dataset; and (c) uncertainty quantification performance of EGP-GM-Var on the High-band open space dataset (RSS values are in dB).

thus achieve superior NMSE performance compared to the $\varepsilon = 1$ case. Similarly, Fig. 14b depicts the NMSE performance of the EGP-GM-Var approach for different values of ε and a given distance budget for the high-band OpenSpace dataset. It is evident that the case of $\varepsilon = 0.3$ enjoys substantially lower NMSE compared to the $\varepsilon = 0$ and $\varepsilon = 1$ cases whereas the AL methods with $\varepsilon = 0$ and $\varepsilon = 1$ have similar performance. This emanates from the fact that although the approach with $\varepsilon = 0$ queries many more locations compared to $\varepsilon = 1$, for a distance budget of 5.2km, these locations are only from two clusters and are not from different regions of the radio map as shown in Fig. 13c. This showcases the importance of small but non-zero values of ε that do not so *heavily* penalize distant locations but instead allow for further exploration that yields improved prediction performance without the need for large traveling costs. Finally, to evaluate the well-quantifiable uncertainty offered by the advocated EGP-AL methods, Fig. 14c shows the predicted RSS values of the EGP-GM-Var ($\varepsilon = 0.3$) method on some test locations of the high-band OpenSpace dataset along with σ -confidence intervals, where it is intuitive that the ground truth RSS values fall inside these intervals.

VII. CONCLUSIONS

This work investigated novel Bayesian AL approaches for sample- and path-efficient 5G radio map reconstruction. The advocated AL approaches judiciously select only a few rep-

resentative locations to collect RSS measurements so that to efficiently and accurately estimate the RSS values at unmeasured locations. To effectively estimate the sought function that predicts RSS values across the coverage area, an ensemble of GPs was utilized. Inspired by the notions of (i) multi-kernel learning; and, (ii) RF approximation, the advocated EGP model offered a more expressive function space compared to a single GP, while enjoying scalable online updates, which are particularly appealing in AL settings where data are processed online. The proposed EGP model not only provided accurate RSS predictions with quantifiable uncertainty, but also offered a suite of pertinent AFs to guide the acquisition of new locations to measure. To further accommodate real-world 5G measurement campaigns, where traveling distances to collect RSS measurements might be limited, these AFs were coupled with a novel path-efficient rule that penalized queries at long distances. Numerical tests on three simulated scenarios with two different radio frequencies and on two real datasets, showcased the significance of the proposed AL methods in the 5G radio map reconstruction problem.

Future research will develop algorithms that account for various practical constraints, such as dynamic and heterogeneous environmental conditions, network traffic, among others, and tailoring EGPs for applications beyond radio map reconstruction.

REFERENCES

- [1] M. I. Rochman, W. Ye, Z.-L. Zhang, and M. Ghosh, "A comprehensive real-world evaluation of 5G improvements over 4G in low-and mid-bands," *arXiv:2312.00957*, 2023.
- [2] W. Ye, X. Hu, S. Sleder, A. Zhang, U. K. Dayalan, A. Hassan, R. A. Fezeu, A. Jajoo, M. Lee, E. Ramadan *et al.*, "Dissecting carrier aggregation in 5G networks: Measurement, QoE implications and prediction," in *Proc. of ACM SIGCOMM*, 2024, pp. 340–357.
- [3] R. A. Fezeu, E. Ramadan, W. Ye, B. Minneci, J. Xie, A. Narayanan, A. Hassan, F. Qian, Z.-L. Zhang, J. Chandrashekar *et al.*, "An in-depth measurement analysis of 5G mmWave PHY latency and its impact on end-to-end delay," in *Intl. Conf. on Passive and Active Netw. Meas.* Springer, 2023, pp. 284–312.
- [4] K. Shaw, "What are the benefits of 5G mmWave?" <https://www.verizon.com/business/resources/articles/what-are-the-benefits-of-5g-mmwave/>, 2023, accessed: 2023-07-29.
- [5] A. Narayanan, E. Ramadan, R. Mehta, X. Hu, Q. Liu, R. A. Fezeu, U. K. Dayalan, S. Verma, P. Ji, T. Li *et al.*, "Lumos5G: Mapping and predicting commercial mmWave 5G throughput," in *Proc. of ACM Intl. Meas. Conf.*, New York, NY, October 27 - 29, 2020, pp. 176–193.
- [6] W. Ye, X. Hu, T. Liu, R. Sun, Y. Li, and Z.-L. Zhang, "5GNN: extrapolating 5G measurements through GNNs," in *Proc. of the 1st Intl. Workshop on Graph Neural Networking*, ser. GNNet '22. New York, NY, USA: Association for Comput. Machinery, 2022, p. 36–41.
- [7] D. Lee and G. B. Giannakis, "A variational bayes approach to adaptive radio tomography," *IEEE Trans. Signal Process.*, vol. 68, pp. 3779–3792, 2020.
- [8] D. Lee, D. Berberidis, and G. B. Giannakis, "Adaptive bayesian radio tomography," *IEEE Trans. Signal Process.*, vol. 67, no. 8, pp. 1964–1977, 2019.
- [9] D. Romero, D. Lee, and G. B. Giannakis, "Blind radio tomography," *IEEE Trans. Signal Process.*, vol. 66, no. 8, pp. 2055–2069, 2018.
- [10] Y. Zhang and S. Wang, "K-nearest neighbors Gaussian process regression for urban radio map reconstruction," *IEEE Commun. Lett.*, vol. 26, no. 12, pp. 3049–3053, 2022.
- [11] M. Hamid and B. Beferull-Lozano, "Non-parametric spectrum cartography using adaptive radial basis functions," in *Intl. Conf. on Acoust., Speech, and Signal Proc.*, New Orleans, LA, March 5-9 2017, pp. 3599–3603.
- [12] B. Khalfi, B. Hamdaoui, and M. Guizani, "Airmap: scalable spectrum occupancy recovery using local low-rank matrix approximation," in *IEEE Global Commun. Conf.*, Abu Dhabi, UAE, Dec. 9-13, 2018, pp. 206–212.
- [13] F. Shen, Z. Wang, G. Ding, K. Li, and Q. Wu, "3D compressed spectrum mapping with sampling locations optimization in spectrum-heterogeneous environment," *IEEE Trans. Wireless Commun.*, vol. 21, no. 1, pp. 326–338, 2021.
- [14] G. Zhang, X. Fu, J. Wang, X.-L. Zhao, and M. Hong, "Spectrum cartography via coupled block-term tensor decomposition," *IEEE Trans. Signal Process.*, vol. 68, pp. 3660–3675, 2020.
- [15] G. Chen, Y. Liu, T. Zhang, J. Zhang, X. Guo, and J. Yang, "A graph neural network based radio map construction method for urban environment," *IEEE Commun. Lett.*, vol. 27, no. 5, pp. 1327–1331, 2023.
- [16] J.-H. Seong and D.-H. Seo, "Selective unsupervised learning-based Wi-Fi fingerprint system using autoencoder and GAN," *IEEE Internet Things J.*, vol. 7, no. 3, pp. 1898–1909, 2019.
- [17] Q. Niu, Y. Nie, S. He, N. Liu, and X. Luo, "RecNet: a convolutional network for efficient radiomap reconstruction," in *IEEE Intl. Conf. on Commun.*, Kansas City, MO, May 20-24, 2018, pp. 1–7.
- [18] H. Zou, C.-L. Chen, M. Li, J. Yang, Y. Zhou, L. Xie, and C. J. Spanos, "Adversarial learning-enabled automatic WiFi indoor radio map construction and adaptation with mobile robot," *IEEE Internet of Things J.*, vol. 7, no. 8, pp. 6946–6954, 2020.
- [19] C. E. Rasmussen and C. K. Williams, *Gaussian processes for machine learning*. MIT press Cambridge, MA, 2006.
- [20] A. Rahimi and B. Recht, "Random features for large-scale kernel machines," *Proc. of Neural Info. Process.*, pp. 1177–1184, Vancouver, B.C., Dec. 8-11, 2008.
- [21] Q. Lu, G. Karanikolas, Y. Shen, and G. B. Giannakis, "Ensemble gaussian processes with spectral features for online interactive learning with scalability," in *Proc. of AISTATS*, Palermo, Italy, June 3-5, 2020, pp. 1910–1920.
- [22] D. Romero, D. Lee, and G. B. Giannakis, "Blind radio tomography," *IEEE Trans. Signal Process.*, vol. 66, no. 8, pp. 2055–2069, 2018.
- [23] W. Liu and J. Chen, "Geography-aware radio map reconstruction for UAV-aided communications and localization," in *IEEE Intl. Conf. Communications*, 2021, pp. 1–6.
- [24] A. A. Bipon, M. S. Islam, A. T. Asyhari, A. Aneiba, and R. Abozariba, "Radio environment maps through spatial interpolation: A web-based approach," in *IEEE Netw. Oper. and Manag. Symp.*, 2024, pp. 1–7.
- [25] V.-P. Chowdappa, C. Botella, J. J. Samper-Zapater, and R. J. Martinez, "Distributed radio map reconstruction for 5G automotive," *IEEE Intel. Transport. Sys. Magazine*, vol. 10, no. 2, pp. 36–49, 2018.
- [26] S. Roger, M. Brambilla, B. C. Tedeschini, C. Botella-Mascarell, M. Cobos, and M. Nicoli, "Deep-learning-based radio map reconstruction for V2X communications," *IEEE Trans. Veh. Technol.*, 2023.
- [27] A. Ivanov, K. Tonchev, V. Poulkov, and A. Manolova, "Deep learning for reduced sampling spatial 3-D REM reconstruction," *IEEE Open J. Commun. Society*, vol. 5, pp. 2287–2301, 2024.
- [28] Q. Chen, M. Huang, and J. Yang, "REM-Net+: 3D radio environment map construction guided by radio propagation model," *Authorea Preprints*, 2024.
- [29] Y. Zhou, A. Wijesinghe, S. Zhang, and Z. Ding, "TiRE-GAN: Task-incentivized generative learning models for radiomap estimation with radio propagation model," *arXiv:2405.02567*, 2024.
- [30] O. Sener and S. Savarese, "Active learning for convolutional neural networks: A core-set approach," *Intl. Conf. on Learn. Rep.*, Vancouver, BC, April 30-May 3, 2018.
- [31] J. T. Ash, C. Zhang, A. Krishnamurthy, J. Langford, and A. Agarwal, "Deep batch active learning by diverse, uncertain gradient lower bounds," *Intl. Conf. on Learn. Rep.*, New Orleans, LA, May 6-9, 2019.
- [32] K. Jedoui, R. Krishna, M. Bernstein, and L. Fei-Fei, "Deep bayesian active learning for multiple correct outputs," *arXiv:1912.01119*, 2019.
- [33] A. Locatelli, A. Carpentier, and S. Kpotufe, "Adaptivity to noise parameters in nonparametric active learning," in *Proc. of Conf. on Learn. Theory*, Amsterdam, Netherlands, July 7-10, 2017.
- [34] A. Kapoor, K. Grauman, R. Urtasun, and T. Darrell, "Active learning with Gaussian processes for object categorization," in *IEEE Intl. Conf. on Computer Vision*, Rio de Janeiro, Brazil, Oct 14-21, 2007, pp. 1–8.
- [35] A. Krause and C. Guestrin, "Nonmyopic active learning of Gaussian processes: an exploration-exploitation approach," in *Proc. of Intl. Conf. on Machine Learn.*, Corvallis Oregon, June 20-24, 2007, pp. 449–456.
- [36] K. Miller and J. Calder, "Poisson reweighted Laplacian uncertainty sampling for graph-based active learning," *SIAM J. Math. Data Science*, vol. 5, no. 4, pp. 1160–1190, 2023.
- [37] J. Huang, Z. Gojcic, M. Atzmon, O. Litany, S. Fidler, and F. Williams, "Neural kernel surface reconstruction," in *Proc. IEEE Conf. Comp. Vision Pat. Recog.*, 2023, pp. 4369–4379.
- [38] K. D. Polyzos, Q. Lu, and G. B. Giannakis, "Weighted ensembles for adaptive active learning," *IEEE Trans. Signal Process.*, vol. 72, pp. 4178–4190, 2024.
- [39] J. Chapman, B. Chen, Z. Tan, J. Calder, K. Miller, and A. L. Bertozzi, "Novel batch active learning approach and its application on the synthetic aperture radar datasets," in *Alg. Synth. Aper. Radar Imag. XXX*, vol. 12520. SPIE, 2023, pp. 96–111.
- [40] G. A. D. Communities, "Cellsignalstrengthnr," <https://developer.android.com/reference/android/telephony/CellSignalStrengthNr>, 2023, accessed: 2023-07-29.
- [41] ACCUVER, "XCAL: PC based advanced 5G network optimization solution," accessed: 2023-07-29. [Online]. Available: <https://accuver.com/sub/products/view.php?idx=6>
- [42] B. Schölkopf, R. Herbrich, and A. J. Smola, "A generalized representer theorem," in *Intl. Conf. Comput. Learn. Theory*. Springer, 2001, pp. 416–426.
- [43] D. Romero, M. Ma, and G. B. Giannakis, "Kernel-based reconstruction of graph signals," *IEEE Trans. Signal Process.*, vol. 65, no. 3, pp. 764–778, 2017.
- [44] J. Hoydis, "Sionna RT: scene creation with Blender using OpenStreetMap," <https://www.youtube.com/watch?v=7xHLDxUaQ7c>, 2023, accessed: 2023-07-10.
- [45] blender.org, "Blender," <https://www.blender.org/>, 2023, accessed: 2023-07-10.
- [46] J. Hoydis, S. Cammerer, F. Ait Aoudia, A. Vem, N. Binder, G. Marcus, and A. Keller, "Sionna: An open-source library for next-generation physical layer research," *arXiv:2203.11854*, Mar. 2022.
- [47] H. Wackernagel and H. Wackernagel, "Ordinary kriging," *Multivariate geostatistics: an introduction with applications*, pp. 79–88, 2003.
- [48] H. Wackernagel, *Universal Kriging*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2003, pp. 300–307.
- [49] R. Christensen, *Advanced linear modeling*. Springer, 2019.