

# Automatic Network Planning with Digital Radio Twin

Xiaomeng Li\*, Yuru Zhang\*, Qiang Liu\*, Mehmet Can Vuran\*,  
Nathan Huynh\*, Li Zhao\*, Mizan Rahman†, Eren Erman Ozguven‡

\* University of Nebraska-Lincoln, † University of Alabama, ‡ Florida State University

**Abstract**—Network planning seeks to determine base station parameters that maximize coverage and capacity in cellular networks. However, achieving optimal planning remains challenging due to the diversity of deployment scenarios and the significant simulation-to-reality discrepancy. In this paper, we propose *AutoPlan*, a new automatic network planning framework by leveraging digital radio twin (DRT) techniques. We derive the DRT by finetuning the parameters of building materials to reduce the sim-to-real discrepancy based on crowdsourced real-world user data. Leveraging the DRT, we design a Bayesian optimization based algorithm to optimize the deployment parameters of base stations efficiently. Using the field measurement from Husker-Net, we extensively evaluate *AutoPlan* under various deployment scenarios, in terms of both coverage and capacity. The evaluation results show that *AutoPlan* flexibly adapts to different scenarios and achieves performance comparable to exhaustive search, while requiring less than 2% of its computation time.

**Index Terms**—Network Planning, Digital Radio Twin, Bayesian Optimization

## I. INTRODUCTION

Network planning constitutes a fundamental stage in the design of cellular networks, wherein critical base station (BS) parameters (e.g., geographic placement, antenna orientation, transmission power, and tilt) are optimized to ensure desired network coverage and performance. With the densification of cellular networks [1], network planning has become increasingly intricate, requiring the joint optimization of heterogeneous base stations under site-specific propagation effects induced by irregular 3D terrain, complex building geometries, and diverse material compositions. Conventional cellular network planning relies on a combination of field measurements (e.g., drive tests) and model-based simulations (e.g., calibrated radio propagation models), which is laborious and time-consuming.

Digital radio twin (DRT) [2] emerges as a promising technique to facilitate network planning. DRT aims to replicate the real-world radio propagation in the scene, with the attributes of fidelity (imitating as closely as possible), tractability (querying as cheaply as possible), and synchronicity (tracking as timely as possible) [3]. By leveraging a DRT, network operators can perform trial-and-error evaluations to optimize base station parameters and assess coverage, capacity, and interference trade-offs, prior to costly field implementation.

However, it is challenging to derive a DRT with all the essential attributes. On the one hand, existing simulators (e.g., Sionna RayTracing [4] and Wireless InSite) suffer from non-trivial simulation-to-reality (sim-to-real) discrepancy [5], necessitating extensive calibration to improve fidelity. On the other hand, the synchronicity attribute is difficult to achieve, as it is not cost-efficient to conduct periodic field measurements for simulator calibration. Note that, given the complex compu-

tation (e.g., raytracing and radio effects), we can only obtain the network performance (e.g., coverage and capacity) after querying the DRT. Hence, it is difficult to derive the optimal base station parameters, under high-dimensional combination spaces, even if the DRT can be queried in seconds.

In this paper, we propose *AutoPlan*, a new automatic network planning framework by leveraging digital radio twin techniques. On the one hand, we derive the DRT with the following steps. First, we crowdsource abundant real-world data from mobile users under existing deployed base stations (e.g., prior deployments or those of other network operators). Second, with the online data, we continuously calibrate the parameters (particularly, the material of buildings) of the existing simulator (i.e., Sionna RT), to reduce the sim-to-real discrepancy. Third, we accelerate the query of the calibrated simulator with GPU parallel computing to improve its tractability. On the other hand, we design a network planning algorithm to derive the optimal base station parameters (particularly, location and power). Specifically, we leverage Bayesian optimization to tackle the non-negligible querying costs of the DRT by creating a Gaussian process as the surrogate model and adopting expected improvement (EI) as the acquisition function. Using the field measurement from Husker-Net (a private 5G network at the University of Nebraska-Lincoln), we extensively evaluate *AutoPlan* under different deployment scenarios, in terms of both coverage, capacity, and efficiency.

Overall, we summarize the contribution of this paper as:

- A digital radio twin (DRT) with the attribute of fidelity, tractability, and synchronicity;
- An automatic network planning algorithm based on Bayesian optimization, that derives optimal base station parameters assisted by the derived DRT;
- An extensive campus-scale evaluation, showing comparable network performance with state-of-the-art works.

## II. DIGITAL RADIO TWIN

In this section, we introduce the design of the DRT by finetuning the parameters of the target scenario in the simulator.

### A. Problem

The goal is to derive the DRT to achieve all the attributes of fidelity, tractability, and synchronicity. We consider the scene given with 3D terrain, including the full set of building objects within the region of interest. Each building is associated with an initial material label, such as marble or concrete, which corresponds to a specific set of material parameters. These material labels serve as the initial configuration for the DRT simulation. Since these default values may not fully reflect the actual propagation environment, a calibration process is subsequently performed to refine the material parameters

using real-world signal measurements, thus improving the fidelity and reliability of the DRT.

In this work, we focus on calibrating two key electromagnetic parameters used in DRT modeling for objects within the region  $\mathcal{A}$ : conductivity (in S/m) and relative permittivity. Specifically, conductivity characterizes how a material attenuates or reflects incident electromagnetic energy. At the same time, relative permittivity affects the bending and delay of waves as they traverse the medium, both playing critical roles in DRT modeling. We use  $\sigma = [\sigma_1, \sigma_2, \dots, \sigma_K]$  and  $\epsilon = [\epsilon_1, \epsilon_2, \dots, \epsilon_K]$  to denote the conductivity and relative permittivity of  $K$  objects, respectively. The parameter set is thus defined as:

$$\Theta = \begin{bmatrix} \sigma_1 & \sigma_2 & \dots & \sigma_K \\ \epsilon_1 & \epsilon_2 & \dots & \epsilon_K \end{bmatrix}. \quad (1)$$

Based on the typical electromagnetic properties of common materials [6], we impose the following constraints on the tunable parameters:  $\sigma_k \in (0, 2)$  and  $\epsilon_k \in (1, 6)$ .

Hence, the goal is further concretized to seek the optimal simulation parameters to minimize the sim-to-real discrepancy.

### B. Challenge

Although DRT systems can simulate real-world wireless propagation, discrepancies often arise between their simulated outputs and actual measurements—referred to as the sim-to-real gap. This gap does not stem from a single source but rather results from the complex interplay of multiple factors, such as multipath effects, material properties, and environmental geometry. Due to this multifactor coupling, the gap cannot be eliminated by simply adjusting a single parameter. The key challenge is developing a systematic methodology for characterizing and mitigating this complex, coupled discrepancy.

### C. Solution

To address the above problem, we develop a material learning method by leveraging the differentiable property of state-of-the-art simulators.

To calibrate  $\Theta$ , we utilize both the measured and simulated reference signal received power (RSRP, in dBm) values. Let  $\hat{\mathbf{r}} = [\hat{r}_1, \hat{r}_2, \dots, \hat{r}_P]$  represent the RSRP collected at  $P$  locations in  $\mathcal{A}$  from existing BSs. The corresponding simulated RSRP vector, computed using the DRT under  $\Theta$ , is denoted by  $\tilde{\mathbf{r}}(\Theta) = [\tilde{r}_1(\Theta), \tilde{r}_2(\Theta), \dots, \tilde{r}_P(\Theta)]$ . The calibration task is defined as

$$\Theta^* = \arg \min_{\Theta} \|\tilde{\mathbf{r}}(\Theta) - \hat{\mathbf{r}}\|, \quad (2)$$

where  $\|\cdot\|$  represents the  $\ell_2$  norm between the simulated and measured RSRP vectors. In practice, we minimize the empirical loss:

$$\mathcal{L}(\Theta) = \frac{1}{P} \sum_{p=1}^P \|\tilde{\mathbf{r}}(\Theta) - \hat{\mathbf{r}}\|. \quad (3)$$

We optimize this loss over  $E$  training epochs using stochastic gradient descent [7]. At each epoch  $e \in \{1, \dots, E\}$ , the parameters are updated as:

$$\Theta^{(e)} \leftarrow \Theta^{(e-1)} - \eta \nabla \mathcal{L}(\Theta^{(e-1)}), \quad (4)$$

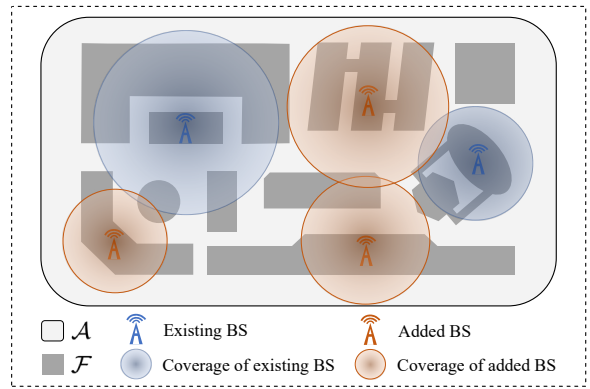


Fig. 1: Illustration of BSs deployment in feasible region  $\mathcal{F}$ .

where  $\eta$  is the learning rate and  $\Theta^{(e)}$  denotes the parameters after the  $e$ -th epoch. In practice, after  $E$  epochs of SGD, we use the checkpoint that yields the lowest training loss as our estimate of the optimizer in (2), denoted by  $\Theta^*$ .

## III. AUTOMATIC NETWORK PLANNING

In this section, we introduce the automatic network planning algorithm by leveraging the aforementioned DRT.

### A. Problem

The goal of network planning is to obtain the optimal base station parameters (e.g., location and orientation) that maximize the specified network performance (e.g., coverage and capacity) in the aforementioned DRT, under a given geographic area.

As illustrated in Fig. 1, we consider a geographic region  $\mathcal{A} \subset \mathbb{R}^2$  with given 3D terrain, such as buildings, trees, and walkways. Here, we denote  $\tilde{\mathcal{B}} = (\hat{x}_m, \hat{y}_m)_{m=1}^M$  as the deployed  $M$  base stations (BSs). To further improve coverage and expand capacity, we aim to deploy additional  $N$  BSs at locations  $\tilde{\mathcal{B}} = (\tilde{x}_n, \tilde{y}_n)_{n=1}^N$ , selected from a feasible region  $\mathcal{F} \subseteq \mathcal{A}$  that meets practical placement constraints (e.g., rooftop availability and height limits). The complete BS set after deployment is  $\mathcal{B} = \tilde{\mathcal{B}} \cup \tilde{\mathcal{B}}$ , and the goal is to determine  $\tilde{\mathcal{B}}$  to maximize the overall performance of the network in  $\mathcal{A}$ .

Next, we introduce two network performance metrics, including coverage and capacity. We discretize  $\mathcal{A}$  into a uniform square grid with resolution  $a \times a$  meters, yielding  $L = \lceil |\mathcal{A}|/a^2 \rceil$  sampling points, where  $|\mathcal{A}|$  is the area of  $\mathcal{A}$ , and  $\lceil \cdot \rceil$  represents the ceiling function (i.e., rounding up to the nearest integer). For each sampling point  $\ell \in \{1, \dots, L\}$ , the characteristics of the received signal are obtained using the aforementioned DRT that reflects the real-world propagation environment.

We evaluate network coverage using  $r_\ell(\Theta)$ , which denotes the received power of the reference signal (RSRP) at the sampling point  $\ell$ , corresponding to the strongest signal received from all BSs in the set  $\mathcal{B}$ . In addition, we evaluate the network capacity using the signal-to-noise ratio (SNR), which captures signal quality in the presence of interference. Let  $s_\ell(\Theta)$  represent the SNR at point  $\ell$ , computed under the current parameter setting  $\Theta$  in the DRT. Given the parameter of the DRT, we can obtain RSRP and SNR values at all

sampling points, and then define the following performance metrics. Given a predefined RSRP threshold  $r_{\text{th}}$ , the coverage rate is defined as:

$$C(\mathcal{B}) = \frac{1}{L} \sum_{\ell=1}^L \mathbb{1}[r_{\ell}(\Theta) > r_{\text{th}}], \quad (5)$$

where  $\mathbb{1}[\cdot]$  is the indicator function, returning 1 if the condition is satisfied and 0 otherwise. The network capacity, reflecting the average spectral efficiency across all sampling points, is expressed by:

$$S(\mathcal{B}) = \frac{1}{L} \sum_{\ell=1}^L \log_2(1 + s_{\ell}(\Theta)). \quad (6)$$

Finally, we combine these two metrics into a composite objective function:

$$T(\mathcal{B}) = \beta \cdot C(\mathcal{B}) + S(\mathcal{B}), \quad (7)$$

where  $\beta > 0$  is a tunable weight controlling the trade-off between coverage and capacity. Here, a larger  $\beta$  places greater emphasis on coverage, whereas a smaller one prioritizes capacity.

Hence, the goal of network planning is further concretized to determine the optimal set of additional deployment locations  $\bar{\mathcal{B}}^*$  that maximizes the overall objective:

$$\bar{\mathcal{B}}^* = \arg \max_{\bar{\mathcal{B}} \subseteq \mathcal{F}} T(\hat{\mathcal{B}} \cup \bar{\mathcal{B}}), \quad (8)$$

where  $\hat{\mathcal{B}}$  denotes the set of existing BSs and  $\mathcal{F} \subseteq \mathcal{A}$  represents the feasible deployment region.

### B. Challenge

The key technical challenge lies in the unknown correlation between the input of base station locations and the output of the network performance. This is attributed to the complex computation (e.g., ray tracing and radio effects) in the DRT, which cannot be mathematically represented. Specifically, we can obtain only the output of network performances by querying the DRT (non-negligible querying costs) under the input base station locations. In other words, the objective function  $T(\mathcal{B})$  is unknown, which falls into the realm of blackbox optimization.

### C. Solution

To address the above problem, we design an automatic network planning method by leveraging Bayesian optimization (BO). BO [8], [9] is the state-of-the-art global searching approach (especially for handling expensive querying costs), generally consisting of a surrogate model and acquisition function. The surrogate model aims to online learn the black box function (e.g., mean and variance), where candidates include Gaussian process [10] and Bayesian neural networks [11]. Based on the updated surrogate model, the acquisition function evaluates the customized utility of different actions (i.e., BS locations) in the action space (i.e.,  $\mathcal{F}$ ). In each iteration, the next action is selected by maximizing the acquisition function (e.g., expected improvement (EI) or probability of improvement (PI)) [12].

To support multiple base station deployment, a naive approach in BO is to use the surrogate model to learn the combinational action space (i.e., all candidate BS locations) and overall network performance. However, this naive approach exponentially expands the action space and practically leads

to a much slower convergence speed and more cumulative querying costs, in terms of time consumption and computation complexity. Hence, we convert the problem to a series of incremental base station deployment problems. In other words, we apply Bayesian optimization only to search for one additional base station deployment at a time. Once the base station location is determined, we put it into the existing deployed base station sets (i.e.,  $\hat{\mathcal{B}}$ ), and then create another Bayesian optimization searching for the next base station. This incremental deployment continues until all  $N$  base stations are deployed.

In BO, the true objective function (7) is replaced by a surrogate model that is cheaper to evaluate but still captures the underlying spatial correlation of the objective. Specifically, we use a Gaussian process (GP) as the surrogate model, which is widely adopted and time-evaluated in multiple application domains. Here, we model the target function as:

$$T(\hat{\mathcal{B}} \cup \bar{\mathcal{B}}) \sim \mathcal{GP}(\mu(\bar{\mathcal{B}}), \sigma^2(\bar{\mathcal{B}})), \quad (9)$$

where  $\mu(\bar{\mathcal{B}})$  and  $\sigma^2(\bar{\mathcal{B}})$  are the posterior mean and variance of the GP given the observations collected in previous iterations. The GP surrogate allows us to predict both the expected performance of a candidate deployment  $\bar{\mathcal{B}}$  and the uncertainty of this prediction, which is crucial for balancing exploration and exploitation.

The next action to query is generally determined by maximizing the selected acquisition function, which defines the utility of all actions in the action space while balancing exploration and exploitation. Specifically, we adopt the EI as the acquisition function, which quantifies the expected gain over the current best performance  $T_{\text{best}}$ :

$$\alpha(\bar{\mathcal{B}}) = \mathbb{E} \left[ \max \left( 0, T(\hat{\mathcal{B}} \cup \bar{\mathcal{B}}) - T_{\text{best}} \right) \right]. \quad (10)$$

This criterion favors actions that are likely to improve the objective while still exploring regions of high uncertainty. At each iteration  $t$ , we select the next set of deployment locations by maximizing the acquisition function over the feasible region  $\mathcal{F}$ :

$$\bar{\mathcal{B}}_{t+1} = \arg \max_{\bar{\mathcal{B}} \subseteq \mathcal{F}} \alpha(\bar{\mathcal{B}}). \quad (11)$$

The newly selected action is then evaluated using the derived DRT model (see Sec. II), and the resulting observation is incorporated into the GP to update its posterior distribution. This iterative process continues until a stopping criterion, such as a fixed iteration limit or convergence threshold, is met.

## IV. THE AUTOPLAN ALGORITHM

In this section, we summarize the workflow of the *AutoPlan* algorithm in Alg. 1. Overall, the algorithm takes the input of 3D terrain of the geographic region and the location of deployed base stations  $\hat{\mathcal{B}}$ , and generates the output of the location of the  $N$  additional base stations. In the digital radio twinning stage, it observes real-world RSRPs  $\hat{r}$  from crowdsourced mobile users and updates the building materials. Specifically, it iteratively calculates the loss function by comparing the real-world RSRPs with the simulated RSRPs from the simulator, and performs multiple gradient descent steps. In the stage of network planning, it iteratively searches

the location of the next base station by 1) selecting a base station location (i.e., maximizing the EI acquisition function); 2) querying the DRT to obtain the network performance; and 3) updating the GP surrogate model with new observations.

---

**Algorithm 1:** The *AutoPlan* Algorithm

---

**Input:** Region  $\mathcal{A}$ , feasible subregion  $\mathcal{F}$ ,  
Existing BS set  $\hat{\mathcal{B}} = \{(\hat{x}_m, \hat{y}_m)\}_{m=1}^M$ ,  
Measured RSRP  $\hat{\mathbf{r}}$ , simulated RSRP  $\tilde{\mathbf{r}}(\Theta)$ ,  
Grid size  $a$ , RSRP threshold  $r_{\text{th}}$ ,  
Number of new BSs  $N$ , training epochs  $E$

**Output:** Calibrated parameters  $\Theta^*$ , new BS locations  
 $\bar{\mathcal{B}} = \{(\bar{x}_n, \bar{y}_n)\}_{n=1}^N$

- 1 **Stage 1: Digital radio twin calibrating;**
  - 2 Initialize  $\Theta$  randomly;
  - 3 **for**  $t \leftarrow 1$  **to**  $E$  **do**
  - 4     Compute simulated RSRP  $\tilde{\mathbf{r}}(\Theta)$  using DRT;
  - 5     Compute loss  $\mathcal{L}(\Theta)$  between  $\hat{\mathbf{r}}$  and  $\tilde{\mathbf{r}}(\Theta)$ ;
  - 6     Update parameters  $\Theta$  according to (4);
  - 7 Keep the checkpoint with the minimum training loss over  $E$  epochs, and denote it as  $\Theta^*$ ;
  - 8 **Stage 2: Automatic network planning;**
  - 9 Initialize  $\bar{\mathcal{B}} \leftarrow \emptyset$ , observation set  $\mathcal{D} \leftarrow \emptyset$ ;
  - 10 **for**  $n \leftarrow 1$  **to**  $N$  **do**
  - 11     Fit GP surrogate model to  $\mathcal{D}$ ;
  - 12     Compute acquisition function  $\alpha(\bar{b})$  using Expected Improvement ( $\bar{b} = (\bar{x}, \bar{y})$ );
  - 13     Select next BS:  $\bar{b}^* \leftarrow \arg \max_{\bar{b} \in \mathcal{F} \setminus (\bar{\mathcal{B}} \cup \hat{\mathcal{B}})} \alpha(\bar{b})$ ;
  - 14     Evaluate  $T(\hat{\mathcal{B}} \cup \bar{\mathcal{B}} \cup \{\bar{b}^*\})$  using calibrated DRT;
  - 15     Update:  $\bar{\mathcal{B}} \leftarrow \bar{\mathcal{B}} \cup \{\bar{b}^*\}$ ;
  - 16     Add  $(\bar{b}^*, T)$  to  $\mathcal{D}$ ;
  - 17 **return**  $\Theta^*$ ,  $\bar{\mathcal{B}}$
- 

## V. PERFORMANCE EVALUATION

In this section, we conduct extensive simulations to evaluate *AutoPlan* by comparing with state-of-the-art works.

**Parameters.** We select a 1210 m  $\times$  1138 m campus area at the University of Nebraska–Lincoln, i.e., the deployment region  $\mathcal{A}$ . There are  $K = 127$  objects (i.e., buildings) in this region. We obtain the scene geometry from OpenStreetMap3D and derive the feasible deployment region  $\mathcal{F} \subseteq \mathcal{A}$  by extracting the union of building footprints. One existing private 5G base station is located near the southwest corner of  $\mathcal{A}$ , yielding the initial BS set  $\hat{\mathcal{B}} = (\hat{x}_m, \hat{y}_m)_{m=1}^M$  with  $M = 1$ . In the region  $\mathcal{A}$ , we collect  $P = 1494$  RSRP measurements  $\hat{\mathbf{r}} = [\hat{r}_1, \hat{r}_2, \dots, \hat{r}_P]$  via drive testing, with a 5G dongle (Quectel RM520N-GL). The accurate location of the receiver is obtained by using RTK-GNSS, providing horizontal accuracy better than 2 cm. The deployed base stations are equipped with a planar antenna array configured as a  $1 \times 2$  horizontal linear array with half-wavelength element spacing. The array adopts the 3GPP TR 38.901 directional antenna pattern and vertical polarization, representing a typical 5G small cell configuration. In the evaluation, we aim to select

$N = 5$  additional BS locations from the feasible region, i.e.,  $\bar{\mathcal{B}} = (\bar{x}_n, \bar{y}_n)_{n=1}^5 \subseteq \mathcal{F}$ . To calculate coverage, we set  $r_{\text{th}} = -90$  dBm and  $a = 0.2$  m. The default maximum transmission power of the base station is 43 dBm. Additionally, we set  $\beta = 10$  to balance the contributions of coverage and capacity in the target objective. All simulations were conducted on a desktop with an Intel Core i7 CPU, an NVIDIA RTX 3080 GPU, and Ubuntu 24.04 OS, using Sionna RT v1.1.0 [4]. Specifically, we calculate the RSRP and SNR by using *RadioMapSolver* on a uniform  $0.2\text{m} \times 0.2\text{m}$  grid. To learn the building materials, we use the Adam optimizer [13] with learning rate  $\eta = 0.01$ , for  $E = 300$  epochs.

We compare *AutoPlan* with the following methods:

- **Random Sampling (RS).** At each selection step, we randomly sample feasible candidate locations from  $\mathcal{F}$  and evaluate  $T(\cdot)$  using the calibrated DRT. For a fair comparison with *AutoPlan*, we use the same pool of 100 candidate BS groups and report the group with the largest objective as the final solution.
- **Exhaustive Search (ES).** ES discretizes the feasible action space  $\mathcal{F}$  using a 5 m step factor, evaluates the objective  $T(\cdot)$  at every candidate point under the calibrated DRT, and selects the location with the highest performance. ES provides a near-optimal result under the given resolution with the highest computational complexity.

### A. Digital Radio Twinning

In this subsection, we evaluate the DRT for all the attributes of fidelity, tractability, and synchronicity, under the *AutoPlan* algorithm. Specifically, we collect the field measurement via a drive test in the City Campus of the University of Nebraska–Lincoln, as shown in Fig. 2. Fig. 4 and 5 show the learned conductivity and permittivity during DRT calibration.

As shown in Fig. 3, the training MSE loss steadily decreases over the 300 gradient descent steps and converges after around 200 steps, with a minimum value of 113.78. In other words, the sim-to-real discrepancy is reduced by 11.23(8.98%), by calibrating the parameters of building materials.

Note that it is impossible to reduce the MSE loss to zero, because the sim-to-real discrepancy is attributed to multiple complex factors, including but not limited to building materials, simulator abbreviation, radio propagation models, etc., losses, circuit noises, and much more. In our hardware setups, the time consumption of a single gradient descent step is 4.97s on average, with the variance of 0.07. This indicates that, with sufficient GPU resources, the DRT can be continuously calibrated and updated in real time.

TABLE I: DRT inference time.

Number of BSs	1	2	3	4	5
Avg. time (s)	0.39	0.72	0.87	1.07	1.24

In addition, Table I reports the average inference time of the DRT for varying numbers of deployed base stations. Notably, it is 0.39s with a single base station and increases to 1.24s with five, highlighting the computational cost of large-scale




Fig. 2: The trajectory of measured RSRP.




Fig. 3: The MSE loss during DRT calibration.




Fig. 4: Conductivity of building materials, during DRT calibration.




Fig. 5: Relative permittivity of building materials, during DRT calibration.




Fig. 6: Network performance under different methods.




Fig. 7: Network performance under different TX powers.

evaluations.

### B. Automatic Network Planning

In this subsection, we evaluate the network performance of network planning under all comparison solutions.

The network performance, as defined in Eq. (7), for all compared network planning methods is illustrated in Fig. 6. As we can see, exhaustive search yields the highest overall performance, while random sampling performs significantly worse than both *AutoPlan* and ES. Notably, *AutoPlan* achieves performance nearly identical to that of ES, demonstrating its effectiveness in accurate and efficient BS deployment. In terms of computation time, ES requires 602 minutes, approximately 58.33 times longer than *AutoPlan*, highlighting the latter’s practical efficiency.

TABLE II: Performance of deployment methods.

	AutoPlan	ES	Random sampling
Coverage (%)	<b>86.04</b>	83.78	80.33
Capacity (bit/s/Hz)	8.04	<b>8.32</b>	6.54
Target	16.64	<b>16.70</b>	14.57

In Table II, we further show the detailed network coverage and capacity under 5 additional base station deployments. We can see that, *AutoPlan* achieves 99.64% overall target performance of ES, including a slightly higher network coverage (2.40%), but a lower network capacity (3.36%). These results demonstrate the effectiveness of *AutoPlan* in selecting near-optimal BS deployment with significantly reduced computational cost.

In Fig. 8-10, we compare the RSRP heatmaps generated by the final BS deployments. Among the three methods, *AutoPlan* achieves the largest coverage area with strong RSRP, while Random Sampling suffers from fragmented coverage and extensive low-signal regions. Moreover, four of the BS

selected by *AutoPlan* are close to the placements in the ES solution, indicating that it can approximate the near-optimal deployment with far fewer evaluations.

In Fig. 11, we show the impact of DRT on the selection of base station location under *AutoPlan*. As observed, the deployment results obtained via the original DRT differ significantly from those based on the calibrated DRT, highlighting the influence of DRT accuracy on BS placement decisions. Quantitatively, the optimal deployment selected using *AutoPlan* under the original DRT achieves a target value of 18.20, while the same method applied on the calibrated DRT yields a target of 16.64—indicating a gap of approximately 10%. This result underscores the necessity of calibrating the DRT to ensure reliable and accurate network planning decisions.

Moreover, we evaluate the sensitivity of coverage to transmit power in Fig. 7, which plots the coverage achieved by *AutoPlan* as a function of the number of newly deployed BSs under different transmit power levels. As expected, coverage improves with increased power; however, the gain becomes marginal beyond 30dBm. For example, when five new BSs are installed, increasing the power from 30dBm to 43dBm only results in an improvement of 2%. Considering the trade-off between performance and power consumption, we adopt 43dBm as the default setting in subsequent experiments. This value yields the highest observed coverage and corresponds to typical configurations of micro-base stations.

## VI. RELATED WORK

Traditional approaches to BS deployment include both manual measurement and optimization-based methods. Early efforts often relied on drive testing or site surveys to evaluate signal quality at predefined locations [14], which are labor-intensive, time-consuming, and inflexible in adapting to changing user demands. To overcome these limitations, optimisation-based methods were introduced to automate the process using mathematical formulations or heuristics. Chiaraviglio et al. [15] formulated BS planning as a mixed-integer linear program under service and EMF constraints, and proposed a heuristic solution named PLATEA. While effective in balancing coverage and EMF compliance, the method relies on static environmental models and sensitive parameters, limiting its adaptability in dynamic scenarios. Similarly, Philip et al.[16] applied meta-heuristic algorithms such as particle swarm optimization (PSO) and genetic algorithms (GA) for BS place-




Fig. 8: RSRP heatmaps of *AutoPlan*.




Fig. 9: RSRP heatmaps of Random Sampling.




Fig. 10: RSRP heatmaps of Extensive Search.




Fig. 11: Base station locations of *AutoPlan*.

ment. Although PSO achieved better coverage and efficiency than GA, both approaches depend on simplified propagation models and static user distributions, making them less suitable for realistic and data-driven deployments.

In contrast, AI/ML-based methods aim to enhance scalability and adaptability by replacing manual engineering and handcrafted models with data-driven learning. *AutoBS* [17] employs reinforcement learning to select BS locations, enabling fast inference during deployment. However, it requires large-scale pretraining on labeled environment performance pairs, which is both computationally expensive and time-consuming to collect. OSSN [18], a Transformer-based framework, unifies radio map estimation and site selection to reduce reliance on exhaustive candidate evaluations. While it improves planning efficiency, its performance remains limited by the availability of training data and a lack of real-world validation. Similarly, Dai and Zhang [19] take a hybrid approach by integrating machine learning with heuristic optimization, using a dataset of over 0.76 million measurements to predict signal strength without explicit propagation modeling.

However, existing AI/ML-based methods require extensive labeled datasets and incur substantial data collection and training costs. More importantly, they tend to overlook the sim-to-real discrepancy introduced by inaccurate or incomplete modeling of physical environments, which hinders generalization in real-world deployments.

## VII. CONCLUSION

In this paper, we propose a new automatic network planning framework (*AutoPlan*) by leveraging digital radio twin techniques. We derive the DRT by finetuning the parameters of building materials to reduce the sim-to-real discrepancy based on crowdsourced real-world user data. Leveraging the DRT, we design a Bayesian optimization-based algorithm to efficiently search for optimal base station parameters. Using the field measurement from Husker-Net, we extensively evaluate *AutoPlan* under various deployment scenarios, in terms of coverage, capacity, and efficiency.

## ACKNOWLEDGMENT

This work is partially supported by the US National Science Foundation under Grant No. 2321699, and No. 2426481.

This material is based upon work supported by the Federal Motor Carrier Safety Administration under grant number: FM-MHP-0888. Any opinions, findings, and conclusions or

recommendations expressed in this publication are those of the author(s) and do not necessarily reflect the view of the Federal Motor Carrier Safety Administration and/or the U.S. Department of Transportation.

## REFERENCES

- [1] S. Andreev, V. Petrov, M. Dohler, and H. Yanikomeroglu, "Future of ultra-dense networks beyond 5g: Harnessing heterogeneous moving cells," *IEEE Communications Magazine*, vol. 57, no. 6, pp. 86–92, 2019.
- [2] L. U. Khan *et al.*, "Digital twin of wireless systems: Overview, taxonomy, challenges, and opportunities," *IEEE Communications Surveys & Tutorials*, vol. 24, no. 4, pp. 2230–2254, 2022.
- [3] P. Almasan *et al.*, "Network digital twin: Context, enabling technologies, and opportunities," *IEEE Communications Magazine*, vol. 60, no. 11, pp. 22–27, 2022.
- [4] J. Hoydis *et al.*, "Sionna rt: Differentiable ray tracing for radio propagation modeling," in *2023 IEEE Globecom Workshops (GC Wkshps)*. IEEE, 2023, pp. 317–321.
- [5] M. Manalastas *et al.*, "From simulators to digital twins for enabling emerging cellular networks: A tutorial and survey," *IEEE Communications Surveys & Tutorials*, 2024.
- [6] P. Series, "Effects of building materials and structures on radiowave propagation above about 100 mhz," *recommendation itu-r*, pp. 2040–1, 2015.
- [7] S. Boyd and L. Vandenberghe, "Convex optimization," Cambridge university press, Cambridge, England, 2010.
- [8] X. Wang, Y. Jin, S. Schmitt, and M. Olhofer, "Recent advances in bayesian optimization," *ACM Computing Surveys*, vol. 55, no. 13s, pp. 1–36, 2023.
- [9] R. Garnett, *Bayesian optimization*. Cambridge University Press, 2023.
- [10] C. E. Rasmussen, "Gaussian processes in machine learning," in *Summer school on machine learning*. Springer, 2003, pp. 63–71.
- [11] Y. L. Li *et al.*, "A study of bayesian neural network surrogates for bayesian optimization," *arXiv preprint arXiv:2305.20028*, 2023.
- [12] P. I. Frazier, "A tutorial on bayesian optimization," *arXiv preprint arXiv:1807.02811*, 2018.
- [13] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [14] A. F. Molisch, *Wireless communications: from fundamentals to beyond 5G*. John Wiley & Sons, 2022.
- [15] L. Chiaraviglio, C. Di Paolo, and N. Blefari-Melazzi, "5g network planning under service and emf constraints: Formulation and solutions," *IEEE Transactions on Mobile Computing*, vol. 21, no. 9, pp. 3053–3070, 2021.
- [16] J. A. Philip, K. M. Mathews, M. Saji, M. Kurian, and M. Vincent, "Cellular base station deployment using pso and genetic algorithm," in *2023 9th International Conference on Smart Computing and Communications (ICSCC)*. IEEE, 2023, pp. 201–206.
- [17] J.-H. Lee and A. F. Molisch, "Autobs: Autonomous base station deployment with reinforcement learning and digital network twins," *arXiv preprint arXiv:2502.19647*, 2025.
- [18] Y. Zheng, C. Liao, J. Wang, and S. Liu, "A transformer-based network for unifying radio map estimation and optimized site selection," in *ICASSPW*. IEEE, 2024, pp. 610–614.
- [19] L. Dai and H. Zhang, "Propagation-model-free base station deployment for mobile networks: Integrating machine learning and heuristic methods," *IEEE Access*, vol. 8, pp. 83 375–83 386, 2020.