

Charting Channels in the Presence of RIS

Mohamed el Mehdi Makhoul^{*}, Yamil Vindas[‡], Anil Kumar^{*}, Maxime Guillaud^{*}, Marco Di Renzo[†]

^{*}Inria, INSA Lyon, CITI Laboratory, UR3720, Villeurbanne, France

[‡] University of Geneva, Switzerland,

[†]Université Paris-Saclay, CNRS, CentraleSupélec, L2S, France and King's College London, United Kingdom

Abstract—This work explores the extension of channel charting, a technique for constructing low-dimensional representations of channel state information (CSI) through self-supervised learning, to wireless environments including reconfigurable intelligent surfaces (RIS). Channel charting has shown promise in pseudo-position based radio resource management applications by capturing spatial relationships between transmitters and receivers. We investigate the challenges and opportunities that RISs present to channel charting, particularly focusing on how RIS reconfiguration affects the learned low-dimensional representations. We focus on a contrastive learning framework with triplet loss to derive a channel chart through distance learning, and incorporate the RIS configuration information into the learning process, aiming to develop channel charts that are robust to changes in the RIS configuration. We also explore the potential for leveraging channel charts to inform RIS configuration, enabling applications where the channel chart is used to determine optimal RIS settings from a given codebook.

I. INTRODUCTION

Channel charting (CC) [1] is a self-supervised learning technique that maps high-dimensional channel state information (CSI) into a low-dimensional latent space, referred to as a *channel chart*. It relies on the spatial correlation structure inherent to wireless channels: users that are physically close typically exhibit similar CSI. Leveraging this property via dimensionality reduction techniques, CC has been successfully applied to tasks such as resource allocation and beamforming [2], [3]. Unlike conventional positioning approaches, CC requires neither explicit location labels nor Global Navigation Satellite System (GNSS) signals.

Radio-map-based network management [4] is a key enabler of intelligent wireless systems, using spatial predictions of performance indicators such as Received Signal Strength (RSS) or blockage probability to support functionalities including handover optimization [5], beam selection and RIS control. Channel charts naturally enable *position-free* radio maps by replacing spatial coordinates with chart coordinates, as demonstrated for the purpose of link adaptation in [6].

Despite these advantages, most existing CC methods rely on a quasi-static propagation environment in which the CSI is assumed to depend solely on the mobile user position. Under this assumption, CSI changes smoothly across space governed by the static geometry of scatterers. This assumption is violated in the presence of reconfigurable intelligent

surfaces (RISs): by altering the phase profile of a large array of passive elements, RIS reconfigurations actively reshape the propagation environment [7]. Consequently, CSI samples collected at the *same* physical location but under different RIS configurations may differ significantly and be mapped to different regions of the chart, thereby degrading spatial consistency.

To address this challenge, the goal of the present work is to **overcome the effect of RIS reconfigurations** and learn a **RIS-agnostic channel chart**, in which embeddings corresponding to the same user location should be mapped to the same channel chart pseudo-location, regardless of the RIS configuration under which the CSI has been obtained. Such a representation makes it possible to compare radio *charts* acquired under different RIS modes, and to perform RIS control in a manner that is robust to RIS reconfiguration dynamics. To the best of our knowledge, no prior work has explored how RIS-induced non-stationarity affects CC or how RIS-aware CC can support RIS control. This paper aims to bridge this gap.

The main contributions of this work are as follows:

- We define a variability metric that quantifies the stability of the learned channel chart across RIS configurations and enables a systematic comparison of training strategies.
- We extend triplet loss based channel charting to RIS-assisted wireless environments, addressing the non-stationarity introduced by RIS reconfiguration.
- We introduce RIS-aware triplet sampling strategies to explicitly incorporate configuration diversity during training and promote cross-mode embedding invariance.
- We evaluate the proposed approach on synthetic CSI generated using Sionna in scenarios with a RIS with a finite number of configurations.
- We demonstrate how radio maps can be used for RIS control based on physical locations or learned chart embeddings, and compare the performance of both approaches.

II. SYSTEM MODEL AND PROBLEM FORMULATION

We consider a wireless system composed of a base station (BS) equipped with M antennas, a reconfigurable intelligent surface (RIS) composed of N passive reflecting elements, and a single-antenna user equipment (UE). The BS and RIS remain fixed, whereas the UE moves along a trajectory that spans both line-of-sight (LoS) and non-line-of-sight (NLoS) regions.

This work was supported by the France's ANR (grant ANR-23-CHR4-0001-01) under the CHIST-ERA project CHASER (CHIST-ERA-22-WAI-01), and by the European Commission under Horizon Europe SNS project INSTINCT (grant 101139161). It was performed using HPC resources from GENCI-IDRIS (Grant 2023-AD010614930).

Assuming a narrowband channel model, for a given user position \mathbf{p} , the effective BS–UE channel under RIS configuration $m = 1 \dots M_{\text{RIS}}$ is modeled as

$$\mathbf{h}^m(\mathbf{p}) = \mathbf{H}_{\text{BS-RIS}} \Phi_m \mathbf{h}_{\text{RIS-UE}}(\mathbf{p}) + \mathbf{h}_{\text{BS-UE}}(\mathbf{p}), \quad (1)$$

where $\mathbf{H}_{\text{BS-RIS}} \in \mathbb{C}^{M \times N}$ models the BS–RIS link, $\mathbf{h}_{\text{RIS-UE}} \in \mathbb{C}^{N \times 1}$ models the RIS–UE link, and $\mathbf{h}_{\text{BS-UE}}$ captures the direct BS–UE component including both LoS and NLoS paths. The RIS applies a phase-shift profile through matrix Φ_m which determines the steering direction or focal point of the reflected radio waves. Without loss of generality, we assume that the RIS operates according to a predefined codebook of size M_{RIS} , defined as $\mathcal{C} = \{\Phi_1, \dots, \Phi_{M_{\text{RIS}}}\}$, with each Φ_m corresponding to a distinct reflection profile.

A. Channel Charting

CC consists in establishing a forward charting function (FCF), denoted by f_θ , which maps a CSI sample to a low-dimensional embedding. f_θ is commonly implemented using an artificial neural network, with learnable parameters denoted by θ . In the classical dimensionality reduction setup, the FCF is optimized such that the low-dimensional representation locally approximates a distance based on CSI features (such as angles of arrival, spatial correlation, impulse responses, angular power density, etc.). In the absence of RIS (letting $\Phi_0 = \mathbf{0}$ in eq. (1)), the CSI $\mathbf{h}^0(\mathbf{p}) = \mathbf{h}_{\text{BS-UE}}(\mathbf{p})$ can be assumed a continuous function of the propagation channel geometry, and it is possible to choose θ such that

$$\|f_\theta(\mathbf{h}(\mathbf{p})) - f_\theta(\mathbf{h}(\mathbf{p}'))\| \propto \|\mathbf{p} - \mathbf{p}'\| \quad (2)$$

when \mathbf{p}' is in the vicinity of \mathbf{p} .

B. RIS-Agnostic Channel Chart

RIS reconfigurations affect the measured CSI, such that in general $\mathbf{h}^m(\mathbf{p}) \neq \mathbf{h}^{m'}(\mathbf{p})$ whenever $m \neq m'$, hence $f_\theta(\mathbf{h}^m(\mathbf{p}))$ and $f_\theta(\mathbf{h}^{m'}(\mathbf{p}))$ are different points in the latent space. In this work, the goal is to learn a FCF f_θ that ensures that spatially close users are represented by nearby points in the embedding space, regardless of the prevailing RIS mode, i.e. $f_\theta(\mathbf{h}^m(\mathbf{p})) = f_\theta(\mathbf{h}^{m'}(\mathbf{p}))$ for all (m, m') . Fig. 1 depicts a *Non RIS-Agnostic FCF* case where embeddings corresponding to the same location \mathbf{p} are widely spread depending on the RIS configuration, and the desired *RIS-agnostic FCF* case in which f_θ maps all embeddings corresponding to \mathbf{p} to a small region—ideally a single point—regardless of the active RIS mode.

In order to capture this objective, we define the variability at position \mathbf{p} for a given FCF f_θ as

$$V_{\mathbf{p}} = \frac{1}{|\mathcal{H}_{\mathbf{p}}|} \sum_{m \in \mathcal{H}_{\mathbf{p}}} \|\mathbf{e}_{\mathbf{p}}^m - \bar{\mathbf{e}}_{\mathbf{p}}\|_2^2, \quad \bar{\mathbf{e}}_{\mathbf{p}} = \frac{1}{|\mathcal{H}_{\mathbf{p}}|} \sum_{m \in \mathcal{H}_{\mathbf{p}}} \mathbf{e}_{\mathbf{p}}^m, \quad (3)$$

where $\mathcal{H}_{\mathbf{p}}$ denotes the RIS configurations for which CSI samples can be acquired at position \mathbf{p} since coverage may vary depending on the RIS configuration, and $\mathbf{e}_{\mathbf{p}}^m = f_\theta(\mathbf{h}^m(\mathbf{p}))$ is the embedding of the CSI at \mathbf{p} obtained under RIS mode m . For each run, the model output is centered and scaled by

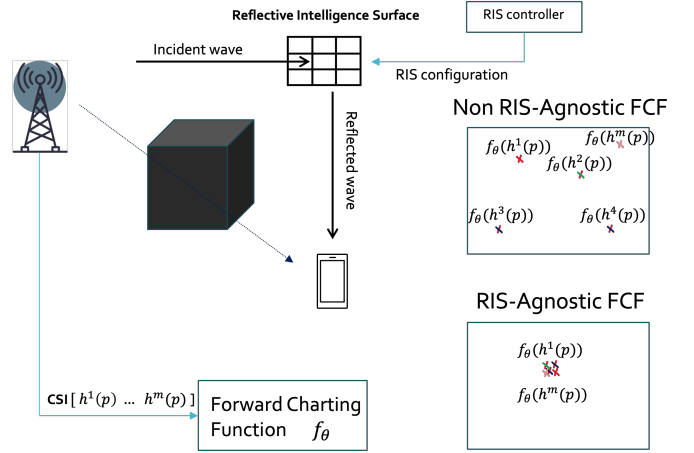


Fig. 1. Illustration of RIS-agnostic channel charting: embeddings of the same physical location corresponding to different RIS configurations remain geometrically consistent in the learned latent space.

the inverse of the training set standard deviation to make the variability metrics comparable across different training runs.

Small values of $V_{\mathbf{p}}$ indicate that the embeddings are robust to RIS reconfiguration. The variability metric above quantifies how strongly RIS reconfiguration affects the learned embeddings. In the next section, we introduce a contrastive training framework that explicitly seeks to minimize it through RIS-aware sampling strategies.

III. RIS-AWARE CONTRASTIVE CHANNEL CHARTING

A. Triplet-Based Distance Learning

In [8], a channel charting method based on contrastive distance learning was introduced; in this approach, CC can be seen as a joint distance-learning and dimensionality-reduction problem. It uses a triplet margin loss of the form

$$\mathcal{L} = \sum_{(A,P,N)} \max(\|f_\theta(A) - f_\theta(P)\|_2^2 - \|f_\theta(A) - f_\theta(N)\|_2^2 + M, 0), \quad (4)$$

where $f_\theta(\cdot)$ denotes the FCF to optimize, M is a margin parameter, and (A, P, N) is a triplet of CSI samples where

- A is the *anchor* CSI sample,
- P is the *positive* CSI sample (physically near the anchor – in [8] it is suggested to use samples acquired within a short time window T_C around the anchor),
- N is the *negative* CSI sample, selected from CSI measurements whose time difference from the anchor exceeds T_C , which implies that the user has likely moved far enough to be considered distant in space.

We now turn to the sampling strategies that determine how the model learns cross-configuration relationships.

B. Proposed Adaptations: RIS-Aware Sampling Strategies

We examine three sampling strategies:

- 1) *No RIS (Baseline)*: The model is trained using CSI obtained in the absence of RIS. This reflects classical CC and serves as a baseline.

2) *Random Mixing*: Triplets are assembled by randomly and uniformly mixing RIS configurations across samples, i.e. random independent RIS states are assumed for A , P and N . This introduces configuration diversity but does not explicitly enforce cross-configuration consistency. This is representative of a case where the RIS configuration evolves over time independently from \mathbf{p} , but its configuration remains unknown to the party performing CC.

3) *RIS-Aware Sampling (Proposed)*: In this case, we assume that the charting entity is aware of the RIS configuration at the time of each CSI sample acquisition. We leverage this information to promote cross-configuration invariance, by enforcing that the RIS configurations associated respectively with A and P differ for all triplets. This encourages the FCF to map samples corresponding to the same location but acquired under different RIS modes to similar embeddings. Note that the existence of such triplets can be guaranteed only if the RIS configuration changes fast enough, such that we can find a positive sample with a different configuration in a small time window T_C . In this work, we assume that this is the case.

IV. EXPERIMENTAL RESULTS

We evaluate the proposed approach using synthetic datasets generated with the Sionna [9] ray tracer (version 0.19.1). The scenario corresponds to an outdoor campus environment comprising one BS operating at 3.5 GHz, a single RIS, and a mobile UE. The RIS has a physical aperture of 5×5 m, discretized into approximately 13,000 reflecting elements with half-wavelength spacing, and operates with a discrete codebook of $M_{\text{RIS}} = 5$ focal points (modes). Each mode steers the reflected beam toward a distinct region of the user area, as illustrated in Fig 3c below. The phase profiles of the RIS are generated based on the method proposed in [10]. The UE follows realistic pedestrian trajectories with an average speed of 1 m/s, covering both LoS and NLoS regions. All channels are modeled as frequency-domain responses over 256 OFDM subcarriers.

A. Training Details

The FCF f_θ is implemented as a feed-forward neural network trained using the triplet margin loss, yielding 2-dimensional CC embeddings. The network architecture follows the design used in [8]. We preprocess the CSI according to the method described in [11]. Triplets are sampled according to the three schemes described in Section III-B, namely *No RIS*, *Random Mixing*, and *RIS-Aware Sampling*. For each strategy, we generate a large pool of triplets based on short-term CSI correlation to select anchor-positive pairs within a

time window parameter $T_C = 1$ s. Each model is trained for 40 epochs using the Adam optimizer with a learning rate of 10^{-2} , a batch size of 2048, and a margin $M = 1$ in eq. (4). We use a training set including around 140k samples and a testing set containing about 18k samples. Training and testing sets contain data from independent simulated user trajectories. After training, we compute the embeddings for each user position and RIS configuration and evaluate the embedding variability metric according to eq. (3).

B. Variability reduction results

Table I summarizes the mean embedding variability across RIS configurations for the three training strategies, averaged over four independent runs for the same set of testing locations. The results indicate a clear improvement in embedding stability when introducing RIS configuration diversity during training. The *Random Mixing* approach already reduces the mean variability by nearly a factor of five compared to the *No-RIS* baseline. This shows that simply exposing the model to multiple RIS configurations during training, even without explicit constraints on cross-configuration pairs, is beneficial for stabilizing the latent representation. The proposed *RIS-Aware Sampling* strategy further enhances this effect, achieving the lowest mean variability and the lowest standard deviation in a consistent manner across multiple runs. Fig. 2 provides a complementary view of the results by showing the distribution of the variability metric over all user positions. The *No-RIS* model exhibits a broader distribution, with a significant fraction of positions experiencing high variability across RIS modes. In contrast, both *Random Mixing* and *RIS-Aware Sampling* yield distributions that are more concentrated near zero, indicating that most positions are embedded consistently across configurations. The *RIS-Aware* strategy shifts the distribution slightly further toward lower variability, which aligns with the quantitative results reported in Table I.

V. RIS CONTROL BASED ON PSEUDO RADIO MAPS

We next investigate the use of the learned channel charts for RIS control through radio maps. The idea is to learn, for each RIS mode, a spatial model of a performance metric (here, the RSS), and then select the RIS configuration that maximizes the predicted performance at a given location or pseudo-location. We compare two control strategies: one that operates directly in the physical coordinate space and one that operates in the learned latent space.




Fig. 2. Distribution of embedding variability for the three training approaches.

TABLE I
EMBEDDING VARIABILITY ACROSS RIS CONFIGURATIONS (LOWER IS BETTER) AVERAGED OVER 4 DIFFERENT RUNS

Approach	Mean Variability	Std Variability
No-RIS Model	1.06×10^{-2}	3.54×10^{-2}
Random Mixing	2.15×10^{-3}	7.78×10^{-3}
RIS-Aware Sampling	2.03×10^{-3}	7.67×10^{-3}

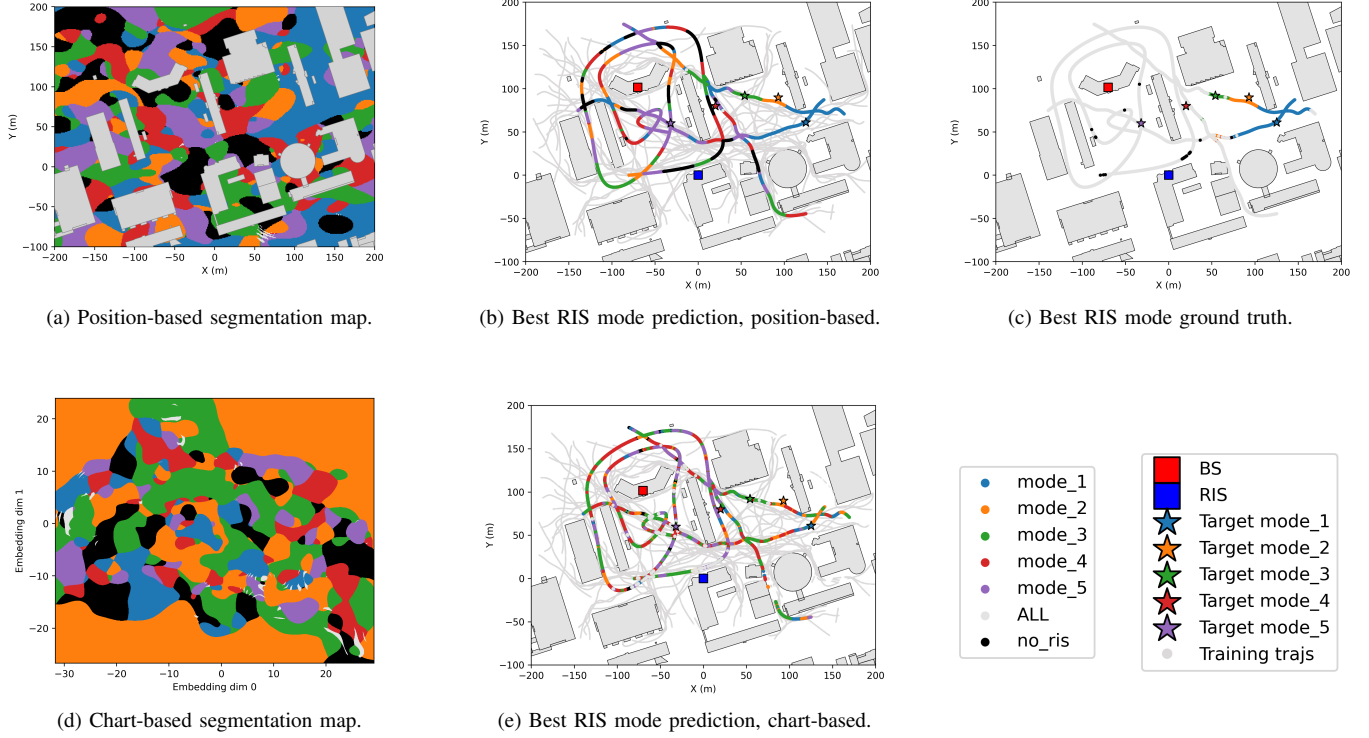


Fig. 3. Segmentation maps and control maps produced by position-based and chart-based RIS control. Here (a) and (d) shows the SVGP-predicted best RIS mode over the physical space and CC embeddings respectively. Figures (b) and (e) shows the predicted best mode for the testing points (colored lines), superimposed over the training set (gray). Figure (c) shows the ground truth labels for the best RIS mode for the testing points.

A. Location-Based RIS Control

In the location-based approach, we assume that user positions are available during the training and testing phases. A Stochastic Variational Gaussian Process (SVGP) regression model [12] is trained to predict RSS for each RIS mode using pairs of user positions and corresponding RSS values using GPyTorch [13]. The use of SVGPs provides a practical compromise between scalability and modeling fidelity: the inducing-point approximation makes GP-based radio maps computationally tractable for our large training set. Each SVGP uses a constant mean and Matérn-5/2 kernel with automatic relevance determination, wrapped in a scale kernel. Up to 4096 inducing points are uniformly sampled from the training inputs to model the RSS spatial distribution. During inference, for each testing point and each RIS configuration, the corresponding SVGP model predicts the expected RSS. The controller then selects, for each location, the RIS mode that maximizes the predicted mean RSS. When all predicted RSS values fall within 1 dB of each other, the system outputs “ALL,” indicating that multiple modes yield essentially identical RSS.

Fig. 3a shows the segmentation map of the physical space into different regions where each color corresponds to a different predicted best mode based on the learned SVGPs. This illustrates that the model can assign a best RIS mode even at locations for which no measurements were available during training. Fig. 3b depicts the test samples only and shows the resulting control decisions. Each sample is color-labeled

with the RIS mode that offers the best predicted RSS (or “ALL” where applicable). It can be observed that the correct RIS mode is predicted whenever the user is in the vicinity of the corresponding RIS focal points. This map serves as an oracle benchmark, as it assumes access to true positions. The real RSS measurements themselves exhibit a high degree of ambiguity: due to the locality of the focusing effect of the RIS, for more than 80% of the testing samples, all RIS modes produce virtually identical RSS (after rounding to the nearest dB), meaning that the “best” configuration is indistinguishable among modes, as shown in Fig. 3c. These ambiguities vanish in the predicted labels probably because the SVGP prediction error ($\approx 3\text{--}4$ dB) dominates the sub-dB differences between modes in the true data, artificially creating separations that lead the model to pick a single mode even though the true RSS values do not meaningfully distinguish between them.

B. Chart-Based RIS Control

In the second approach, we remove the dependency on physical locations and instead operate in the learned channel-chart space. The same SVGP regression framework is used, but now the input to each model consists of 2D chart embedding obtained from the pre-trained RIS-Aware model, rather than physical coordinates. Importantly, this approach only requires CSI and the learned FCF $f_{\theta}(\cdot)$; no position labels are needed. Training proceeds analogously to the location-based case: for each RIS mode, we train an SVGP on (embedding, RSS) pairs. At inference time, we first embed the CSI of the testing




Fig. 4. Regret comparison different approaches to control the RIS .

points—measured in the absence of RIS—using the FCF, then infer the SVGP models at the resulting pseudo-positions to obtain the expected RSS for each RIS configuration, and finally select the mode with the highest predicted RSS. The decision rule from Sec. V-A is used again to select the best RIS mode, including the 1 dB threshold for the “ALL” output. As in the location-based case, Fig. 3d shows the segmentation obtained when control is performed in the learned latent space. Although expressed in embedding coordinates rather than physical positions, the resulting regions exhibit a spatial organization analogous to the location-based case and offer the possibility to predict the best RIS mode even for pseudo-positions never seen at training. Fig. 3e illustrates the resulting control for the same discrete testing points used in the location based control projected into the real-world domain, where we can see that the controller can predict the correct best RIS mode, especially for points near to focal points of the RIS (depicted by stars in Fig. 3), confirming that the chart preserves the geometric information required for downstream control.

C. Discussion

In our evaluations, complete invariance across RIS modes could not be achieved, since each mode purposefully alters the propagation geometry, however the proposed *Random mixing* and *RIS-aware Sampling* methods managed to keep the variability low enough for downstream tasks to remain reliable. A regret-based comparison (Fig. 4) provides a concise assessment of RIS control performance. Here, regret is defined as the performance loss with respect to the best achievable configuration, i.e., the difference between the RSS of the true optimal mode (the mode yielding the highest measured RSS) and that of the selected mode, and is computed only over non-ambiguous points (those for which one mode is strictly better after rounding). The location-based controller reaches the lowest regret, while the chart-based controller remains competitive and clearly outperforms spatially agnostic approaches such as random mode selection.

VI. CONCLUSION

Overall, the results indicate that:

- RIS-aware channel charting effectively mitigates RIS-induced distortions in the latent space.

- The resulting embeddings enable meaningful RIS control without requiring physical positions.
- Chart-based control performs worse than location-based control because its decisions depend on the quality of the learned chart: residual embedding variability and imperfections in the forward charting function propagate thorough the SVGP radio maps and amplify RSS predictions errors.

More broadly, the chart-based controller highlights the potential of self-supervised CSI embeddings for position-free network optimization, providing a practical alternative when location labels are unavailable. Future work will explore stronger invariance-promoting mechanisms, such as explicit regularizers, and extend the analysis to real-world measurements to assess robustness under hardware impairments and non-ideal RIS behavior.

Despite these limitations, this study provides a first step toward *RIS-agnostic channel charting*. The combination of self-supervised learning, RIS-aware sampling, and radio-map regression forms a promising foundation for future intelligent and reconfigurable wireless networks.

REFERENCES

- [1] C. Studer, S. Medjkouh, E. Gonultaş, T. Goldstein, and O. Tirkkonen, “Channel charting: Locating users within the radio environment using channel state information,” *IEEE Access*, vol. 6, 2018.
- [2] P. Ferrand, M. Guillaud, C. Studer, and O. Tirkkonen, “Wireless channel charting: Theory, practice, and applications,” *IEEE Communications Magazine*, vol. 61, no. 6, pp. 124–130, 2023.
- [3] L. Le Magoarou, T. Yassine, S. Paquelet, and M. Crussière, “Channel charting based beamforming,” in *Proc. Asilomar Conference on Signals, Systems, and Computers*, 2022, pp. 1185–1189.
- [4] S. Bi, J. Lyu, Z. Ding, and R. Zhang, “Engineering radio maps for wireless resource management,” *IEEE Wireless Communications*, vol. 26, no. 2, pp. 133–141, 2019.
- [5] S. Lembo, S. Horsmanheimo, and M. Laukkanen, “Positioning based intra-frequency handover in indoor cellular network for ultra reliable communications assisted by radio maps,” in *Int. Congress on Ultra Modern Telecommunications and Control Systems and Workshops*, 2019.
- [6] T. Kallehauge, A. E. Kalør, P. Ramírez-Espinosa, M. Guillaud, and P. Popovski, “Delivering ultra-reliable low-latency communications via statistical radio maps,” *IEEE Wireless Communications*, vol. 30, no. 2, pp. 14–20, 2023.
- [7] M. Di Renzo, A. Zappone, M. Debbah, M.-S. Alouini, C. Yuen, J. de Rosny, and S. Tretjakov, “Smart radio environments empowered by reconfigurable intelligent surfaces: How it works, state of research, and the road ahead,” *IEEE Journal on Selected Areas in Communications*, vol. 38, no. 11, pp. 2450–2525, 2020.
- [8] P. Ferrand, A. Decurninge, L. G. Ordoñez, and M. Guillaud, “Triplet-based wireless channel charting: Architecture and experiments,” *IEEE Journ. Sel. Areas in Comm.*, vol. 39, no. 8, pp. 2361–2373, 2021.
- [9] NVIDIA Corporation. Sionna: An open-source library for next-generation physical layer research. Accessed: 2024-07-22. [Online]. Available: <https://github.com/NVlabs/sionna>
- [10] E. M. Vitucci, M. Albani, S. Kodra, M. Barbiroli, and V. Degli-Esposti, “An efficient ray-based modeling approach for scattering from reconfigurable intelligent surfaces,” *IEEE Transactions on Antennas and Propagation*, vol. 72, no. 3, pp. 2673–2685, 2024.
- [11] S. Taner, V. Palhares, and C. Studer, “Channel charting in real-world coordinates,” in *Proc. Global Communications Conference (GLOBECOM)*, 2023, pp. 3940–3946.
- [12] J. Hensman, A. Matthews, and Z. Ghahramani, “Scalable variational Gaussian process classification,” in *Artificial intelligence and statistics*. PMLR, 2015, pp. 351–360.
- [13] J. Gardner, G. Pleiss, K. Q. Weinberger, D. Bindel, and A. G. Wilson, “GPpyTorch: Blackbox matrix-matrix Gaussian process inference with GPU acceleration,” *Advances in neural information processing systems*, vol. 31, 2018.