Proceedings of the 4th International Conference on Mechatronics and Smart Systems
DOI: 10.54254/2755-2721/2026.33393

Cross-Scenario Few-Shot Wireless Channel Modeling with Analogical Adaptation

Jiabei Wang

Sino-European Institute of Engineering and Technology, Shanghai University, Shanghai, China
yukii0523@shu.edu.cn

Abstract. Learning-based wireless channel models often perform well in the scenario used for training, but their performance can drop when the deployment region changes and only a few labeled samples are available. This paper studies this few-shot cross-scenario setting using channel data from the DeepMIMO dataset framework for millimeter-wave (mmWave) and massive multiple-input multiple-output (MIMO) research. This research builds a pipeline that combines dual-view channel state information (CSI) and channel impulse response (CIR) pre-training, representative source-scenario organization, and a lightweight support-sample-driven adapter. The adapter uses the labeled support samples in the target region as local references, allowing each query representation to be adjusted according to its relation to the available target examples. The experiments show a clear pattern. Adapter-based methods reduce the error substantially compared with no-adapter transfer, while the gap between pretrained and scratch adapter variants is small under the current frozen-backbone protocol. This suggests that the main benefit in this setting comes from explicit support-query adaptation rather than from pre-training alone. Additional experiments further clarify the role of each component. Source selection remains competitive but brings only a small performance margin. Support coverage has a stronger influence, as spatially concentrated support samples consistently hurt adaptation. Lightweight configuration tuning gives limited but stable improvement. Overall, the adapter is the most reliable part of the current framework, whereas the benefit of pre-training is not yet fully separated from the adapter effect.

Keywords: Wireless channel modeling, Cross-scenario transfer, Few-shot learning, Dual-view pre-training, Analogical adaptation.

1. Introduction

In recent years, wireless channel learning has moved from task-specific supervised models to foundation models and self-supervised representation learning. DeepMIMO provides a reproducible dataset framework for learning-based millimeter-wave (mmWave) and massive multiple-input multiple-output (MIMO) research, while channel charting shows that CSI can also reveal useful spatial structure through self-supervised representation learning [1,2]. WirelessGPT explored large-scale pre-training for multiple wireless tasks [3]. Jiang et al. treated channel state information (CSI) and channel impulse response (CIR) as two aligned views of the same wireless channel, and used

© 2026 The Authors. This is an open access article distributed under the terms of the Creative Commons Attribution License 4.0 (https://creativecommons.org/licenses/by/4.0/).

DOI: 10.54254/2755-2721/2026.33393

their consistency to learn more robust channel representations [4]. WavesFM further extended wireless foundation modeling to communication, sensing, and localization tasks [5]. WiFo viewed channel prediction as a unified reconstruction problem under a foundation model framework [6]. Fontaine et al. also pointed out that wireless physical-layer models cannot simply follow the design of text or image models. This is because wireless data, such as in-phase and quadrature (IQ) samples, CSI, and CIR, are heterogeneous time-series signals with different structures and physical meanings [7]. More recent studies have started to handle this heterogeneity directly. HeterCSI uses adaptive batching and dual masking to deal with cross-scale and cross-scenario CSI pre-training [8]. ContraWiMAE combines masked reconstruction with contrastive learning, so that the model can preserve channel structure and learn more discriminative features at the same time [9].

However, there are still two problems. First, many wireless foundation models use a large amount of available source-scene data for pre-training. They rarely discuss which source scenes are really useful for transfer. As a result, some weakly related data may increase the training cost but bring limited benefit. Second, when the target scene provides only a small support set, many methods adapt the model by direct transfer or ordinary fine-tuning. In such a design, the support samples are used mainly for parameter updates, rather than as local references for each query sample. They do not clearly model how the support samples should help the query samples in the target scene. Recent studies on fast adaptation in wireless communications show that few-shot adaptation depends not only on a good initialization, but also on how well the limited target-scene reference samples are used [10]. Similar ideas have also been shown in few-shot orthogonal frequency division multiplexing (OFDM) receiver and mmWave beamforming adaptation, and in transfer-learning and meta-learning based fast beamforming adaptation [11,12].

To address these issues, this paper studies a more complete technical chain for few-shot cross-scenario wireless channel modeling. The proposed framework first organizes representative source scenarios to reduce redundancy while preserving cross-scenario coverage; it then learns shared channel representations by dual-view pre-training on CSI and CIR; after that, it performs support-sample-driven analogical adaptation in the target scenario; finally, it applies lightweight strategy optimization at deployment time.

The study focuses on four connected aspects of few-shot cross-scenario wireless channel modeling: representative source-scenario organization, dual-view pre-training, support-sample-driven analogical adaptation, and lightweight deployment optimization. Together, these components form a practical low-cost adaptation pipeline for new wireless scenarios.

2. Method

2.1. Dual-view pre-training and problem formulation

This research considers a source-to-target setting. The model first learns from several source scenarios and is then evaluated in a target region where only K labeled support samples are available. The support samples are used for adaptation, while the remaining target samples are used for testing.

Each wireless sample is represented by two aligned views. The frequency-domain CSI keeps the subcarrier- and antenna-dependent channel response, while the CIR view is obtained from the same channel and emphasizes delay-domain multipath structure. These two views describe the same propagation process from different domains. Using both of them during pre-training encourages the encoder to retain information that is useful for both channel reconstruction and later target-region adaptation.

DOI: 10.54254/2755-2721/2026.33393

The training loss contains three groups of terms: reconstruction losses for the two channel views, an alignment loss between the CSI and CIR representations, and spatial regularization terms:

$$\mathscr{L}_{\text{pre}} = \lambda_f L_f + \lambda_t L_t + \lambda_a L_a + \lambda_p L_p + \lambda_s L_s \tag{1}$$

where L_f and L_t denote the frequency-domain CSI reconstruction loss and the time-domain CIR reconstruction loss, respectively. L_a is the dual-view alignment loss between the CSI and CIR representations. L_p and L_s are auxiliary spatial terms used to preserve position-sensitive and local-geometry information in the latent space. The coefficients λ_f, λ_t, λ_a, λ_p, and λ_s control the relative weights of these terms.

The reconstruction losses prevent the encoder from discarding useful channel information in either domain. The alignment term encourages the CSI and CIR branches to form a consistent latent structure, while the spatial terms make the representation more suitable for the downstream position-regression task. Thus, pre-training gives the later support-sample-driven adapter a more stable feature space, rather than only reconstructing the input channel.

The pre-training stage should therefore be viewed as a shared representation-learning step. It does not solve the target adaptation problem by itself. Instead, it provides a stable encoder before the model observes the few labeled samples in the target region.

2.2. Representative source-scenario organization

In cross-scenario pre-training, simply concatenating all available source scenarios does not necessarily improve transferability. Highly similar scenarios may introduce redundancy, and large but weakly relevant scenarios may dilute the truly transferable knowledge. Therefore, source selection is treated here as a representative organization problem rather than a passive data collection step. Source scenarios therefore serve as structured sources of transferable experience.

For each candidate source scenario, statistical channel descriptors, geometric propagation descriptors, and latent representation prototypes are jointly constructed, and the inter-scenario distance is defined as

$$d(S_i, S_j) = \alpha \|\mu_i - \mu_j\|_2 + \beta \|g_i - g_j\|_2 + \gamma \|h_i - h_j\|_2 \tag{2}$$

where the statistical term reflects the overall channel distribution, the geometric term describes the propagation structure, and the latent prototype term measures representational similarity from the model perspective. Rather than formulating source selection as a separate optimization problem, this research uses it as a data-organization step to construct a representative source pool under a limited source budget.

Transferability is hard to describe with one criterion alone. Two scenarios may have similar coarse channel statistics but different geometric layouts. Conversely, scenarios with similar propagation geometry may still produce different latent-channel patterns because of local scattering structure or user distribution. The proposed source-organization strategy therefore combines statistical, geometric, and representation-level cues, instead of relying on a single measure of source relevance. Explicit source organization is also useful from a deployment perspective. In practice, it may be unrealistic to use all available source scenarios with equal priority, because storage, preprocessing cost, and experiment budget can all limit full aggregation. A representative source pool is therefore useful even when its performance gain is moderate, since it provides a more controlled way to construct the pre-training data.

DOI: 10.54254/2755-2721/2026.33393

Source organization and target-scene adaptation operate at different stages of the pipeline. Source organization is applied before deployment to shape the transferable information learned from source scenarios. The adapter is applied after deployment and uses the few labeled target samples to correct query representations under local distribution mismatch.

2.3. Support-sample-driven analogical adaptation

After pre-training, only a small number of labeled support samples are available in the target scenario. Because updating the full backbone with such limited supervision can easily lead to overfitting, this research freezes the encoder and update only the lightweight adapter and regression head.

This setup matches the intended deployment context. In a new target region, there are usually not enough labels for reliable end-to-end retraining. The support samples are therefore used as local references to guide query prediction under the new regional distribution, rather than merely as a small fine-tuning set.

For a query sample q with representation z_q, this research computes its relation to the i-th support sample using the similarity weight

$$a_i = \frac{\exp(sim(z_q, z_i)/\tau)}{\sum_{j=1}^{K} \exp(sim(z_q, z_j)/\tau)} \tag{3}$$

Here, $sim(z_q, z_i)$ measures the similarity between the query representation z_q and the i-th support representation z_i, z_j denotes the j-th support representation in the normalization term, and τ controls the sharpness of the weighting. The support features are then aggregated into a local context vector and injected into the query representation through a gated residual correction. The adapted feature is finally mapped to the downstream prediction head. The key idea is not to continue large-scale learning in the target scene, but to let the limited support set explicitly participate in building a local reference frame for target-scene modeling.

2.4. Lightweight strategy optimization

The final performance is also affected by deployment choices such as support selection, learning rate, training duration, and adaptation scope. These choices influence accuracy, adaptation stability, and update cost. Instead of using a large-scale hyper-parameter search for every target region, this research treat deployment as a lightweight configuration-selection problem. The goal is to obtain a reasonable balance between adaptation accuracy, update cost, and parameter efficiency under deployment constraints.

This part of the framework is used as an auxiliary deployment module. The main performance gain still comes from shared pre-training and support-sample-driven analogical adaptation. Once the main pipeline is fixed, strategy-level tuning can further reduce the risk of using an unstable manually chosen setting.

In practice, this module has three roles. It reduces dependence on a single hand-tuned configuration, avoids exhaustive tuning for every target scenario, and makes it possible to compare how different configuration choices affect deployment quality. It is therefore used as a deployment-oriented refinement of the adaptation pipeline, rather than as the main source of performance improvement. As shown later in the experiments, this module offers mild but reproducible gains and

DOI: 10.54254/2755-2721/2026.33393

is therefore treated as an auxiliary deployment enhancement rather than the dominant source of performance.

3. Experimental setup

The experiments are built on DeepMIMO scenarios and use dual-view CSI/CIR representations as model input, while the downstream task is position regression in the target region. Evaluation follows a target-region few-shot protocol. Data are divided into source and target regions by spatial coordinates, and the support set and test set are both drawn from the target region. This setting reflects deployment in a new region where only a few labeled target samples can be collected. Unlike random train-test splitting, the protocol introduces a spatial distribution shift between the source region used for training and the target region used for deployment.

The main benchmark compares four methods: Pretrained + Adapter, Scratch + Adapter, Pretrained + NoAdapter (Frozen), and Pretrained + NoAdapter (Unfrozen). Scratch + Adapter uses the same architecture and few-shot pipeline as the proposed method, but starts from random initialization. It is included to test whether pre-training brings an additional benefit under the same adapter-based adaptation protocol. The two no-adapter variants are used to test whether a pre-trained backbone alone is sufficient when explicit support-query relation modeling is removed. This benchmark separates the effect of pre-training from the effect of adaptation, rather than mixing the two in a single comparison.

This research report mean squared error (MSE), mean absolute error (MAE), and root mean squared error (RMSE), with RMSE used as the primary metric because it measures position error on the original scale. To avoid relying on a single case, the main benchmark is summarized across multiple target regions and support settings. Additional multi-seed evidence is reported in the strengthened 16-shot experiment.

Besides the main benchmark, this research include separate experiments for source-scenario comparison, support-coverage analysis, strengthened 16-shot evaluation, and lightweight strategy optimization. Each experiment changes one main factor, so that the corresponding claim can be checked with a more direct comparison.

The representative source-scenario package evaluates whether explicit source-domain organization remains competitive under source-budget constraints. The support-coverage experiment studies the effect of support-sample distribution on final adaptation. The strengthened 16-shot experiment examines the most data-limited setting in more detail. The lightweight strategy experiment evaluates deployment-level configuration choices.

4. Results and analysis

4.1. Main benchmark results

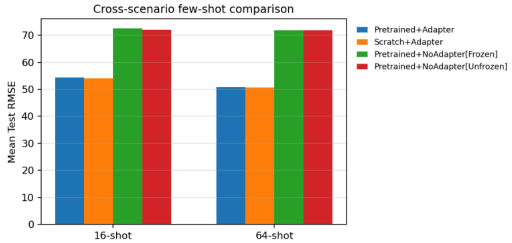


Figure 1. Main cross-scenario few-shot comparison

DOI: 10.54254/2755-2721/2026.33393

Table 1. Main cross-scenario few-shot benchmark

Method	Support Size	Mean RMSE	RMSE Std.	Runs
Pretrained + Adapter	16-shot	53.94	23.38	3
Scratch + Adapter	16-shot	53.86	22.99	3
Pretrained + NoAdapter (Frozen)	16-shot	72.88	21.99	3
Pretrained + NoAdapter (Unfrozen)	16-shot	72.27	22.50	3
Pretrained + Adapter	64-shot	51.19	23.30	3
Scratch + Adapter	64-shot	51.00	23.47	3
Pretrained + NoAdapter (Frozen)	64-shot	72.81	22.03	3
Pretrained + NoAdapter (Unfrozen)	64-shot	72.19	22.55	3

Table 1 reports the main cross-scenario few-shot benchmark. Figure 1 presents the comparison among the four main methods. In Table 1, the three runs correspond to the three target regions used in the compact cross-scenario benchmark. The strengthened 16-shot study further expands this setting with additional seeds and paired comparisons. For both 16-shot and 64-shot settings, the two adapter-based methods perform much better than the no-adapter baselines. At 16-shot, Pretrained + Adapter and Scratch + Adapter obtain mean RMSE values of 53.94 and 53.86, while Pretrained + NoAdapter (Frozen) and Pretrained + NoAdapter (Unfrozen) reach 72.88 and 72.27. The same trend appears at 64-shot: the adapter-based methods stay close to 51, whereas the no-adapter methods remain around 72.

These results show that the adapter is necessary in this few-shot cross-scenario setting. A pretrained backbone alone does not adapt well to the target region when support-query relation modeling is removed. The comparison between Pretrained + Adapter and Scratch + Adapter should be interpreted more carefully. Their RMSE values are almost identical, so the current results should not be used to claim a strong standalone pre-training gain. Under the frozen-backbone protocol, the main improvement comes from the explicit use of target support samples.

The no-adapter comparison gives the same message. Unfreezing the pretrained backbone only slightly improves the no-adapter result and still leaves a large gap to the adapter-based methods. This suggests that simply updating more parameters with a few target labels is not enough. The model also needs a mechanism to use the support samples as local references for query prediction.

For deployment, this result is useful because the best-performing methods do not require full backbone retraining in the target scene. They keep the shared encoder frozen and update only the lightweight adapter and regression head. This keeps the adaptation process relatively cheap while still allowing the model to use the limited target support set.

DOI: 10.54254/2755-2721/2026.33393

4.2. Representative source-scenario organization

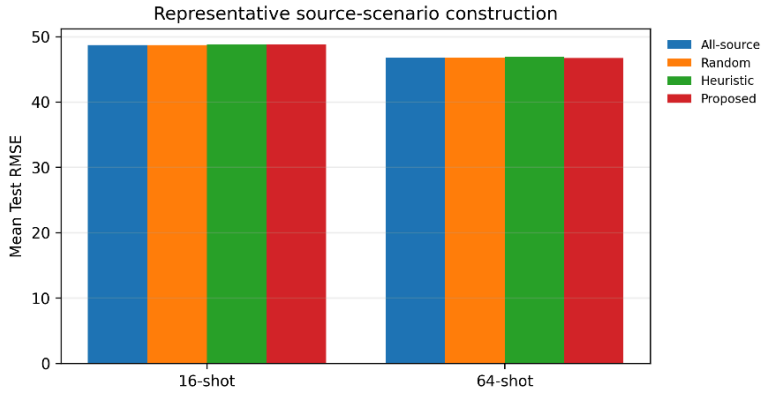


Figure 2. Representative source-scenario organization results

Figure 2 reports the representative source-scenario experiments. Overall, the performance gap among different source-organization strategies is moderate rather than large. In the 16-shot setting, All-source, Random, Proposed, and Heuristic obtain average Test RMSE values of 48.72, 48.71, 48.74, and 48.76, respectively. In the 64-shot setting, the corresponding values are approximately 46.78, 46.76, 46.73, and 46.87.

These results suggest that representative source-scenario organization should not be interpreted as a source of large accuracy gains. Its value is more limited but still useful: it keeps transfer performance competitive under a limited source budget. In this sense, the proposed source organization changes multi-scenario pre-training from simple source aggregation to a more controlled source-pool construction step. Even when the numerical gain is small, the module makes the source construction process clearer and easier to justify.

The 64-shot result is the most favorable case for the proposed strategy. It achieves the best average RMSE among the compared source-organization methods, although the margin is still small. This may indicate that source organization becomes more useful when the target support set is large enough to make use of the transferred representation, but still small enough for transfer to matter. When the support size is very small, the quality of local adaptation may dominate the composition of the source pool.

The methodological implication is also important. This experiment allows source-scenario organization to be included in the framework without making an overly strong performance claim. The evidence supports a restrained conclusion: representative source organization is a reasonable and practical design under limited source budgets, but it is not the main driver of the final performance. This positioning is more credible than claiming that source selection alone drives the main performance advantage.

4.3. Support coverage and adaptation quality

Table 2 and Figure 3 report the main support-coverage results. Under both 16-shot and 64-shot settings, clustered support selection performs worst, while random, diverse, and proposed selection achieve much lower RMSE. The proposed mode obtains the lowest average RMSE in the 16-shot setting. In the 64-shot setting, diverse selection performs best, and the proposed mode remains close to it.

DOI: 10.54254/2755-2721/2026.33393

Table 2. Summary of support-coverage experiments

Support Mode	16-shot Mean RMSE	64-shot Mean RMSE	Runs per shot
clustered	74.12	74.16	9
random	53.05	51.40	9
diverse	54.69	50.40	9
proposed	52.53	50.78	9

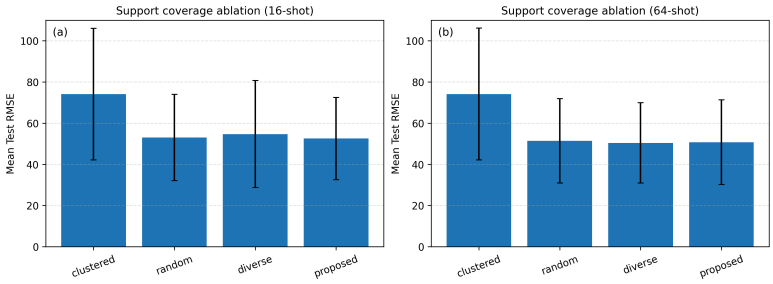


Figure 3. Support-coverage ablation under 16-shot and 64-shot

These results show that few-shot adaptation is affected not only by backbone initialization, but also by the spatial coverage of the support set.

Figure 4 further examines the relationship between support coverage and prediction error. The clustered setting shows very low support_bbox_ratio values and is generally associated with higher RMSE. By comparison, random, diverse, and proposed support sets provide broader spatial coverage and tend to produce lower prediction errors.

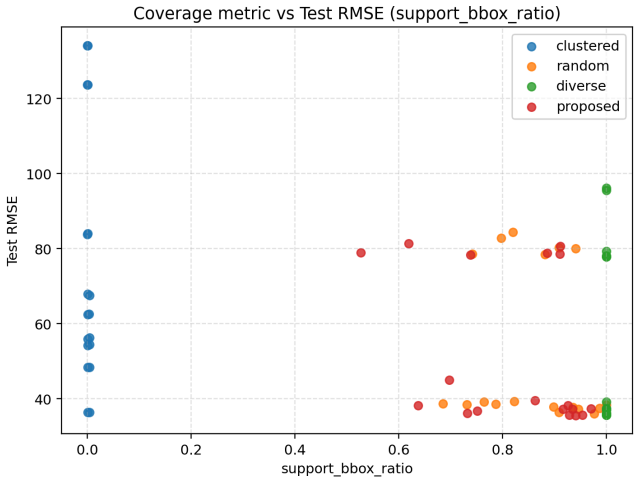


Figure 4. Relationship between support_bbox_ratio coverage and RMSE

This result indicates that the adapter depends not only on the number of support samples, but also on how well these samples represent the target region. When support samples are confined to a small local area, the model can still form support-query relations, but these relations provide limited guidance for query points outside the covered region. Consequently, the adaptation becomes less reliable, leading to larger final errors.

DOI: 10.54254/2755-2721/2026.33393

This finding links support coverage to adaptation quality. The comparison is not only about which support-selection mode achieves lower error. It also explains the mechanism behind the difference: broader support coverage provides more reliable local references for the adapter, while concentrated support samples weaken support-query matching in uncovered regions. This strengthens the credibility of the adapter-based narrative and makes the support-coverage package more than an auxiliary visualization result.

The same pattern appears under both shot settings. Although the exact ranking among random, diverse, and proposed changes slightly between 16-shot and 64-shot, the overall pattern remains stable: clustered support is clearly harmful, while wider spatial coverage is consistently beneficial. This means that support quality is not a marginal issue that only appears in one specific setting, but a general factor shaping few-shot transfer quality in this framework.

4.4. Strengthened 16-shot results and lightweight strategy

The strengthened 16-shot experiments, summarized in Figure 5, further confirm the necessity of adapter-based adaptation. Across Arizona State University (ASU), New York, and Los Angeles, both adapter-based methods consistently achieve lower RMSE than the frozen no-adapter baseline. The average RMSE values of the main methods are 53.69 for Pretrained + Adapter, 53.60 for Scratch + Adapter, and 69.71 for Pretrained + NoAdapter (Frozen). Consistent with the main benchmark, the two adapter-based methods remain very close, but both are substantially better than the no-adapter baseline. Pairwise comparisons also show 30 wins and 0 losses for both Pretrained + Adapter and Scratch + Adapter against the no-adapter baseline, while Pretrained + Adapter versus Scratch + Adapter yields 12 wins and 18 losses. This again suggests that the major gain comes from the adapter mechanism itself rather than from a dominant pre-training margin alone.

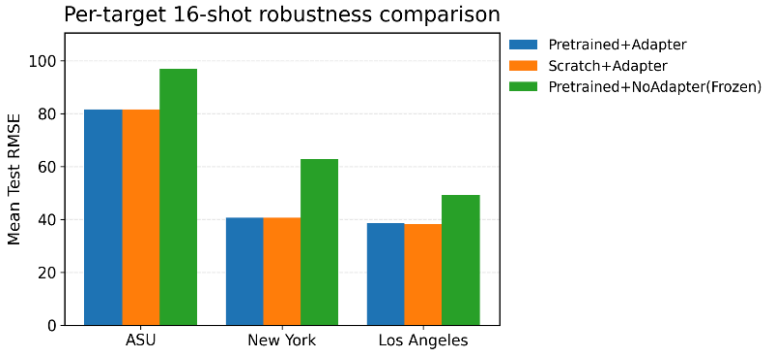


Figure 5. Per-target strengthened 16-shot robustness comparison

The value of the strengthened 16-shot package is that it stresses the practically most challenging and relevant low-shot condition. In real deployment, the extremely small-support setting is often the hardest case, because data collection is limited and local mismatch is large. The fact that the same pattern appears across different target regions increases confidence that the paper's central conclusion is not an artifact of a particular seed or target split.

The lightweight strategy experiments show a mild deployment-level effect. Across 18 cases, proposed tuner slightly improves over the default setting and performs close to random search, with average RMSE values around 51.95 for the tuned strategies and 52.29 for the default configuration. The gain is small, so this module is not treated as the main performance driver. Its role is to reduce

DOI: 10.54254/2755-2721/2026.33393

dependence on a single manually chosen setting and to provide a more stable configuration-selection process.

4.5. Discussion

Taken together, four observations can be drawn. First, support-sample-driven analogical adaptation is the key enabling factor in cross-scenario few-shot modeling. Without the adapter, the pretrained backbone alone cannot handle the local distribution shift in the target region. Therefore, pre-training should be viewed as a shared representation foundation, not as a guarantee that the pretrained model will always outperform scratch training. Source-scenario organization and support-set organization also act at different stages. The former shapes transferable information before deployment, while the latter controls the quality of local reference construction during adaptation. Lightweight strategy selection can reduce the reliance on manual configuration, but its gain remains auxiliary.

The contribution of this work is therefore a complete adaptation pipeline rather than a single dominant module. The pipeline includes representative source-scenario organization, dual-view pre-training, support-driven analogical adaptation, and lightweight deployment optimization. This interpretation is consistent with the experimental results: the framework is most effective when it explicitly organizes and uses limited cross-scenario information, especially the few labeled support samples in the target region.

5. Conclusion

This paper investigated a unified framework for few-shot cross-scenario wireless channel modeling, integrating representative source-scenario organization, dual-view pre-training, support-sample-driven analogical adaptation, and lightweight strategy optimization.

The results consistently show that reliable few-shot transfer across scenarios depends on explicit adapter-based analogical adaptation, which consistently outperforms the no-adapter baselines in both the main benchmark and the strengthened 16-shot evaluation. Under the current frozen-backbone protocol, pre-training mainly provides a shared feature space for later adaptation, and its value becomes more visible when target-side support information is explicitly used. Representative source-scenario organization also remains useful. Although its gain over alternative source-selection strategies is limited, it preserves competitive transfer performance under constrained source budgets and makes multi-scenario pre-training more systematic. Adaptation quality is also closely tied to support-set coverage: spatially concentrated support samples lead to weaker results, while broader and more representative coverage improves final performance. Lightweight strategy optimization further provides small but repeatable gains with little extra deployment cost. Overall, the framework forms a coherent and practical pipeline for low-cost wireless channel modeling in new scenarios.

References

[1] A. Alkhateeb, "DeepMIMO: A Generic Deep Learning Dataset for Millimeter Wave and Massive MIMO Applications, " arXiv: 1902.06435, 2019.
[2] P. Ferrand, M. Guillaud, C. Studer, and O. Tirkkonen, "Wireless Channel Charting: Theory, Practice, and Applications, " IEEE Commun. Mag., vol. 61, no. 6, pp. 124–130, 2023.
[3] T. Yang, P. Zhang, M. Zheng, Y. Shi, L. Jing, J. Huang, and N. Li, "WirelessGPT: A Generative Pre-Trained Multi-Task Learning Framework for Wireless Communication, " IEEE Netw., vol. 39, no. 5, pp. 58–65, 2025.
[4] J. Jiang, W. Yu, Y. Li, Y. Gao, and S. Xu, "A MIMO Wireless Channel Foundation Model via CIR-CSI Consistency, " in Proc. IEEE Int. Conf. Mach. Learn. Commun. Netw. (ICMLCN), 2025, doi: 10.1109/ICMLCN64995.2025.11140262.

DOI: 10.54254/2755-2721/2026.33393

[5] A. Aboulfotouh, E. Mohammed, and H. Abou-Zeid, "6G WavesFM: A Foundation Model for Sensing, Communication, and Localization, " IEEE Open J. Commun. Soc., vol. 6, pp. 6792–6807, 2025.
[6] B. Liu, S. Gao, X. Liu, X. Cheng, and L. Yang, "WiFo: Wireless Foundation Model for Channel Prediction, " Sci. China Inf. Sci., vol. 68, no. 6, Art. no. 162302, 2025.
[7] J. Fontaine, A. Shahid, and E. De Poorter, "Towards a Wireless Physical-Layer Foundation Model: Challenges and Strategies, " in Proc. IEEE Int. Conf. Commun. Workshops (ICC Workshops), 2024, pp. 1566–1572.
[8] C. Zhang, X. Lyu, C. Ren, S. Liu, Q. Cui, and X. Tao, "HeterCSI: Channel-Adaptive Heterogeneous CSI Pretraining Framework for Generalized Wireless Foundation Models, " arXiv: 2601.18200, 2026.
[9] B. Guler, G. Geraci, and H. Jafarkhani, "A Multi-Task Foundation Model for Wireless Channel Representation Using Contrastive and Masked Autoencoder Learning, " arXiv: 2505.09160, 2025.
[10] O. Wang, H. He, S. Zhou, Z. Ding, S. Jin, K. B. Letaief, and G. Y. Li, "Fast Adaptation for Deep Learning-Based Wireless Communications, " IEEE Commun. Mag., vol. 63, no. 10, pp. 158–164, 2025.
[11] O. Wang, S. Zhou, and G. Y. Li, "New Environment Adaptation with Few Shots for OFDM Receiver and mmWave Beamforming, " arXiv: 2310.12343, 2023.
[12] Y. Yuan, G. Zheng, K.-K. Wong, B. Ottersten, and Z.-Q. Luo, "Transfer Learning and Meta Learning-Based Fast Downlink Beamforming Adaptation, " IEEE Trans. Wireless Commun., vol. 20, no. 3, pp. 1742–1755, 2021.