

Consistency-based Test Time Adaptation for Deep Learning-based Parametric Channel Estimation

Wonseok Shin, Hyeongjun Kim and Byonghyo Shim
Seoul National University, Seoul, Republic of Korea
Email: {wsshin, hjkim}@islab.snu.ac.kr, bshim@snu.ac.kr

Abstract—To support extremely high data rates in 6G wireless networks, acquiring accurate channel state information (CSI) is critical, yet pilot-based channel estimation incurs massive overhead. Parametric channel estimation (PCE), in which a deep neural network is trained to estimate a small set of geometric channel parameters from received pilots, has emerged as a promising solution to reduce this overhead. However, the pretrained PCE network suffers from a significant performance drop when deployed in an environment whose channel parameter distribution differs from that of the training set. In this paper, we propose a test time adaptation scheme that adapts a pretrained PCE network to the target domain using only unlabeled received signals. The key idea is to model the domain shift as a covariate shift, where only the input distribution changes while the mapping from the input to the channel parameters remains unchanged across domains. Under this view, we adapt the network using a consistency loss against additive perturbations of the received signal, together with auxiliary terms that prevent ordering confusion among the paths and drift from the source weights. Simulation results demonstrate that the proposed scheme substantially mitigates the degraded performance of the unadapted network on the target scenario.

Index Terms—channel estimation, test time adaptation, domain shift, deep learning.

I. INTRODUCTION

To support extremely high data rates in 6G wireless communications, the use of high-frequency bands (e.g., mmWave and sub-THz) along with massive multiple-input multiple-output (MIMO) antenna arrays has received a great amount of attention in recent years [1]. To exploit such resources, accurate channel state information (CSI) at the base station (BS) is essential. However, the conventional pilot-based channel estimation incurs a heavy overhead that scales with the bandwidth and the antenna dimension, making the accurate CSI acquisition a major challenge in practice [2].

Over the years, deep learning (DL) has emerged as a powerful tool to address this challenge and has been applied to channel estimation [3]. Among DL-based channel estimation methods, parametric channel estimation (PCE) reduces pilot overhead by estimating a small set of geometric parameters, i.e., path gains, AoAs, and delays, from pilot measurements and reconstructing the full channel response from them [4].

This research was partly supported by the MSIT (Ministry of Science and ICT), Korea, under the ITRC (Information Technology Research Center) support program(IITP-2026-2021-0-02048) supervised by the IITP(Institute for Information & Communications Technology Planning & Evaluation) and IITP under the artificial intelligence semiconductor support program to nurture the best talents (IITP-2023-RS-2023-00256081) grant funded by the Korea government (MSIT).

One crucial drawback of DL-based PCE is that the pretrained network may suffer from a significant performance degradation due to *domain shift*. Specifically, the PCE network is a regression model whose output is closely tied to the surrounding propagation environment, so the distribution of channel parameters varies substantially across deployment scenarios. As a result, a network trained on a given channel distribution, referred to as the *source domain*, may not generalize well to a different distribution, the *target domain*, where the changes in environment geometry, SNR level, user mobility, or carrier frequency severely affect the estimation accuracy.

To handle this issue, various adaptation approaches have been studied for DL-based wireless modules. In [5], a domain adaptation scheme bridging the gap between a simulated (quasi-static) channel model and a map-based (ray-traced) channel model has been proposed. In [6], a meta-learning approach has been proposed to enable fast beamforming adaptation to new wireless environments from only a limited amount of labeled target data. While effective, these schemes require perfect channels as ground-truth, which are typically unavailable during operation since the underlying channel is not directly observable. On the other hand, test time adaptation (TTA) techniques that adapt the pretrained model online using only unlabeled target data have been developed in the machine learning community [7], [8]. However, since these methods are mostly designed for classification tasks (e.g., entropy minimization [8] or feature distribution alignment [9]), they are not directly applicable to regression problems such as PCE.

An aim of this paper is to propose a TTA scheme that can be applied to DL-based PCE. Our key observation is that environmental variations primarily affect the statistical distribution of the received signal rather than the underlying mapping from the input to the channel parameters, since the channel is generated from the same few parameters regardless of the domain. We therefore model the domain shift as a *covariate shift*, where the marginal input distribution changes while the conditional relation between input and output is preserved. Under this view, the pretrained network already encodes the correct mapping and only needs to be adjusted to the new input distribution. To this end, we adopt a consistency loss, which has been widely used in semi-supervised learning and TTA to align a model with the target input distribution while preserving the mapping learned in the source domain [10]. The proposed scheme employs an additive consistency term that forces the network to produce consistent predictions

under perturbations of the received signal, augmented by a physical consistency term that enforces a valid ordering of the estimated delays and a regularization term that anchors the network to the source weights. From the simulation result, we demonstrate that the proposed scheme substantially improves the degraded performance of the network.

II. SYSTEM MODEL AND PROBLEM FORMULATION

We consider an uplink single-input multiple-output (SIMO) OFDM system in which the BS equipped with a uniform linear array (ULA) of N_{Rx} antennas serves a single-antenna user equipment (UE). The system employs N_s subcarriers with subcarrier spacing Δf . The uplink channel vector at subcarrier s is expressed as

$$\mathbf{h}_s = \sum_{l=1}^L \alpha_l e^{-j2\pi f_s \tau_l} \mathbf{a}(\theta_l), \quad s = 1, \dots, N_s, \quad (1)$$

where $\alpha_l \in \mathbb{C}$, τ_l , and θ_l denote the complex gain, delay, and AoA of the l -th path, respectively, and $f_s = f_c + (s - \frac{N_s+1}{2})\Delta f$ is the frequency of the s -th subcarrier with carrier frequency f_c . The array steering vector $\mathbf{a}(\theta_l) \in \mathbb{C}^{N_{\text{Rx}}}$ of the ULA is

$$\mathbf{a}(\theta_l) = \left[1, e^{-j\frac{2\pi d}{\lambda} \sin \theta_l}, \dots, e^{-j\frac{2\pi d}{\lambda} (N_{\text{Rx}}-1) \sin \theta_l} \right]^T, \quad (2)$$

where d is the antenna spacing and λ is the signal wavelength. Such a geometric multipath representation is widely adopted in high-frequency systems, where the propagation is dominated by a small number of specular paths [11].

With this channel, the received pilot signal at subcarrier s is given by

$$\mathbf{y}_s = \mathbf{h}_s x + \mathbf{n}_s, \quad (3)$$

where x is the known pilot symbol transmitted by the UE and $\mathbf{n}_s \sim \mathcal{CN}(\mathbf{0}, \sigma^2 \mathbf{I})$ is the additive white Gaussian noise (AWGN). Stacking the received signals over all N_s subcarriers, the received signal matrix is given by

$$\mathbf{Y} = [\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_{N_s}] \in \mathbb{C}^{N_{\text{Rx}} \times N_s}. \quad (4)$$

The channel in (1) is represented by the parameter set $\mathbf{p} = \{\alpha_l, \theta_l, \tau_l\}_{l=1}^L$, and the goal of PCE is to recover the estimate $\hat{\mathbf{p}}$ from $\mathbf{Y}$ and to reconstruct the channel $\hat{\mathbf{h}}_s$ through (1). To this end, we employ a DL network $f(\cdot; \Theta)$ that learns the nonlinear mapping from the received signal to the channel parameters,

$$\hat{\mathbf{p}} = f(\mathbf{Y}; \Theta). \quad (5)$$

Since the geometric parameters are estimated rather than the per-subcarrier coefficients, the channel response over the entire band can be reconstructed from $\hat{\mathbf{p}}$, which substantially reduces the pilot overhead.

III. PROPOSED TEST TIME ADAPTATION

The main goal of the proposed scheme is to adapt a pretrained PCE network to the target domain using only unlabeled target signals. We view the domain shift as a covariate shift, where the input distribution changes across domains while the mapping from the input to the channel parameters is preserved. Under




Fig. 1. Overall framework of the proposed test time adaptation scheme.

this view, the adaptation is reduced to aligning the network with the target input distribution while maintaining the mapping learned in the source domain, which we realize through a consistency loss against additive perturbations of the received signal. To prevent the unsupervised objective from drifting to invalid configurations or away from the mapping learned in the source domain, we further introduce a physical consistency term that enforces a valid ordering of the estimated delays and a regularization term that anchors the network to the source weights.

A. Adaptation Setup

During source training, the network parameters $\Theta = \{\Theta_e, \Theta_h\}$, consisting of the feature extractor parameters Θ_e and the MLP head parameters Θ_h , are learned by supervision on labeled source data, yielding the source weights $\Theta_s = \{\Theta_{s,e}, \Theta_{s,h}\}$. During target adaptation, we freeze the MLP head, $\Theta_h = \Theta_{s,h}$, and adapt only the feature extractor to the target input distribution. The adaptation therefore minimizes $\mathcal{L}_{\text{TTA}}$ over Θ_e alone, i.e.,

$$\Theta_e \leftarrow \Theta_e - \eta \nabla_{\Theta_e} \mathcal{L}_{\text{TTA}}, \quad (6)$$

where η is the learning rate.

B. Network Architecture

We design the PCE network $f(\cdot; \Theta)$ as a Transformer-based architecture [12] consisting of a feature extractor and a multi-layer perceptron (MLP) head. This is motivated by the effectiveness of self-attention in capturing the structure of high-dimensional channel measurements for parameter acquisition [13]. As an input, the received pilot signal $\mathbf{Y}$ is first projected through an input embedding layer and combined with positional encodings to preserve the resource grid structure of the input. The embedded sequence is then passed through a stack of N_B Transformer encoder blocks, each consisting of the multi-head self-attention layer and the feed-forward layer. By the self-attention mechanism, the network captures the correlations across the antenna and subcarrier dimensions. The extracted features are then passed through the MLP head to produce the estimated channel parameters $\hat{\mathbf{p}}$.




Fig. 2. Network Architecture.

Fig. 1 illustrates the overall framework of the proposed scheme. The same network architecture shown in Fig. 2 is used in both the source training and the target adaptation phases, while the training objective changes from a supervised loss in the source training to the proposed unsupervised $\mathcal{L}_{\text{TTA}}$ in the target adaptation. By doing so, the proposed scheme preserves the architectural backbone learned in the source domain while adapting the network to the target distribution.

C. Adaptation Loss

The proposed objective is a weighted sum of three complementary terms,

$$\mathcal{L}_{\text{TTA}} = \lambda_{\text{con}} \mathcal{L}_{\text{con}} + \lambda_{\text{phys}} \mathcal{L}_{\text{phys}} + \lambda_{\text{reg}} \mathcal{L}_{\text{reg}}, \quad (7)$$

where λ_{con} , λ_{phys} , and λ_{reg} are the weighting factors balancing the three contributions, and are tuned empirically.

1) Noise Consistency. At each iteration, random perturbations are added to the received signal and the network is forced to produce the same prediction on the original and perturbed inputs. This encourages the network output to be stable under input variation, and thus to remain consistent across the gap between the source and target distributions. The objective is expressed as

$$\mathcal{L}_{\text{con}} = \frac{1}{M} \sum_{m=1}^M \|f(\mathbf{Y}; \Theta) - f(\mathbf{Y} + \mathbf{E}_m; \Theta)\|_2^2, \quad (8)$$

where $\mathbf{E}_m$ is the m -th random perturbation matrix and M is the number of perturbations.

2) Physical Consistency. The estimated channel parameters must remain physically admissible after the adaptation. Since the propagation paths are indexed in ascending order of delay, the estimated delays should satisfy $\hat{\tau}_1 \leq \hat{\tau}_2 \leq \dots \leq \hat{\tau}_L$. We penalize any violation of this ordering through

$$\mathcal{L}_{\text{phys}} = \sum_{l=1}^{L-1} \max(0, \hat{\tau}_l - \hat{\tau}_{l+1}), \quad (9)$$

Algorithm 1 Proposed Test Time Adaptation

Require: Source weights $\Theta_s = \{\Theta_{s,e}, \Theta_{s,h}\}$; unlabeled target batches $\{\mathbf{Y}^{(b)}\}_{b=1}^B$; weights $\lambda_{\text{con}}, \lambda_{\text{phys}}, \lambda_{\text{reg}}$; learning rate η ; perturbation count M and variance σ_p^2 .

- 1: Initialize $\Theta \leftarrow \Theta_s$ and freeze the MLP head $\Theta_h = \Theta_{s,h}$.
- 2: **for** each target batch $\mathbf{Y}$ **do**
- 3: **for** $m = 1, \dots, M$ **do**
- 4: Draw $\mathbf{E}_m \sim \mathcal{CN}(\mathbf{0}, \sigma_p^2 \mathbf{I})$ and evaluate $f(\mathbf{Y} + \mathbf{E}_m; \Theta)$.
- 5: **end for**
- 6: Compute $\mathcal{L}_{\text{con}}, \mathcal{L}_{\text{phys}}, \mathcal{L}_{\text{reg}}$ from (8)–(10).
- 7: $\mathcal{L}_{\text{TTA}} \leftarrow \lambda_{\text{con}} \mathcal{L}_{\text{con}} + \lambda_{\text{phys}} \mathcal{L}_{\text{phys}} + \lambda_{\text{reg}} \mathcal{L}_{\text{reg}}$.
- 8: $\Theta_e \leftarrow \Theta_e - \eta \nabla_{\Theta_e} \mathcal{L}_{\text{TTA}}$.
- 9: **end for**
- 10: **return** Adapted weights $\Theta = \{\Theta_e, \Theta_{s,h}\}$.

which prevents the adaptation from collapsing into physically invalid configurations such as out-of-order delays caused by confusion among the propagation paths.

3) Regularization. Adapting with an unsupervised objective alone may cause the network to drift away from the mapping acquired during source training. To preserve it, we anchor the adapted weights to the source model [14], expressed as

$$\mathcal{L}_{\text{reg}} = \|\Theta_e - \Theta_{s,e}\|_2^2, \quad (10)$$

which mitigates catastrophic forgetting and stabilizes the adaptation process.

The three loss terms serve complementary roles: $\mathcal{L}_{\text{con}}$ makes the estimate insensitive to domain-induced variations, $\mathcal{L}_{\text{phys}}$ prevents confusion in the ordering of the paths, and $\mathcal{L}_{\text{reg}}$ preserves the mapping learned during source training.

D. Adaptation Procedure

Starting from the source-trained weights Θ_s , for each batch of unlabeled target signals, we draw M random perturbations, evaluate the three loss components, and update Θ by gradient descent on $\mathcal{L}_{\text{TTA}}$. The procedure requires only the unlabeled target signals and the source weights, and incurs only a small computational overhead beyond the standard forward pass, making the proposed scheme suitable for online deployment.

IV. SIMULATIONS AND DISCUSSIONS

A. Simulation Setup

We consider a domain-shift scenario from a synthetic source domain to a ray-traced target domain. In the source domain, the channel parameters $\{\alpha_l, \theta_l, \tau_l\}_{l=1}^L$ are randomly generated with $L = 3$ paths including a dominant line-of-sight (LoS) component. In the target domain, the channel parameters are instead obtained from a ray-tracing simulation of an outdoor environment. The BS is equipped with an $N_{\text{Rx}} = 16$ ULA and the input is taken over $N_s = 16$ subcarriers, i.e., $\mathbf{Y} \in \mathbb{C}^{16 \times 16}$.

The PCE network uses a Transformer-based feature extractor with $N_B = 4$ encoder blocks, 4 attention heads, and an embedding dimension of 128, followed by an MLP head. For source training, we use received pilot signals from 200,000 UEs at different positions. Training is done with Adam optimizer




Fig. 3. Channel NMSE as a function of SNR.

at a learning rate of 10^{-3} and a batch size of 128 over 50 epochs. For target adaptation, 10,000 unlabeled target signals are used, the adaptation runs for 1 step per batch at a learning rate of 10^{-4} , $M = 4$ perturbations per sample are drawn with variance σ_p^2 corresponding to the operating SNR, and the loss weights λ_{con} , λ_{phys} , λ_{reg} are 0.5, 0.3 and 0.1, respectively.

As the performance metric, we adopt the channel normalized mean square error (NMSE) defined as

$$\text{NMSE} = \mathbb{E} \left[\frac{\|\mathbf{h} - \hat{\mathbf{h}}\|_2^2}{\|\mathbf{h}\|_2^2} \right], \quad (11)$$

where $\hat{\mathbf{h}}$ is reconstructed from the estimated channel parameters. We compare the proposed scheme with the following baselines: 1) least squares (LS) estimator, 2) the unadapted pretrained network, and 3) the network fine-tuned with labeled target data.

B. Simulation Results

In Fig. 3, we plot the channel NMSE as a function of the received SNR over the range of 0–10 dB. We observe that the proposed scheme consistently achieves a lower NMSE than the unadapted network across the entire SNR range. For example, when SNR is 0 dB, the proposed scheme achieves an NMSE reduction of more than 1.2 dB over the unadapted network, confirming that the pretrained network can be effectively adapted to the target domain without any target labels.

The improvement is most pronounced in the low-SNR regime, where the source-trained network is most vulnerable to the domain shift. LS estimator exhibits high NMSE at low SNR and improves as the SNR increases. Unlike model-based [15] or DL-based estimators, it relies solely on the noisy pilots without exploiting any prior knowledge of the channel structure. From these observations, we can see that the DL-based methods offer the largest advantage in the low-SNR, where accurate channel estimation is most challenging, by adapting the pretrained network through the noise consistency constraint, while the ordering and regularization terms prevent confusion among the paths and drift from the source weights. The fine-tuned network, which exploits labeled target data, serves as a performance

upper bound, and the proposed scheme narrows a large portion of this gap without using any target labels.

V. CONCLUSION

In this paper, we addressed the performance degradation of DL-based parametric channel estimation under domain shift. The key insight is that environmental variations primarily affect the input distribution while leaving the mapping from the input to the channel parameters unchanged. We therefore proposed a test time adaptation scheme that aligns the pretrained network with the target input distribution through a consistency loss under additive perturbations of the received signal, together with a term enforcing a valid ordering of the estimated delays and a regularization term that anchors the network to the source weights. The simulation results showed that the proposed scheme effectively restores the performance without using any target labels.

REFERENCES

- [1] T. S. Rappaport et al., “Wireless communications and applications above 100 GHz: Opportunities and challenges for 6G and beyond,” *IEEE Access*, vol. 7, pp. 78729–78757, 2019.
- [2] S. Kim, J.W. Choi, B. Shim, “Downlink pilot precoding and compressed channel feedback for FDD-based cell-free systems,” *IEEE Trans. Wireless Commun.*, vol. 19, no. 6, pp. 3658–3672, Jun. 2020.
- [3] W. Kim, Y. Ahn, J. Kim and B. Shim, “Towards deep learning-aided wireless channel estimation and channel state information feedback for 6G,” *J. Commun. Netw.*, vol. 25, no. 1, pp. 61–75, Feb. 2023.
- [4] J. Kim, Y. Ahn, S. Kim, and B. Shim, “Deep learning-aided parametric sparse channel estimation for terahertz massive MIMO systems,” *IEEE Trans. Cogn. Commun. Netw.*, vol. 10, no. 6, pp. 2136–2148, Dec. 2024.
- [5] T. H. Hoang, T. N. Do, and G. Kaddoum, “Domain adaptation-enabled realistic map-based channel estimation for MIMO-OFDM,” *arXiv:2507.08974*, 2025.
- [6] 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, Mar. 2021.
- [7] J. Liang, D. Hu, and J. Feng, “Do we really need to access the source data? Source hypothesis transfer for unsupervised domain adaptation,” in *Proc. Int. Conf. Mach. Learn. (ICML)*, 2020, pp. 6028–6039.
- [8] D. Wang, E. Shelhamer, S. Liu, B. Olshausen, and T. Darrell, “Tent: Fully test time adaptation by entropy minimization,” in *Proc. Int. Conf. Learn. Represent. (ICLR)*, 2021.
- [9] B. Sun and K. Saenko, “Deep CORAL: Correlation alignment for deep domain adaptation,” in *Proc. Eur. Conf. Comput. Vis. (ECCV) Workshops*, 2016, pp. 443–450.
- [10] A. Tarvainen and H. Valpola, “Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results,” in *Proc. Adv. Neural Inf. Process. Syst. (NeurIPS)*, 2017, pp. 1195–1204.
- [11] J. Wu, S. Kim, and B. Shim, “Parametric sparse channel estimation for RIS-assisted terahertz systems,” *IEEE Trans. Wireless Commun.*, vol. 71, no. 9, pp. 5503–5518, Sep. 2023.
- [12] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, “Attention is all you need,” in *Proc. Adv. Neural Inf. Process. Syst. (NeurIPS)*, 2017, pp. 5998–6008.
- [13] S. Kim, A. Lee, H. Ju, K. Ngo, J. Moon and B. Shim, “Transformer-based channel parameter acquisition for terahertz ultra-massive MIMO systems,” *IEEE Trans. Veh. Technol.*, vol. 72, no. 11, pp.15127–15132, Nov. 2023.
- [14] X. Li, Y. Grandvalet, and F. Davoine, “Explicit inductive bias for transfer learning with convolutional networks,” in *Proc. Int. Conf. Mach. Learn. (ICML)*, 2018, pp. 2825–2834.
- [15] S. Park, J. W. Choi, J.-Y. Seol, and B. Shim, “Expectation-maximization based channel estimation for multiuser MIMO systems,” *IEEE Trans. Commun.*, vol. 65, no. 6, pp. 2397–2410, Jun. 2017.