

Environment-Conditioned Generative Channel Modeling: A Diffusion-Based Approach for Site-Specific Wireless Dataset Synthesis

The Imperative for High-Fidelity Synthetic Channel Data

The design and optimization of next-generation wireless communication systems, including 5G-Advanced and the forthcoming 6G, are increasingly dependent on sophisticated machine learning (ML) and artificial intelligence (AI) methodologies. These data-driven techniques are being applied across the protocol stack, from physical layer tasks like channel estimation and beamforming to network-level challenges such as resource allocation and interference management.¹ However, the efficacy of these ML models is fundamentally constrained by the availability of large, diverse, and realistic datasets. This has created a significant bottleneck in the field, as the primary method for acquiring such data—extensive real-world field measurements—is exceptionally resource-intensive, demanding significant investments in time, cost, and specialized equipment.⁴ The challenge of data scarcity represents a fundamental barrier to the rapid prototyping, validation, and deployment of AI-enabled wireless technologies.

The Data Scarcity Bottleneck in Wireless ML

The physical wireless channel is a complex, high-dimensional, and stochastic medium. Its characteristics are determined by a multitude of factors, including the physical geometry of the environment, the materials of surrounding objects, the mobility of transmitters and receivers, and atmospheric conditions. Accurately capturing this complexity in a dataset requires comprehensive measurement campaigns across a wide range of scenarios. For tasks such as massive Multiple-Input Multiple-Output (MIMO) beamforming or millimeter-wave (mmWave) channel estimation, the required datasets must contain thousands or millions of channel state information (CSI) instances to train deep neural networks effectively. The process of collecting this data is not only expensive but also difficult to scale and replicate,

hindering the progress of ML-driven wireless research and creating a high barrier to entry for academic and industrial researchers alike. This data-centric challenge necessitates a paradigm shift away from sole reliance on physical measurements towards scalable and cost-effective data synthesis solutions.

Generative AI as a Paradigm Shift for Data Synthesis

Generative Artificial Intelligence (GAI) has emerged as a transformative solution to the data scarcity problem in wireless communications and numerous other fields.¹ GAI models are designed to learn the underlying probability distribution of a given dataset and can subsequently generate new, synthetic data samples that exhibit the same statistical properties as the original data. This capability allows researchers to augment limited real-world datasets, creating vast repositories of high-fidelity synthetic data for training and testing ML algorithms.

The application of generative models to wireless channel modeling has evolved significantly over the past several years. Initial explorations leveraged architectures such as Variational Autoencoders (VAEs) and Generative Adversarial Networks (GANs).⁷ VAEs proved effective in learning compressed latent representations of channel data, while GANs demonstrated a powerful ability to generate highly realistic channel matrices through an adversarial training process.⁹ To address the inherent training instability of early GAN architectures, more advanced variants like the Wasserstein GAN with Gradient Penalty (WGAN-GP) were developed, leading to improved sample quality and more stable convergence when modeling channel distributions.¹⁰

More recently, the field of generative modeling has been dominated by the rise of Diffusion Models (DMs). DMs have established a new state-of-the-art in generating high-quality, high-dimensional data across various domains, including image synthesis and audio generation. Their application to wireless communications is particularly promising due to several key advantages: superior sample fidelity, stable and robust training dynamics, and an inherent ability to model complex, multi-modal data distributions.¹² This progression from GANs and VAEs to Diffusion Models is not merely an algorithmic substitution but reflects a strategic shift within the research community. The stringent fidelity requirements of physical layer modeling, where unrealistic artifacts can invalidate synthetic data for downstream engineering tasks, have driven the adoption of models that prioritize sample quality and training robustness above all else. The known issues of earlier models, such as mode collapse in GANs or the characteristic blurriness of VAE-generated samples, were found to be unacceptable for the high-precision domain of wireless channel synthesis. Consequently, building upon a diffusion-based framework aligns with the field's trajectory towards achieving the highest possible fidelity in synthetic data generation.

The Unaddressed Frontier: From User-Centric to Environment-Aware

Generation

Despite the rapid progress in generative channel modeling, a critical limitation persists in many state-of-the-art approaches. Models are often user-centric or scenario-specific, meaning they are trained to generate data for a fixed, predefined set of users or environmental conditions. For instance, the Conditional Denoising Diffusion Implicit Model (cDDIM) for wireless channels learns to generate channel matrices conditioned on a discrete user index, effectively learning a separate distribution for each user in the training set.¹⁵ While this is useful for data augmentation within a static scenario, it offers no generalizability; the model is incapable of generating a channel for a new user or a new physical location not encountered during training.

This limitation points towards the next frontier in generative channel modeling: creating models that are environment-aware and possess true generalization capabilities. A recurring and valid critique of purely data-driven generative models is that their outputs may not correspond to physically or geometrically viable channels.⁴ An ideal generative model should not simply memorize statistical patterns but should learn the underlying relationship between the physical environment and the resulting channel characteristics. This would enable the generation of valid channel data for any arbitrary location within a given environment, transforming the model from a simple data augementer into a powerful simulation tool or a "digital twin" of the wireless propagation space.

From User-Specific to Environment-Aware Channel Generation: A Literature Synthesis

To define a clear path forward, it is essential to analyze the two dominant but divergent paradigms in conditional channel generation. The first, exemplified by the cDDIM baseline, focuses on powerful data-driven generation conditioned on discrete identifiers. The second pursues physical realism by integrating deterministic channel models directly into the generative pipeline. Contrasting these approaches reveals a significant and pragmatic research opportunity that synthesizes their respective strengths while mitigating their weaknesses.

User-Conditioned Generation: The cDDIM Baseline

The work "Generating High Dimensional User-Specific Wireless Channels using Diffusion Models" and its associated cDDIM codebase provide a strong starting point for analysis.¹⁵ This framework employs a conditional diffusion model to generate user-specific channel matrices. The core mechanism involves conditioning the denoising process on a specific user's identity.

Mathematically, the model learns to approximate the conditional probability distribution $P(H|c)$, where H is the channel matrix and c is a discrete class label representing a unique user index (e.g., `user_id`). During training, the model is fed pairs of channel matrices and their corresponding user IDs. During inference, a user ID is provided as a condition, and the model generates a new channel matrix drawn from that specific user's learned distribution. The primary critique of this approach lies in its fundamental lack of generalizability. The model treats each user location as a disconnected, categorical entity. It learns a distinct channel distribution for each `user_id` but learns nothing about the spatial relationships between them. As a result, it is incapable of generating a channel for a user at a new location that was not part of the original training set. This severely limits its utility for applications that require exploring the channel characteristics of an entire area, such as network planning, mobility management, or the development of location-aware communication protocols. The model excels at data augmentation for a fixed set of users but fails as a general-purpose environmental simulator.

Physics-Informed Generation: The PBGC-VAE Approach

An alternative and increasingly influential paradigm is that of physics-informed generative modeling. A prominent example is the work on "Physics-Informed Generative Approaches for Wireless Channel Modeling," which integrates a parametric, Physics-Based Geometric Channel (PBGC) model directly into a VAE framework.⁴ In this approach, the generative model (the VAE's decoder) does not generate the channel matrix directly. Instead, it learns to generate the underlying physical parameters that govern the channel, such as the number of propagation paths, their respective gains, angles of arrival, and angles of departure. These generated parameters are then fed into the deterministic PBGC model, which calculates the final channel matrix according to established geometric and electromagnetic principles. This method offers a powerful advantage: by construction, every generated channel matrix is guaranteed to be physically consistent and geometrically viable. It grounds the generative process in the laws of physics, providing a strong inductive bias that can lead to better generalization. However, this approach introduces significant implementation complexity. The authors of the PBGC-VAE work note substantial challenges related to the non-convex loss landscape and poor gradient flow through the highly non-linear PBGC model. To make the end-to-end system trainable, they were required to develop a "linearized reformulation" of the PBGC model, a non-trivial research contribution in its own right.¹⁶ This indicates that integrating explicit physics-based models into deep generative frameworks is a high-risk, time-intensive endeavor that requires deep domain expertise in both machine learning and propagation modeling.

Defining the Opportunity: A Pragmatic Synthesis

The current literature presents a dichotomy between two compelling but flawed extremes. On one side are powerful, relatively easy-to-train, but non-generalizable data-driven models like cDDIM. On the other are generalizable, physically-grounded, but complex and difficult-to-train frameworks like the PBGC-VAE. This division highlights a clear and valuable research gap for a "third way"—an approach that achieves environmental awareness and generalizability without the prohibitive engineering complexity of a full physics-based integration.

The technical difficulty of integrating explicit, differentiable physics models has inadvertently created an opening for what can be termed "soft" physics-informed methods. The complexity of the PBGC-VAE approach acts as a significant barrier to entry, particularly for rapid-prototyping scenarios such as a two-week research competition. Attempting to replicate and integrate a complex, non-differentiable physical model into a diffusion framework under such constraints would be infeasible. However, the *goal* of the physics-informed approach—to make channel generation sensitive to the physical environment—remains critically important.

A pragmatic solution lies in leveraging modern, high-fidelity ray-tracing datasets like DeepMIMO, which provide a direct link between a channel matrix and its corresponding physical context, such as the precise 3D coordinates of the user equipment.²⁰ This allows for a reformulation of the problem. Instead of learning to generate intermediate physical parameters for an explicit model, a generative model can be trained to learn the direct, implicit mapping from geometric coordinates to the resulting channel matrix distribution: $f:(x,y,z) \rightarrow P(H)$. This approach bypasses the complex physics modeling by conditioning the powerful generative capabilities of a diffusion model on continuous physical parameters. It represents a much more tractable machine learning problem that is perfectly suited for rapid implementation while still directly addressing the core limitation of non-generalizable, user-specific models. It achieves a similar outcome—environment-aware generation—with a far simpler and more direct architectural modification.

Proposed Framework: Geometrically-Conditioned Diffusion Models (Geo-cDDIM)

The proposed research introduces the Geometrically-Conditioned Diffusion Model (Geo-cDDIM), a novel framework for generating site-specific wireless channel data. This framework modifies the state-of-the-art cDDIM architecture to enable conditioning on continuous geometric parameters, thereby transforming it from a user-specific data augementer into a generalizable digital twin of a wireless environment.

Problem Formulation: From Categorical to Continuous Conditioning

The fundamental limitation of the baseline cDDIM model is its reliance on categorical conditioning. The model is trained to learn the conditional distribution $P(H|c)$, where $H \in \mathbb{C}^{N_r \times N_t}$ is the complex channel matrix, N_r and N_t are the number of receive and transmit antennas, and c is a one-hot encoded vector representing a discrete user_id from a finite set $\{1, \dots, N\}$. This formulation prevents the model from generating channels for any location not explicitly included in the training set.

The proposed Geo-cDDIM framework reformulates this problem by replacing the discrete condition with a continuous, geometrically meaningful vector. The new objective is to learn the conditional distribution $P(H|c_{\text{geo}})$, where $c_{\text{geo}} \in \mathbb{R}^D$ is a vector of geometric parameters. For a typical wireless scenario, this vector can be defined as $c_{\text{geo}} = [x_{\text{user}}, y_{\text{user}}, z_{\text{user}}, \text{BS_id}]$, where $(x_{\text{user}}, y_{\text{user}}, z_{\text{user}})$ are the 3D Cartesian coordinates of the user equipment and BS_id is an identifier for the serving base station. This change fundamentally alters the model's task from memorizing a discrete set of channel distributions to learning a continuous function that maps spatial coordinates within an environment to their corresponding channel distributions.

Dataset: DeepMIMO 'O1' Scenario

The successful implementation of Geo-cDDIM is contingent on the availability of a dataset that provides both high-fidelity channel matrices and their associated geometric metadata. The DeepMIMO dataset is uniquely suited for this purpose.²⁰ Unlike datasets generated by stochastic channel models, DeepMIMO is constructed from accurate, physics-based ray-tracing simulations, ensuring that the channel data realistically captures the complex dependencies on the surrounding environment.

For this project, the 'O1' outdoor scenario from DeepMIMO will be utilized. This scenario provides a rich and complex environment ideal for training and evaluating the proposed model. The key parameters of the 'O1' scenario are as follows²⁰:

- **Environment:** An outdoor urban setting consisting of two intersecting streets, lined with buildings of varying heights. The main street is 600m long, and the second street is 440m long.
- **Base Stations (BS):** The scenario includes 18 active base stations positioned at a height of 6m along both sides of the streets.
- **Users:** A dense grid of over 1.1 million user locations is defined, providing extensive spatial coverage of the environment. The user locations are organized into three distinct grids with user-to-user spacing of 10 cm or 20 cm.
- **Frequency and Materials:** The simulation is conducted at a carrier frequency of 60 GHz. The materials for the ground and buildings (e.g., ITU dry earth, ITU layered drywall) are specified, ensuring realistic reflection and scattering properties.
- **Data Format:** The ray-tracing output for each BS-user link provides detailed channel parameters, including path gains, phases, delays, and angles of arrival/departure. The DeepMIMO generation code processes this raw data to construct the final complex channel matrices. Crucially, the dataset provides the precise (x, y, z) coordinates for

every user location.

For this project, the data will be parsed to create a set of training and validation samples. Each sample will consist of a pair: the complex channel matrix H for a link between a specific user and BS, and the corresponding conditioning vector c_{geo} containing the user's normalized coordinates and the BS identifier.

Model Adaptation: Engineering the Geo-cDDIM

The architectural modification required to transform the baseline cDDIM into the proposed Geo-cDDIM is targeted and low-complexity, making it highly feasible for rapid implementation. The cDDIM codebase, being derived from the minDiffusion repository for conditional MNIST generation, utilizes a standard approach for categorical conditioning.¹⁵ A discrete class label (the `user_id`) is passed through a `torch.nn.Embedding` layer, which converts the integer index into a dense vector embedding. This embedding is then added to the time step embedding and injected into the intermediate layers of the U-Net denoising network, guiding the generation process.

The engineering of Geo-cDDIM involves the following precise steps:

1. **Remove the Embedding Layer:** The `torch.nn.Embedding` layer, which is designed for discrete, categorical inputs, is removed from the model architecture.
2. **Insert a Conditioning MLP:** In its place, a small Multi-Layer Perceptron (MLP) is introduced. This MLP will serve as the conditioning network. Its input will be the continuous geometric vector c_{geo} . The MLP will consist of a few linear layers with non-linear activations (e.g., ReLU or SiLU).
3. **Map to Embedding Dimension:** The output of the conditioning MLP will be a vector of the same dimension as the original class embedding produced by the `nn.Embedding` layer. This ensures seamless integration with the rest of the U-Net architecture.
4. **Inject Conditioning:** The output vector from the MLP is then added to the time step embedding and fed into the U-Net at each step of the reverse diffusion process, exactly as the original class embedding was.

This modification is powerful because it replaces a simple lookup table (the embedding layer) with a function approximator (the MLP). The MLP learns the complex, non-linear mapping from a continuous spatial coordinate to the appropriate conditional guidance for the diffusion model. This targeted change fundamentally alters the model's capability from discrete memorization to continuous spatial interpolation and extrapolation, enabling true environmental generalization.

Positioning the Contribution

The Geo-cDDIM framework represents a significant and pragmatic advancement over existing

methods for generative channel modeling. Its position relative to the state-of-the-art is summarized in Table 1. The proposed model combines the high sample fidelity and training stability of diffusion models with a novel, continuous conditioning mechanism that enables spatial generalizability, a feature critically lacking in prior data-driven approaches. It achieves this without incurring the high implementation complexity and training challenges associated with models that rely on explicit, differentiable physics engines.

Table 1: Comparison of Conditional Generative Channel Models

Model	Generative Core	Conditioning Information	Generalizability	Key Limitation
ChannelGAN ¹⁰	GAN	Implicit (Learned from data distribution)	Low (to seen scenarios)	Training instability; no explicit control.
PBGC-VAE ¹⁸	VAE + Physics Model	Latent vector z	Moderate (to new physical parameters)	High implementation complexity; gradient flow issues.
cDDIM (Baseline)	Diffusion Model	User Index (Categorical)	None (to new users/locations)	Cannot generate channels for unseen locations.
Geo-cDDIM (Proposed)	Diffusion Model	Geometric Vector (e.g.,)	High (to any location within the environment)	Fidelity depends on density of training locations.

A Two-Week Plan for Prototyping and Evaluation

To ensure the feasibility of this research within the demanding two-week timeframe of the competition, a detailed, day-by-day project plan is outlined. This plan breaks the project into three distinct phases: (1) Data Curation and Model Re-engineering, (2) Training and Inference, and (3) Evaluation and Analysis.

Week 1: Data Curation and Model Re-engineering (Days 1-7)

The first week is dedicated to setting up the computational environment, acquiring and processing the necessary data, and implementing the core modifications to the generative model.

- **Day 1-2: Environment Setup & Data Acquisition:**
 - The first priority is to establish a functional and reproducible software

environment. This will be accomplished by creating a Conda environment using the environment.yml file provided in the public cDDIM repository.¹⁵ This step ensures all required dependencies, such as PyTorch and other scientific computing libraries, are correctly installed.

- Concurrently, the DeepMIMO 'O1' scenario dataset will be downloaded from the official DeepMIMO project website.²³ As this dataset is large, consisting of multiple .mat files containing the raw ray-tracing outputs, this download process should be initiated immediately to avoid delays.

- **Day 3-4: Data Parsing and Preprocessing:**

- A Python script will be developed to parse the downloaded DeepMIMO .mat files. This script will utilize the scipy.io.loadmat function to load the data into memory.
- The primary function of this script is to iterate through the vast number of user locations and base stations within the 'O1' scenario and extract the necessary data pairs: the complex channel matrix (H) and its corresponding geometric conditioning vector (cgeo).
- The conditioning vector will be constructed as [user_x,user_y,user_z,bs_id]. To ensure stable training, the coordinate values will be normalized to a standard range, such as [-1,1].
- The processed data pairs will be saved to disk in a format optimized for efficient loading by PyTorch, such as individual .pt files or a consolidated HDF5 file. A critical step in this phase is the creation of distinct training and validation data splits. To rigorously test the model's spatial generalization capability, the validation set will consist of a contiguous block of user locations that is entirely held out from the training process. This ensures that the model is evaluated on its ability to interpolate and generate channels for a geographic region it has never seen before.

- **Day 5-7: Model Implementation (Geo-cDDIM):**

- The official cDDIM repository will be forked to serve as the codebase for the project.¹⁵
- The core coding task will involve modifying the main training script (script_channel_ddim.py) and the underlying model definition.
- The architectural change detailed in Section 3.3 will be implemented: the nn.Embedding layer responsible for categorical conditioning will be replaced with a small MLP designed to process the continuous geometric conditioning vector.
- A custom PyTorch Dataset class and DataLoader will be implemented to efficiently load the preprocessed DeepMIMO data, feeding batches of (channel matrix, conditioning vector) pairs to the model during training.

Week 2: Training, Inference, and Analysis (Days 8-14)

The second week focuses on training the model, using it to generate synthetic data, and performing a rigorous quantitative and qualitative evaluation of the results.

- **Day 8-11: Model Training:**

- The training process for the Geo-cDDIM model will be initiated on a machine with a high-performance GPU. The training loop logic from the original cDDIM script can be largely reused, with modifications to handle the new data loading and conditioning mechanism.
- Training progress will be closely monitored using a tool like TensorBoard or Weights & Biases to track the loss curve and other relevant metrics. Given the limited timeframe, extensive hyperparameter optimization is not feasible. The primary goal is to achieve stable convergence and train the model for a sufficient number of epochs to learn the complex mapping from geometry to channel structure. Key parameters such as the learning rate and batch size will be fine-tuned as necessary based on initial training dynamics.

- **Day 12-13: Generation and Evaluation:**

- Once the model is trained, it will be used for inference. The ultimate test of the proposed framework is its ability to generate high-fidelity channels for the held-out validation set of locations. This will be done by feeding the model the geometric conditioning vectors from the validation set and generating the corresponding synthetic channel matrices.
- **Quantitative Evaluation:** The statistical properties of the generated channels will be compared against the ground-truth channels from the DeepMIMO validation set. To ensure a rigorous comparison, metrics cited in related state-of-the-art literature will be employed. These include distribution-level metrics such as the **2-Wasserstein distance** and **Maximum Mean Discrepancy (MMD)**, which were used in the evaluation of the physics-informed generative model and provide a robust measure of similarity between two distributions of high-dimensional data.¹⁶ Additionally, summary statistics of the channel matrices, such as the distribution of singular values, power delay profiles, or delay spreads, will be compared.²⁴
- **Qualitative Evaluation:** To provide an intuitive understanding of the generation quality, visualizations will be created. For a few sample locations from the validation set, key channel characteristics like the channel power delay profile or the angular power spectrum will be plotted for both the real and the generated channels, allowing for a direct visual comparison.

- **Day 14: Finalization:**

- The final day will be dedicated to compiling all results, tables, and figures into a coherent narrative. The abstract, introduction, methodology, results, and conclusion sections of the final research paper will be written.
- The code will be cleaned, commented, and organized. A comprehensive README.md file will be created to explain the project setup, data preprocessing steps, and how to train the model and run the evaluation scripts, ensuring the work is reproducible for submission.

Anticipated Contributions and Future Directions

This research is poised to make a significant contribution to the field of wireless communications by introducing a novel, generalizable framework for generative channel modeling. Beyond the immediate results, this work opens up several promising avenues for future investigation and application.

Primary Contribution: A Generalizable Channel Digital Twin

The core contribution of this project is the development of a framework capable of generating high-fidelity, site-specific wireless channels for *any arbitrary location* within a learned environment. This moves beyond the paradigm of user-specific data augmentation, where models are limited to generating data for a fixed set of pre-registered users. The Geo-cDDIM learns the continuous relationship between physical space and channel characteristics, effectively creating a functional "digital twin" of the wireless propagation environment. The broader implication of this contribution is substantial. A generative digital twin can revolutionize how wireless networks are designed, tested, and optimized. Network operators could use such a tool to perform virtual site surveys, predicting network coverage and performance in a new deployment area without conducting costly and time-consuming physical measurements. Algorithm developers could test and validate novel communication protocols, such as location-aware beamforming or mobility management schemes, in a realistic, simulated environment. This tool democratizes access to large-scale, site-specific channel data, accelerating the research and development cycle for next-generation wireless systems.

Future Work

The proposed Geo-cDDIM framework serves as a strong foundation for numerous future research directions.

- **Richer Conditioning:** The current proposal conditions the model on basic geometric information (3D coordinates and BS ID). A natural extension is to incorporate richer environmental parameters available in datasets like DeepMIMO. This could include line-of-sight (LoS)/non-line-of-sight (NLoS) flags, information about building materials, antenna orientations, or even atmospheric conditions. Adding this information to the conditioning vector c_{geo} could further improve the fidelity and realism of the generated channels.
- **Application to Emerging Technologies:** The scarcity of high-quality channel data is even more acute for emerging wireless technologies. The Geo-cDDIM framework could be adapted and applied to generate channel datasets for systems where real data is

exceptionally difficult or expensive to obtain. This includes channels for Reconfigurable Intelligent Surfaces (RIS)²⁵, which involve complex interactions with programmable metasurfaces, and Terahertz (THz) communications²⁶, which operate at extremely high frequencies with unique propagation characteristics.

- **Inference Acceleration:** A well-known limitation of diffusion models is their slow iterative sampling process, which can be a bottleneck for real-time applications.²⁷ Future work could focus on integrating faster sampling techniques, such as Denoising Diffusion Implicit Models (DDIM), consistency models, or one-step generation methods, with the Geo-cDDIM framework. Accelerating the inference speed would be a critical step towards making the channel digital twin usable for applications that require low-latency channel generation, such as real-time network optimization or dynamic spectrum access.
- **Transformer-based Architectures:** The U-Net architecture has been the de facto standard for the denoising network in diffusion models. However, recent research has demonstrated the significant potential of Transformer architectures for various tasks in wireless communications, owing to their ability to capture long-range dependencies via self-attention mechanisms.²⁸ An exciting avenue for future research would be to replace the convolutional U-Net backbone of the Geo-cDDIM with a Transformer-based architecture. This could potentially improve the model's ability to learn complex spatial correlations between different locations in the environment, leading to even more accurate and globally consistent channel generation.

Works cited

1. Applications of Generative AI (GAI) for Mobile and Wireless Networking: A Survey - arXiv, accessed August 19, 2025, <https://arxiv.org/html/2405.20024v2>
2. AI-enhanced channel estimation and signal processing for MIMO systems in 5G/6G radio frequency networks - Global Journal of Engineering and Technology Advances, accessed August 19, 2025, <https://gjeta.com/sites/default/files/GJETA-2024-0253.pdf>
3. A Survey of Artificial Intelligence Enabled Channel Estimation Methods: Recent Advance, Performance, and Outlook - ResearchGate, accessed August 19, 2025, https://www.researchgate.net/publication/390704827_A_Survey_of_Artificial_Intelligence_Enabled_Channel_Estimation_Methods_Recent_Advance_Performance_and_Outlook
4. [2503.05988] Physics-Informed Generative Approaches for Wireless Channel Modeling, accessed August 19, 2025, <https://arxiv.org/abs/2503.05988>
5. COST CA20120 INTERACT Framework of Artificial Intelligence Based Channel Modeling | Request PDF - ResearchGate, accessed August 19, 2025, https://www.researchgate.net/publication/385921286_COST_CA20120_INTERACT_Framework_of_Artificial_Intelligence_Based_Channel_Modeling
6. Applications of Generative AI (GAI) for Mobile and Wireless Networking: A Survey, accessed August 19, 2025, https://www.researchgate.net/publication/381006048_Applications_of_Generativ

[e_AI_GAI_for_Mobile_and_Wireless_Networking_A_Survey](#)

7. Generative AI at DeepSig and its Role in Shaping the Future of Wireless Communications, accessed August 19, 2025, <https://www.deepsig.ai/generative-ai-at-deepsig-and-its-role-in-shaping-the-future-of-wireless-communications/>
8. End-to-end wireless communication architecture consists of the VAE and classifier DNNs., accessed August 19, 2025, https://www.researchgate.net/figure/End-to-end-wireless-communication-architecture-consists-of-the-VAE-and-classifier-DNNs_fig4_365811541
9. Generative adversarial networks for generating synthetic features for Wi-Fi signal quality - PMC - PubMed Central, accessed August 19, 2025, <https://pmc.ncbi.nlm.nih.gov/articles/PMC8610258/>
10. (PDF) Attention-Guided Wireless Channel Modeling and Generating - ResearchGate, accessed August 19, 2025, https://www.researchgate.net/publication/389800154_Attention-Guided_Wireless_Channel_Modeling_and_Generating
11. Attention-Guided Wireless Channel Modeling and Generating - MDPI, accessed August 19, 2025, <https://www.mdpi.com/2076-3417/15/6/3058>
12. Diffusion Models for Future Networks and Communications: A Comprehensive Survey | Request PDF - ResearchGate, accessed August 19, 2025, https://www.researchgate.net/publication/394293667_Diffusion_Models_for_Future_Networks_and_Communications_A_Comprehensive_Survey
13. [2508.01586] Diffusion Models for Future Networks and Communications: A Comprehensive Survey - arXiv, accessed August 19, 2025, <https://www.arxiv.org/abs/2508.01586>
14. Diffusion Models for Future Networks and Communications: A Comprehensive Survey - arXiv, accessed August 19, 2025, <https://arxiv.org/html/2508.01586v1>
15. taekyun/cDDIM - GitHub, accessed August 19, 2025, <https://github.com/taekyun/cDDIM>
16. Physics-Informed Generative Approaches for Wireless Channel Modeling - ResearchGate, accessed August 19, 2025, https://www.researchgate.net/publication/389715027_Physics-Informed_Generative_Approaches_for_Wireless_Channel_Modeling
17. Physics-Informed Generative Approaches for Wireless Channel Modeling - arXiv, accessed August 19, 2025, <https://arxiv.org/html/2503.05988v3>
18. PHYSICS-INFORMED GENERATIVE APPROACHES FOR WIRELESS CHANNEL MODELING - OpenReview, accessed August 19, 2025, <https://openreview.net/pdf?id=CSkBnuaMkC>
19. PHYSICS-INFORMED GENERATIVE APPROACHES FOR WIRELESS CHANNEL MODELING | OpenReview, accessed August 19, 2025, <https://openreview.net/forum?id=CSkBnuaMkC>
20. (PDF) DeepMIMO: A Generic Deep Learning Dataset for Millimeter ..., accessed August 19, 2025, https://www.researchgate.net/publication/331195696_DeepMIMO_A_Generic_Deep_Learning_Dataset_for_Millimeter_Wave_and_Massive_MIMO_Applications

21. Deep MIMO Data Set | PDF | Mimo | Machine Learning - Scribd, accessed August 19, 2025, <https://www.scribd.com/document/419552639/Deep-MIMO-data-set>
22. DeepMIMO - OpenDataLab, accessed August 19, 2025, <https://opendatalab.com/OpenDataLab/DeepMIMO>
23. Use DeepMIMO dataset to generate samples for wireless power allocation - GitHub, accessed August 19, 2025, <https://github.com/Haoran-S/DeepMIMO>
24. Partial source code for paper "GAN-based massive MIMO channel model trained on measured data" - GitHub, accessed August 19, 2025, <https://github.com/Jeija/GAN-Wireless-Channel-Model>
25. [2211.13487] Sensing Aided Reconfigurable Intelligent Surfaces for 3GPP 5G Transparent Operation - arXiv, accessed August 19, 2025, <https://arxiv.org/abs/2211.13487>
26. Wireless Channel Modeling and Prediction Technology Based on Deep Learning, accessed August 19, 2025, https://www.researchgate.net/publication/384646975_Wireless_Channel_Modeling_and_Prediction_Technology_Based_on_Deep_Learning
27. Faster Diffusion: Rethinking the Role of the Encoder for Diffusion Model Inference | OpenReview, accessed August 19, 2025, [https://openreview.net/forum?id=ca2mABGV6p&referrer=%5Bthe%20profile%20of%20jian%20Yang%5D\(%2Fprofile%3Fid%3D~jian_Yang14\)](https://openreview.net/forum?id=ca2mABGV6p&referrer=%5Bthe%20profile%20of%20jian%20Yang%5D(%2Fprofile%3Fid%3D~jian_Yang14))
28. Transformer Based Next Generation Wireless Communication Receiver | LUP Student Papers - Lund University Publications, accessed August 19, 2025, <https://lup.lub.lu.se/student-papers/search/publication/9198959>
29. Transformer-Driven Neural Beamforming with Imperfect CSI in Urban Macro Wireless Channels - arXiv, accessed August 19, 2025, <https://arxiv.org/pdf/2504.11667>