



Department of Electrical and Computer Engineering

電機與計算機工程系

Site-Specific Learning via Digital Twins: The Next-Generation AI for Wireless

Dr. Yuanwei Liu

Professor, IEEE Fellow, AAIA Fellow,
Web of Science Highly Cited Researcher,
IEEE Communication Society Distinguished Lecturer,
IEEE Vehicular Technology Distinguished Lecturer,

e-mail: yuanwei@hku.hk

website: <https://www.eee.hku.hk/~yuanwei/>



THE UNIVERSITY OF HONG KONG 香港大學 • FACULTY OF ENGINEERING 港大工程學院



Outline

- ❑ Background: From AI-native 6G to Site-Specific Learning
- ❑ Fundamentals of Site-Specific Beamforming
- ❑ Case Studies
 - ❑ SSBF via diffusion
 - ❑ Site-specific channel probing codebook
 - ❑ SSBF via flow matching
 - ❑ SSBF for high-accuracy type-II feedback
- ❑ Future Directions



Outline

- ❑ **Background: From AI-native 6G to Site-Specific Learning**
- ❑ Fundamentals of Site-Specific Beamforming
- ❑ Case Studies
 - ❑ SSBF via diffusion
 - ❑ Site-specific channel probing codebook
 - ❑ SSBF via flow matching
 - ❑ SSBF for high-accuracy type-II feedback
- ❑ Future Directions



□ 6G is envisioned to be AI-native

Nvidia and Nokia Team Up

Nvidia and Nokia have announced a **US\$1 billion** partnership to develop **AI-native 6G** technology, leveraging the **Nvidia Aerial RAN Computer** platform with **GPU-accelerated** processing.



<https://investor.nvidia.com/news/press-release-details/2025/NVIDIA-and-Nokia-to-Pioneer-the-AI-Platform-for-6G--Powering-Americas-Return-to-Telecommunications-Leadership/default.aspx>

6G and AI

6G Base Station is named "aNB"

In the 131st RAN3 meeting held by **3GPP** in Gothenburg, the official name of the 6G base station/RAN node has been determined as **aNB**, where **"a" stands for "AI" and "Advanced"**.



<https://mp.weixin.qq.com/s/eRUOv-4Txea2lf3PUwZp2A>



6G and AI

□ 6G is envisioned to be AI-native

AI-for-RAN

AI for the
enhancement of
RAN



Spectral Efficiency

AI-and-RAN

AI and RAN
sharing the same
infrastructure



Asset Utilization

AI-on-RAN

AI applications
enabled by RAN



New Applications



6G and AI

□ 6G is envisioned to be AI-native

AI-for-RAN

Neural Transceiver,
Site-Specific
Learning, etc.



Spectral Efficiency

AI-and-RAN


Mobile Edge
Generation, etc.



Asset Utilization

AI-on-RAN

Self-driving,
Embodied AI, etc.



New Applications



6G and AI

□ 6G is envisioned to be AI-native

AI-for-RAN

Neural Transceiver,
Site-Specific Learning, etc.



Spectral Efficiency

AI-and-RAN


Mobile Edge
Generation, etc.



Asset Utilization

AI-on-RAN

Self-driving,
Embodied AI, etc.





New Applications



Wireless Digital Twin

❑ Classical Digital Twin - Ray Tracing

- Wireless propagation is **extremely complex**
- Requires detailed **site geometry** and **dielectric properties** --- *hard to model!*
 - Easier than Maxwell (geometry vs. differential equations)
- Commercial software often used
 - closed-source, costly, computationally expensive



Wireless InSite[®]



Wireless Digital Twin

□ Digital Twin is now accelerating with AI

NVIDIA Sionna RT


NVIDIA's platform simulated 5G coverage across the US in **under five minutes** by tracing more than **35 trillion rays** on 96 NVIDIA L40S GPUs, a task computationally infeasible for traditional methods.



<https://developer.nvidia.com/blog/powering-ai-native-6g-research-with-the-nvidia-sionna-research-kit/?ncid=so-link-383155-vt26&linkId=100000388985859>

Northeastern Uni. & VIAVI AIRMap

AIRMap generated site-specific path-gain radio maps in **~4 ms** per inference on an NVIDIA L40S GPU, achieving **7000x speedup** compared with GPU-accelerated ray-tracing radio-map simulation.



<https://arxiv.org/abs/2511.05522>



Department of Electrical and
Computer Engineering

電機與計算機工程系

Wireless Digital Twin

□ HKU Wireless Digital Twin Platform






□ What can we do with Digital Twin?

Site-Specific Learning for Wireless

From Blind Estimation to Knowledge-Informed Prediction

- **Old Way:** "**One-size-fits-all**" algorithms that treat every wireless channel as a **random statistical** event.
- **New Way:** "**Site-specific**" paradigms where AI models are **tailored** to the unique geometry and physics of the local environment.

Site-Specific Learning





Site-Specific Learning

□ Why site-specific learning is useful?

Conventional Learning

Training Data: Massive, generalized datasets (but often limited and biased in practice)

Core Assumption: **"One-size-fits-all"**; The model will generalize to any new environment.

Handling of Channel:

1. **Learns statistical averages;**
2. Cannot break through conventional designs.

Real-World Performance: Accuracy drops significantly in new environments.

Overhead: High upfront training cost; low adaptation cost

Site-Specific Learning

Training Data: Small-scale, local datasets (much easier to obtain, e.g., via digital twins)

Core Assumption: **"Every site is unique"**; The model learns the site-specific RF fingerprint.

Handling of Channel:

1. **Learns the exact physical layout;**
2. Enable environment-aware predictive design.

Real-World Performance: High accuracy; perfectly tailored to local features.

Overhead: Lower initial training, but requires continuous data collection for the site changes.




Site-Specific Learning

□ Why site-specific learning is useful?

Optimal BS settings vary site to site ➡ *Site-Specific BS Tuning*

- Many BS parameters can be configured: bandwidths, subcarrier spacings, antenna configurations, handover thresholds, etc.
- These parameters depends on **site-specific** propagation environment, user locations, density, mobility, etc.



Higher channel estimation overhead ➡ *Site-Specific Beamforming/Localization*

- **Large-dimensional channel:** 64 antennas (5G) → 512+ antennas (6G)
- **High user density:** $10^6 \rightarrow 10^7 \sim 10^8$ devices/km²
- **Stringent latency requirement:** low-altitude networks, high-speed train

And more...




Outline

- ❑ Background: From AI-native 6G to Site-Specific Learning
- ❑ **Fundamentals of Site-Specific Beamforming**
- ❑ Case Studies
 - ❑ Beam-Brainstorm: SSBF via diffusion
 - ❑ Site-specific channel probing codebook
 - ❑ SSBF via flow matching
 - ❑ SSBF for high-accuracy type-II feedback
- ❑ Future Directions

Beamforming

□ Beamforming in wireless communications



Use **antenna arrays** to

- Focus most energy at users (**beamforming gain**)
- Combat multi-path fading (**diversity gain**)
- Serve multiple users (**multiplexing gain**)


The reality: **Channel estimation must be accurate!**

- Computational complexity is not the key burden with more advanced chips
- **Time overhead is the key burden**
- Can take **8~15%** time in typically 5G MU-MIMO, which increases with more antennas and users



Beamforming

□ A Typical Working Flow of Beamforming



Stage I: Channel Estimation/Sensing

Acquire channel information from environment
(Pilot-based / Codebook-based)



Stage II: Beam Forming/Selection


Optimize/select the best beam for UEs
(Optimization-based / Codebook-based)



Beamforming

❑ Existing channel estimation method – Pilot-based

Pilot-based



Uplink channel
estimation at BS

$$\underset{\text{Receive signal at BS}}{y_{UL}(t)} = \overset{\text{Channel to estimate}}{\mathbf{h}} \underset{\text{Pilot reference signal}}{s_{\text{ref}}(t)} + \overset{\text{Noise}}{\mathbf{n}(t)}$$

- Rely on **uplink-downlink reciprocity**, applied only to **TDD** systems
- **Theoretical perfect channel estimation** with sufficient pilot length $t \rightarrow$ High performance limit.
- **High computational and time overhead**, especially with large number of antenna arrays and users.



Beamforming


❑ Existing channel estimation method – Codebook-based

Channel

$$y_{\text{DL},n}(t) = \mathbf{h}^H \mathbf{w}_n s_{\text{ref}}(t) + \mathbf{n}(t)$$

Receive
signal at user

Beam n in
codebook



- Sweeping the entire codebook; User selects the best one(s) **based on the power of receive signals**.
- Apply to both **TDD** and **FDD** systems.
- Lower time overhead than the pilot method, but can still be high in practice due to the **need for very large codebooks**.
- **Low DoFs**, due to the usage of discrete codebooks.
- **Hard to handle multi-path propagation**, since each beam typically exhibits a single mainlobe.



Beamforming

❑ Challenges in future wireless networks (e.g., Low-altitude network)

Highly Dynamic 3D Channels

High-DoF user trajectories combined with urban environments (high rises, street canyons) create rapidly changing signal paths.

Need High-DoF

Beamforming

Frequent Linkage & Loss

The 3D environment (buildings, cranes, foliage) causes sudden, unpredictable signal occlusions.

Beam Misalignment

Highly dynamic user motions are essential but easily misaligned by the user motion.


Need Low Channel

Estimation

Strict Low-Latency Budgets

Critical control & command and collision avoidance signals cannot be delayed.

Overhead





Low-Altitude Beamforming

□ Revisit Working Flow of Beamforming



Stage I: Channel Estimation/Sensing

**As low-overhead as possible
to reduce latency**



Stage II: Beam Forming/Selection


**As accurate as possible
to guarantee performance
in complex environment**

This is a dilemma !!



Site-Specific Beamforming

□ Site-Specific Beamforming with Digital Twin



- Wireless channel is **not as random as we imagine.**
- Many environment objects (e.g., building geometry and persistent scatterer) change very slowly.

- **Prior** site-specific knowledge reduces the sensing overhead
- Use **AI model** to learn implicit site-specific propagation pattern


Z. Wang, Z. Zhou, C. J. Zhao, and Y. Liu, "Generative Site-Specific Beamforming for Next-Generation Spatial Intelligence," *arXiv preprint arXiv:2601.02301*, Jan. 2026



Site-Specific Beamforming



□ Components of Site-Specific Beamforming

Codebook-based/ Multimodality Channel Sensing




Provide input
"Wireless
Prompt"

Beamforming Inference Network



Provide dataset

Site-Specific Environment Database



- Use **small-size** DFT codebook or site-specific optimized codebook; user feedback signal powers of each beam
- **Ultra-low** time overhead
- **Very coarse** channel information
- Potential **multi-modal** fusion (e.g., assistant with GPS, camera, LiDAR)
- Learn **site-specific propagation feature** from database.
- Map **very coarse channel information** to **accurate beamformers**
- Generate **multipath-aware** beams
- Mostly from **digital twin**
- Continuously updated and calibrated by **real-world data**


Z. Wang, Z. Zhou, C. J. Zhao, and Y. Liu, "Generative Site-Specific Beamforming for Next-Generation Spatial Intelligence," *arXiv preprint arXiv:2601.02301*, Jan. 2026



Site-Specific Beamforming


□ Advantages of Site-Specific Beamforming

Codebook-based/ Multimodality
Channel Sensing




Provide input
"Wireless
Prompt"

Beamforming
Inference Network




Provide dataset

Site-Specific
Environment Database



Low-overhead codebook-
based channel sensing

Apply to both **TDD**
and **FDD** systems



Environment-aware
high-DoF beamforming, adapting
to multi-path propagation


1. Z. Wang, Z. Zhou, C. J. Zhao, and **Y. Liu**, "Generative Site-Specific Beamforming for Next-Generation Spatial Intelligence," *arXiv preprint arXiv:2601.02301*, Jan. 2026
2. Y. Heng, Y. Zhang, A. Alkhateeb and J. G. Andrews, "Site-Specific Beam Alignment in 6G via Deep Learning," *IEEE Communications Magazine*, vol. 62, no. 8, pp. 162-168, August 2024



Site-Specific Beamforming

❑ Challenges – Multimodal Prediction

Multimodality: one observation implies multiple plausible beams.







Site-Specific Beamforming

❑ Challenges – Multimodal Prediction

Site-specific beamforming is essentially to learn a **multimodal conditional distribution**



multimodal conditional distribution



- A single coarse observation is mapped to an infinite number of beams with different probabilities.
- Some beams have peak probabilities, which are of our interest.




Site-Specific Beamforming

❑ Challenges – Multimodal Prediction


Conventional Discriminative Models with Point Output

Regression Model



- **Averaging** the whole distribution into a single point
- May lead to a **completely misaligned beam**

Classification Model



- **Discretized** distribution into a set of labelled classes, and predict the best class
- We cannot set good labels in advance since the distribution is unknown




Site-Specific Beamforming

□ Generative Site-Specific Beamforming (GenSSBF)

Generative models are needed

Generative Model



- Learn the **entire distribution**
- Sampling the beams with high probability from the distribution

Generate-and-Select Procedure

$$\{w_1, w_2, \dots, w_Q\} \sim p_\theta(w|x)$$

- **Sampling multiple beams** from the learned distribution to form a **new codebook**
- User select **the best one**



Convert unavoidable ambiguity into a controlled exploration step.



Site-Specific Beamforming


□ Compatibility with 3GPP

Synchronization Signal Block

1. Used for initial access
2. Broast to all users
3. Send periodically
4. **Send channel sensing beams**

Reference Signal Received Power

1. **Low-dimensional
coarse observation**



Channel State Information Reference Signal

1. Used for fine-tuning the link
2. Tailored for individual users
3. **Send generated beams**



Outline

- ❑ Background: From AI-native 6G to Site-Specific Learning
- ❑ Fundamentals of Site-Specific Beamforming
- ❑ **Case Studies**
 - ❑ SSBF via diffusion
 - ❑ Site-specific channel probing codebook
 - ❑ SSBF via flow matching
 - ❑ SSBF for high-accuracy type-II feedback
- ❑ Future Directions



Outline


- ❑ Background: From AI-native 6G to Site-Specific Learning
- ❑ Fundamentals of Site-Specific Beamforming
- ❑ **Case Studies**
 - ❑ **SSBF via diffusion**
 - ❑ Site-specific channel probing codebook
 - ❑ SSBF via flow matching
 - ❑ SSBF for high-accuracy type-II feedback
- ❑ Future Directions



Beam-Brainstorm: SSBF via diffusion

□ A General GenSSBF Framework

1. The element of the beamforming vector should be carefully designed;
2. Practical constraints should be considered during generation;
3. Quality vs Diversity vs Speed vs Overhead (4D trade-off)





The Principle of BBS

□ The Cornerstone of BBS


beam generator f

$$\begin{aligned} \max_f \quad & |\mathbf{h}_k^H f(\mathbf{c}_k)|^2 \\ \text{s.t.} \quad & \|f(\mathbf{c}_k)\|^2 = 1, \\ & f(\mathbf{c}_k) \in \mathcal{S}, \end{aligned}$$

Particle CSI observations

Forward diffusion

$$q(\mathbf{X}_t^k | \mathbf{X}_0^k)$$



$$p_g(\mathbf{X}_{t-1}^k | \mathbf{X}_t^k)$$

Backward denoise

Forward diffusion:

$$\mathbf{X}_T^k = \sqrt{\prod_{t=1}^T (1 - \beta_t)} \mathbf{X}_0^k + \sqrt{1 - \prod_{t=1}^T (1 - \beta_t)} \mathbf{Z}_{0 \rightarrow T},$$

Backward Beam Generation:

$$\mathbf{U}'_t = \frac{1}{\sqrt{1 - \beta_t}} \left(\mathbf{X}_t^k - \frac{\beta_t}{\sqrt{1 - \prod_{i=1}^t (1 - \beta_i)}} \epsilon_\vartheta(\mathbf{X}_t^k, \mathbf{c}_k, t) \right)$$

Sweeping the probing beams
from DFT codebook




BS



UE

Diffusion model



RSRP as
wireless
prompts


Cleaner beam



The Performance of BBS


□ Numerical Results of BBS

I2_28B, DeepMIMO




(a) Indoor Blockage

O1B_28, DeepMIMO




(b) Outdoor Blockage


Boston5G_28, DeepMIMO



(c) City




(a) I2_28B, Q=15




(b) I2_28B, Optimal

denoising →




(c) I2_28B, Q=32



(d) I2_28B, 64-beam DFT

- The generated beam has **multiple main lobes** compared to conventional DFT beams, thus effectively exploiting **multipath propagation**.






generated

optimal



Performance of BBS

□ Beamforming Gain Performance of BBS






The proposed model **with mild brainstorm** (red/blue lines) can achieve best beamforming gain while significantly reducing the beam sweeping overhead. This is because for the trained generative model, its parameters constitute a latent **beam space** capable of generating (without explicit storage) a virtually unlimited set of high-quality and site-specific beams on-demand for any UE. Consequently, utilizing a set of such customized beams for sweeping typically leads to a superior performance.



Performance of BBS

□ Performance Under Imperfect Wireless Prompts



- Mild brainstorm is also beneficial to **improve resistance to noise**;
- At low SNR case, especially when the noise power exceeds the received power, the useful information within the vector is dominated by noise, resulting in performance degradation




Outline

- ❑ Background: From AI-native 6G to Site-Specific Learning
- ❑ Fundamentals of Site-Specific Beamforming
- ❑ **Case Studies**
 - ❑ SSBF via diffusion
 - ❑ **Site-specific channel probing codebook**
 - ❑ **SSBF via flow matching**
 - ❑ SSBF for high-accuracy type-II feedback
- ❑ Future Directions




Probing codebook design & SSBF via flow matching

□ Motivation



Two ways of improving SSBF performance

- Make RSRPs **include more CSI information** $\Rightarrow$ **Codebook design**



More informative RSRPs lead to **more concentrated distribution**;
Easier to learn

- Make AI model to learn the residual uncertainty with **low latency and overhead**
 $\Rightarrow$ **Conditional flow matching (CFM)**



Information Maximizing Codebook

□ An Information-Theoretical Perspective

Objective: Site-Information-Maximizing

Maximizing the **mutual information between RSRP measurements and beamformer**, i.e., channel state information.


Constraints:

(i) **Orthogonal beams:** Avoid redundant measurements

$$\mathbf{G}_C = \mathbf{I}_K \quad [\mathbf{G}_C]_{i,j} = \frac{\mathbf{c}_i^H \mathbf{c}_j}{\|\mathbf{c}_i\| \|\mathbf{c}_j\|}$$

(i) **Coverage guarantee: codebook must cover the whole site region**

$$\max_k y_k \geq y_{\text{th}} \quad \text{for all users}$$



RSRP Channel

$$\max_{\mathbf{C}} I(\mathbf{y}; \mathbf{h}) = h(\mathbf{y}) - h(\mathbf{y}|\mathbf{h})$$

$$\begin{aligned} h(\mathbf{y}|\mathbf{h}) &= \frac{1}{2} E_{\mathbf{h}} \left[\log \left((2\pi e)^K \det \Sigma_y(\mathbf{h}, \mathbf{C}) \right) \right] \\ &= \frac{1}{2} E_{\mathbf{h}} \left[\sum_{k=1}^K \frac{1}{2} \log \left(2\pi e \sigma_{y, \mathbf{c}_k}^2(\mathbf{h}) \right) \right] \end{aligned}$$

$$h(\mathbf{y}) \leq \frac{1}{2} \log((2\pi e)^K \det \mathbf{R}_y)$$

$$I(\mathbf{y}; \mathbf{h}) \leq \frac{1}{2} \log \det \mathbf{R}_y - I_c$$

$$\max_{\mathbf{C}} \log \det \mathbf{R}_y$$

Maximizing a more trackable upper bound



Information Maximizing Codebook

□ Proposed Solution

Original problem:

$$\begin{aligned} \max_{\mathbf{C}} \quad & \log \det \mathbf{R}_y \\ \text{s.t.} \quad & \mathbf{G}_{\mathbf{C}} = \mathbf{I}_K, \max_k y_k \geq y_{\text{th}}, |\mathbf{C}|_{i,j} = 1, \forall i, j, k \end{aligned}$$

Penalized version:

$$\max_{\mathbf{C}} \log \det \mathbf{R}_y^{(N)} - \lambda_1 L_{\text{orth}} - \lambda_2 \hat{L}_{\text{cov}}$$

$$L_{\text{orth}} = \|\hat{\mathbf{C}}^H \hat{\mathbf{C}} - \mathbf{I}_K\|_F^2 \quad \hat{L}_{\text{cov}} = \sum_{k=1}^K e^{\beta y_k}$$

- **Directly acquiring $\mathbf{R}_y$ is intractable**

Approximate $\mathbf{R}_y$ by the samples from dataset:

$$\hat{\mathbf{R}}_y^{(N)} \leftarrow \sum_{j=1}^N (\mathbf{y}_j - \bar{\mathbf{y}}^{(N)})(\mathbf{y}_j - \bar{\mathbf{y}}^{(N)})^T$$

- **Constraints are intractable**

$$\mathbf{G}_{\mathbf{C}} = \mathbf{I}_K \rightarrow \|\hat{\mathbf{C}}^H \hat{\mathbf{C}} - \mathbf{I}_K\|_F^2$$



$$\max_k y_k \geq y_{\text{th}} \rightarrow \left[y_{\text{th}} - \frac{1}{\beta} \log \sum_{k=1}^K e^{\beta y_k} \right]_+$$

Penalized version is equivalent
to the original if $N, \lambda_1, \lambda_2 \rightarrow \infty$.



SSBF via Flow Matching

□ From Diffusion to Flow Matching



Flow Matching provides the fastest and most smooth mapping from the source to the target distribution




CFM-based Beamforming

□ Training Target with RSRP Condition and Linear Path

$$L_{CFM}(\theta) = E_{\mathbf{z}_0, \mathbf{z}_1, \mathbf{y}, t} [\|\mathbf{v}_\theta(\mathbf{z}_t, t; \mathbf{y}) - (\mathbf{z}_1 - \mathbf{z}_0)\|^2]$$

Conditional distribution of
beamformer $p(\mathbf{w}|\mathbf{y})$ is
captured by the CFM

Inferring the CFM corresponds to sampling $p(\mathbf{w}|\mathbf{y})$
Selecting the strongest candidate approximates MAP




C. J. Zhao, Z. Wang, and Y. Liu,
"Generative Site-Specific Beamforming via
Information-Maximizing Codebook," *arXiv
preprint arXiv:2602.12552*, Feb. 2026




Simulation Results

□ Optimized Codebook



O1_28, DeepMIMO



Site-specific channel probing codebooks





O1B_28, DeepMIMO





Simulation Results

□ Performance v.s. Transmit Power




- **With the proposed codebook**, the conventional Discriminative Model is enhanced, even outperforming the the Generative Model without brainstorm ($M=1$).
- **With the brainstorm ($M>1$)**, the Generative Model exhibit much better performance



Simulation Results

□ Performance v.s. Codebook Size



Model Advantage

Achieves near-optimal performance outperforming the conventional PMI-based method with significantly less overhead

Codebook Advantage

Improves performance of all methods in terms of both beamforming gain and overhead





SSBF via Mean-Flow (F-BBS)

□ Flexible One or Few-step Beam Generation

Diffusion Transformer

**Flow field matching still requires many generation steps.
Can we further accelerate beam generation?**



*The core idea for
fast beam
generation is from
instantaneous
velocity to average
velocity!*



SSBF via Mean-Flow (F-BBS)

Flexible One or Few-step Beam Generation

**Flow field matching still requires many generation steps.
Can we further accelerate beam generation?**

Bridge the relationship between instantaneous velocity and average velocity:


$$\mathbf{u}(\mathbf{X}_r, r, t) \triangleq \frac{\mathbf{X}_t - \mathbf{X}_r}{t - r} = \frac{1}{t - r} \int_r^t \mathbf{v}(\mathbf{X}_\tau, \tau) d\tau, \quad 0 \leq r < t \leq 1.$$

Construct the consistency for average velocity:

$$(t - r)\mathbf{u}(\mathbf{X}_r, r, t) = (s - r)\mathbf{u}(\mathbf{X}_r, r, s) + (t - s)\mathbf{u}(\mathbf{X}_s, s, t).$$


Build a accessible loss function to learn average velocity:

$$\mathcal{L}_{av}(\vartheta) = \mathbb{E}_{r,s,t,\mathbf{X}_0,\mathbf{X}_1} [\|\mathbf{u}^\vartheta(\mathbf{X}_r, r, t) - \text{sg}[\bar{\mathbf{u}}(\mathbf{X}_r, r, t)]\|^2],$$





Flexible One or Few-step Beam Generation






Z. Zhou, Z. Wang, and Y. Liu, "Fast Beam-Brainstorm: Few-Step Generative Site-Specific Beamforming with Flexible Probing," *arXiv preprint arXiv:2603.17622*, Apr. 2026



The Performance of F-BBS

□ Numerical Results of F-BBS

HKU






Z. Zhou, Z. Wang, and **Y. Liu**, "Fast Beam-Brainstorm: Few-Step Generative Site-Specific Beamforming with Flexible Probing," *arXiv preprint arXiv:2603.17622*, Apr. 2026



The Performance of F-BBS

□ Numerical Results of F-BBS



- F-BBS can achieve **better Overhead-Gain tradeoff**;
- Compared with diffusion-based beam generation or even standard flow matching algorithm, F-BBS can effectively **lower the number of inference** during beam generation;
- F-BBS demonstrated good **extrapolation** performance across different wireless prompt lengths.



Outline

- ❑ Background: From AI-native 6G to Site-Specific Learning
- ❑ Fundamentals of Site-Specific Beamforming
- ❑ **Case Studies**
 - ❑ SSBF via diffusion
 - ❑ Site-specific channel probing codebook
 - ❑ SSBF via flow matching
 - ❑ **SSBF with high-accuracy type-II feedback**
- ❑ Future Directions



Site-Specific Type-II Codebook

□ Motivation

- Previous method generate **all beam parameters at the BS**
 - a) Uncertainty caused by the **BS-side generation error**
- Can we further improve the SSBF performance through the **online calibration?**


Our method: **SSBF + Type-II Feedback**

- **BS** generates only the **beam subspace** instead of the final beam
- **UE** determine **the beam parameters in the subspace** via **Type-II feedback**



Site-Specific Type-II Codebook

□ Proposed Offline Training Workflow




- **Input:**
Site-specific channel data generated by digital twin/measurement
- **Joint learn:**
SSB probing codebook + subspace inference network
- **Objective:**
Maximizing the end-to-end channel capture efficiency



Site-Specific Type-II Codebook


❑ Proposed Online Workflow v.s. Conventional Type-II Feedback





Site-Specific Type-II Codebook

□ End-to-End Joint Design Convergence



Key Observations

- **Stable convergence within around 100 epochs**
- Small train-validation gap across all settings
- No obvious overfitting

K (Codebook size)/Q (Subspace dimension) Effect

- $K \uparrow Q \uparrow$, performance $\uparrow$, but overhead $\uparrow$

Design Insight

- **Increase K/Q to improve performance**




Effective spectral efficiency



Site-Specific Type-II Codebook

Overhead-Performance Tradeoff



Key Observations

- Capture efficiency increases with overhead
- Performance saturates after moderate overhead
- Pareto points show the best

Impact of K/Q

- Larger K: more informative RSRP fingerprint
- Larger Q: stronger subspace representation
- **Increasing only one of them is inefficient**


Design Insight

- Balance K and Q
- Maximum capture efficiency is not always best
- **Effective spectral efficiency depends on both gain and overhead**




Site-Specific Type-II Codebook

Comparison with conventional codebooks



Site-specific Type-II consistently outperforms conventional codebooks




Higher and more stable UE-level performance



Site-Specific Type-II Codebook

Physical Interpretation of Proposed Scheme



Key Observations

- Proposed response aligns with dominant physical paths
- Response is close to Type-II
- Learned SSI recovers meaningful angular structure




Future Directions

- ❑ Fast and robust wireless digital twin construction
- ❑ Site-specific learning for
 - ❑ Sensing and localization
 - ❑ Multi-user, multi-cell beamforming and resource allocation
 - ❑ Pinching antenna system
 - ❑ Adaptive modulation and coding
 - ❑ User tracking
 - ❑ Distributed multi-site training
 - ❑ ...




HKU Demo


6G Massive MIMO BS
1024 antennas, 128 ports



5G Multi-Cell Pico BS
4 antennas



NVIDIA
Digital Twin
Server



6G Test UE



OpenClaw
Server






Department of Electrical and
Computer Engineering

電機與計算機工程系

HKU Demo

Agent-Driven Data Collector





Department of Electrical and Electronic Engineering
電 機 電 子 工 程 系

Thanks for Your Attention



HKUEEE



www.eee.hku.hk

THE UNIVERSITY OF HONG KONG 香港大學 • FACULTY OF ENGINEERING 港大工程學院