Digital Twin Aided Millimeter Wave MIMO: Site-Specific Beam Codebook Learning

Hao Luo and Ahmed Alkhateeb

Wireless Intelligence Laboratory (WI-Lab)
School of Electrical, Computer, and Energy Engineering
Arizona State University

IEEE International Conference on Communications (ICC), 2026

Wireless Intelligence Lab

IRA A. FULTON SCHOOLS OF engineering
ARIZONA STATE UNIVERSITY

mmWave MIMO communications

- mmWave MIMO can deliver high data rates and support emerging data-intensive applications

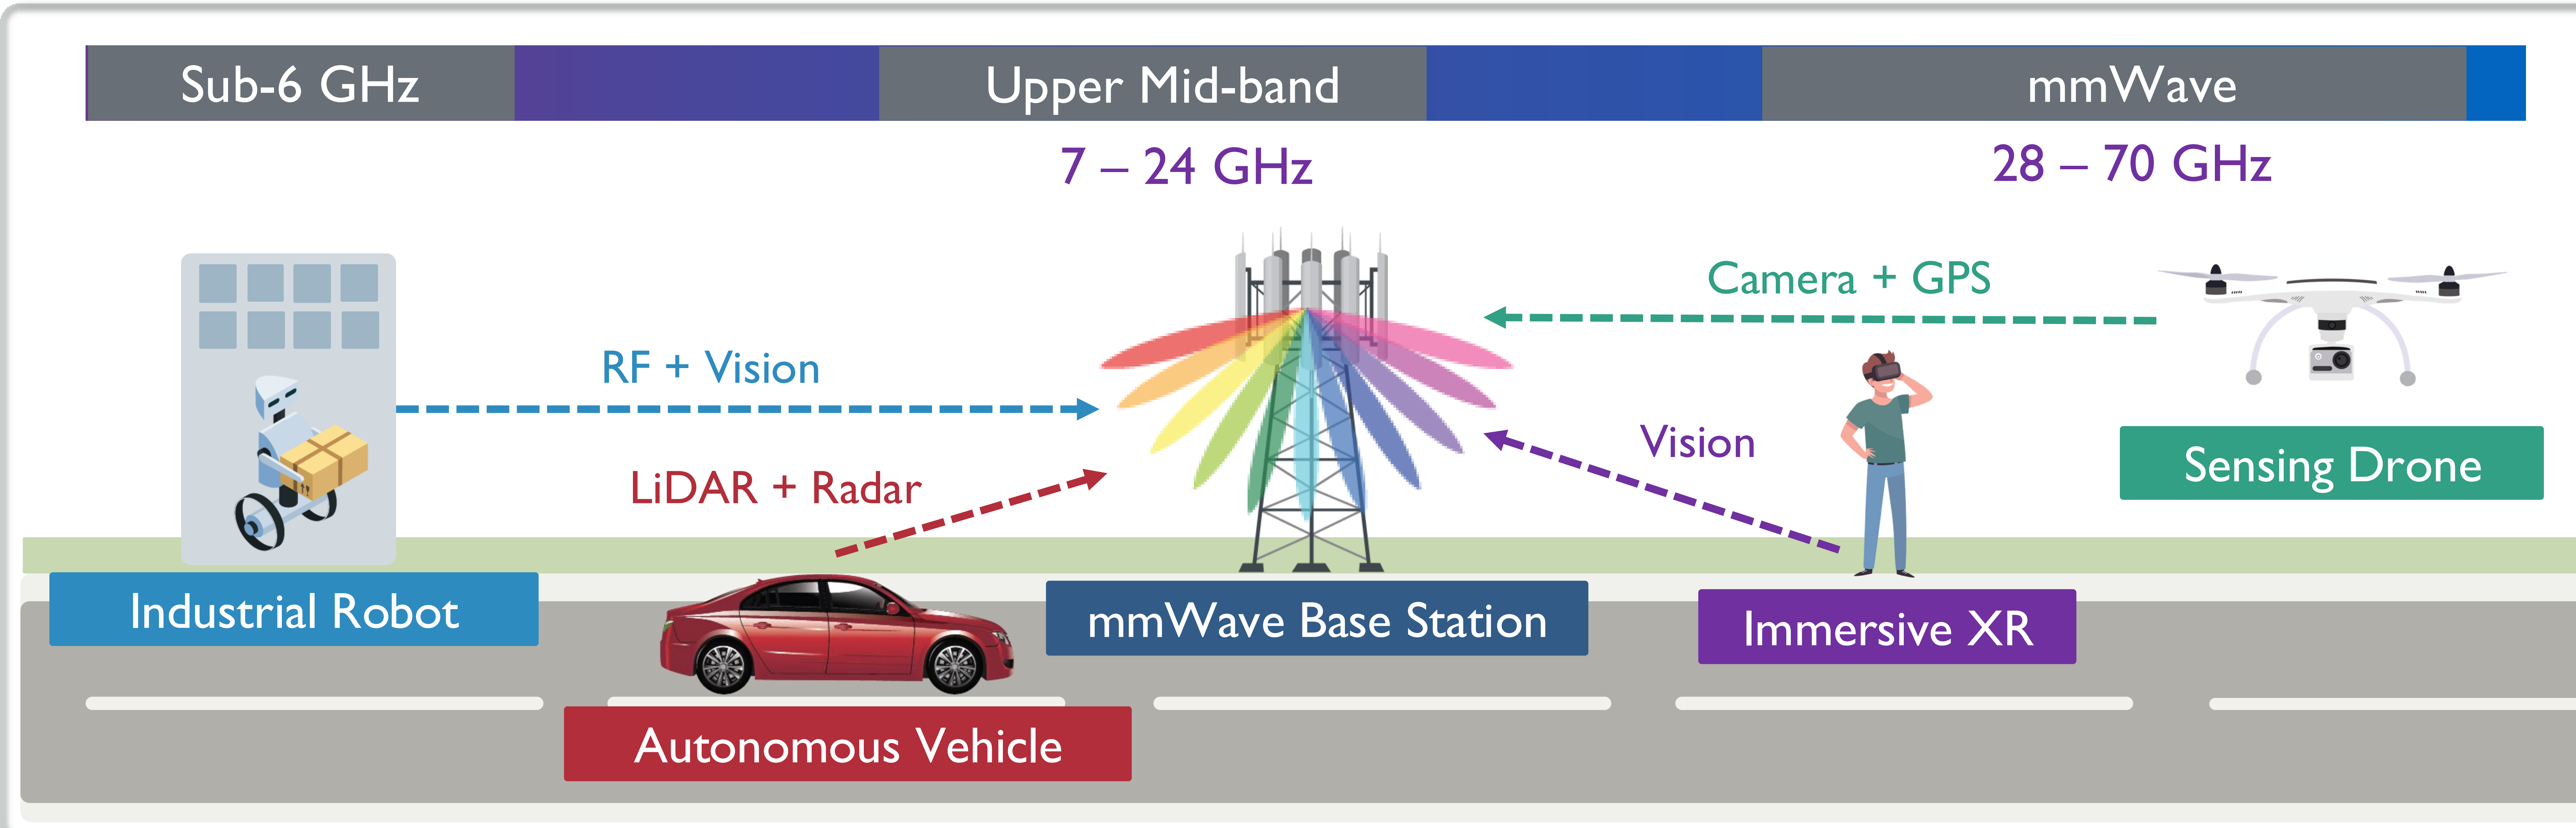

- Analog/hybrid architecture challenge in beamforming design
  - Fewer RF chains than antennas
  - Limited access to channel state information

Efficient beamforming codebook design is critical in mmWave communications

Data-driven site-specific codebook learning

- Existing approaches

Classical codebooks
- DFT or fixed beam grids
- Simple and robust
- Ignore site geometry and user distribution

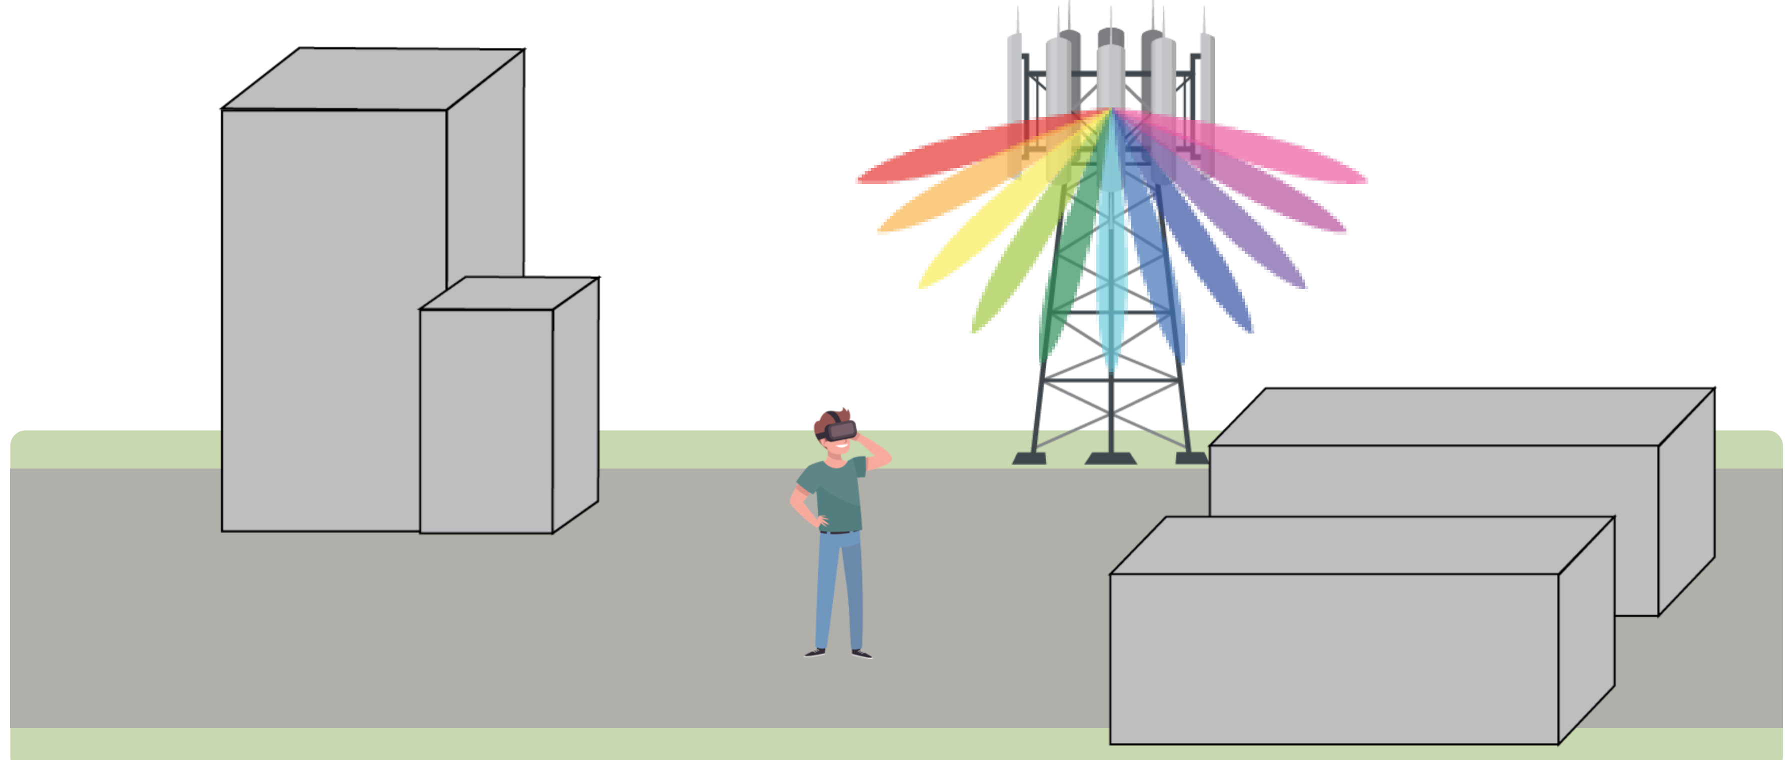

Data-driven codebook design
- Learn beams from measurements
- Can adapt to deployment and hardware
- Needs many online iterations to converge

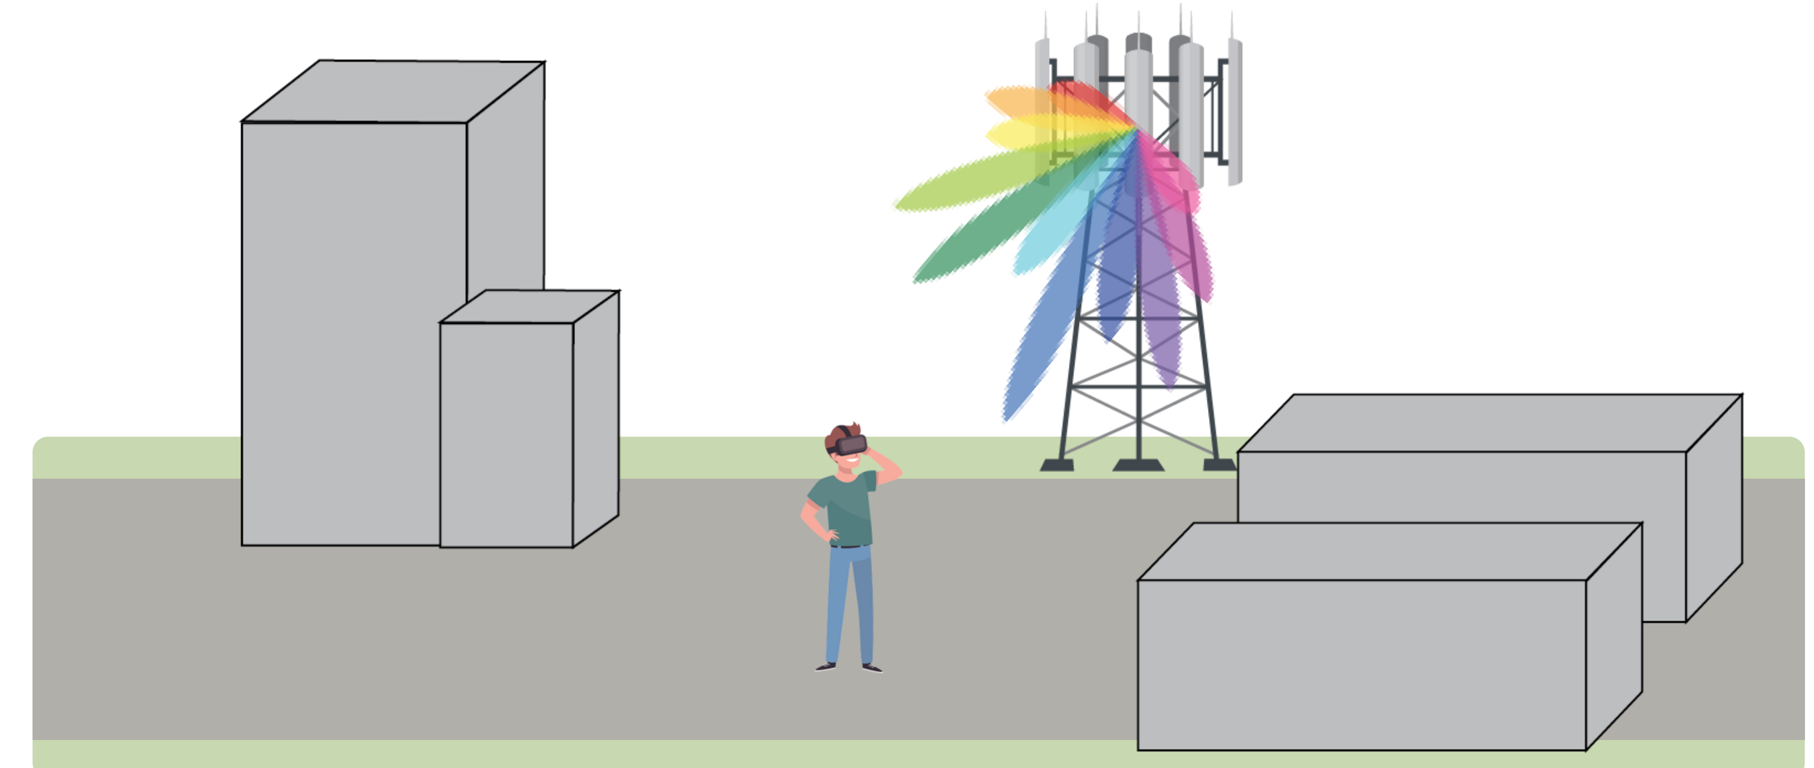

- Our proposal

Use a site-specific digital twin to generate synthetic channels

Train beams offline with the digital twin, and learn separate LoS/NLoS codebooks

[Zhang'22] Y. Zhang, M. Alrabeiah and A. Alkhateeb, "Reinforcement Learning of Beam Codebooks in Millimeter Wave and Terahertz MIMO Systems," in *IEEE Transactions on Communications*
[Alrabeiah'22] M. Alrabeiah, Y. Zhang and A. Alkhateeb, "Neural Networks Based Beam Codebooks: Learning mmWave Massive MIMO Beams That Adapt to Deployment and Hardware," in *IEEE Transactions on Communications*

System model

- Base station (BS)
  - M antennas
  - Serving single-antenna users
  - Analog-only architecture
  - One RF chain with r-bit phase shifters
  - Codebook of combining vectors $\mathcal{W}$

$$\mathbf{w} = \frac{1}{\sqrt{M}}[e^{j\theta_1}, \ldots, e^{j\theta_M}]^T$$

$$\theta_m \in \boldsymbol{\Theta}$$

Set of phase shift values that contains 2^r uniformly spaced values drawn from $(-\pi, \pi]$

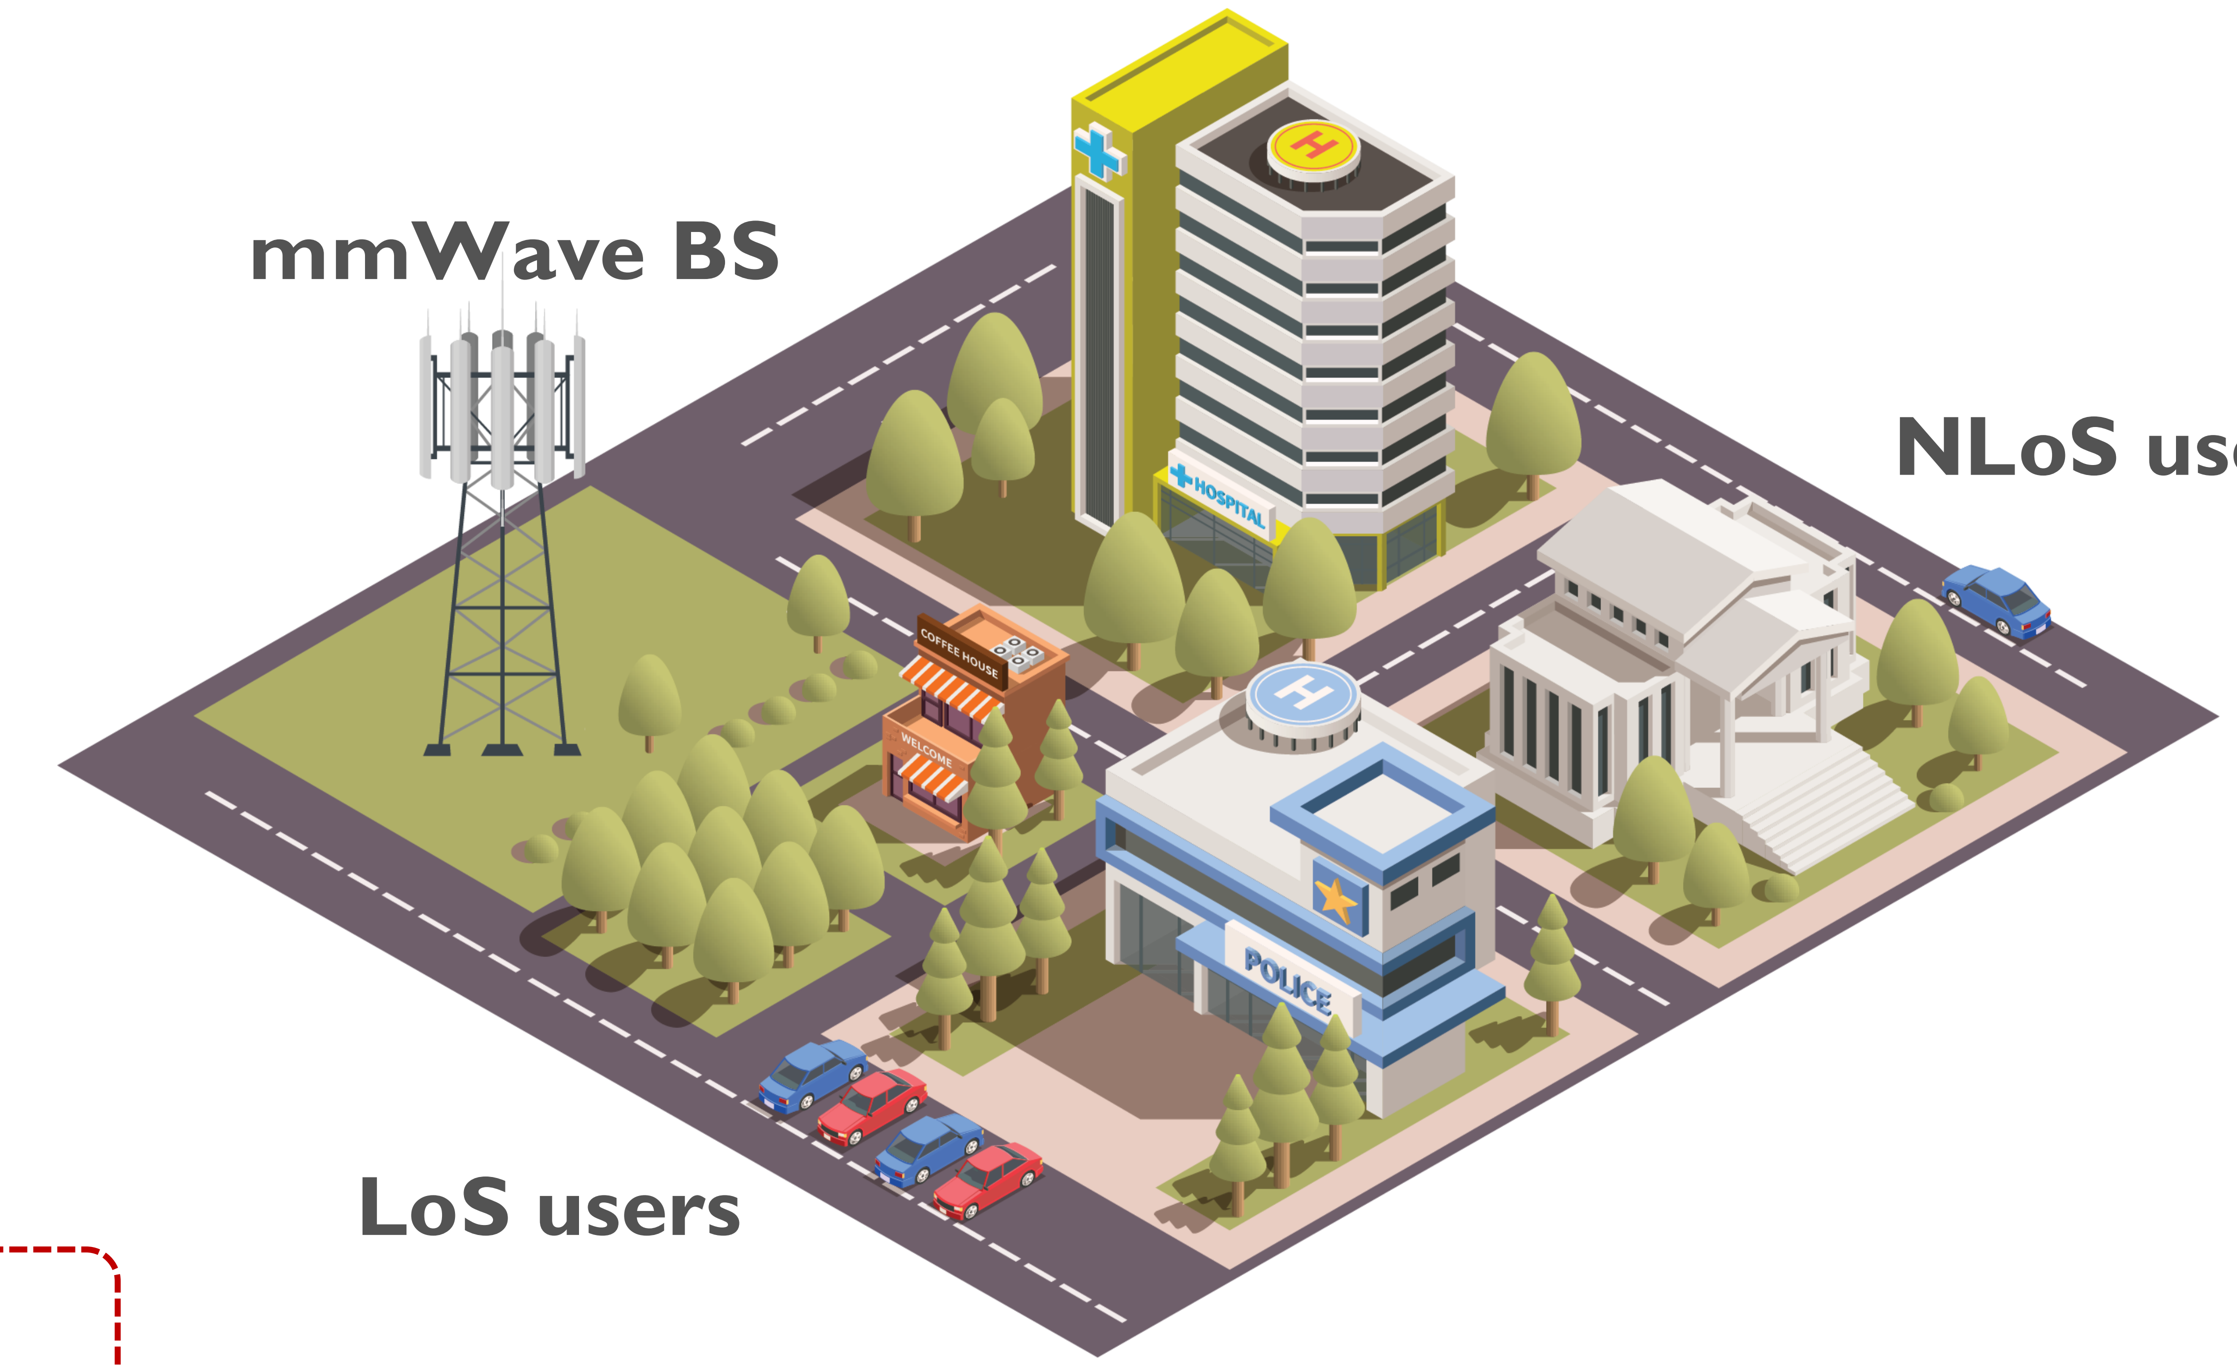

- Uplink received signal at the BS

$$y = \mathbf{w}^H \mathbf{h} x + \mathbf{w}^H \mathbf{n}$$

x: Transmitted symbol

$\mathbf{n}$: $\sim \mathcal{CN}(0, \sigma^2 \mathbf{I})$

- Geometric channel model

$$\mathbf{h} = \sum_{l=1}^{L} \alpha_l \mathbf{a}(\phi_l)$$

$\sum_{l=1}^{L}$: Superposition of paths

α_l: Complex-valued path gain

$\mathbf{a}(\phi_l)$: Array response vector

Problem formulation

- Objective: design a digital twin aided codebook learning framework
- Received SNR after combining at the based station

$$\mathsf{SNR} = \rho \max_{\mathbf{w}\in\mathcal{W}} |\mathbf{w}^H \mathbf{h}|^2, \qquad \rho = P/\sigma^2$$

- Problem formulation

Avg. over user channels

$$\boldsymbol{\mathcal{W}}_{\mathsf{opt}} = \arg\max_{\boldsymbol{\mathcal{W}}} \boxed{\frac{1}{|\boldsymbol{\mathcal{H}}|}\sum_{\mathbf{h}\in\boldsymbol{\mathcal{H}}}} \left(\max_{\mathbf{w}_n\in\boldsymbol{\mathcal{W}}} |\mathbf{w}_n^H \mathbf{h}|^2\right),$$

$$\text{s.t.} \quad w_{mn} = \frac{1}{\sqrt{M}} e^{j\theta_{mn}}, \ \theta_{mn} \in \boldsymbol{\Theta}, \ \forall m, n, \Big\} \text{ Quantized phase shift}$$

m^{th} element of the
n^{th} combining vector

- This is a non-convex combinatorial optimization problem
  - Existing ML approaches: learn over the air → high training overhead
  - Require a large amount of real-world measurements

Key Idea: digital twin aided codebook learning framework

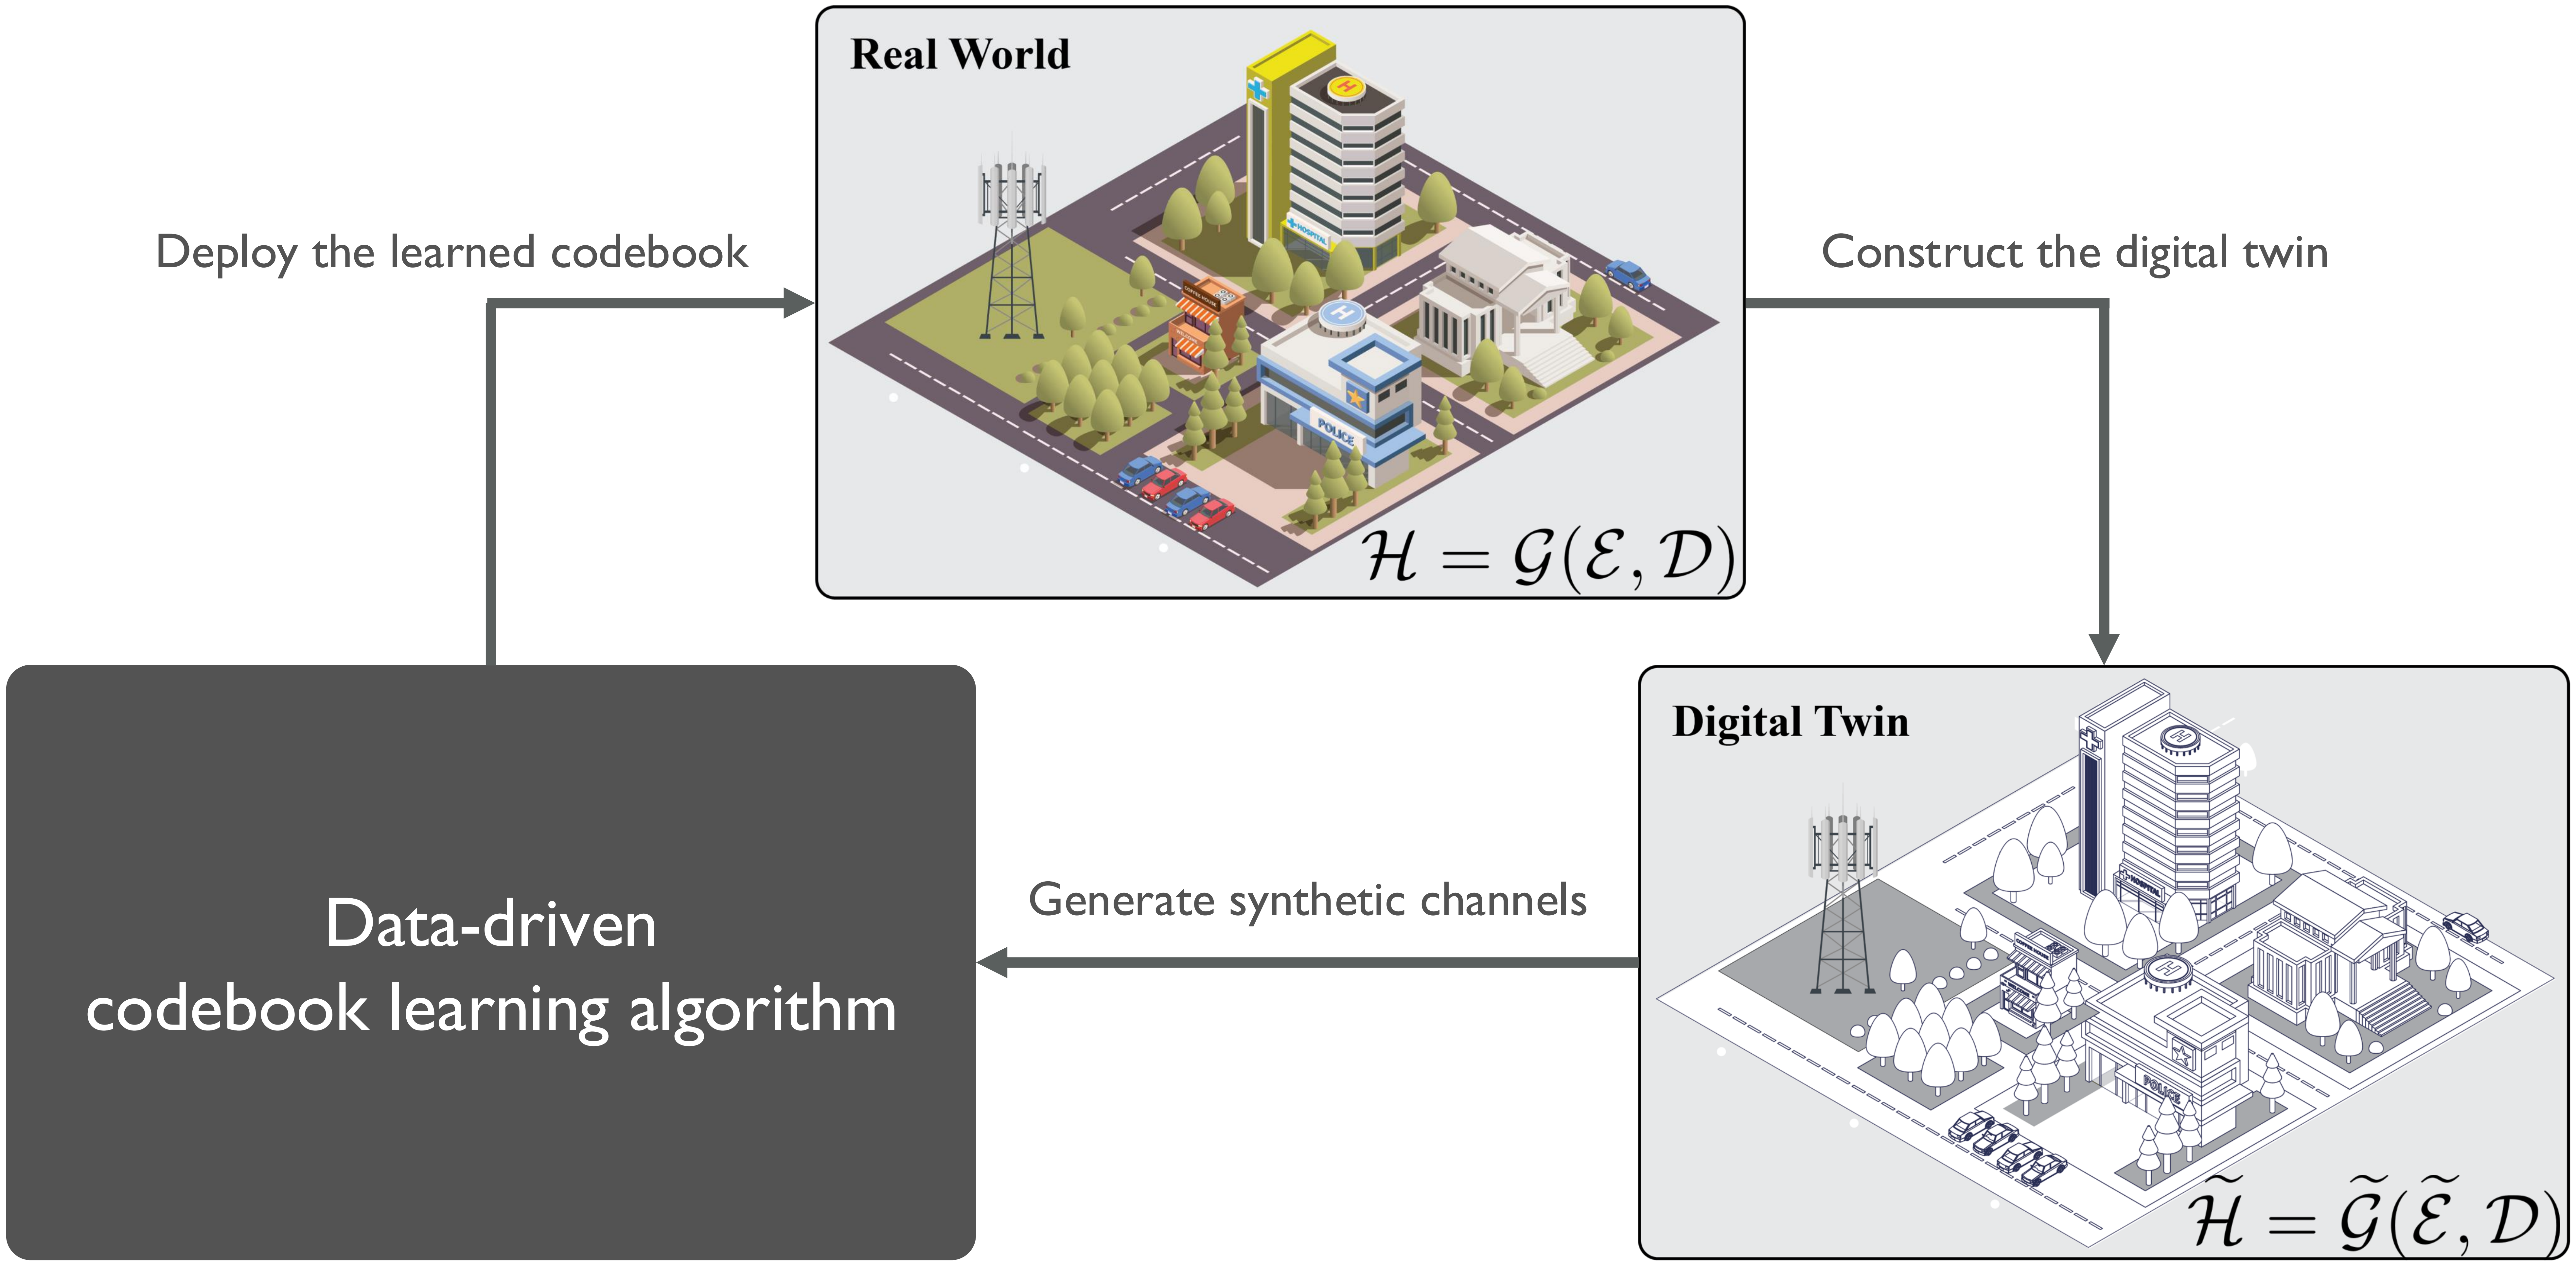

Leverage site-specific digital twins to generate synthetic channels for codebook learning

Case study: reinforcement learning (RL) based codebook learning

- Decouple codebook design into N independent beam subproblems

Avg. over digital twin synthetic channels

$$\mathbf{w}_{\text{opt}} = \arg\max_{\mathbf{w}} \frac{1}{|\tilde{\mathbf{H}}_s|} \sum_{\mathbf{h}\in\tilde{\mathbf{H}}_s} |\mathbf{w}^H \mathbf{h}|^2,$$

$$\text{s.t.} \quad w_m = \frac{1}{\sqrt{M}} e^{j\theta_m},\ \theta_m \in \Theta,\ \forall m$$

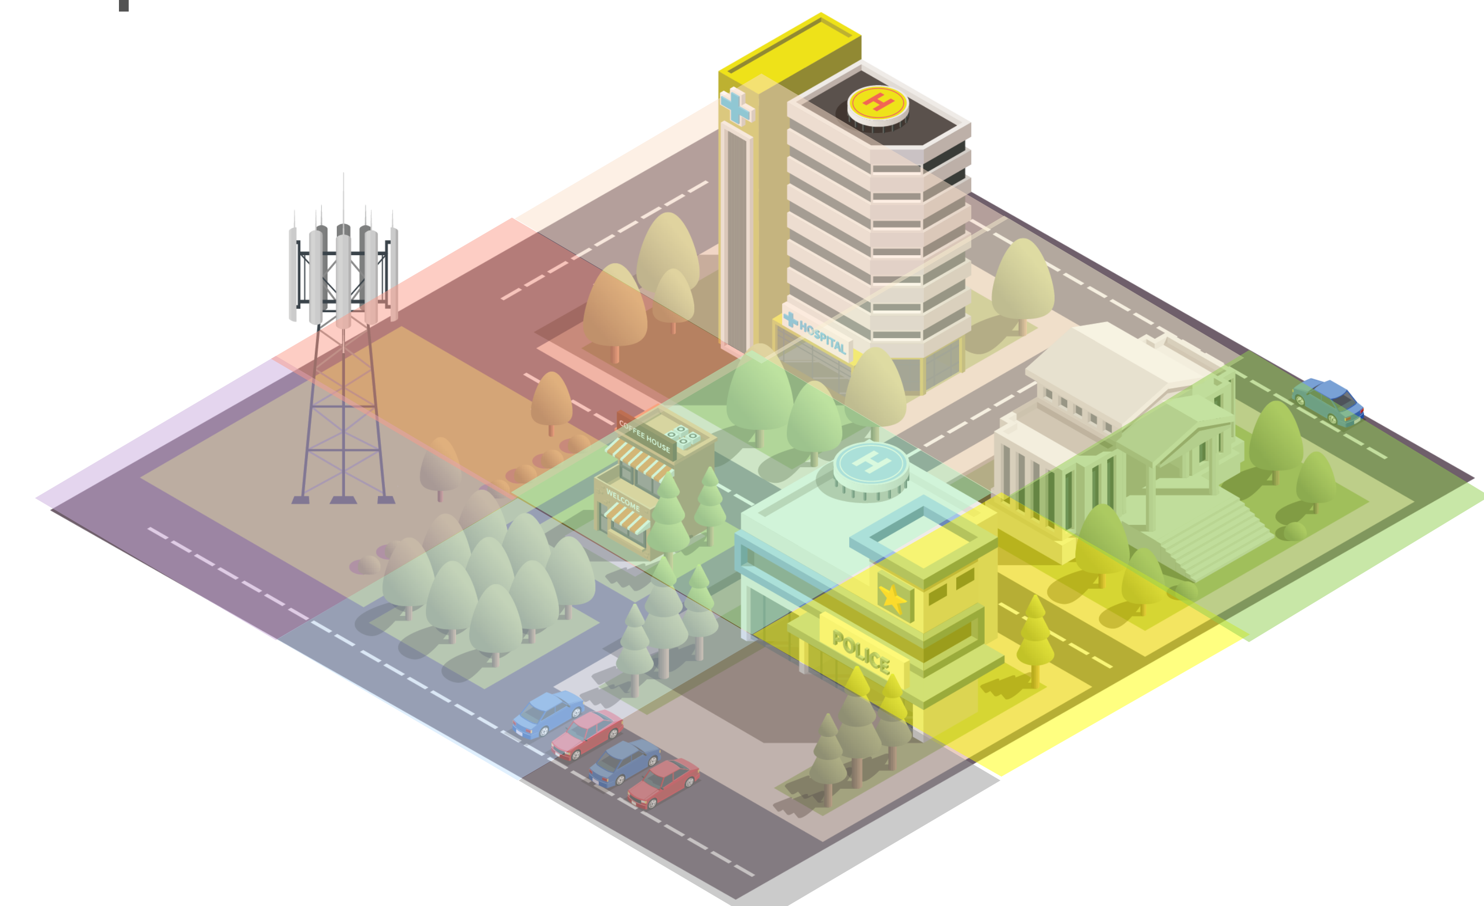

- SNR-based clustering of digital twin synthetic channels
  - Construct S random sensing beam $\mathcal{F} = \{\mathbf{f}_1, \dots, \mathbf{f}_S\}$
  - Building sensing matrix for K user positions
  - Apply K-means on the columns of P → partition into N clusters

Capture spatial characteristics of each user position

$$\mathbf{P} = \begin{bmatrix} |\mathbf{f}_1^H \tilde{\mathbf{h}}_1|^2 & \cdots & |\mathbf{f}_1^H \tilde{\mathbf{h}}_K|^2 \\ \vdots & \ddots & \vdots \\ |\mathbf{f}_S^H \tilde{\mathbf{h}}_1|^2 & \cdots & |\mathbf{f}_S^H \tilde{\mathbf{h}}_K|^2 \end{bmatrix}$$

Each cluster is solved independently for its own beamforming vector

[Zhang'22] Y. Zhang, M. Alrabeiah and A. Alkhateeb, "Reinforcement Learning of Beam Codebooks in Millimeter Wave and Terahertz MIMO Systems," in *IEEE Transactions on Communications*

Case study: reinforcement learning (RL) based codebook learning (cont'd)

- RL formulation for the beamforming design problem

State	Action	Reward
$\mathbf{s}_t = [\theta_1, \ldots, \theta_M]^T$ Phases of the M phase shifters at iteration t	$\mathbf{a}_t$ Selected new phase values $\mathbf{s}_{t+1} = \mathbf{a}_t$	Combining gain $r_t = \begin{cases} +1, & g_t > \beta_t \\ +1, & g_t \leq \beta_t \text{ and } g_t > g_{t-1} \\ -1, & \text{otherwise} \end{cases}$

- Deep deterministic policy gradient (DDPG) with quantized actions
  - Actor output: proto-action $\hat{\mathbf{a}}_t = \mu_\varphi(\mathbf{s}_t)$
  - Quantized action $[\mathbf{a}_t]_m = \arg\min_{\theta \in \Theta} |\theta - [\hat{\mathbf{a}}_t]_m|$
  - Critic estimate Q-value; MSE loss with target network
  - Beam interacts with DT and get avg. combing gain

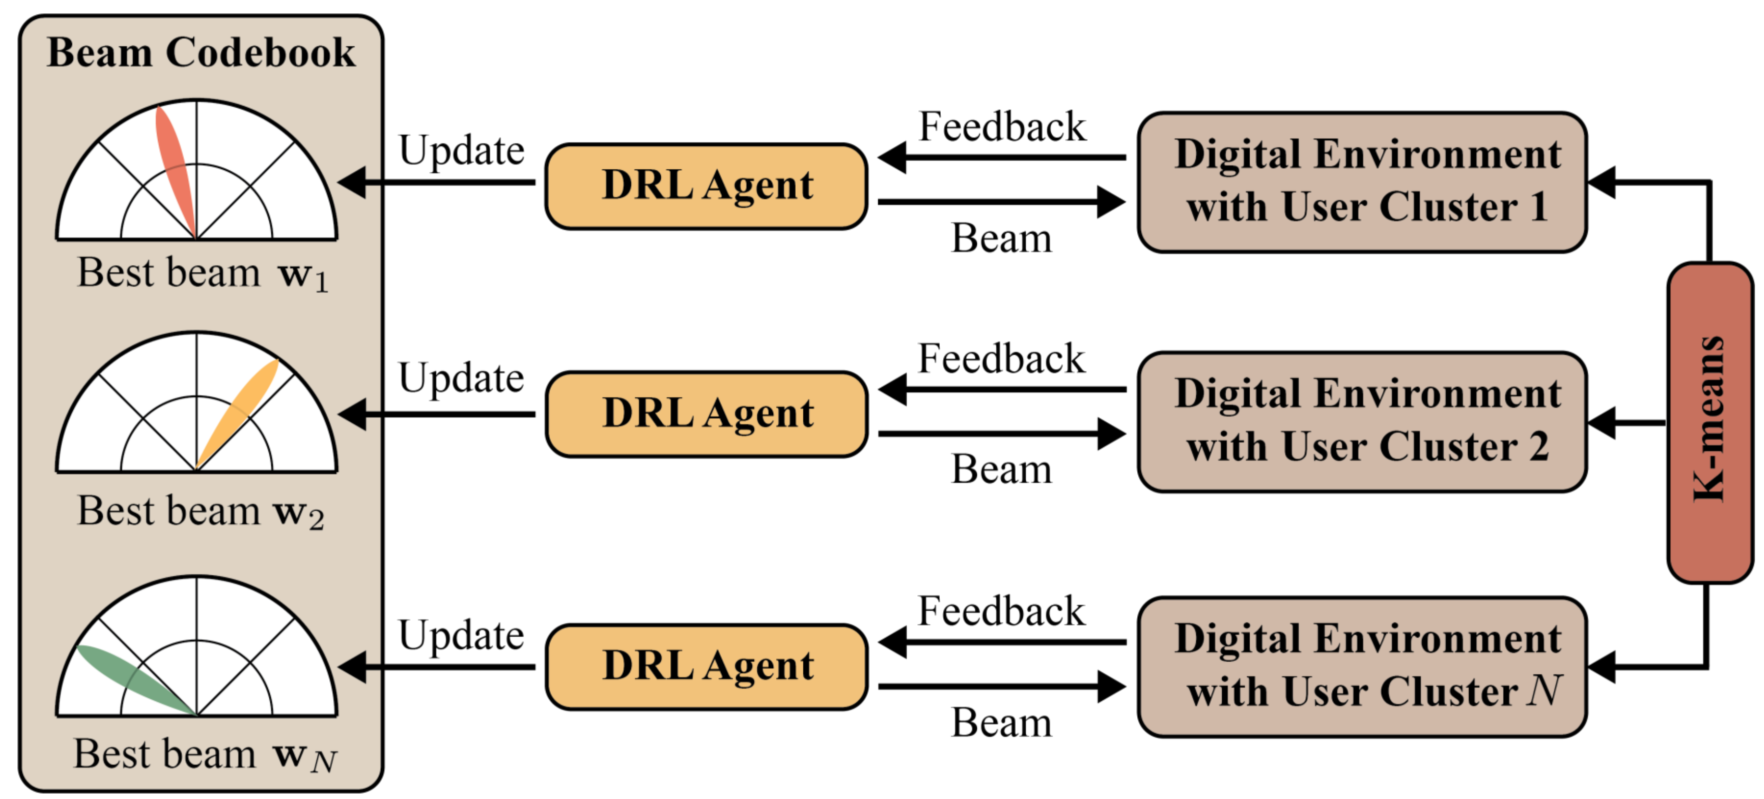

- Threshold β_t updates to the max gain so far → corresponding beam is stored

After training, the codebook holds the best beam for each user cluster

[Zhang'22] Y. Zhang, M. Alrabeiah and A. Alkhateeb, "Reinforcement Learning of Beam Codebooks in Millimeter Wave and Terahertz MIMO Systems," in *IEEE Transactions on Communications*

Separate codebooks for LoS and NLoS users

- With digital twins, we have new opportunities

LoS users
• Need more beams for coverage • Strong direct path dominates • Dedicated codebook capture LoS angles

NLoS users
• Signal arrives through reflections • Channels can be similar across positions • Few angles can cover most users

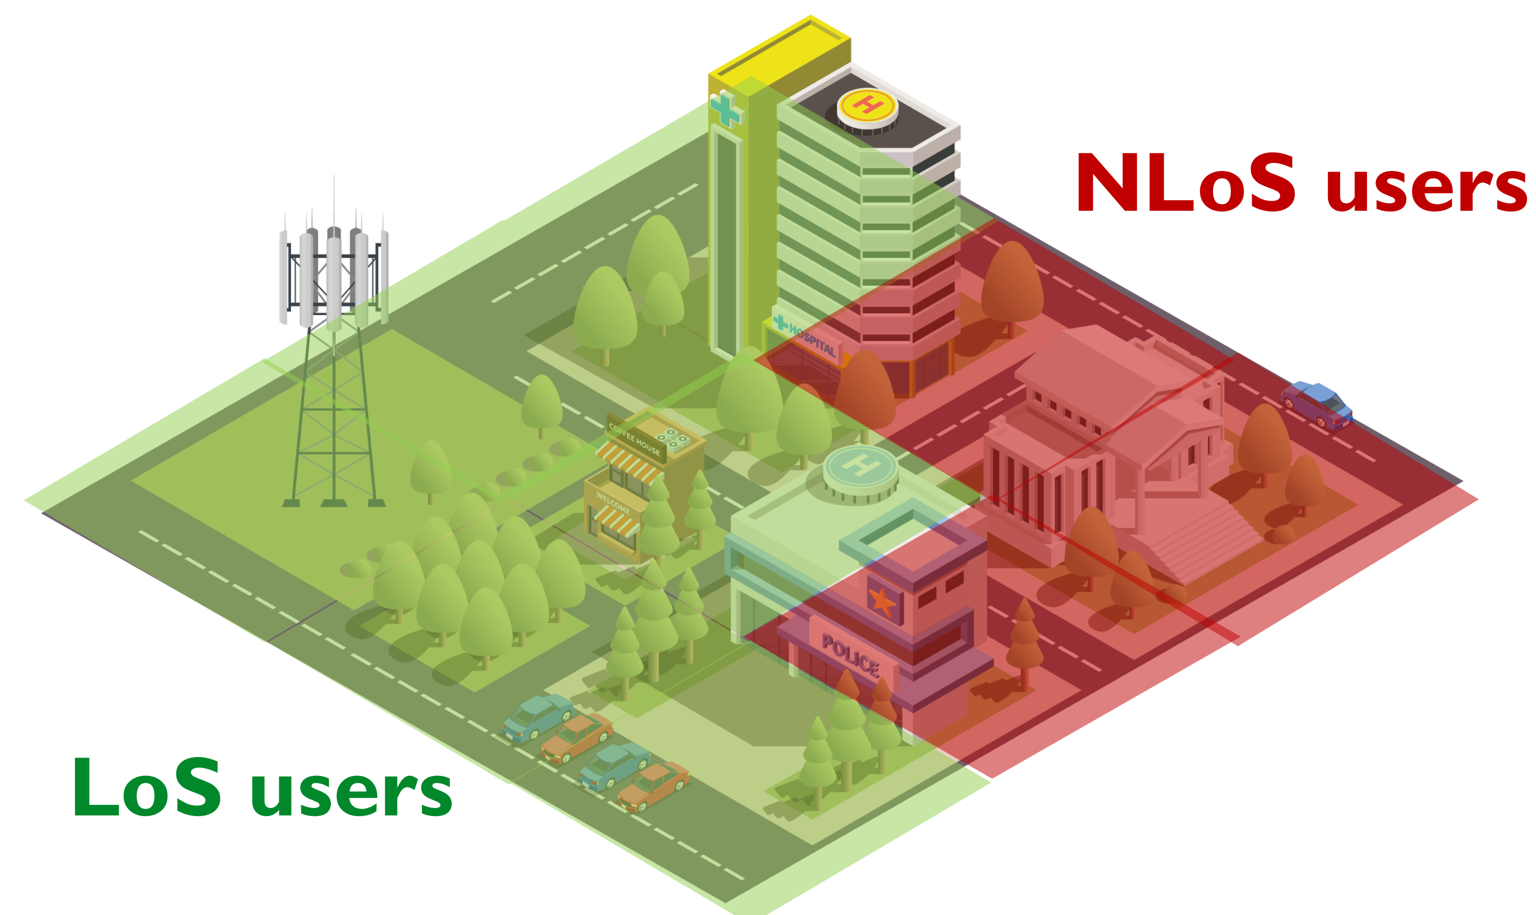

Separate codebooks tailor beamforming to LoS/NLoS – no extra beam training overhead

- However, we also have new challenges …
  - The accuracy of the digital twin model can affect the quality of the synthetic channel data
  - We will study the impact of digital twin fidelity through sensitivity analysis

Simulation setup

Times Square Scenario

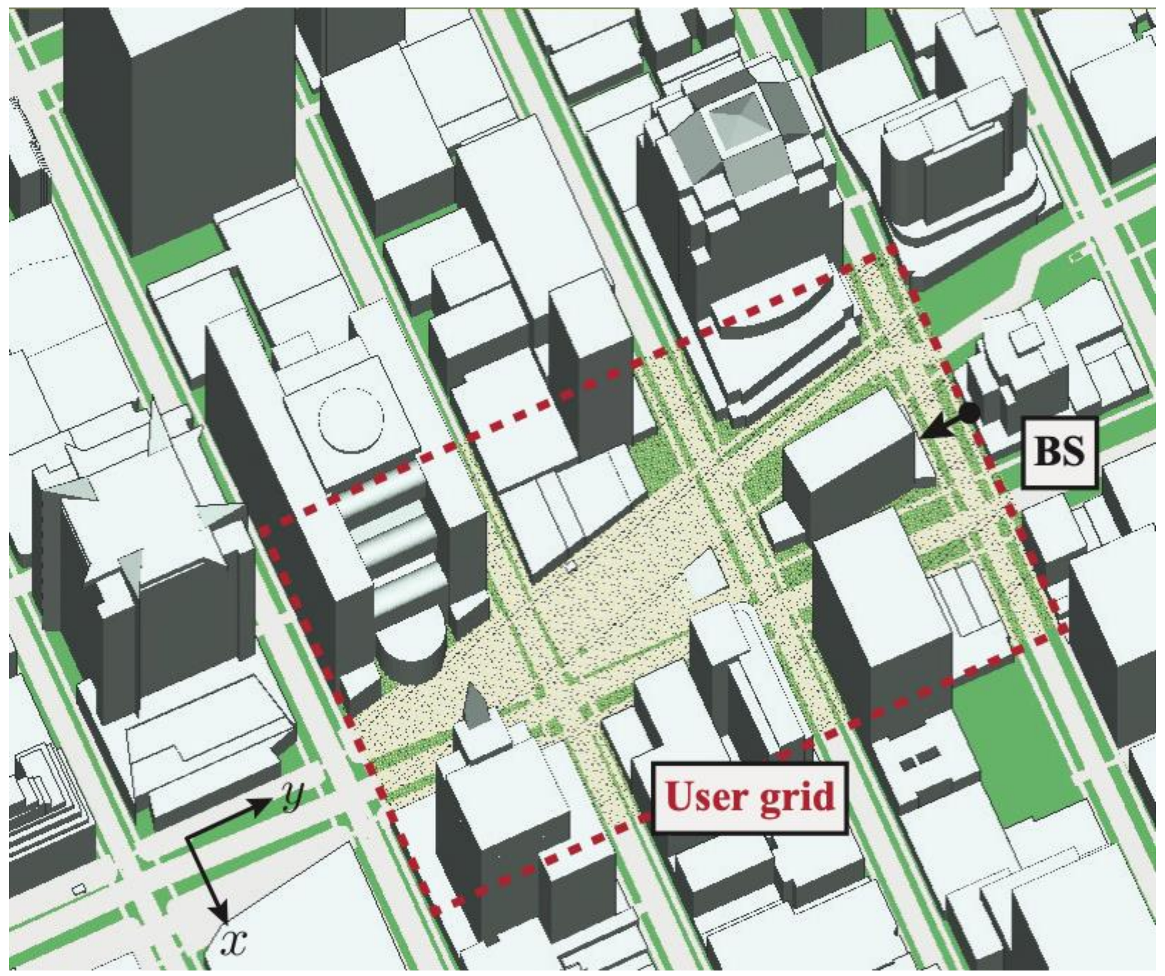

Digital Twin Fidelity Example

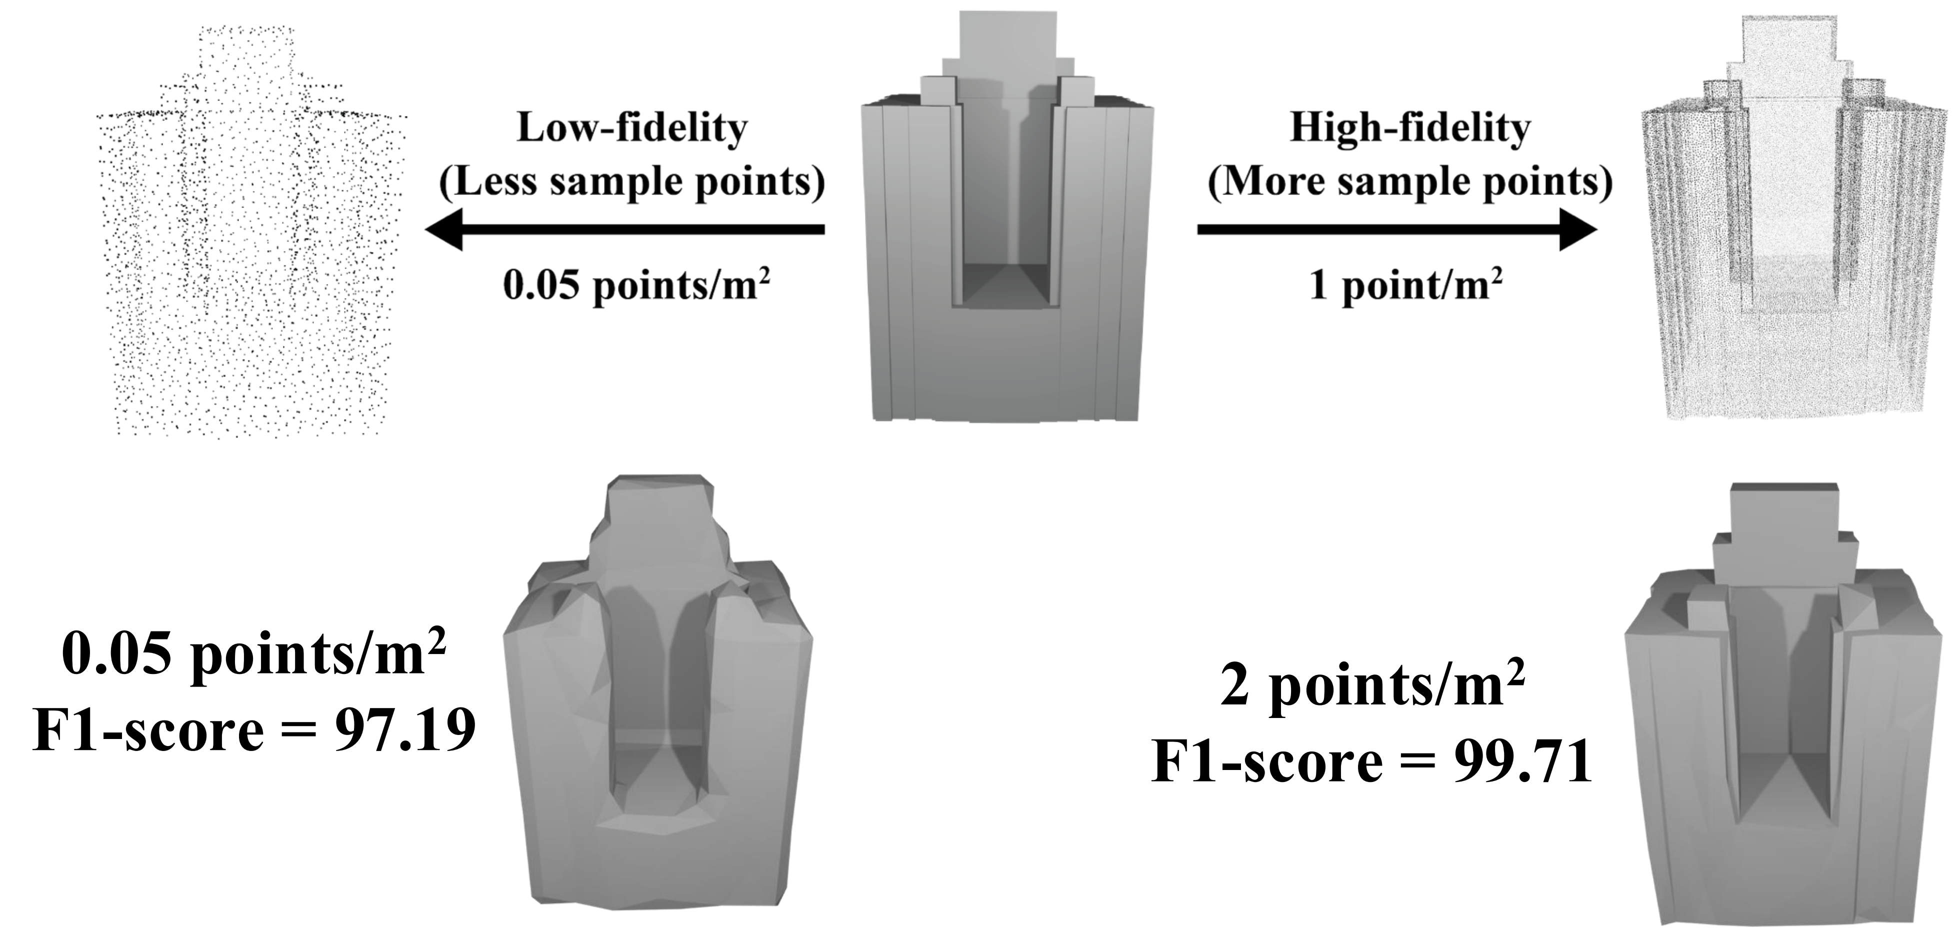

- Target (real-world) scenario
  - Build based on Times Square
  - BS with 4x8 uniform planar array (UPA)
  - Single-antenna users
- Dataset generation
  - Wireless Insite ray-tracing simulator
  - DeepMIMO channel generator
- Digital twin scenario
  - **Geometry fidelity**: sampling density = 0.5 points/m^2
  - **EM material fidelity**: all concrete
  - **Ray-tracing fidelity**: limited to 1st order reflection
- Deep learning architecture
  - Actor network is a fully connected network
  - Critic network is same as actor with diff I/O

[Luo'25] H. Luo, S. Jiang, S. R. Khosravirad and A. Alkhateeb, "Digital Twin Aided Massive MIMO CSI Feedback: Exploring the Impact of Twinning Fidelity," in *IEEE Transactions on Communications*.

What is the benefit of using separate codebooks for LoS and NLoS users?

LoS User

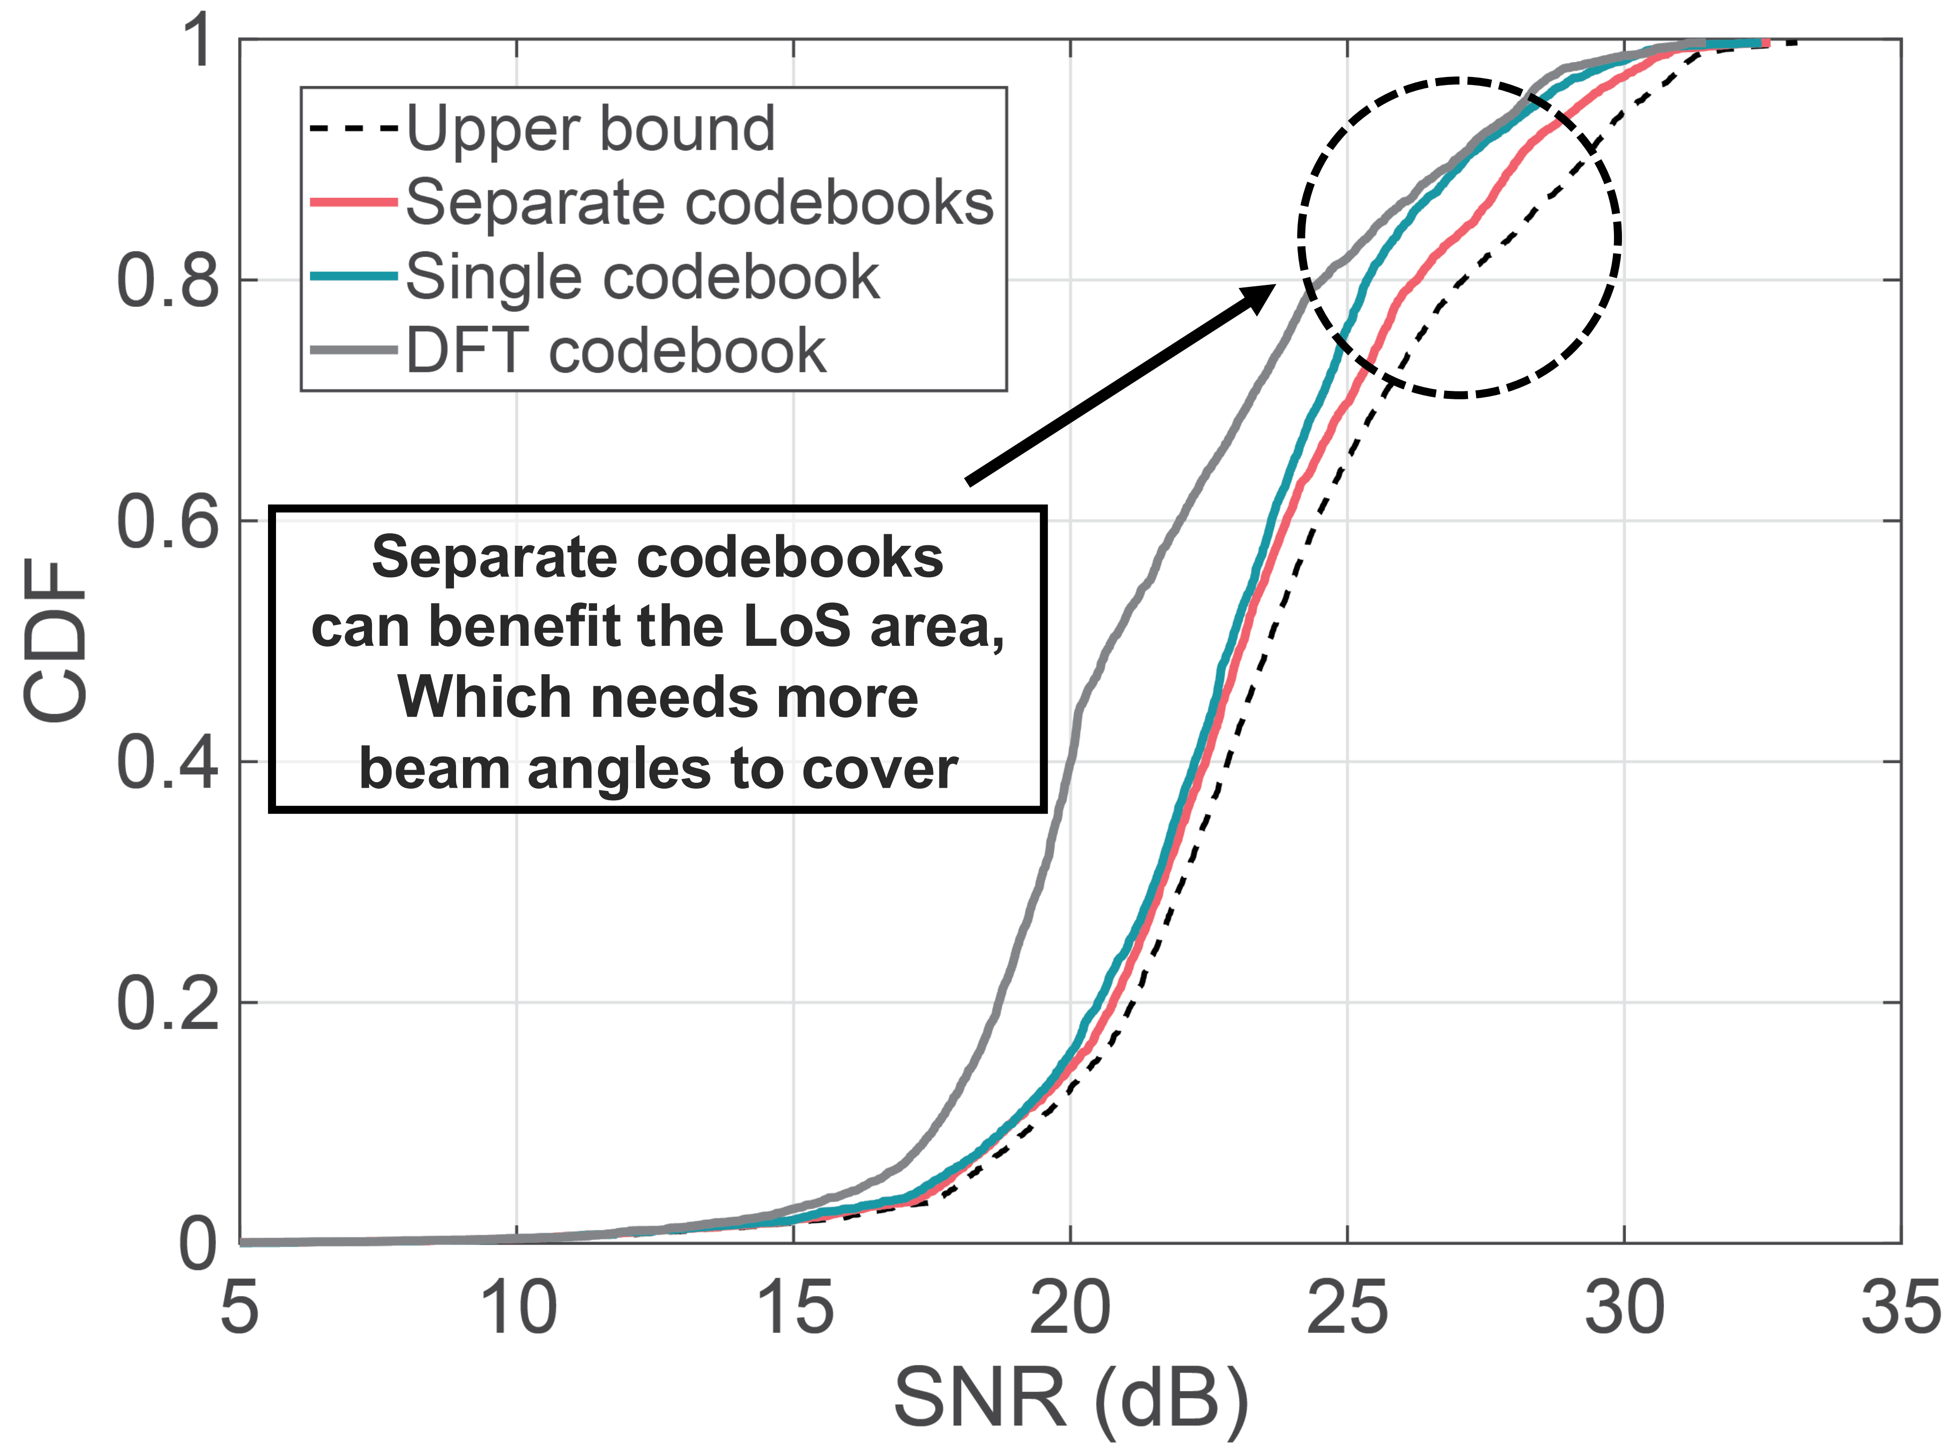

NLoS User

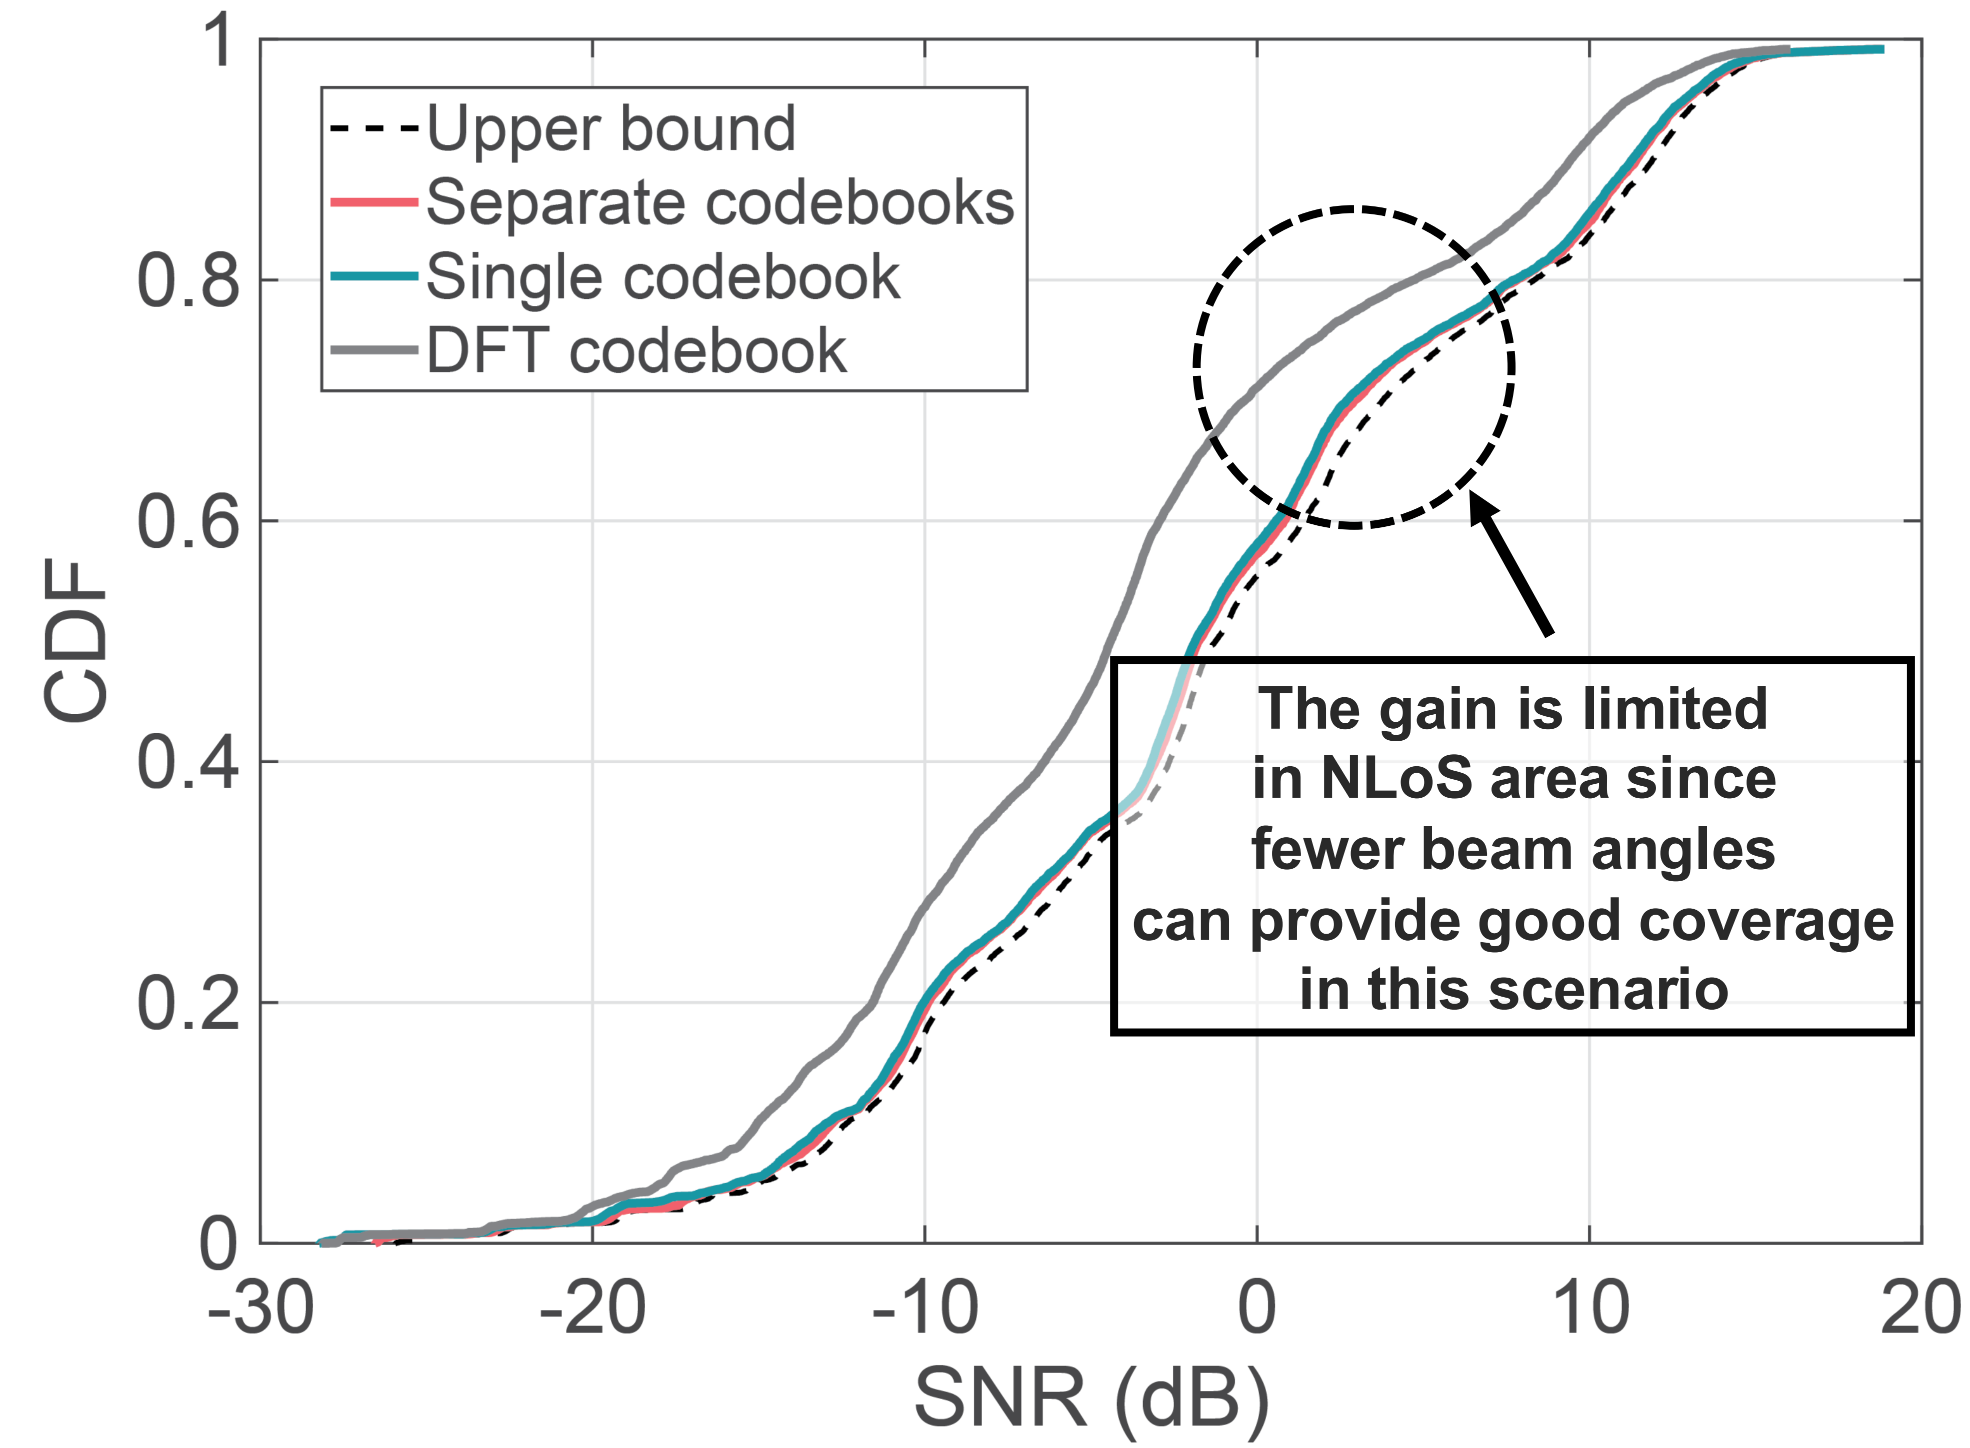

Separate codebooks enhances performance in regions that has wider angular spread

The improvement is attributed to the geometric information provided by digital twins

Can the codebook learned from the digital twin work effectively?

LoS User

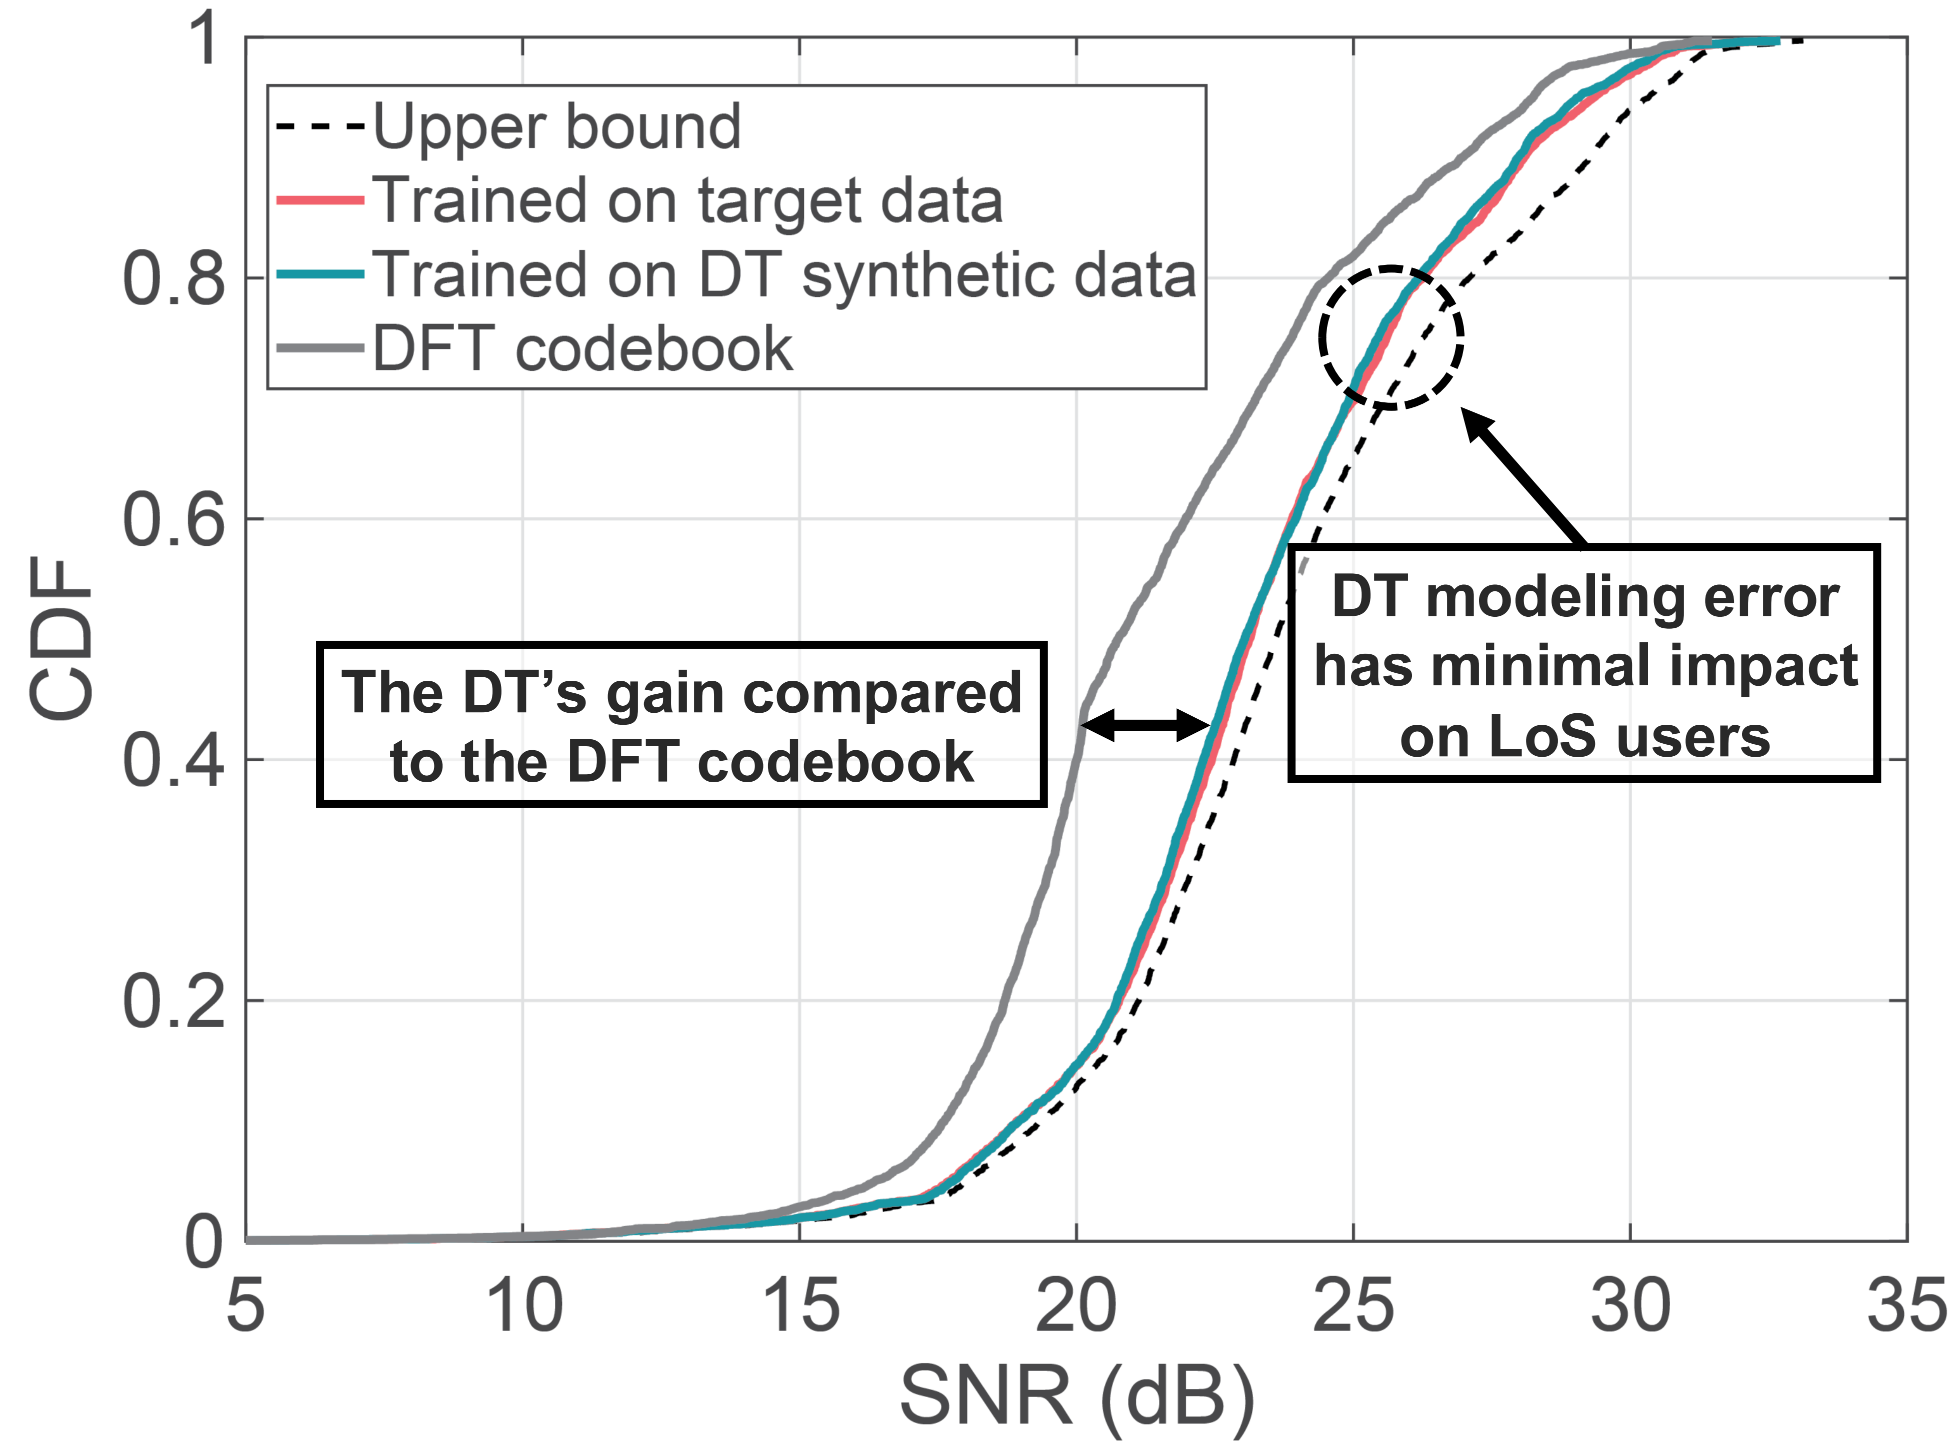

NLoS User

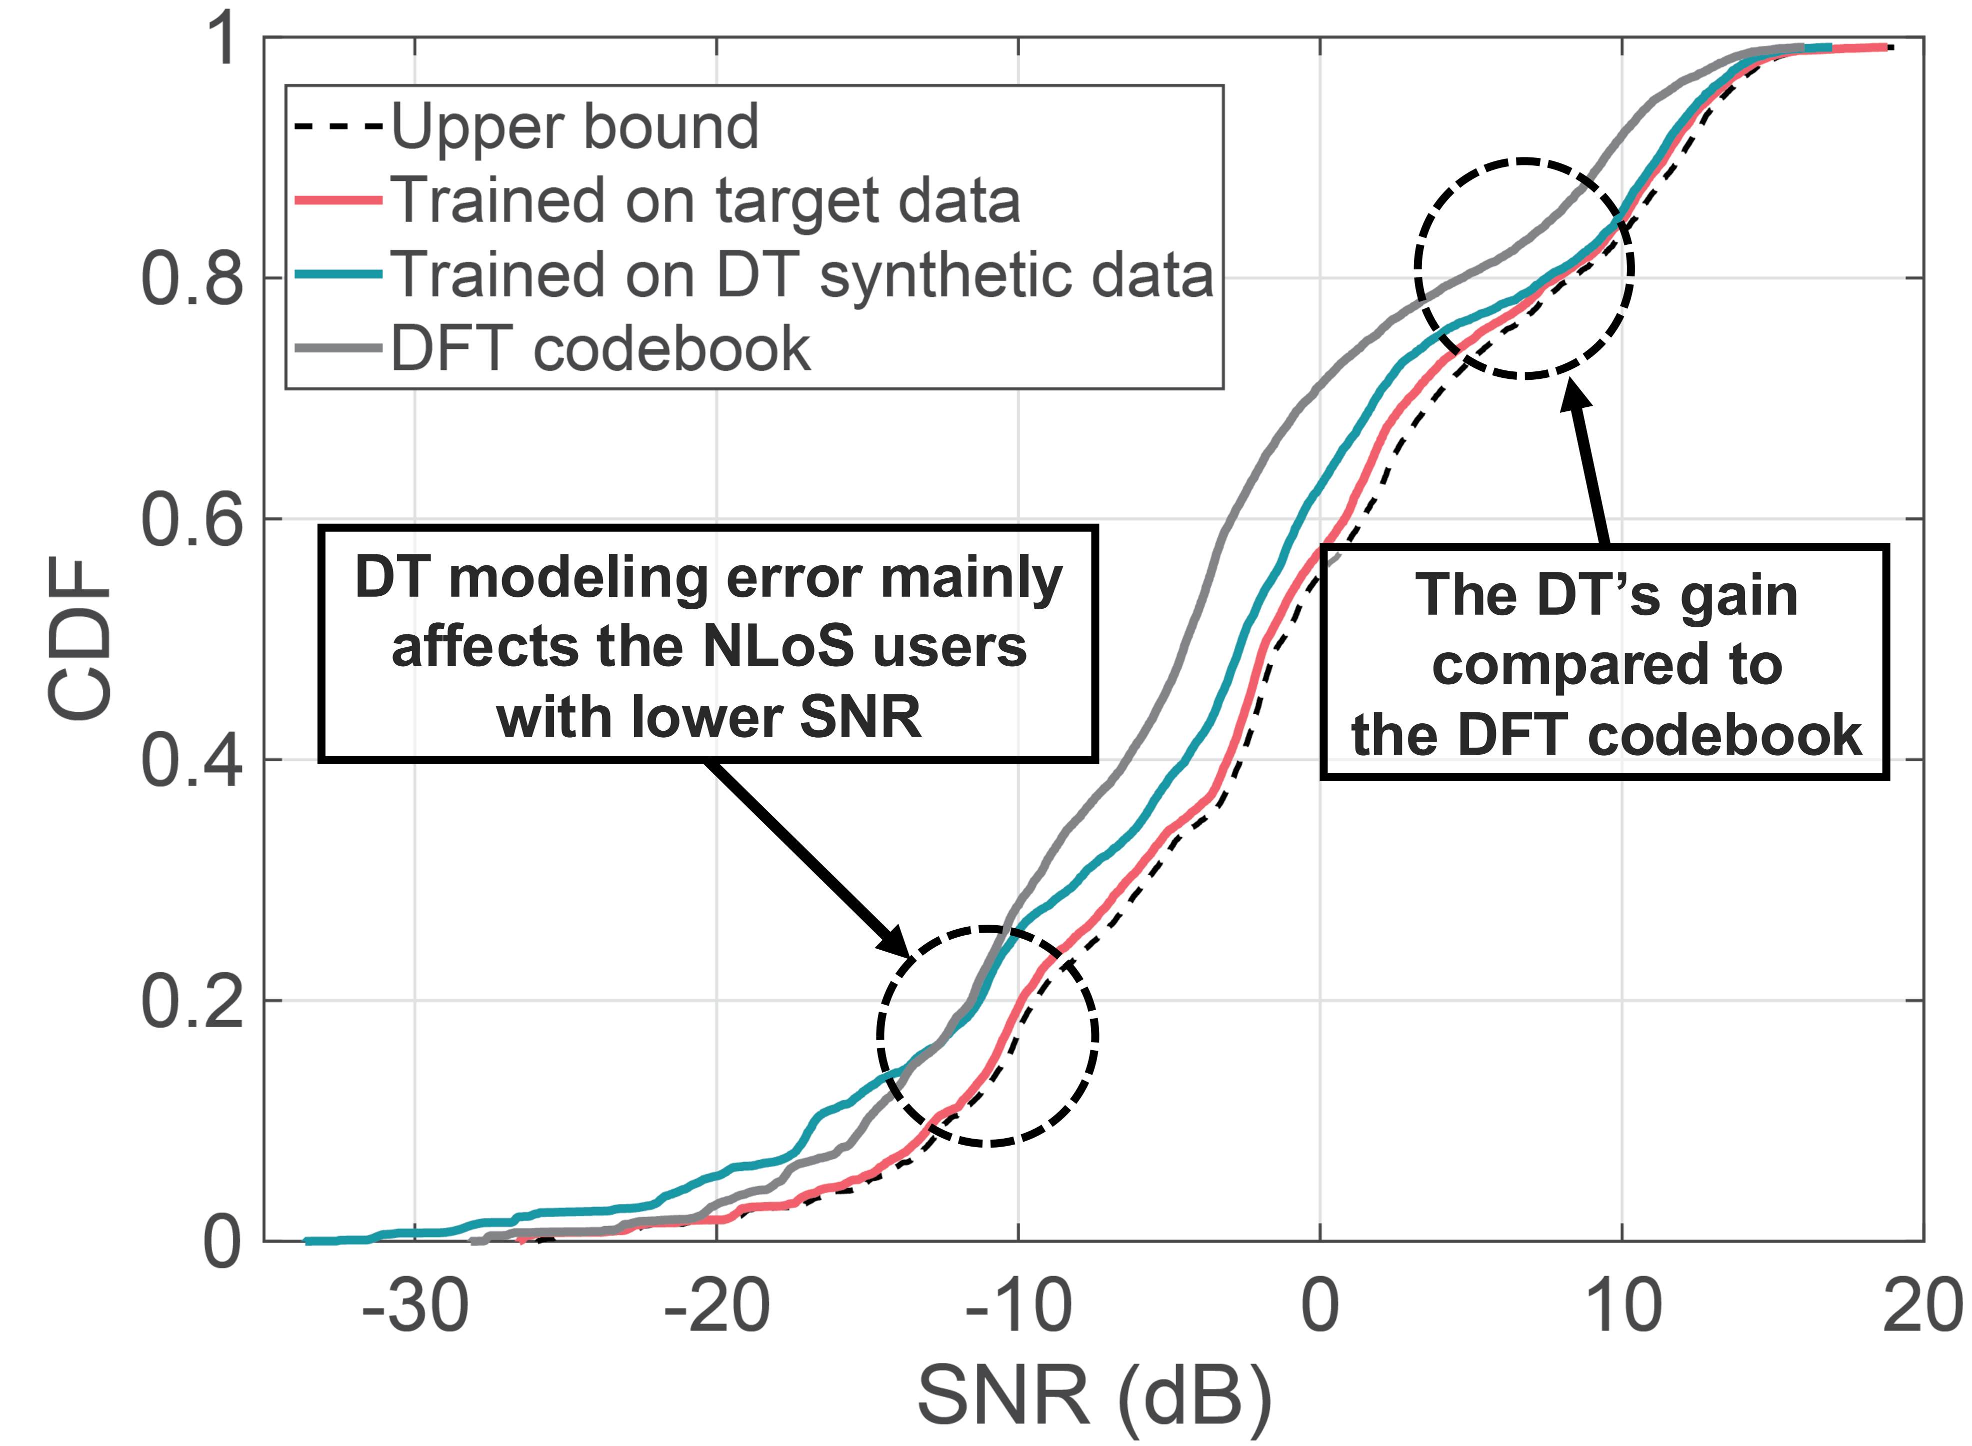

The proposed digital twin based approach outperforms the DFT codebook

Digital twin imperfections primarily cause performance loss

What is the key twinning fidelity factor?

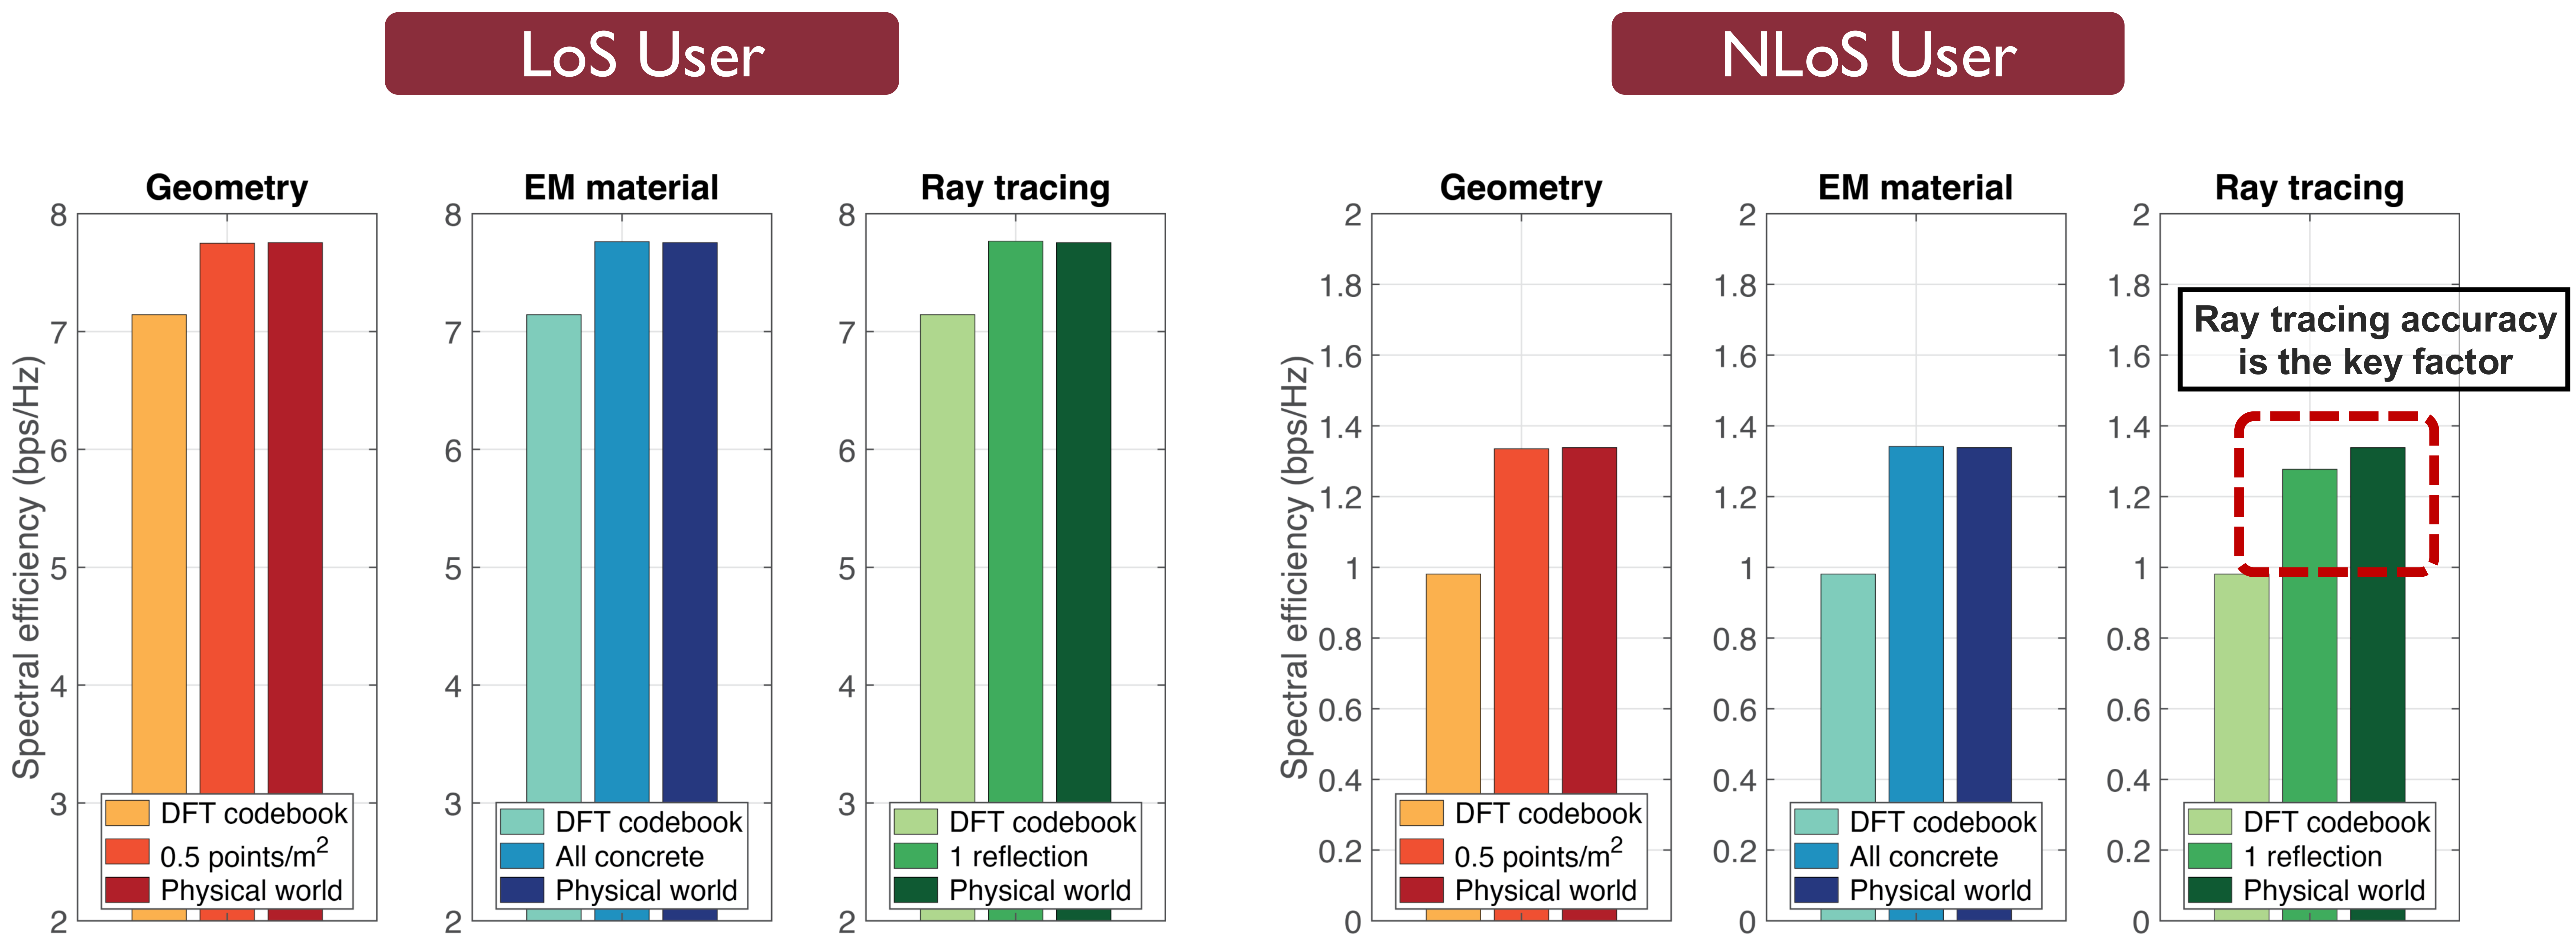

The codebook learning performance is sensitive to the ray-tracing accuracy.

Codebook performance remains stable despite 3D geometry and material inaccuracies

[Luo'25] H. Luo, S. Jiang, S. R. Khosravirad and A. Alkhateeb, "Digital Twin Aided Massive MIMO CSI Feedback: Exploring the Impact of Twinning Fidelity," in *IEEE Transactions on Communications*.

Can the learned codebook adapt to the geometry and user distribution?

- Compare digital twin based and DFT beams using the beam pattern and received SNR (dB)

DFT beam

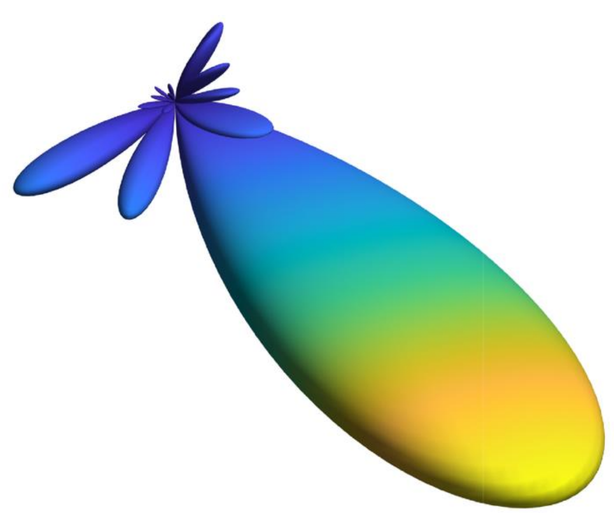

Pre-defined, single-lobe beam

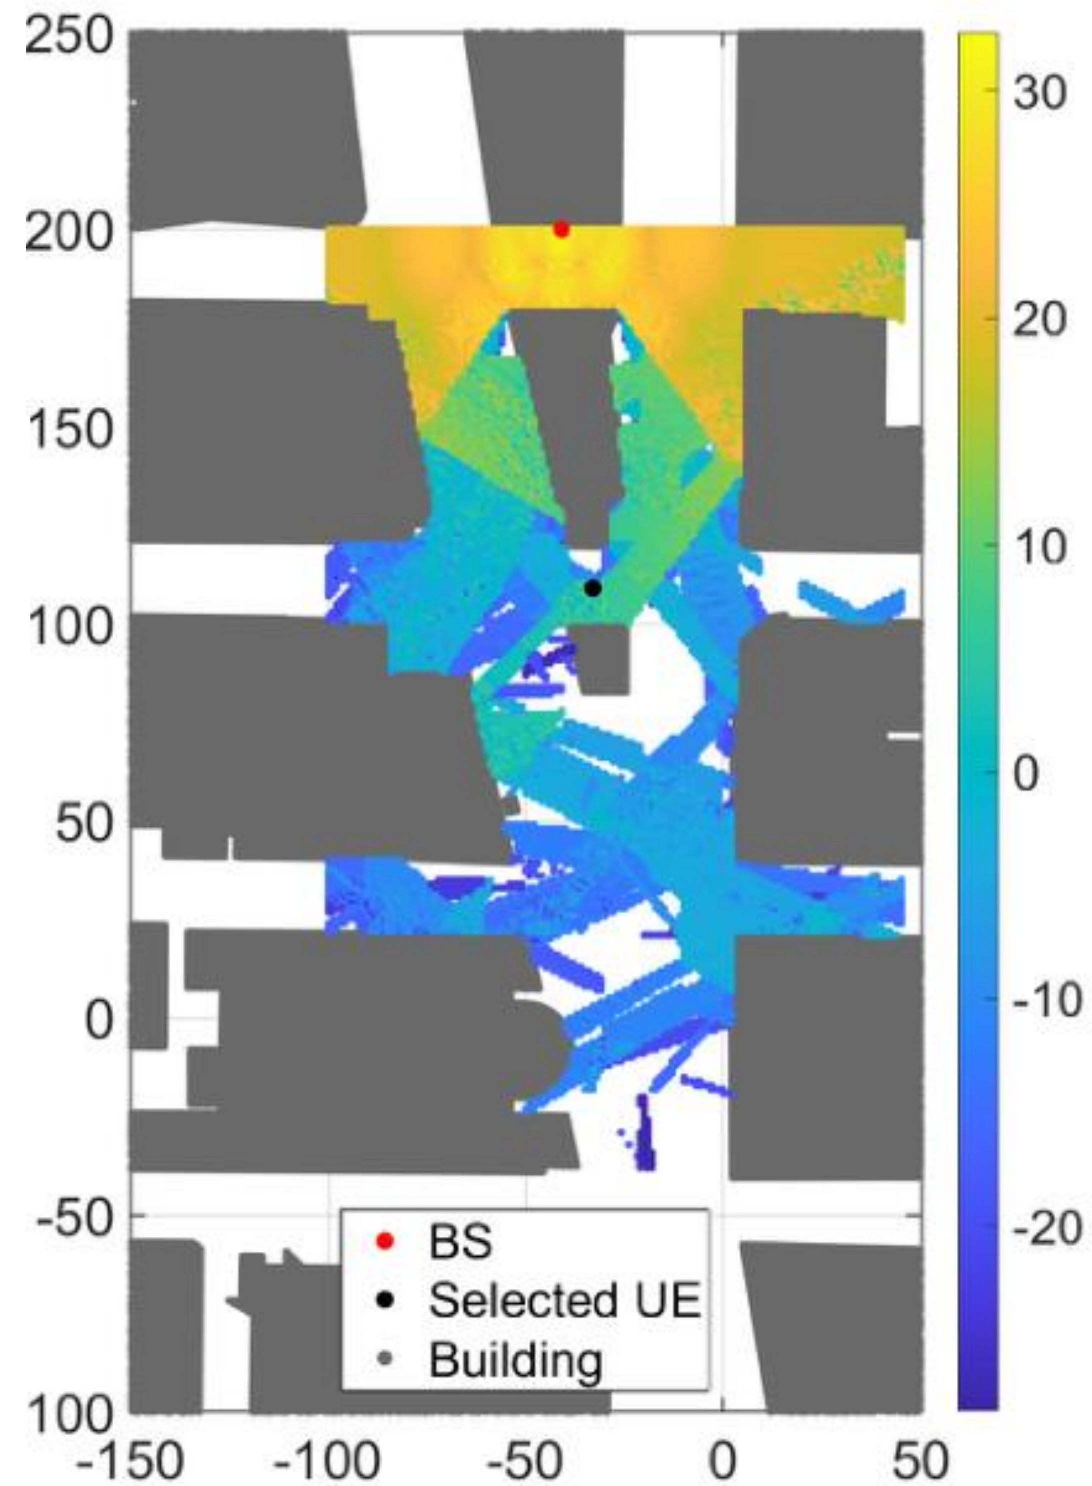

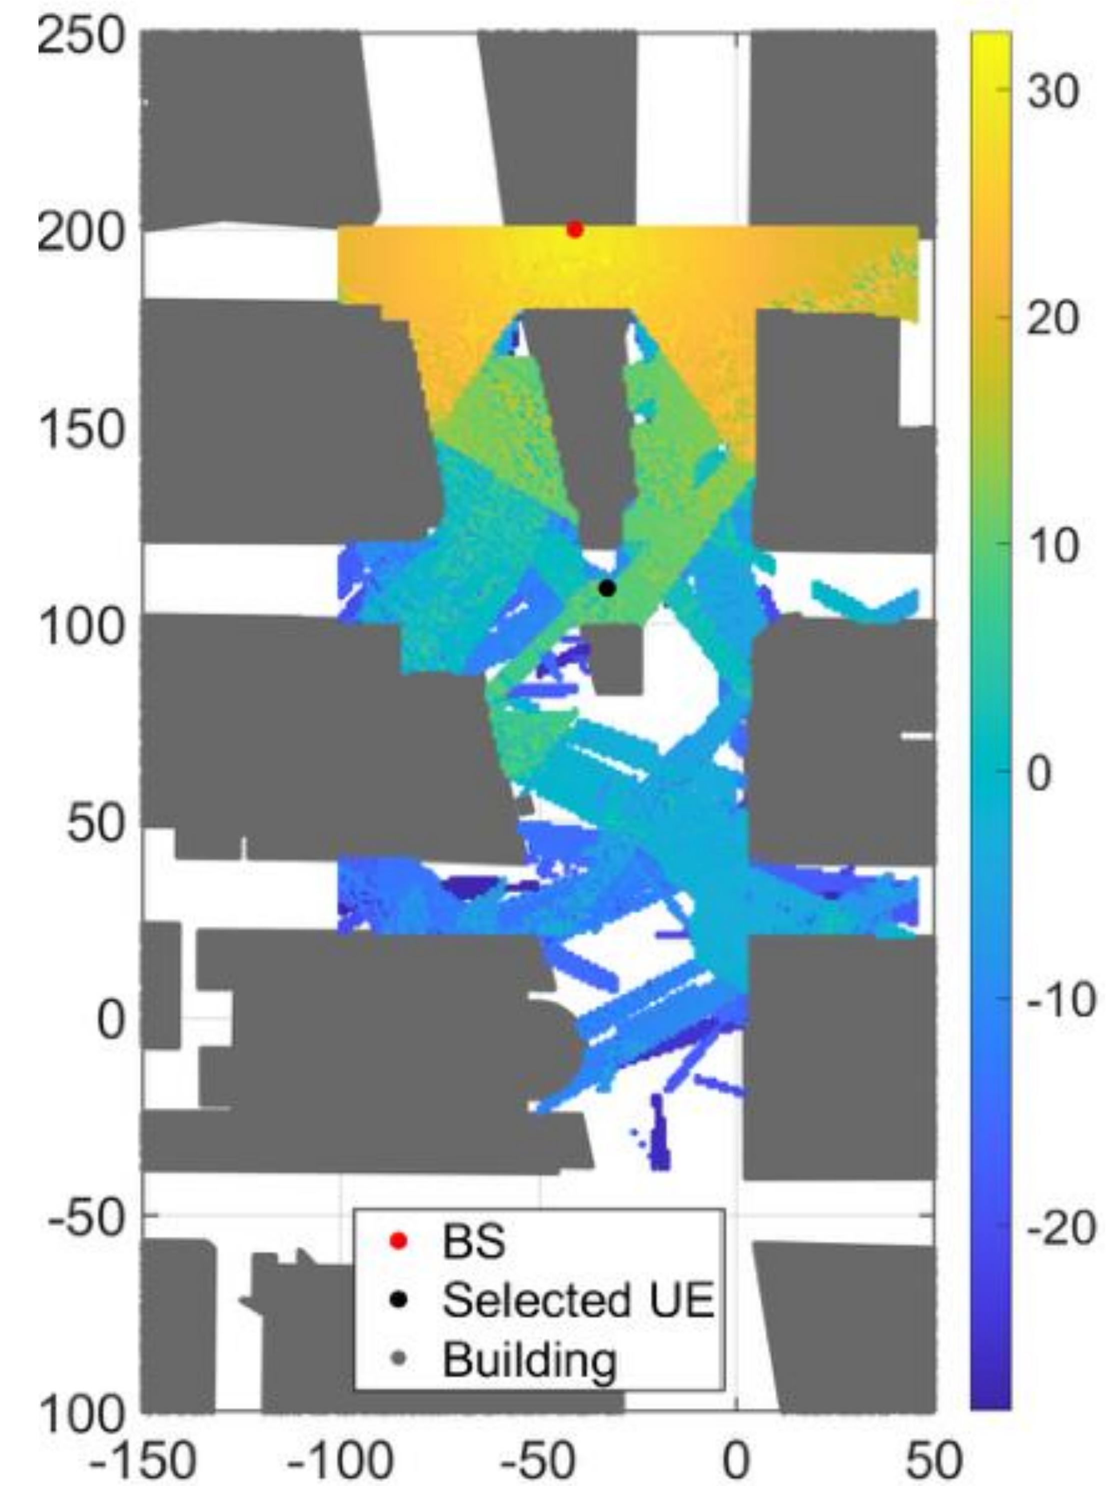

Learned beam

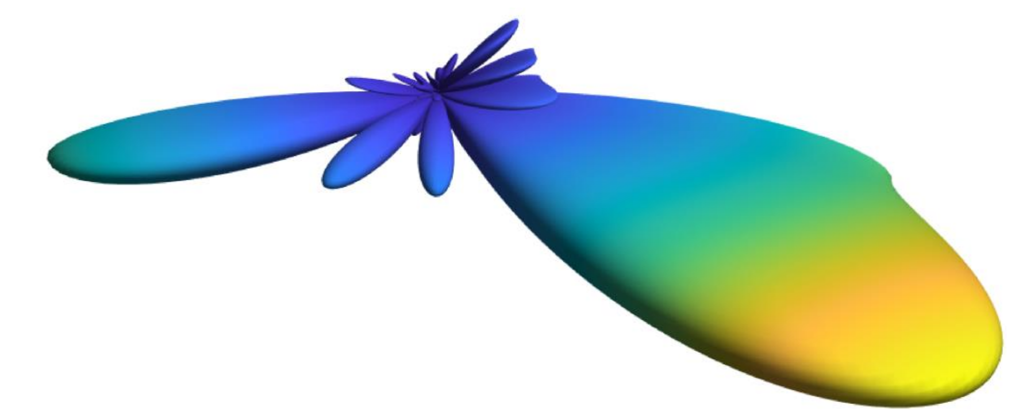

Adaptive, multi-lobe beam

Standard single-lobe beams are less efficient and fail to navigate site-specific blockages.

The learned beams form multi-lobe shapes to exploit specific urban propagation paths.

Conclusion and future work

- We propose a digital twin aided codebook learning framework for mmWave MIMO
  - Generate synthetic channel data from a site-specific digital twin
  - Apply synthetic channels to data-driven codebook learning approaches
  - Leverage digital twin geometry to learn separate LoS/NLoS codebooks
- Key findings
  - Learned codebooks significantly outperform DFT in both LoS and NLoS regions
  - Digital twin based codebook adapt to environment geometry and user distribution
  - Ray-tracing accuracy is the critical fidelity factor
- Future work
  - Fidelity study: systematically vary geometry, material, ray-tracing, hardware modeling fidelity level
  - Refinement strategy: fine-tune the digital twin codebook with a small number of real-world measurements

The code and dataset files of this paper will be available at www.wi-lab.net

Q&A

Thank you!

Email: h.luo@asu.edu