

# A Physics-Informed Digital Twin Framework for Calibrated Sim-to-Real FMCW Radar Occupancy Estimation

Huy Trinh  
University of Waterloo  
h3trinh@uwaterloo.ca

Sebastian Ratto V  
University of Waterloo  
srattovalderrama@uwaterloo.ca

Elliot Creager  
University of Waterloo  
creager@uwaterloo.ca

George Shaker  
University of Waterloo  
gshaker@uwaterloo.ca

## Abstract

*Learning robust radar perception models directly from real measurements is costly due to the need for controlled experiments, repeated calibration, and extensive annotation. This paper proposes a lightweight simulation-to-real framework that can enable reliable Frequency Modulated Continuous Wave (FMCW) radar occupancy detection and people counting using only a physics-informed geometric simulator plus a small unlabeled real calibration set. A calibrated domain randomization (CDR) step is introduced that aligns the global noise-floor statistics of simulated range-Doppler (RD) maps to those observed in real environments, while preserving discriminative micro-Doppler structure. Across real-world evaluations, ResNet18 models trained purely on CDR-adjusted simulation achieve 97% accuracy for occupancy detection and 72% accuracy for people counting, outperforming ray tracing baseline simulation and conventional random domain-randomization baselines.*

## 1. Introduction

High-fidelity simulators and radar digital twins based on full-wave or ray-tracing electromagnetic solvers promise scalable synthetic data generation in various perception applications. Given a geometric and electromagnetic model of the scene, a simulator can synthesize large quantities of labeled radar data under controlled conditions, including difficult or safety-critical scenarios [6, 12]. Recent works have leveraged full-physics simulation [3, 4], motion-capture driven micro-Doppler synthesis [19], or Generative Adversarial Network (GAN) based generation [7, 11]. In addition, various studies focus on bridging the gap between simulation and reality on object detection, gesture, and human activity recognition [2, 8, 10]. These works often depend on

proprietary electromagnetic toolchains, extensive parameter sweeps, and specific domain knowledge tuning. This paper proposes a one-shot, calibration-based noise-floor alignment between simulation and reality to enable efficient sim-to-real transfer for FMCW radar occupancy detection and people counting. To support this, we employ a modular physics-informed digital twin framework that integrates human animation, electromagnetic modelling, and automated radar data generation. The remainder of this paper is organized as follows. Section 2 details the digital twin simulated radar generation, real data collection, calibration formulation, and learning setup. Section 3 presents results, and Section 4 concludes and discusses future research.

## 2. Methodology

### 2.1. Physics Informed Digital Twin for Radar Data Generation

A physics-informed digital twin of the monitored corridor was developed using an electromagnetic simulation platform that combines Shooting and Bouncing Rays (SBR) and Physical Optics (PO) methods [14, 18]. The corridor environment, including walls, ceiling, and floor, was reconstructed with measured dielectric properties to capture realistic multipath propagation and clutter effects. A virtual 60 GHz FMCW radar modeled after the Infineon BGT60TR13C sensor [9] replicated the antenna configuration, polarization, and chirp parameters of the physical hardware. The radar was emulated with a carrier frequency of 60 GHz and bandwidth of 882.35 MHz. The sampling rate and chirp duration were set to 1 MHz and 0.256 ms, respectively, yielding 256 samples per chirp and a chirp slope of approximately  $3.44 \times 10^{12}$  Hz/s. These parameters correspond to a theoretical range resolution of 0.17 m and a maximum unambiguous range of 21.7 m, consistent with the operating characteristics of the real device. Human motion

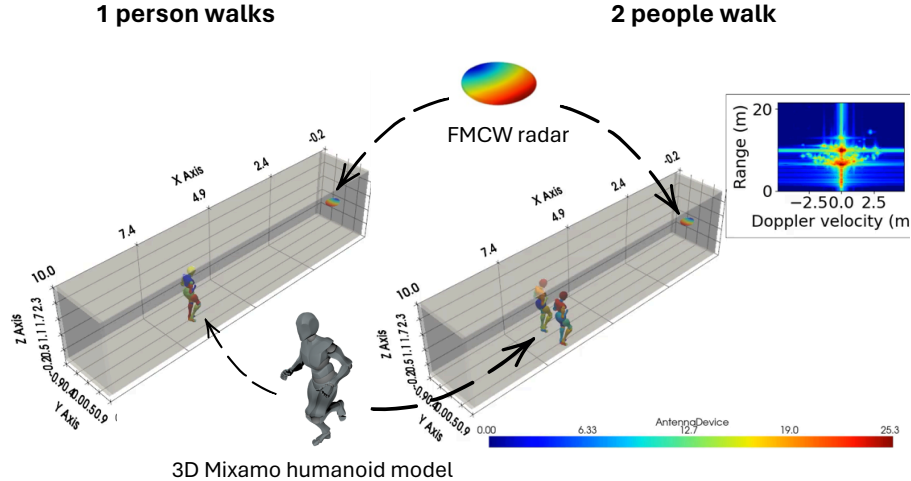


Figure 1. Simulation setup showing the created radar scenes. A 3D Mixamo humanoid model is animated to represent walking motion and imported into the virtual corridor environment. The FMCW radar observes different scenarios—(left) a single walking subject and (right) two subjects walking at different speeds, producing corresponding range–Doppler signatures.

dynamics were introduced using animated 3D humanoid meshes derived from Mixamo and CMU Motion Capture datasets [1, 5]. These animated models provided time-varying radar cross sections and realistic trajectories for walking and interaction behaviours. The electromagnetic solver iteratively computed scattered fields through multiple ray interactions until the residual reflected power dropped below a predefined threshold. The simulated radar outputs were organized as complex data cubes containing amplitude and phase information across fast time, slow time, and frequency dimensions. Range Doppler maps were then generated for representative scenes, including static clutter, a single walking subject, and two subjects moving toward the radar, as Fig. 1. All simulated outputs were exported as floating-point arrays for subsequent machine learning analysis and calibration with real radar measurements. In the remainder of the paper, we refer to this method as **ray-tracing baseline (RT baseline)** or **RT simulation**.

## 2.2. Real Data Collection and Processing Pipelines

A real-world FMCW radar dataset is collected in the corridor using the same Infineon BGT60TR13C radar [9] and configuration specified above. We record three occupancy states: empty room, one person walking, and two people walking the radar sensor as seen in Fig. 2. An additional short unlabeled calibration sequence (10 seconds or 100 frames) of an empty corridor is acquired once and used only to estimate the real noise-floor statistics for CDR (Sec. 2.3). Both simulated and real data are applied with a 2D FFT applied across fast-time and slow-time dimensions to obtain a complex Range-Doppler map, and their magnitudes are converted to dB [17]. The real data processing stage



Figure 2. Experiment setup for 2 people walking scenario.

does not go through any randomization and is used exclusively for testing purposes. They are then clipped to a fixed dB range  $[v_{\min}, v_{\max}]$  estimated from simulation, normalized to  $[0, 1]$ , and mapped via a perceptually uniform colormap (viridis) to form 3-channel images for the input of the neural network model (ResNet18 with random initialization and a task-specific classifier head). The final fully connected layer is replaced with a task-specific classifier: 2 units for binary occupancy or 3 units for people counting applications. Training is performed on simulated RD maps: either (i) RT simulation, (ii) RT simulation with Random Domain Randomization (Random DR), or (iii) RT simulation with the proposed Calibrated Domain Randomization (CDR). The real labeled test set is held out entirely for sim-to-real evaluation. Real RD maps at test time are passed through the same clip-and-normalize pipeline as the simulated inputs (Fig. 3).

## 2.3. Noise-Floor Domain Randomization and CDR

Inspired by domain randomization (DR) in computer vision and robotics [13, 15, 16], our approach introduces synthetic noise to bridge this gap. Both our baseline Random DR and a proposed Calibrated DR (CDR) method apply a per-frame, per-cell clamp in the dB domain, before clipping and normalization:

$$\tilde{S}(i, j) = \max(S(i, j), N(i, j)). \quad (1)$$

This operation raises the overly clean simulated background cells  $S(i, j)$  to a sampled noise-floor level  $N(i, j)$  with range-Doppler cell  $(i, j)$ . The two methods differ only in how they sample the noise map  $N(i, j)$ . For the Random DR baseline, we emulate generic domain randomization by sampling:

$$N(i, j) \sim \mathcal{N}(m, s^2) \quad (2)$$

with mean  $m$  from  $[-130, -100]$  dB and a standard deviation  $s$  from  $[1, 4]$  dB. These ranges approximate typical background levels and variability observed in our quantitative indoor FMCW experiment. CDR keeps the same clamping form (1) but replaces heuristic sampling with a one-shot, data-driven calibration from the unlabeled real empty-room calibration set by aggregating all RD magnitudes and computing their empirical mean  $m_t$  and standard deviation  $s_t$ . CDR then samples

$$N_t(i, j) \sim \mathcal{N}(m_t, s_t^2) \quad (3)$$

i.i.d. over  $(i, j)$  and applies the same clamp  $\tilde{S}(i, j) = \max(S(i, j), N_t(i, j))$ . Fig. 4 shows that CDR shifts the simulated histogram toward the real empty-room distribution, while baseline simulation underestimates the noise floor. Fig. 5 further illustrates how CDR transforms overly clean simulated RD maps into clutter patterns that visually match real measurements while maintaining Doppler walking signatures. Fig. 6 confirms that CDR moves the ratio of low-Doppler (clutter) energy to total energy for simulated samples closer to real data. Finally, to better understand how CDR alters the learned representation, Fig. 7 visualizes channel-averaged activations from the first residual block of ResNet-18 for representative simulated and real RD inputs. When applied to real measurements, the baseline model produces diffuse responses dominated by corridor clutter and simulator-specific artifacts. In contrast, the CDR-trained model concentrates its responses on the range-Doppler cells associated with human motion while suppressing background structure in both domains.




Figure 7. Qualitative visualization of early feature maps from the first residual block of ResNet-18 for representative simulated and real radar range-Doppler inputs. (a,d) inputs; (b,e) features from the baseline RT model trained only on unmodified simulation; (c,f) features from the CDR-trained model.

## 3. Results

On the real corridor test set, the baseline RT model trained solely on unmodified simulation struggles to generalize. Fig. 8 shows that it misclassifies many occupied frames as empty due to mismatched clutter and noise, resulting in a balanced accuracy close to 50% (chance-level for a binary task). The introduction of uncalibrated Random DR improves robustness: the model becomes less sensitive to background mismatch and achieves balanced accuracy around 83%. However, the resulting RD samples sometimes exhibit unrealistic structure, and performance is still limited by residual sim-to-real discrepancies. CDR provides a further and substantial improvement. By aligning simulated noise floors to real statistics while preserving micro-Doppler content, the CDR-trained model achieves balanced accuracy  $\approx 97\%$  and high overall accuracy. The confusion matrices (Fig. 8) show that CDR dramatically reduces both false negatives and false positives compared to the baselines. The people counting task is more challenging due to the overlap between one- and two-person signatures and limited data. Simulated training remains feasible by using same pipeline, while the real test set is used exclusively for evaluation. As shown in Fig. 9, the baseline RT model again fails to transfer, achieving balanced accuracy near 33% (random guessing) Random DR improves performance to balanced accuracy  $\approx 61\%$ , indicating that generic noise-floor perturbation offers some regularization but does not fully align domains, and heavily confuses among 3 classes. CDR yields the best performance with balanced accuracy  $\approx 72\%$  and improved confusion matrices, particularly for distinguishing one vs. two-person cases. Misclassifications mainly occur in ambiguous scenes with partial overlap or occlusion in Doppler.

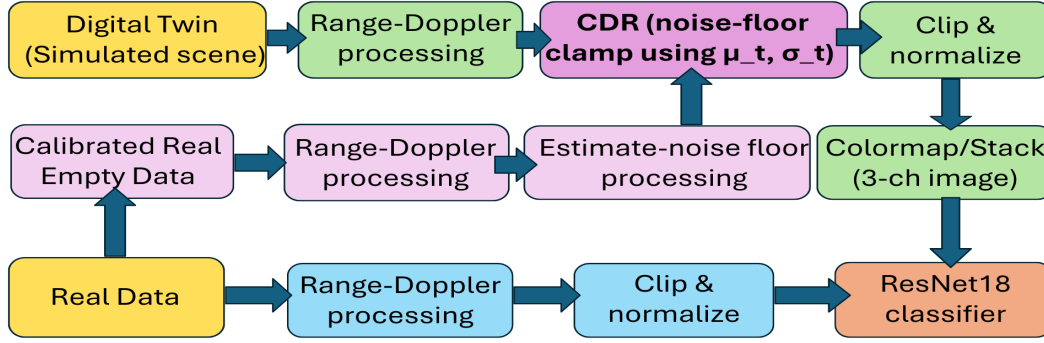


Figure 3. Overview of the proposed sim-to-real pipeline. **Yellow blocks:** Start of simulated data and real data preprocessing **Green path:** simulated radar frames from the digital twin are converted to range–Doppler (RD) maps and passed through the noise-floor CDR block, then clipped, normalized, and colour-mapped for network training. **Purple path:** CDR uses 10 seconds of unlabeled real empty-room data to estimate the target noise-floor statistics that drive the calibration. **Blue path:** real radar frames for evaluation are converted to RD maps and only clipped and normalized, without any randomization or CDR modification.




Figure 4. Effect of CDR on global RD magnitude statistics. Baseline simulation (blue) has quite a different low noise floor. CDR (green) shifts the simulated histogram toward the real empty-room distribution (orange).




Figure 5. Range-Doppler maps for (a,d) baseline simulation, (b,e) simulation after CDR, and (c,f) real measurements for empty room and one-person walking.

#### 4. Conclusion and Future Work

This work introduced calibrated domain randomization as a simple but effective method for sim-to-real transfer in FMCW radar occupancy detection and people counting. By calibrating simulated range–Doppler noise floors to




Figure 6. Static-to-total energy ratio in RD space. Each point is one frame. CDR moves simulated points toward the real cluster, indicating that the global clutter level has been calibrated.

	Ray-Tracing Baseline		Random DR		CDR (proposed)	
empty_room	0	427	425	2	417	10
1_person	0	184	63	121	8	176
	empty_room	1_person	empty_room	1_person	empty_room	1_person

Count

Figure 8. Confusion matrices for real binary occupancy detection across 3 methods.

match global statistics of unlabeled real measurements, CDR significantly improves balanced accuracy on real test data compared to both baseline RT simulation-only training and uncalibrated domain randomization. The approach is lightweight, interpretable, and deployment constraints, making it practical for radar-based human sensing systems built from modern simulators. Future work will focus on enabling angle-of-arrival estimation for people counting and conducting studies on different activity recognition tasks.




Figure 9. Confusion matrices for real people counting across 3 methods.

## References

- [1] Adobe. Mixamo: Characters and animations, 2021. Available: <https://www.mixamo.com/>, Accessed: Aug. 19, 2025. 2
- [2] Oded Bialer and Yuval Haitman. Radsimreal: Bridging the gap between synthetic and real data in radar object detection with simulation. *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 15407–15416, 2024. 1
- [3] Ushemadzero Chipengo, Peter Krenz, and Shawn Carpenter. From antenna design to high fidelity, full physics automotive radar sensor corner case simulation. *Modelling and Simulation in Engineering*, 2018. 1
- [4] Ushemadzero Chipengo, Arien P. Sligar, Stefano Mihai Canta, Markus Goldgruber, Hen Leibovich, and Shawn Carpenter. High fidelity physics simulation-based convolutional neural network for automotive radar target classification using micro-doppler. *IEEE Access*, 9:82597–82617, 2021. 1
- [5] CMU Graphics Lab Motion Capture Database. CMU motion capture database. <http://mocap.cs.cmu.edu/>, 2003–2007. 2
- [6] Abdelrahman Elbadrawy, Hajar Abedi, Omar Ramahi, and George Shaker. Severity of fall detection using digital twin for radar systems. In *2023 IEEE Microwaves, Antennas, and Propagation Conference (MAPCON)*, pages 1–3, 2023. 1
- [7] Barış Filiz Erol, Sevgi Zubeyde Gurbuz, and Moeness G. Amin. Gan-based synthetic radar micro-doppler augmentations for improved human activity recognition. *2019 IEEE Radar Conference (RadarConf)*, pages 1–5, 2019. 1
- [8] Reda El Hail, Pouya Mehrjousesht, Dominique M. M.-P. Schreurs, and Peter Karsmakers. Radar-based human activity recognition: A study on cross-environment robustness. *Electronics*, 2025. 1
- [9] Infineon Technologies AG. Bgt60tr13c rf microelectromechanical sensor datasheet, 2023. [Online; accessed: Jul. 1, 2025]. 1, 2
- [10] Nicolai Kern, Julian Aguilar, Timo Grebner, Benedikt Meinel, and Christian Waldschmidt. Learning on multistatic simulation data for radar-based automotive gesture recognition. *IEEE Transactions on Microwave Theory and Techniques*, 70:5039–5050, 2022. 1
- [11] Mohammed Mahbubur Rahman, Sevgi Z. Gurbuz, and Moeness G. Amin. Physics-aware generative adversarial networks for radar-based human activity recognition. *IEEE Transactions on Aerospace and Electronic Systems*, 59(3): 2994–3008, 2023. 1
- [12] Sebastian Ratto, Ahmed N. Sayed, Abdelrahman Elbadrawy, Arien Sligar, Omar Ramahi, and George Shaker. Digital twin-driven radar for human monitoring optimization. In *Extended Reality*, pages 33–45, Cham, 2026. Springer Nature Switzerland. 1
- [13] Fereshteh Sadeghi and Sergey Levine. CAD2RL: real single-image flight without a single real image. In *Robotics: Science and Systems XIII, Massachusetts Institute of Technology, Cambridge, Massachusetts, USA, July 12-16, 2017*, 2017. 3
- [14] V. Sebastian Ratto, Ahmed N. Sayed, Arien Sligar, Omar M. Ramahi, and George Shaker. Radar-based people presence estimation using digital twin simulation framework. In *2025 International Telecommunications Conference (ITC-Egypt)*, pages 164–166, 2025. 1
- [15] Joshua Tobin, Rachel Fong, Alex Ray, Jonas Schneider, Wojciech Zaremba, and P. Abbeel. Domain randomization for transferring deep neural networks from simulation to the real world. *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 23–30, 2017. 3
- [16] Jonathan Tremblay, Aayush Prakash, David Acuna, Mark Brophy, V. Jampani, Cem Anil, Thang To, Eric Cameracci, Shaad Boochoon, and Stan Birchfield. Training deep networks with synthetic data: Bridging the reality gap by domain randomization. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 1082–10828, 2018. 3
- [17] Huy Trinh, Elliot Creager, and George Shaker. Doppler-domain respiratory amplification for semi-static human occupancy detection using low-resolution simo fmcw radar, 2026. 2
- [18] Sebastian Ratto V, Ahmed N. Sayed, Abdelrahman Elbadrawy, Arien Sligar, Omar M. Ramahi, and George Shaker. Digital twin system for realistic radar optimization in healthcare monitoring. In *2025 IEEE 20th International Symposium on Antenna Technology and Applied Electromagnetics (ANTEM)*, pages 182–185, 2025. 1
- [19] Sahil Waqar and Matthias Pätzold. A simulation-based framework for the design of human activity recognition systems using radar sensors. *IEEE Internet of Things Journal*, 11(8):14494–14507, 2024. 1