

WiP: Super-Resolution mmWave Normal Estimation

Laura Dodds¹, Nazish Naeem¹, Jizheng He¹, Matt Reynolds², Fadel Adib^{1,3}

ldodds@mit.edu, nazishn@mit.edu, jizhengh@mit.edu, msreynol@uw.edu, fadel@mit.edu

¹ Massachusetts Institute of Technology ² University of Washington ³ Cartesian Systems

Abstract

Recent research has introduced a novel method for millimeter-wave (mmWave) imaging that leverages surface normal estimation for 3D object reconstruction. Unlike traditional voxel-based or ray-tracing methods, this approach estimates surfaces directly, demonstrating dramatically improved resolution. However, despite its impressive performance for reconstructing relatively smooth surfaces, this new approach suffers from normal-based interference, where closely spaced surfaces with different normals, such as the stitches on a baseball, create interference in the normal estimation process, resulting in inaccurate reconstructions.

In this work-in-progress paper, we present a mmWave super-resolution reconstruction method. Our method introduces normal vectors as a new dimension of separability, allowing us to isolate reflections from closely spaced surface components through their different orientations. To do so, since surface components with different normal vectors will reflect mmWave signals in different directions, we separately analyze portions of the aperture to isolate each surface component's reflections. This allows us to accurately estimate normal vectors in high-interference regimes, with our preliminary results showing that this leads to strong reconstruction performance compared to our state-of-the-art baseline, with applications spanning robotics, shipping and logistics, and augmented reality.

CCS Concepts

• Computer systems organization → Sensor networks.

ACM Reference Format:

Laura Dodds, Nazish Naeem, Jizheng He, Matt Reynolds, Fadel Adib. 2026. WiP: Super-Resolution mmWave Normal Estimation. In *The 32nd Annual International Conference on Mobile Computing and Networking (MobiCom '26)*, October 26–30, 2026, Austin, TX, USA. ACM, New York, NY, USA, 4 pages. <https://doi.org/10.1145/3795866.3844760>

1 Introduction

Recent years have witnessed growing interest in using millimeter wave (mmWave) signals to reconstruct fully occluded 3D objects [1, 3, 4, 7, 8]. Since mmWave signals can traverse




Figure 1: Normal-Based Imaging a) For smooth surfaces, our recent work enables high-fidelity 3D reconstruction compared to prior mmWave backprojection methods. b) Fine surface features (e.g., baseball seams) can create interference and distort reconstructions.

many everyday occlusions (e.g., cardboard, plastic, fabric) and reflect off objects hidden behind them, they can enable 3D reconstruction of objects entirely hidden from view. This capability can enable novel applications in domains spanning robotics, augmented reality, and shipping and logistics.

However, most mmWave reconstruction methods fail to achieve high-accuracy reconstruction due to their reliance on occupancy-based estimation methods. These methods, such as backprojection [4] or range-migration algorithm [6, 10], produce a 3D image by estimating the reflected power from each voxel. This image is then converted to a 3D point cloud by selecting all voxels above a power threshold [2, 4, 5]. Unfortunately, their accuracy is fundamentally limited by the commercially-available bandwidth, which causes reflection power to spread along the depth dimension and produce significant artifacts above and below the object (Fig. 1a).

To overcome this challenge, our recent work, mmNorm [3], not only estimates occupancy, but also the orientation of the surface. It does so by estimating normal vectors at each voxel, then using these normals to reconstruct the final surface, resulting in much higher fidelity reconstructions (Fig. 1a).

However, one of the key limitations of this approach (as well as occupancy-based methods) arises from normal-based interference, where closely spaced surfaces with different orientations, such as the stitches on a baseball, create interference within an image voxel. In this case, mmNorm is unable to separate their reflections, resulting in inaccurate normal estimations and surface reconstructions (Fig. 1b).

In this work-in-progress paper, we introduce a super-resolution mmWave normal estimation method to resolve such interference. Our key insight is that while both occupancy-based methods and mmNorm resolve interference through the location of reflectors,¹ we add a new dimension of separability to the imaging problem: normal vectors. Specifically, even when two surfaces are located arbitrarily close

¹Note that while mmNorm performs normal estimation, it still requires surfaces of different orientations to be spatially separated.



This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License.

MobiCom '26, Austin, TX, USA

© 2026 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-2505-0/2026/10

<https://doi.org/10.1145/3795866.3844760>




Figure 2: Antennas mostly receive reflections from points with normals pointed to antenna.




Figure 3: Super-Resolution Normal Estimation. a) A 2D slice of a mmWave image for two surfaces that do not interfere, where b) our prior normal estimation produces accurate normals. c) A mmWave image for two interfering surfaces, where d) our prior normal estimation fails but e) our proposed super-resolution normal estimation succeeds.

in 3D space, they may have different normal vectors. Since mmWave signals exhibit primarily specular reflections, these different normal vectors will cause the two surfaces to reflect signals in different directions. By exploiting this directional diversity in the received reflections, we can isolate reflections from individual surface components and compute more accurate normal estimates despite spatial overlap.

2 Primer

mmWave Imaging: The majority of mmWave imaging methods operate by collecting frequency-modulated continuous wave (FMCW) measurements across either a physical or synthetic aperture, and then forming a 3D image. Formally, the reflected power, $P(v)$, at each voxel v is [4]:

$$P(v) = \sum_{a=1}^A \sum_{t=1}^T h(a, t) e^{j2\pi d_{v,a}/\lambda_t} \quad (1)$$

where $h(a, t)$ is the t^{th} sample from the a^{th} antenna location (p_a). $d_{v,a} = 2||p_a - v||$ is the round trip distance from v to p_a .

Then, to construct a 3D point cloud, existing methods typically select all voxels in the 3D mmWave image above a power threshold [5]. However, the reconstruction accuracy of this approach is directly limited by the commercially available ISM bandwidth of 4 GHz (at the 77 GHz band). This results in a limited depth resolution of ~ 4 cm, leading to amorphous 3D shapes as shown in Fig. 1b.

Normal-Based Reconstruction: To overcome this limitation, our recent work, mmNorm [3], estimates surface normal vectors to enable higher accuracy 3D reconstruction. To do so, mmNorm relies on the fact that mmWave signals experience primarily specular ("mirror-like") reflections [9]. As shown in Fig. 2, an antenna (black triangle) will primarily receive reflections from the portion of the object's surface (light blue) which has a normal (red arrow) pointing directly towards the antenna, while other portions will primarily reflect signal away (purple).

mmNorm leverages this relationship to perform surface normal estimation. At a high level, for a given voxel, it estimates antenna location(s) which receive the strongest reflection from this voxel, and estimates the normal pointing towards that location. For more details, please refer to [3].

Normal Interference: One primary limitation of mmNorm is normal-based interference. Consider the illustrative example shown in Fig. 3, with two surfaces shown in black and the synthetic aperture shown in pink. In Fig. 3a, we show




Figure 4: We use direction-dependent weighting to resolve interference.

a simulated mmWave image where red and blue represents areas of high and low reflection power, respectively. Since the two surfaces are spaced far apart, their respective red images do not overlap, and mmNorm successfully estimates normal vector fields for both surfaces (Fig. 3b). Now consider a second scenario in Fig. 3c, with two surface components located such that their images overlap, where many voxels in the red region result from reflections from both surfaces. In this case, mmNorm is unable to estimate the direction of highest reflection power for each voxel, resulting in inaccurate normal estimates (red circle in Fig. 3d). This leads to inaccurate reconstructions similar to Fig. 1b.

3 System Design

Our approach aims to resolve normal interference in two stages: multi-normal estimation and surface reconstruction.

3.1 Multi-Normal Estimation

At a high level, we aim to resolve interference by leveraging normal vectors as a new dimension of separability. While existing approaches (backprojection and mmNorm) only attempt to separate surfaces based on their location, we instead separate surfaces through both their location and orientation. To do so, we rely on the fact that, due to specularity, two interfering surfaces with different normals (such as those shown in Fig. 3c) will reflect signals towards different portions of an aperture. Thus, by separately analyzing portions of the aperture, we can isolate reflections from different surfaces, even when they originate from the same image voxel.

To do so, we explain our approach through the illustrative examples in Fig. 4. We define our image as a function of not only voxel location (purple), but also surface orientation (defined by a normal vector in blue). Then, for each voxel location and orientation pair, instead of naively combining responses from the entire aperture, we apply a spatial weighting function (shown by the shade of the triangle antennas) which emphasizes radar locations in the direction of the normal vector. Then, if the same voxel location is affected by two different surfaces, we can separate their individual images through two different surface orientations. Formally,




Figure 5: Super-Resolution Surface Reconstruction. Our preliminary pipeline a) computes a multi-normal vector field, b) clusters the normals, c) proposes valid surfaces, and d) selects a final surface.

our normal-based image $P(v, n)$ at voxel v and normal n is:

$$P(v, n) = \sum_{a=1}^A \sum_{t=1}^T h(a, t) \text{sinc}(2\pi L_c \sin(\theta_{v,a,n})/\lambda_t) e^{j2\pi d_{v,a}/\lambda_t} \quad (2)$$

where the $\text{sinc}(2\pi L_c \sin(\theta_{v,a,n})/\lambda_t)$ term is the spatial weight function that emphasizes antennas in the n direction.

Once we compute this image, we can use it to derive a normal vector field which contains multiple normal vectors for each voxel (corresponding to different surface components), which we call our *multi-normal vector field* (MNVF). We do so by selecting up to K normals for each voxel which produce the strongest responses (e.g., using `scipy's find_peaks` function). For example, Fig. 3e shows our MNVF, with normal vectors colored based on their orientation. Now, many points now contain two vectors, where one points left (shown in teal) and corresponds to the right surface, and another points right (shown in purple) and corresponds to the left surface. Thus, our approach allows us to successfully isolate the interference of the two surfaces and estimate accurate normal vectors in a scenario where mmNorm failed.

3.2 (In Progress) Surface Reconstruction

Building on our super-resolution normal estimation, we aim to enable high-resolution 3D surface reconstruction from mmWave measurements. Unlike our prior work [3], where each voxel was assigned a single normal, our approach can estimate multiple surface normals for each voxel (corresponding to multiple interfering surface component). While this allows us to separate interfering surfaces, it also introduces additional complexity when inverting these normal vectors to recover the object's shape, since there are many more surfaces consistent with the more complex field.

To overcome this challenge, our super-resolution reconstruction follows three main steps, shown in Fig. 5:

- (1) First, we estimate our MNVF as described in §3.1 (Fig. 5a).
- (2) Second, we cluster the normals within the MNVF (using HDBSCAN) to form separate vector fields associated with each surface component (different color normals in Fig. 5b).
- (3) Finally, we apply the normal field inversion (Fig. 5c) and isosurface optimization (Fig. 5d) similar to that proposed in [3] to find surface estimates for each cluster of normals.

4 Preliminary Results

We apply a standard free-space path loss simulation to estimate reflected signals from simulated object geometries, then apply our super-resolution reconstruction method. Fig. 6a




Figure 6: Preliminary Results for a) mmNorm and b) our method.

shows the standard backprojection image computed for the surface shown in black, as well as the surface reconstruction (white) from our baseline, mmNorm [3]. Fig. 6b shows results from our approach, where the heatmap shows the highest power across all normals ($\max_n(P(v, n))$), and the white line is our final reconstruction. We note that while mmNorm suffers from normal interference, resulting in an inaccurate surface reconstruction, our approach can resolve the interference and reconstruct the surface with high accuracy. Furthermore, this simulated object geometry is representative of surface with a small dent (~ 2 cm wide), demonstrating the ability of our proposed method to enable novel capabilities such as through-occlusion quality control.

5 Ongoing Work & Conclusion

While this work-in-progress paper presents preliminary results for our super-resolution normal estimation and reconstruction, we are pursuing several directions to improve upon this initial success. First, we will continue to build on our surface reconstruction pipeline (§3.2) to enable more robust surface estimation, especially for real-world non-idealities. For example, we will add constraints to the isosurface selections to ensure that neighboring surface components are selected to form a continuous surface, without discrete jumps. We will also investigate various clustering methods to ensure robustness to different object sizes, geometries, and noise levels. Second, we will investigate various weighting functions, and evaluate their impact on reconstruction accuracy and robustness. Finally, we are developing the theoretical foundations of this method. Due to our directional weighting function, our approach introduces a tradeoff between normal and cross-range resolutions. By deriving these resolutions, we can formalize this tradeoff and guide future deployments.

Conclusion: In this work-in-progress paper, we presented a novel super-resolution mmWave normal estimation method, which enables higher accuracy 3D reconstruction of fully occluded objects. Our method introduces surface orientation as a new dimension of separability, allowing it to resolve interfering surfaces with different orientations. Looking forward, we believe this method will enable new applications in logistics, robotics, and augmented reality.

Acknowledgments. We thank the anonymous reviewers and Signal Kinetics group for their help and feedback. This research is sponsored by NSF (Award #2313234), Amazon MIT Science Hub, and the MIT Media Lab. Laura Dodds is supported by the Amazon Robotics Thriving Stars Fellowship.

References

- [1] Daniel Arnitz, Claire M Watts, Andreas Pedross-Engel, and Matthew S Reynolds. 2022. A portable K-band (24 GHz) 3-D millimeter wave imaging system for detecting and dimensioning hidden objects. In *Passive and Active Millimeter-Wave Imaging XXV*, Vol. 12111. SPIE, 38–42.
- [2] Pingping Cai and Sanjib Sur. 2023. Millipcd: Beyond traditional vision indoor point cloud generation via handheld millimeter-wave devices. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 6, 4 (2023), 1–24.
- [3] Laura Dodds, Tara Boroushaki, Kaichen Zhou, and Fadel Adib. 2025. *Non-Line-of-Sight 3D Object Reconstruction via mmWave Surface Normal Estimation*. Association for Computing Machinery, New York, NY, USA, 445–458. <https://doi.org/10.1145/3711875.3729138>
- [4] Laura Dodds, Hailan Shanbhag, Junfeng Guan, Saurabh Gupta, and Haitham Hassanieh. 2024. Around the Corner mmWave Imaging in Practical Environments. In *Proceedings of the 30th Annual International Conference on Mobile Computing and Networking*. 1–15.
- [5] Junfeng Guan, Sohrab Madani, Suraj Jog, Saurabh Gupta, and Haitham Hassanieh. 2020. Through fog high-resolution imaging using millimeter wave radar. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 11464–11473.
- [6] Qijia Guo, Jie Liang, Tianying Chang, and Hong-Liang Cui. 2019. Millimeter-wave imaging with accelerated super-resolution range migration algorithm. *IEEE Transactions on Microwave Theory and Techniques* 67, 11 (2019), 4610–4621.
- [7] Jiachen Lu, Hailan Shanbhag, and Haitham Al Hassanieh. 2025. GeRaF: Neural Geometry Reconstruction from Radio Frequency Signals. In *Conference on Neural Information Processing Systems (NeurIPS)*.
- [8] Jiachen Lu, Hailan Shanbhag, and Haitham Al Hassanieh. 2026. Seeing through boxes: Non-Line-of-Sight 3D Reconstruction from Radar Signals. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [9] Jonathan S Lu, Patrick Cabrol, Daniel Steinbach, and Ravikumar V Pragada. 2013. Measurement and characterization of various outdoor 60 GHz diffracted and scattered paths. In *MILCOM 2013-2013 IEEE Military Communications Conference*. IEEE, 1238–1243.
- [10] Mou Wang, Shunjun Wei, Jiadian Liang, Xiangfeng Zeng, Chen Wang, Jun Shi, and Xiaoling Zhang. 2021. RMIST-Net: Joint range migration and sparse reconstruction network for 3-D mmW imaging. *IEEE Transactions on Geoscience and Remote Sensing* 60 (2021), 1–17.