Demo: High Resolution Point Clouds from mmWave Radar

Akarsh Prabhakara[★], Tao Jin[★], Arnav Das[†], Gantavya Bhatt[†], Lilly Kumari[†], Elahe Soltanaghai[#], Jeff Bilmes[†], Swarun Kumar[★] and Anthony Rowe[★$]
[★]Carnegie Mellon University [†]University of Washington
[#]University of Illinois, Urbana-Champaign [$]Bosch Research

ABSTRACT

Millimeter wave radars can perceive through occlusions like dust, fog, smoke and clothes. But compared to cameras and lidars, their perception quality is orders of magnitude poorer. RadarHD [3] tackles this problem of poor quality by creating a machine learning super resolution pipeline trained against high quality lidar scans to mimic lidar. RadarHD ingests low resolution radar and generates high quality *lidar-like* point clouds even in occluded settings. RadarHD can also make use of the high quality output for typical robotics tasks like odometry, mapping and classification using conventional lidar workflows. Here, we demonstrate the effectiveness of RadarHD's point clouds against lidar in occluded settings.

1 INTRODUCTION

Cameras and lidars, which are the standard sensors for most robotic tasks suffer in the presence of occlusions from particulate matter like dust, smoke and fog. Millimeter wave (mmWave) sensing systems have shown great potential for perceiving through these occlusions [1]. Compact mmWave radars are becoming more ubiquitous for portable, robotic use cases where sensing through occlusions is essential [2].

While mmWave radars can see through occlusions, compact radars yield only poor quality perception. This is due to the compact form factor, which results in poor angular resolution – that is, the ability to tell apart two nearby objects. Poorer the resolution, coarser the perception and performance of all downstream application tasks like odometry, mapping and classification degrades. For comparison, while cameras and lidars have an angular resolution of 0.01° and 0.1°, mmWave radars are at 15°. This results in coarse grained, blobby perception. We tackled this problem in RadarHD [3].

ACM MobiCom '23, October 2–6, 2023, Madrid, Spain

ACM ISBN 978-1-4503-9990-6/23/10.
https://doi.org/10.1145/3570361.3614077

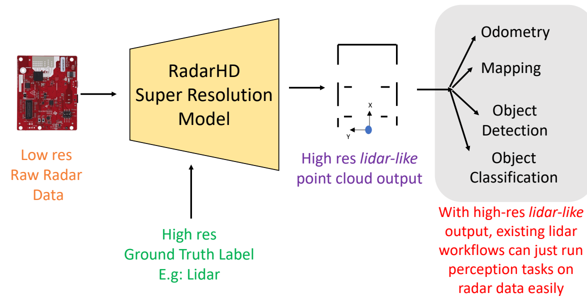

Figure 1: Overview: RadarHD obtains high resolution point clouds from low resolution, compact, mmWave radars to enable high quality perception in occlusions.

RadarHD upsamples the low resolution, compact mmWave radar to obtain *lidar-like* point clouds even in occluded environments. In an effort to bring our robotics work to the wireless audience, we are presenting RadarHD as a demo.

As a motivating scenario, we show in [3] that we can enable a futuristic fire fighting robot to perceive and navigate through thick smoke using compact, mmWave radars. Unlike past work which uses accurate trajectory for synthetic aperture super resolution [4, 5] we make no assumptions on the motion, so the radar can remain static/move unconstrained. This allows RadarHD to be broadly applicable to any freely moving robotic system that needs to see past occlusions.

Demo: https://youtu.be/me8ozpgyy0M
GitHub: https://github.com/akarsh-prabhakara/RadarHD

2 RADARHD OVERVIEW

Before we describe our demo setup, we will briefly describe how RadarHD tackles the resolution challenge.

RadarHD performs machine learning driven super resolution to obtain a *general purpose*, high resolution *point cloud*. We achieve this by training low resolution radar I/Q with high resolution ground truth lidar (Fig. 1) using a custom designed super resolution model. Learning to obtain a *general purpose* representation strives to replace lidar in lidar denied scenes but still obtain *lidar-like* point clouds that can go through any lidar processing workflows. This makes it easy for robotic system integration.

The key challenges encountered in designing the supervised learning framework are all to do with the nature of radar data being different from vision data where super resolution has been extensively studied. First, we have a preprocessing step that decides how to input the raw I/Q data into the network as "images". Our solution is to perform a low level thresholding on range-azimuth radar heatmap that filters out noise but allows most of the information rich, radar artifacts arising due to poor resolution to be captured in the radar input image.

Second, we enable the model to understand these *sinc-like* artifacts and remove them. We achieve this by choosing a polar representation of range-azimuth for the input image, as opposed to Cartesian. This makes the artifacts occur along a row (along azimuth) instead of spreading angularly (in Cartesian). This helps learning blocks like convolutional layers to capture these artifact patterns more effectively.

Lastly, in contrast to camera images where neighboring pixels are similar, visualizing a time series of low resolution radar images reveals that they are sparse and suffer from specular noise from frame to frame. RadarHD tackles this by choosing an encoder-decoder U-net based architecture. The encoder learns the radar artifacts and removes them to create a semantic understanding of the scene. The decoder upsamples the artifact-free semantic understanding to yield *lidar-like* images. We account for specular noise by making U-net input a history of frames to look for persistent reflections and ignore fluttering reflections. To get lidar-like images with sharp edges and boundaries, in addition to Binary Cross Entropy, we also make use of Dice loss that promotes sharpness. [3] describes our methodology in detail.

3 DEMO SETUP

Sensor Hardware: We use TI AWR1843BOOST as the low resolution radar. We deploy it along with DCA1000EVM to get raw radar I/Q data. We use Ouster OS-0 64 beam lidar for ground truth. The lidar and radar assembly will be mounted on a rigid body that is portable.

Software: We obtain raw I/Q radar and lidar data over ethernet. We then process the packets as radar and lidar frames and take it through RadarHD pipeline. For demo purposes, lidar frames are visualized as is. Radar frames go through the inference pipeline on a pre-trained model to obtain *lidar-like* images which are then thresholded to convert to *lidar-like* point clouds for visualization. The raw radar, lidar and RadarHD point clouds will be visualized similar to this.

Scenes: We will show the effectiveness of RadarHD inferences on the following different types of scenes (Fig. 2).

- Static radar-lidar assembly overlooking a corner of a room as a baseline.
- Static radar-lidar assembly covered with a cloth curtain to show a proof of concept occluded setting.

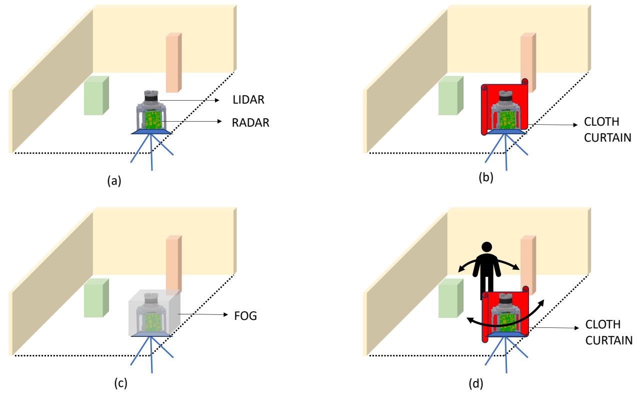

Figure 2: We demo RadarHD against lidar in (a) static (b) static but occluded by cloth curtain (c) static but occluded by fog (d) allow people to move and the radar-lidar assembly to move around in the region.

- Static radar-lidar assembly covered with a self contained chamber with dense fog. We will generate the fog using fog machines. We will fully comply with venue policies on the use of fog machines by coordinating with the organizers.
- Radar-lidar assembly with curtain/fog overlooking a corner where people can walk by and test the system.
- Moving the radar-lidar assembly over a region to show that the point clouds from walls and anything else in the region moves accordingly.

Requirements: We would like to get a portion of room (5m x 5m, maybe even a corner of a room) where we could setup our radar-lidar assembly and add objects like chairs and tripods to create an interesting scene. Besides this, generic requirements include power, WiFi (that allows NTP), and tables (to setup display monitors for visualization).

4 CONCLUSION

We present a demonstration of RadarHD - a solution to obtain high quality point clouds from a cheap, compact, mmWave radar. RadarHD achieves this by designing a machine learning super resolution model trained against high quality lidar. We show the effectiveness of RadarHD against lidar in occluded settings such as with a cloth curtain and fog. RadarHD paves the way to engineer sensing applications that require high quality perception even in occluded environments.

Acknowledgements: This research was supported by NSF (2030154, 2106921, 2007786, 1942902, 2111751), Bosch, CONIX Center, DARPA-TRIAD, ARL and CMU-MFI.

REFERENCES

[1] J Guan, S Madani, S Jog, S Gupta, and H Hassanieh. 2020. Through fog high-resolution imaging using millimeter wave radar. In *Proceedings of CVPR*. 11464–11473.

[2] C Lu, M Saputra, P Zhao, Y Almalioglu, P de Gusmao, C Chen, K Sun, N Trigoni, and A Markham. 2020. milliEgo: single-chip mmWave radar

aided egomotion estimation via deep sensor fusion. In *Proceedings of SenSys*. 109–122.
[3] A Prabhakara, T Jin, A Das, G Bhatt, L Kumari, E Soltanaghai, J Bilmes, S Kumar, and A Rowe. 2023. High Resolution Point Clouds from mmWave Radar. In *IEEE ICRA*. 4135–4142.
[4] C Watts, P Lancaster, A Pedross-Engel, J Smith, and M Reynolds. 2016. 2D and 3D millimeter-wave synthetic aperture radar imaging on a PR2 platform. In *IEEE IROS*. 4304–4310.
[5] M Yanik and M Torlak. 2019. Near-field MIMO-SAR millimeter-wave imaging with sparsely sampled aperture data. *IEEE Access* 7 (2019), 31801–31819.