PanoRadar: Enabling Visual Recognition at Radio Frequency

Richard Beattie
MAS.S61 Wireless & Mobile Sensing
Nov 17th, 2025

Objectives

Learn about mmWave 3D sensing

- In what ways is RF sensing better than optical sensing / LiDAR?
- What are the challenges with current RF sensing methods?
- How to improve azimuth radar resolution with SAR
- Can we mount radars to moving platforms
- Improving elevation resolution with machine-learning

What do we need to sense?

Robots

Self-driving cars

Humans

What’s wrong with cameras / LiDAR?

What's wrong with cameras / LiDAR?

Environmental resilience

Darkness

Glass

Image by freepik

What’s wrong with cameras / LiDAR?

Environmental resilience

Darkness

Glass

& more

Image by freepik

The solution?

mmWave has **longer wavelength.** Doesn’t scatter off obscurants

No dependency on lighting conditions

RF Sensing

Glass is opaque to mmWave

Ok? So why isn't it being used

Poor spatial resolution

- iPhone 17 camera has 48 million pixels on it's camera sensor
  - But these measure **intensity** not phase and frequency
  - For that need **antennas**

- **Rayleigh criterion** – Angular resolution inversely proportional to aperture size

- Antennas spaced λ / 2 apart

- Most RF sensors therefore have small aperture → poor spatial resolution

Ok? So why isn't it being used

Consequences

- Objects in close proximity appear smeared
- Inability to capture fine-grained environmental details
    - Can't do downstream tasks such as semantic segmentation or object detection

How is this currently mitigated?

SAR

Category Priors

Robot Motion

Category Priors

Use prior knowledge to improve resolution

Example: Human body

Issue: Generalization

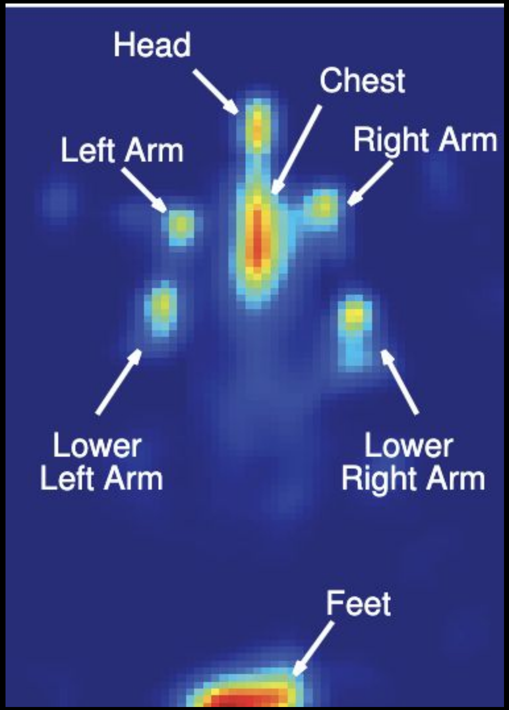

Image: RF-Capture; F. Adib, et. al

Synthetic Aperture Radar (SAR)

Move a radar through space and coherently combine the results to form array of “virtual antennas”

Often put on a linear rail

- Minimal drift or rotation
- Encoders are very accurate

Issues:

- Physical large
- Slow – device has to move; settle; then scan

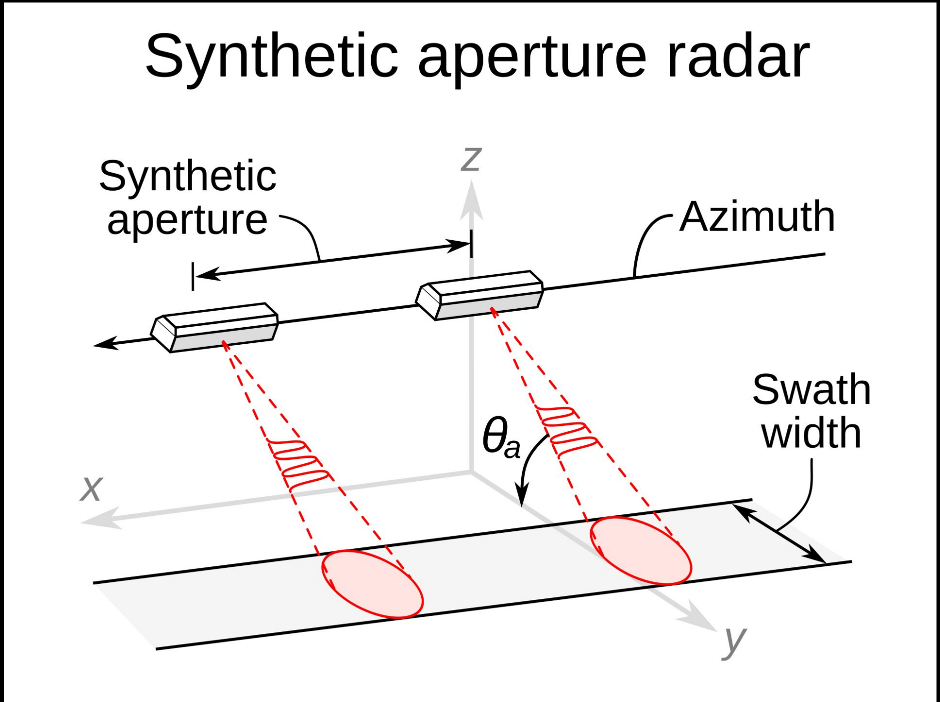

Robot Motion

Use robot’s own motion to form synthetic aperture

Issues:

- Only improves resolution in moving direction
- Only works when robot is moving

PanoRadar

Signal Processing

Hardware Design

3D reconstruction
on mobile robots

Machine Learning

Hardware Design: Cylindrical Array Imaging

Rotates 8 virtual antennas 360 degrees

Scans at 1200 positions

Creates 8x1200 virtual antenna array

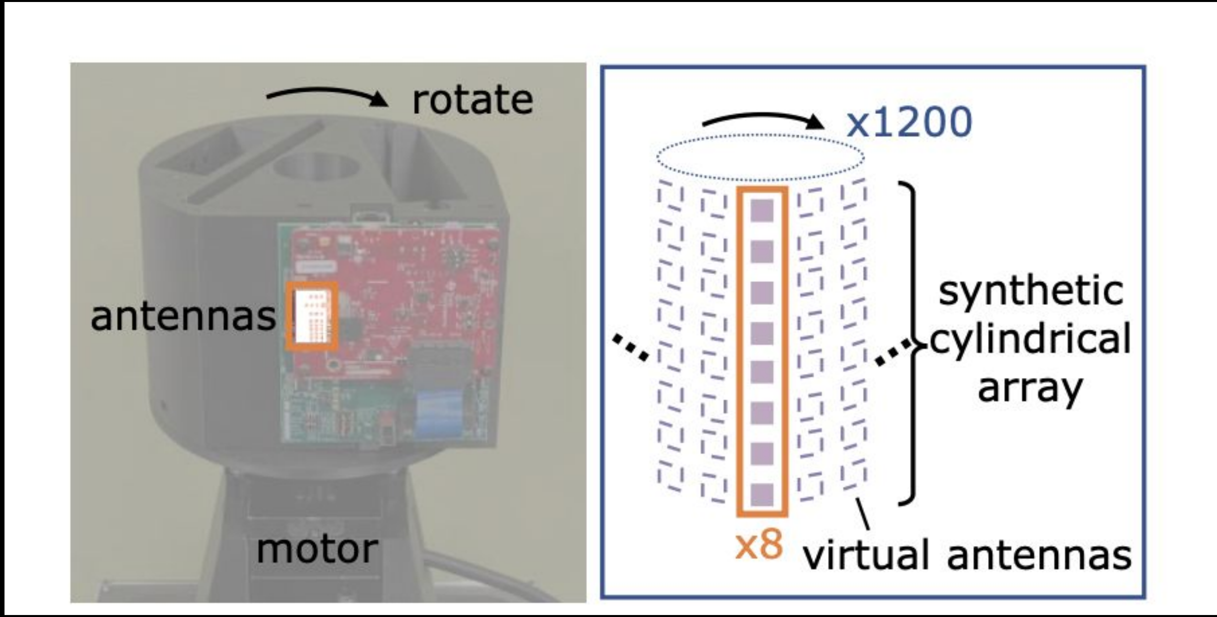

Hardware Design: Cylindrical Array Imaging

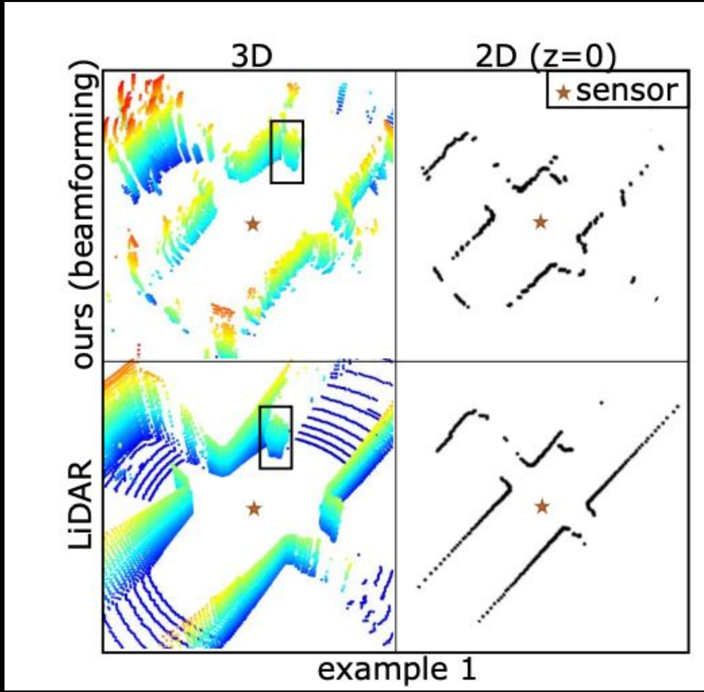

One-shot measurement comparison RF vs. LiDAR

Hardware Design: Cylindrical Array Imaging

Design Advantages?

Hardware Design: Cylindrical Array Imaging

- Azimuth resolution: 0.96 degrees
  - TI AWR1843 without any other hardware normally has ~15 degree res
- 3D imaging: Vertical placement of the linear array enables beamforming along the elevation axis
- Panoramic: Radar has limited 30-60 degree field-of-view. Rotation provides 360 panoramic sensing
- Low cost: ~$500 for COTS radar

Mounting to a robot

With SAR signals arriving at virtual antennas have **phase shift**

- Geometry-induced – contains range + angle information (good)
- Platform-motion induced (bad)

To remove bad phase shifts must know radar location at $\lambda/2$ resolution

- For 79 Ghz wavelength is ~3.8mm so 1.9mm

IMUs / wheel odometers are not accurate enough

Need another way to get estimate the robot's motion

Doppler Effect

As a reflector moves towards or away from radar, the returned waves gets **compressed** or **stretched**, causing a frequency shift

Can use this to estimate the speed of the robot

Doppler Effect

Issue: Doppler gives us the **speed** but not the **direction**

For direction need **Angle-of-Arrival** (AoA)

However, AoA is also determined by frequency shift across measurements

Signal Processing: Untangling Doppler and AoA

PanoRadar introduces signal processing to compensate

Consider a reflector on the x,y-plane with the range R_n and azimuth θ_n

$$d(t) = R_n - \underbrace{r\cos(\omega t - \theta_n)}_{\text{radar rotation}} - \underbrace{vt\cos(\theta_v - \theta_n)}_{\text{robot motion}},$$

Signal Processing: Untangling Doppler and AoA

Rotation introduces **non-linearity**

But this is known so can remove it

(a) original spectrogram **(b)** compensated spectrogram

Machine Learning: Improving the elevation resolution

Now have good **range** and **azimuth** resolution

Still only have **8 virtual antennas** – poor elevation resolution

- ~14.2 degrees

Leads to smearing

Machine Learning: Improving the elevation resolution

Idea: Spatial dimensions are not independent in 3D environments

Constant Depth

Gravity

Repetitive Patterns

Machine Learning: Improving the elevation resolution

Use 2D CNNs by treating the range information as the channel

- Computational faster than trying to do 3D CNNs

Trained with RF and LiDAR data pairs as inputs

Consequences of training with LiDAR:

- Multipath resistant
- Have to perform glass masking

Machine Learning: Improving the elevation resolution

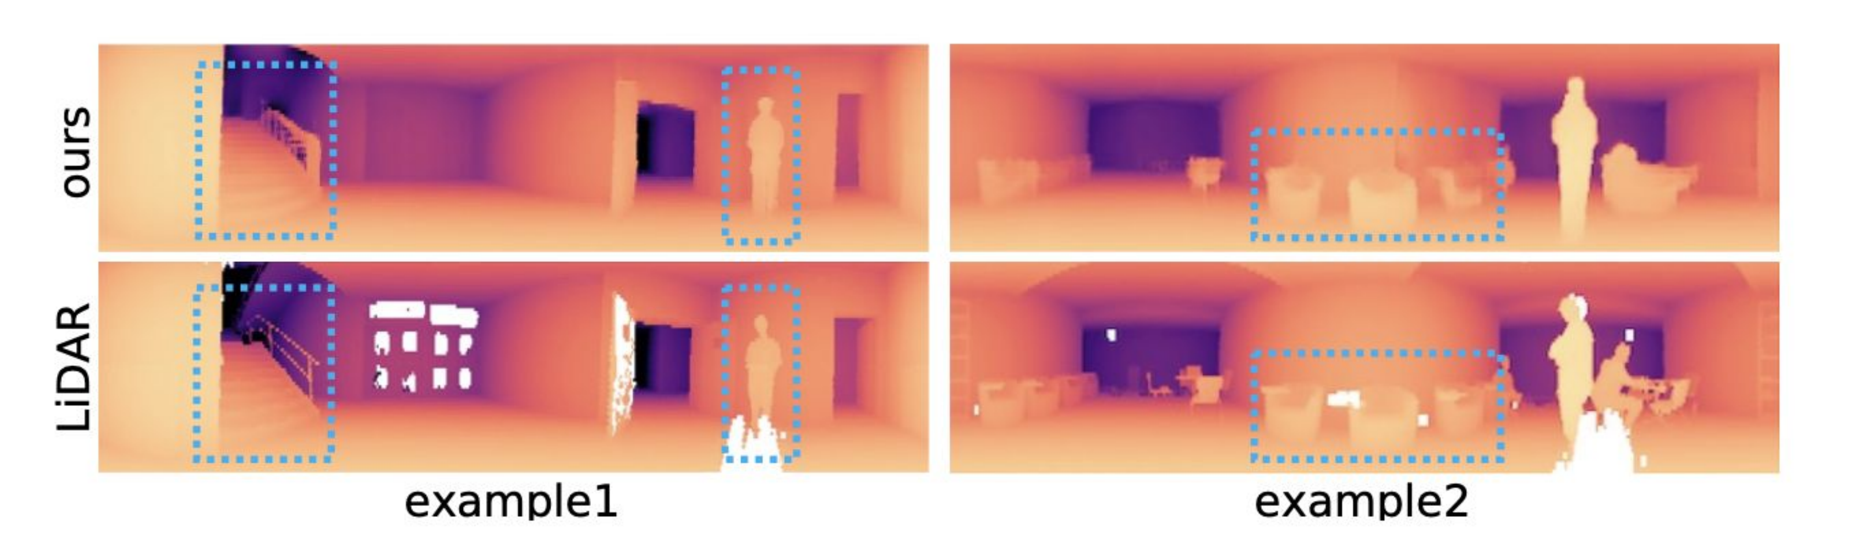

Experimental Validation

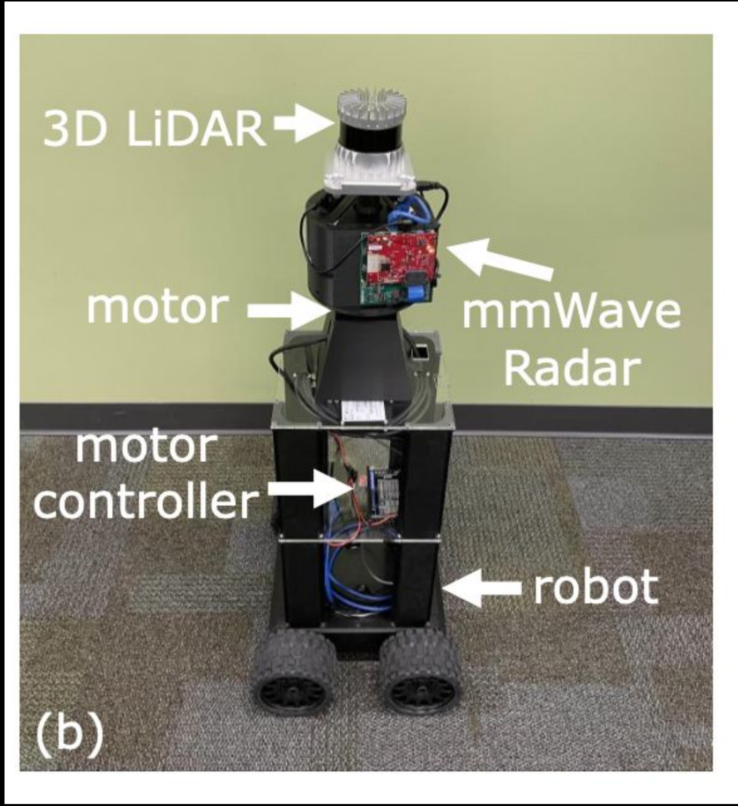

Training

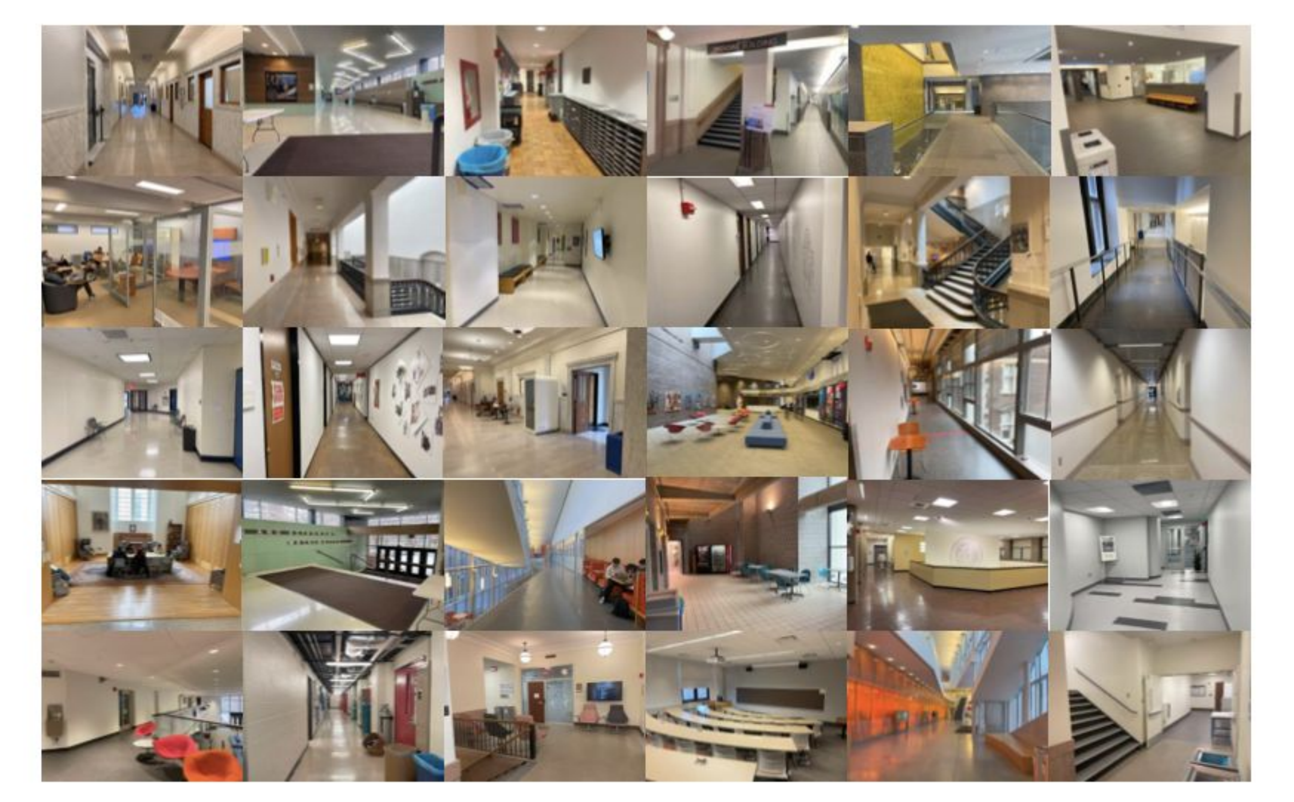

Results

Mean Absolute Error: 15.76 cm
Median: 3.39 cm

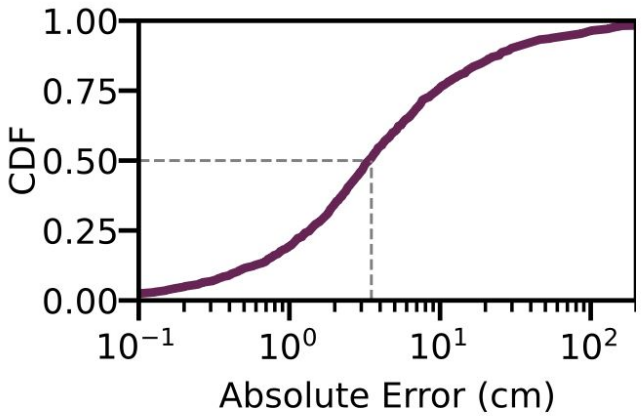

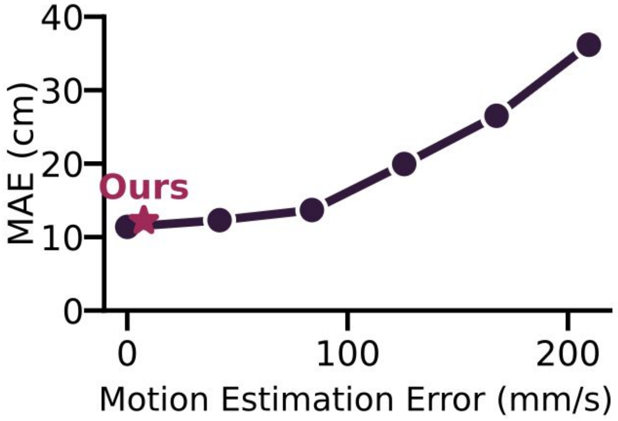

Figure 18: The CDF for absolute error of range image estimation.

Figure 19: The effect of motion errors to imaging performance.

Downstream Applications

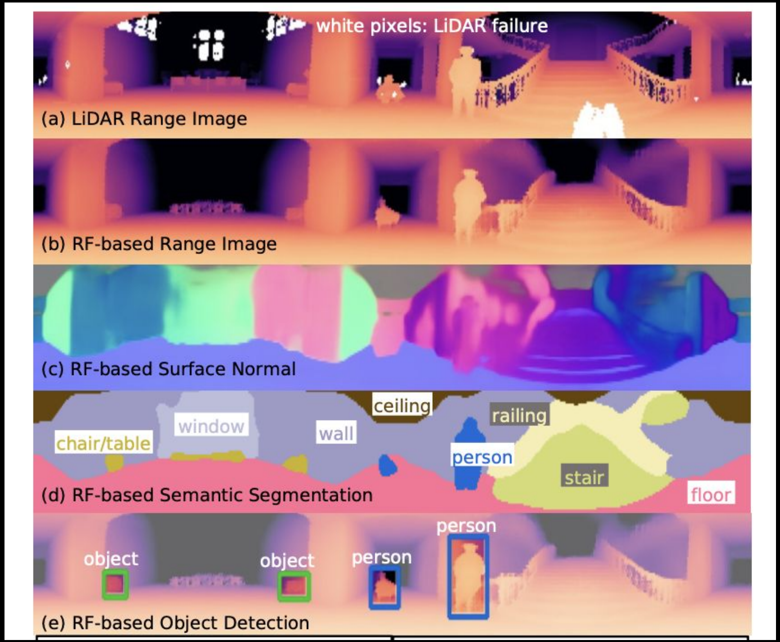

Human Localization

Object Detection

Surface Normal Estimation

Limitations

Runtime performance:

- Desktop GPU takes 51ms to perform range estimation (20Hz) and 95ms for all downstream tasks (10Hz)
- Nvidia Jetson board takes 726ms

Ignores multipath reflections (good or bad?)

Long 0.5-second scanning time – No discussion of if system can handle external motion in environment

Discussion

- **Are there dangers of relying on ML models to "fill in the blanks" of the 3D environment?**
- **How could the system be extended to allow motion in the environment?**
- **Are you convinced by the authors experiments?**