

Extending Computer Vision to Hidden Objects

A Tutorial on Millimeter-Wave Imaging and Reconstruction of Occluded Scenes



Mingmin Zhao

Assistant Professor

University of Pennsylvania



Laura Dodds

PhD Student

Massachusetts Institute of Technology



Hailan Shanbhag

PhD Student

École Polytechnique Fédérale de Lausanne



Jiachen Lu

PhD Student

École Polytechnique Fédérale de Lausanne



Penn
UNIVERSITY of PENNSYLVANIA

EPFL

How many of you have heard of wireless signals (WiFi, Bluetooth, mmWave)?

How many of you have heard of mmWave (5G, 6G, Next-Gen WiFi)?

How many of you have prior knowledge on mmWave sensing (Imaging, FMCW, AoA)?

How many of you have prior research experience in mmWave sensing?

No prior mmWave experience necessary!

There have been incredible advances in computer vision



There have been incredible advances in computer vision



There have been incredible advances in computer vision

Home icon | Change objects | Share icon

Review objects


Adjust results with clicks or boxes.

Search: box

+ Add object

- Object 1
- Object 2
- Object 3

Continue to effects →



The image displays a computer vision application interface. On the left, a 'Review objects' panel allows users to adjust results with clicks or boxes. It features a search bar containing the text 'box', an 'Add object' button, and a list of three objects labeled 'Object 1', 'Object 2', and 'Object 3'. A 'Continue to effects' button is located at the bottom of the panel. The main image shows a conveyor belt with four boxes. Each box is highlighted with a colored bounding box: a yellow box, a light blue box, a green box, and a pink box. The boxes are labeled '4/1 GALLON' and 'Tampico Mango'. The interface also includes a home icon, a 'Change objects' button, and a share icon at the top.

There have been incredible advances in computer vision

What is inside
the boxes?



What if we could build imaging & perception systems
which can see hidden objects?

High-Accuracy Non-Line-of-Sight 3D Reconstruction has Many Applications



Shipping & Logistics



Robotics



Augmented Reality

Wireless signals can travel through everyday occlusions

Wireless signals can travel through everyday occlusions

Millimeter-Wave (mmWave)

5G, 6G, Next-Gen WiFi

Wireless signals can travel through everyday occlusions

Millimeter-Wave (mmWave)

5G, 6G, Next-Gen WiFi

Cardboard


Plastic

Fabric

Fog

Smoke

Wireless signals can travel through everyday occlusions
We can use them to image, reconstruct, & perceive the hidden environment




Wireless signals can travel through everyday occlusions
We can use them to image, reconstruct, & perceive the hidden environment

Autonomous Driving In adverse conditions




Visual sensors fail!



mmWave Radar

Penetrates fog and rain
Robust under adverse weather

Wireless signals can travel through everyday occlusions
We can use them to image, reconstruct, & perceive the hidden environment



Wireless signals can travel through everyday occlusions
We can use them to image, reconstruct, & perceive the hidden environment







PanoRadar uses radio signals to enable robots to accurately perceive its surroundings...

Is mmWave imaging & reconstruction a solved problem?

Top View

Side View







Is mmWave imaging & reconstruction a solved problem?

Top View


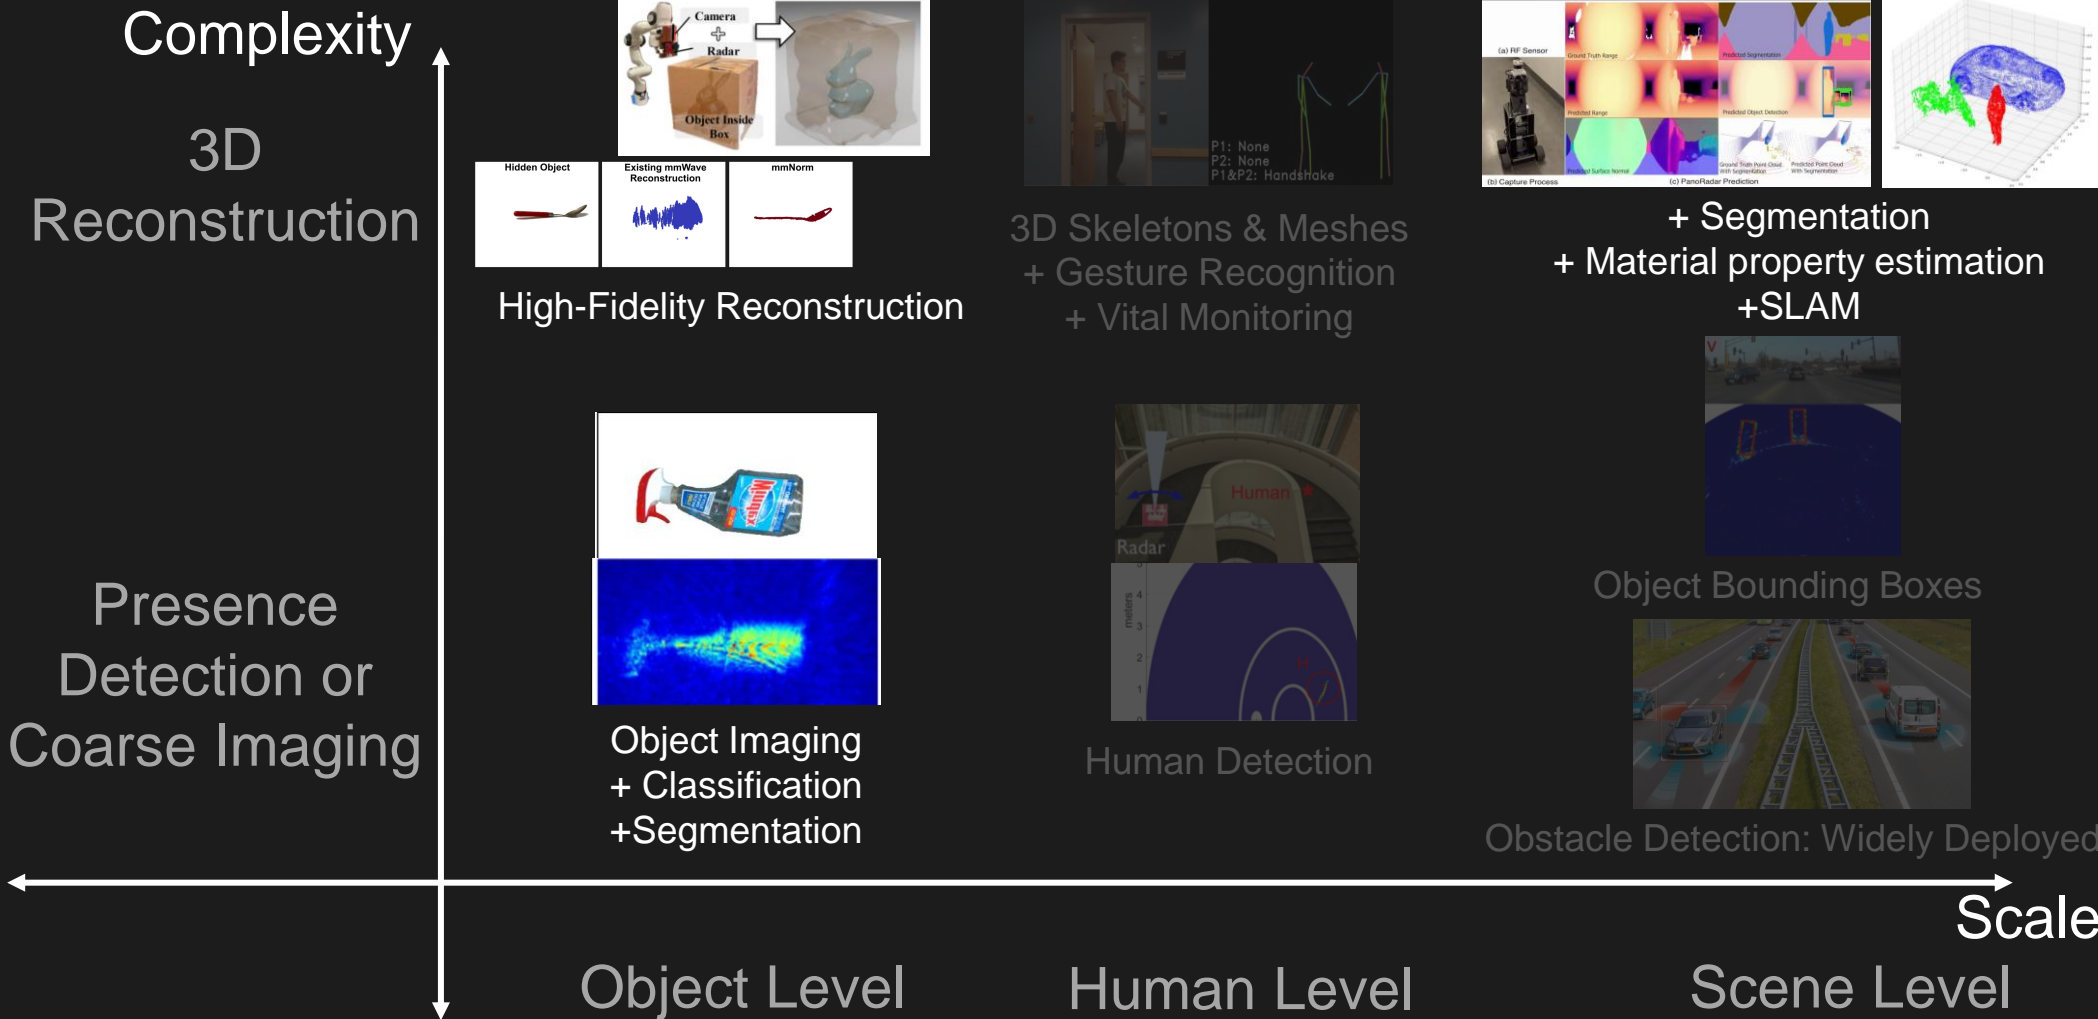


Side View




Large
Power
Spread


What are most common mmWave imaging research domains?




High-Fidelity Reconstruction




3D Skeletons & Meshes + Gesture Recognition + Vital Monitoring




+ Segmentation + Material property estimation + SLAM




Object Imaging + Classification + Segmentation



Human Detection



Object Bounding Boxes



Obstacle Detection: Widely Deployed

Extending Computer Vision to Hidden Objects

A Tutorial on Millimeter-Wave Imaging and Reconstruction of Occluded Scenes

mmWave 101: Basics of mmWave Imaging, Limitations

State-of-the-art Reconstruction Methods



mmWave Neural Radiance Fields

Coffee Break

Vision-quality mmWave Scene Understanding

Open Problems; Tour of Existing Datasets & Tools

Hands-On Tutorial: mmWave Imaging & Reconstruction



Extending Computer Vision to Hidden Objects

A Tutorial on Millimeter-Wave Imaging and Reconstruction of Occluded Scenes

mmWave 101: Basics of mmWave Imaging, Limitations

State-of-the-art Reconstruction Methods


mmWave Neural Radiance Fields

Coffee Break

Vision-quality mmWave Scene Understanding

Open Problems; Tour of Existing Datasets & Tools

Hands-On Tutorial: mmWave Imaging & Reconstruction



How does mmWave compare to other through-occlusion modalities?

XRay



✗ Unsafe for extended exposure

Thermal




✗ Requires heat
Cannot see through walls

Ultrasound



✗ Requires gel to minimize mismatch

Reflected Lasers




✗ Cannot travel through occlusions

mmWave is safe, operates through occlusions, and can image everyday objects


What is a mmWave sensor?

“Radar”



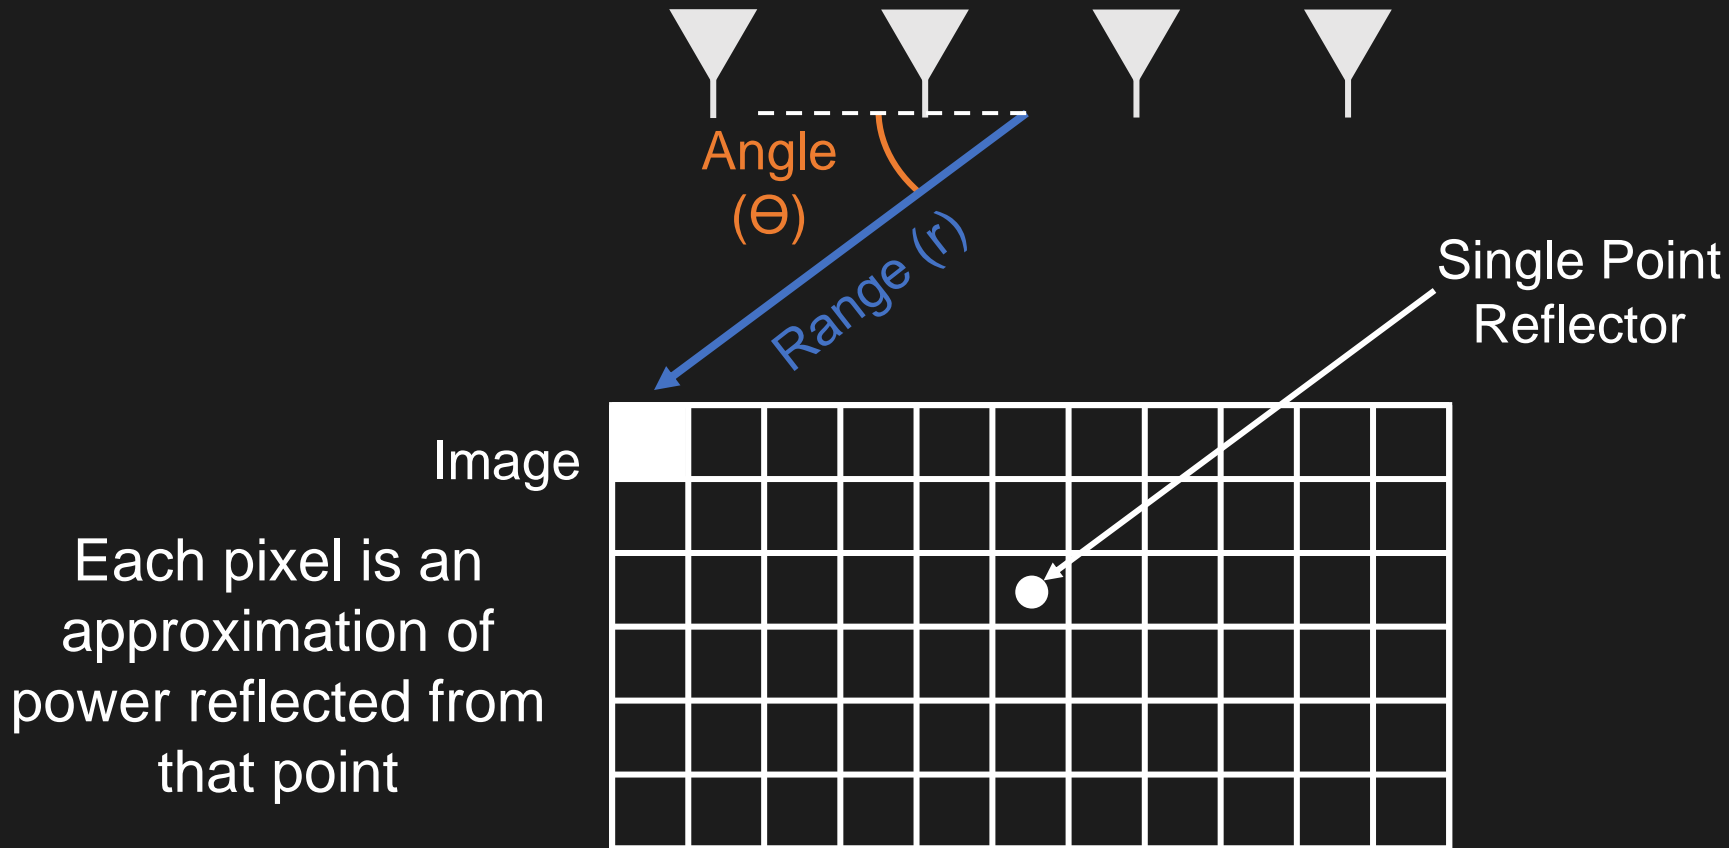
What is a mmWave sensor?

“Radar”




A mmWave radar transmits and receives signals
It does not directly capture an image

How do we measure a mmWave Image?



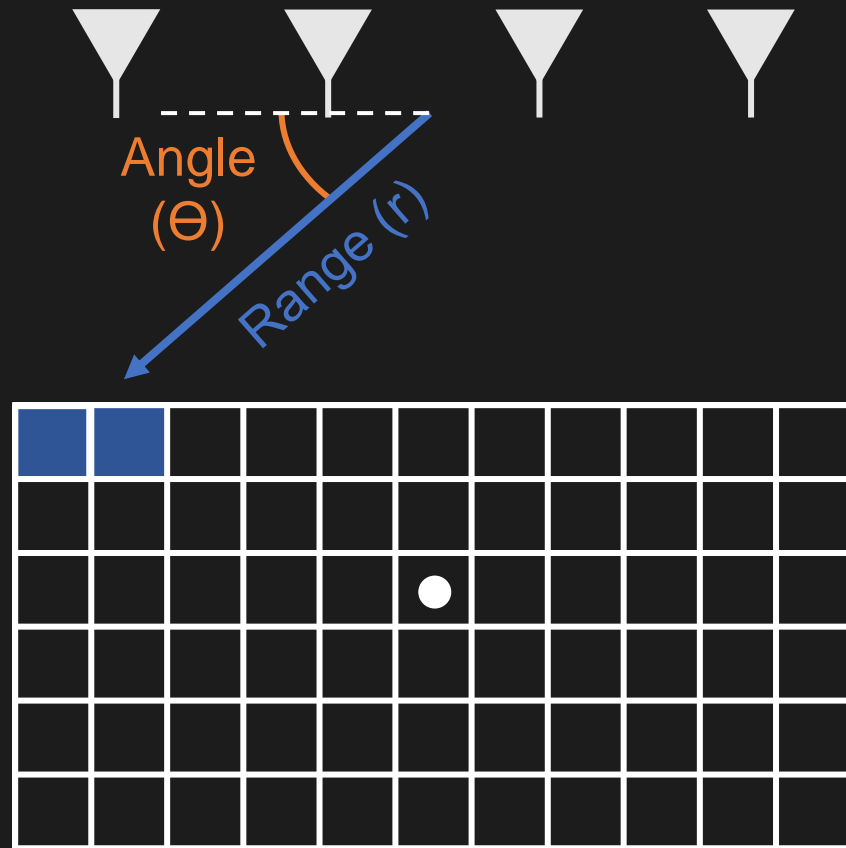
Estimate the **power** reflected from this **point (range, angle)** in space

How do we measure a mmWave Image?



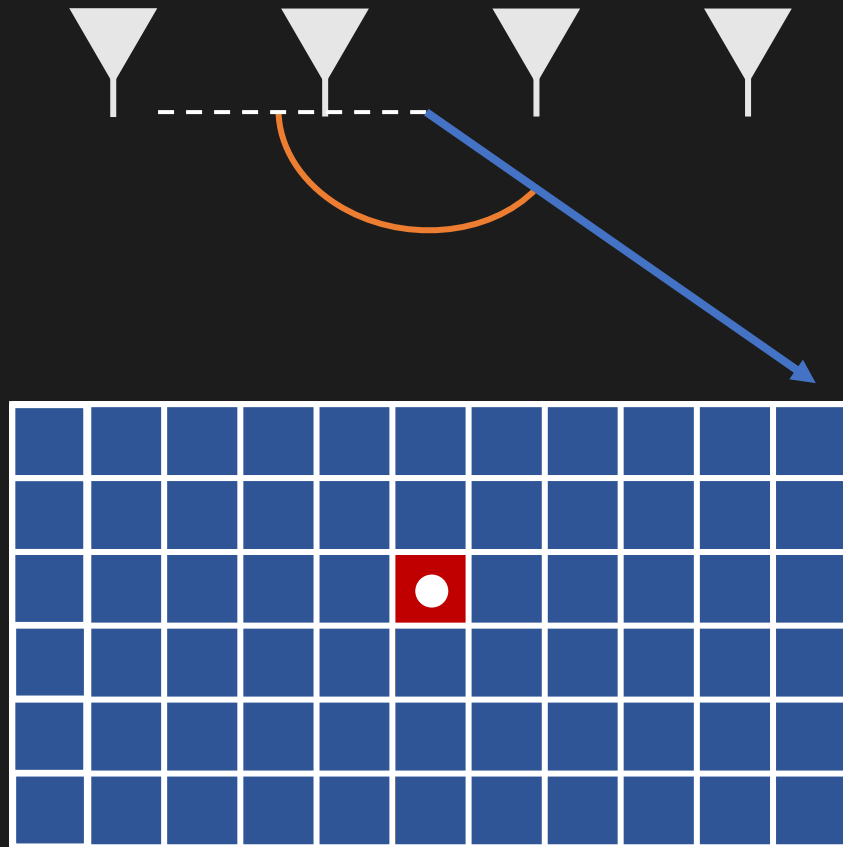
Estimate the **power** reflected from this **point (range, angle)** in space

How do we measure a mmWave Image?



Estimate the **power** reflected from this **point (range, angle)** in space

How do we measure a mmWave Image?



Estimate the **power** reflected from this **point (range, angle)** in space

How do we measure a mmWave Image?

Estimate the **power** reflected from this **point (range, angle)** in space

How do we estimate **range** with mmWave signals?

How do we estimate **angle** with mmWave signals?

How do we estimate range with mmWave signals?



Measure time it takes
signal to travel to
object and back

“Time of Flight”



How do we estimate range with mmWave signals?



Measure time it takes
signal to travel to
object and back

“Time of Flight”



How do we estimate range with mmWave signals?



Measure time it takes
signal to travel to
object and back

“Time of Flight”



How do we estimate range with mmWave signals?



Measure time it takes
signal to travel to
object and back

“Time of Flight”

Measuring time of flight directly is very difficult
> \$10k




Range = Speed of Light x Time of Flight

$$r = c \times \tau$$

How do we estimate range with mmWave signals?

Frequency Modulated Continuous Wave (FMCW)


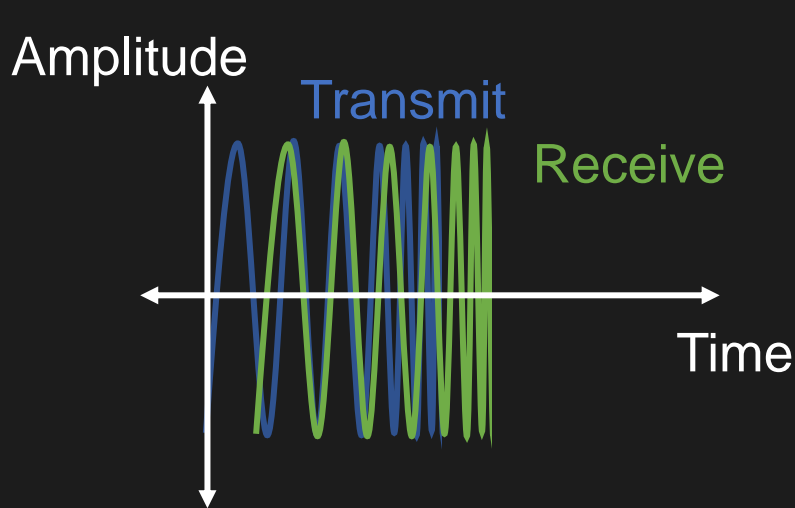
Create a relationship between time of flight and frequency



How do we estimate range with mmWave signals?

Frequency Modulated Continuous Wave (FMCW)

Create a relationship between time of flight and frequency



How do we estimate range with mmWave signals?

Frequency Modulated Continuous Wave (FMCW)

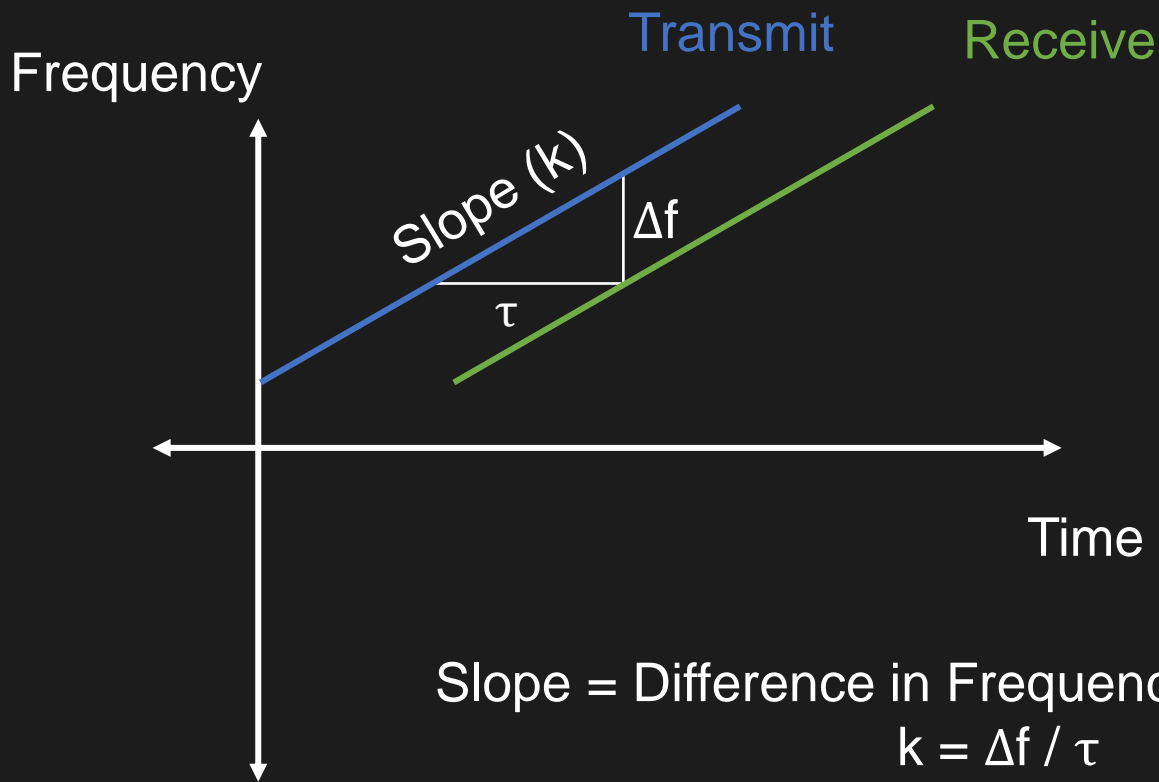
Create a relationship between time of flight and frequency



How do we estimate range with mmWave signals?

Frequency Modulated Continuous Wave (FMCW)

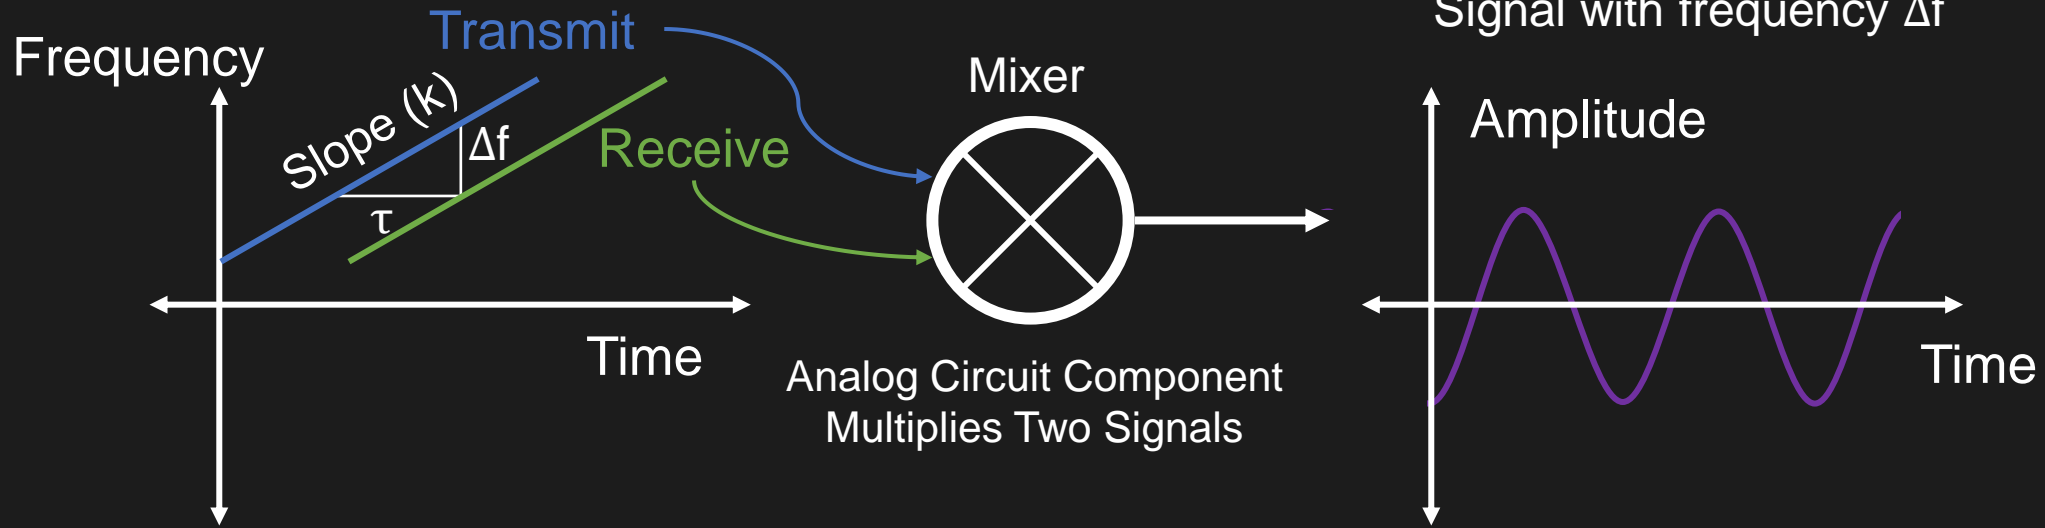
Create a relationship between time of flight and frequency



How do we estimate range with mmWave signals?

Frequency Modulated Continuous Wave (FMCW)

Create a relationship between time of flight and frequency



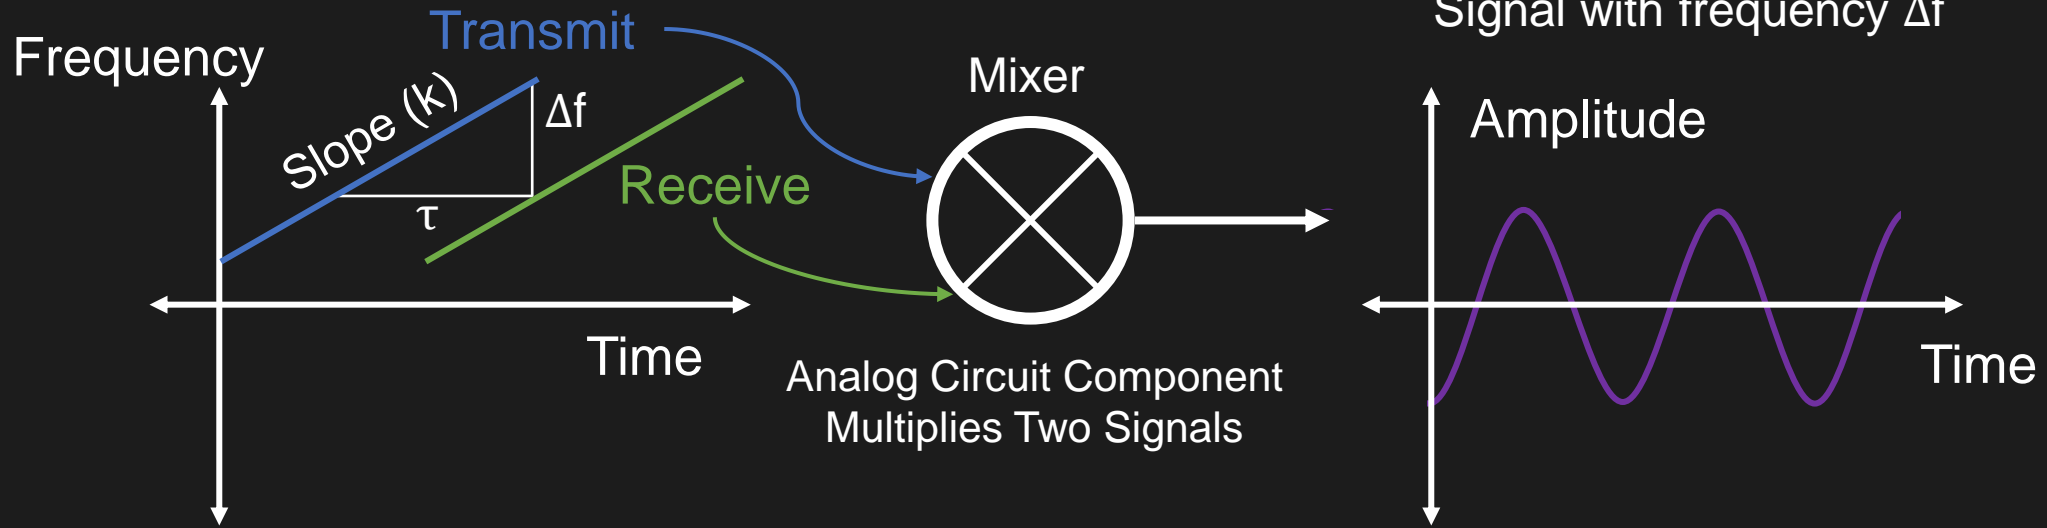
Slope = Difference in Frequency / Time of Flight

$$k = \Delta f / \tau$$

How do we estimate range with mmWave signals?

Frequency Modulated Continuous Wave (FMCW)

Create a relationship between time of flight and frequency



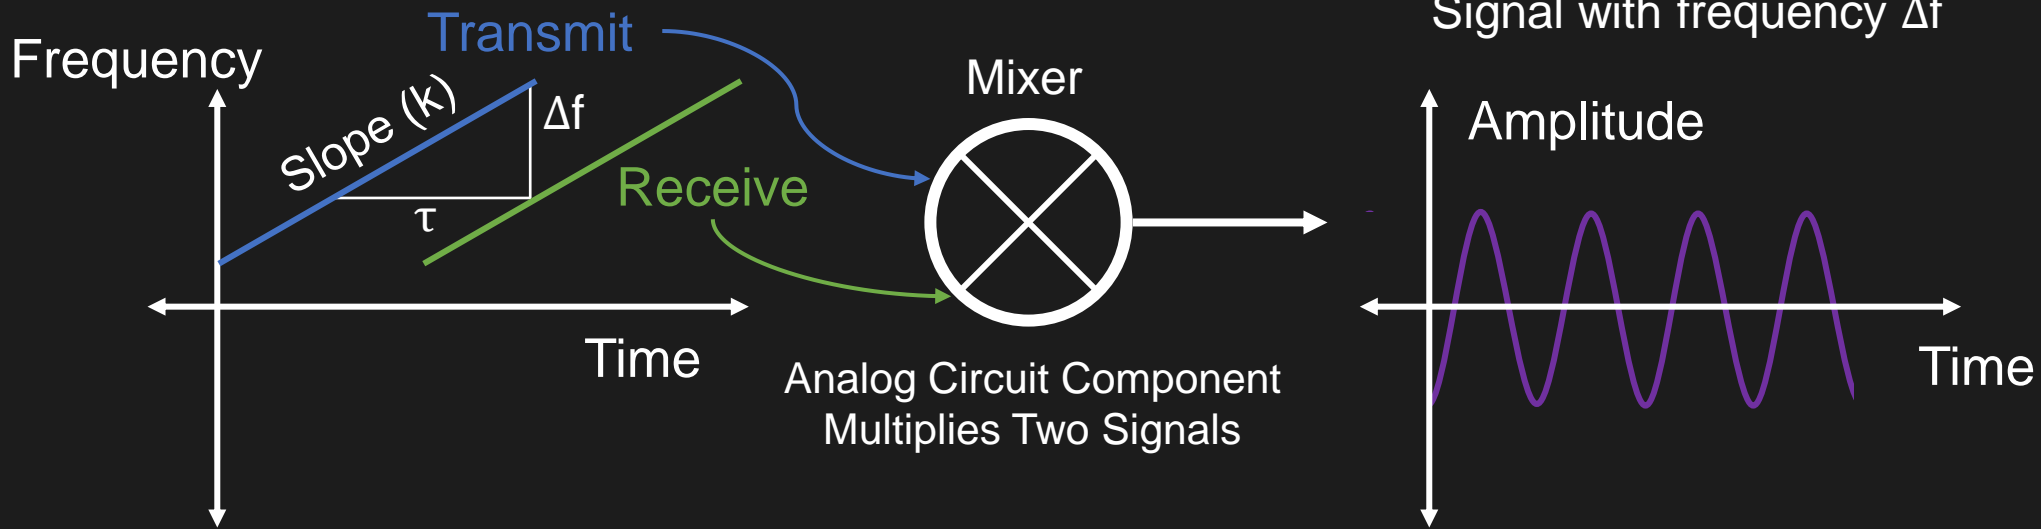
Slope = Difference in Frequency / Time of Flight

$$k = \Delta f / \tau$$

How do we estimate range with mmWave signals?

Frequency Modulated Continuous Wave (FMCW)

Create a relationship between time of flight and frequency



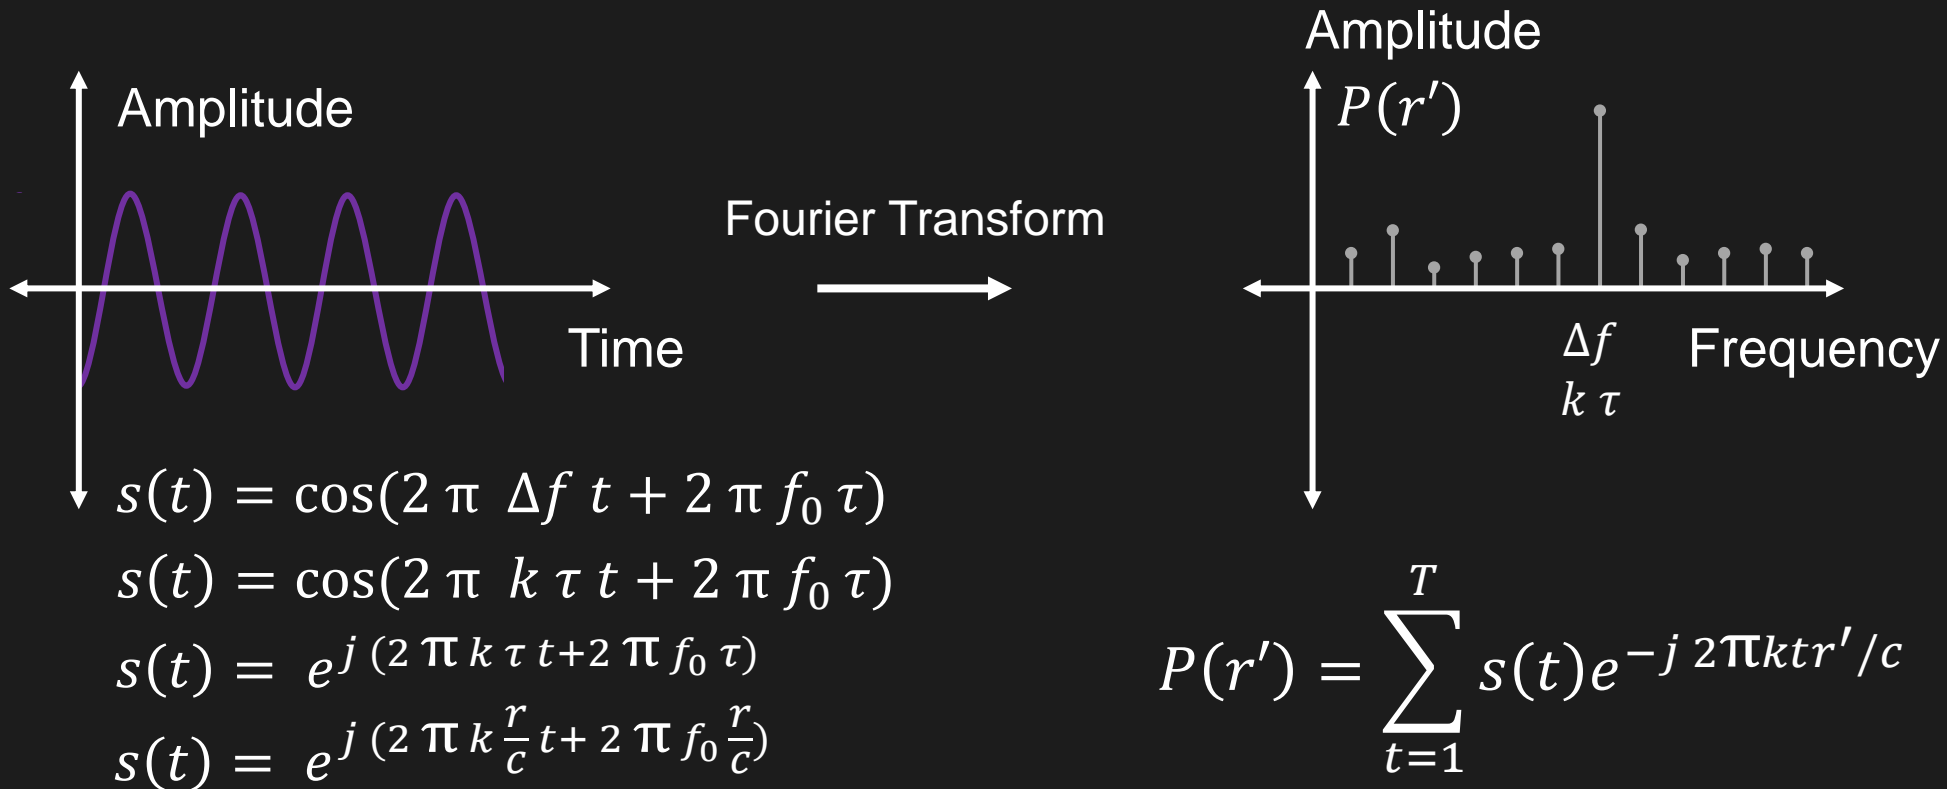
Slope = Difference in Frequency / Time of Flight
 $k = \Delta f / \tau$

We need to find Δf to find range

How do we estimate range with mmWave signals?


Frequency Modulated Continuous Wave (FMCW)

Create a relationship between time of flight and frequency



We need to find Δf to find range

How do we estimate range with mmWave signals?

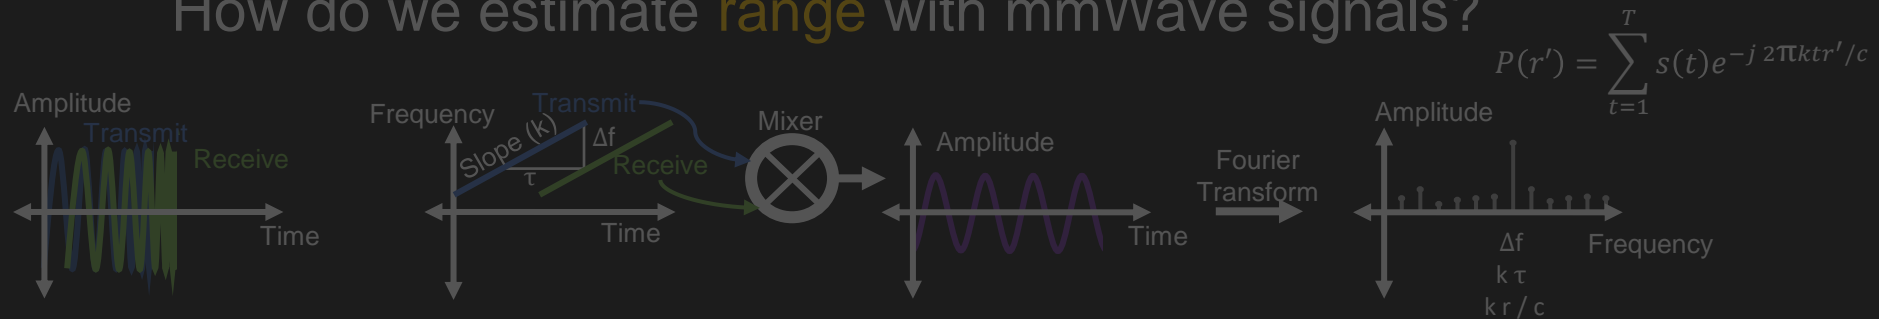


$$P(r') = \sum_{t=1}^T s(t) e^{-j 2\pi k t r' / c}$$

How do we measure a mmWave Image?

Estimate the **power** reflected from this **point (range, angle)** in space

How do we estimate **range** with mmWave signals?



How do we estimate **angle** with mmWave signals?

How do we estimate angle with mmWave signals?



How do we estimate angle with mmWave signals?



Compare difference
between signal at two
locations



How do we estimate angle with mmWave signals?



Compare difference
between signal at two
locations



How do we estimate angle with mmWave signals?



Compare difference
between signal at two
locations



How do we estimate angle with mmWave signals?



Compare difference
between signal at two
locations




How do we estimate angle with mmWave signals?



Far Field Assumption
Reflector is far away that
the lines become parallel

How do we estimate angle with mmWave signals?




Far Field Assumption
Reflector is far away that
the lines become parallel

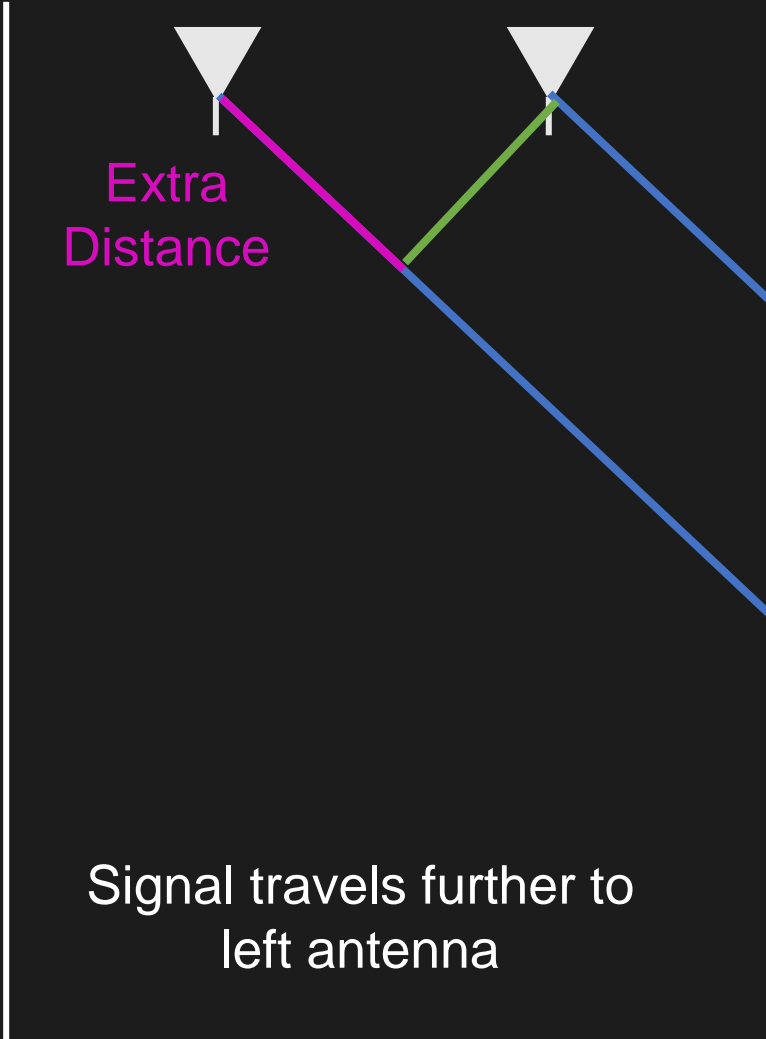

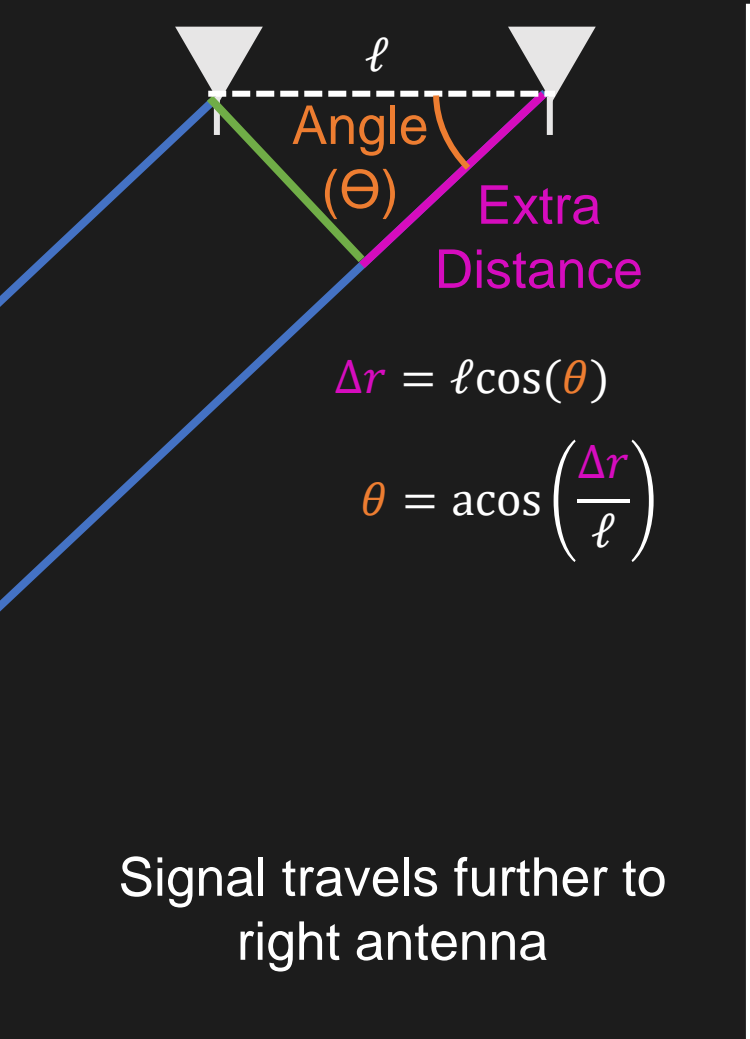
How do we estimate angle with mmWave signals?



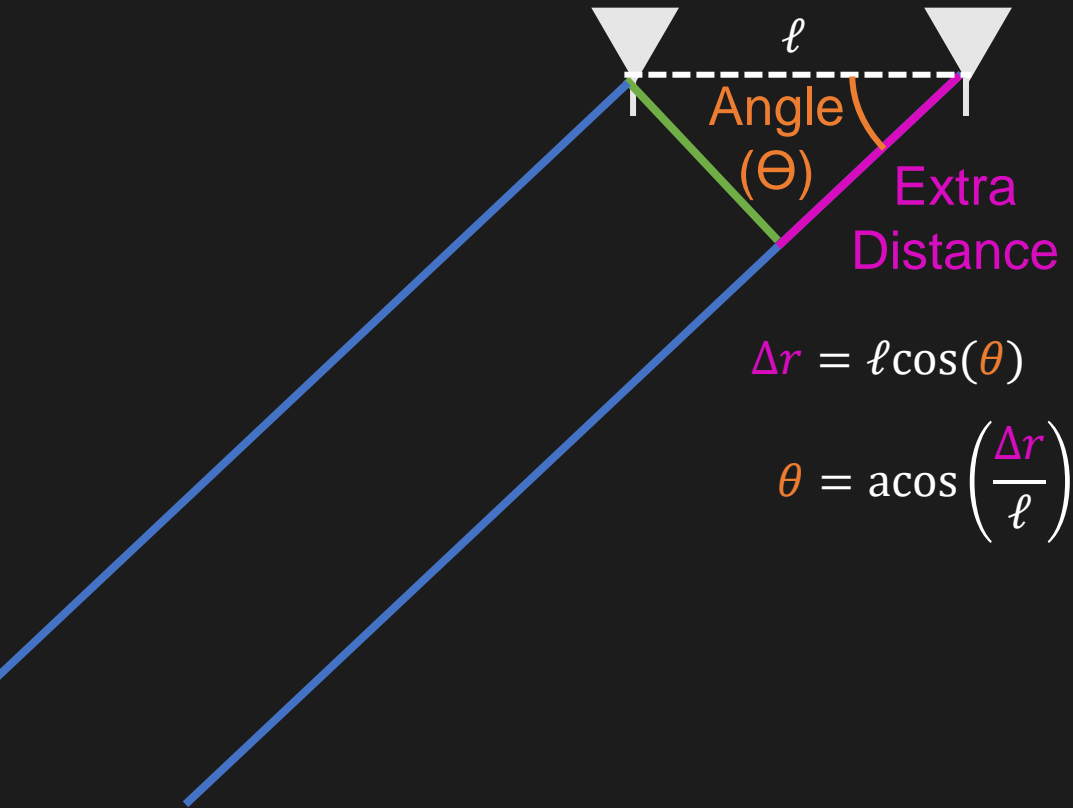
How do we estimate angle with mmWave signals?



How do we estimate angle with mmWave signals?



How do we estimate angle with mmWave signals?



$$s(t) = e^{j2\pi k \frac{r}{c} t + j2\pi f_0 \frac{r}{c}}$$

$$\angle s_1 = 2\pi f_0 \frac{r_1}{c}$$

$$\angle s_2 = 2\pi f_0 \frac{r_2}{c}$$

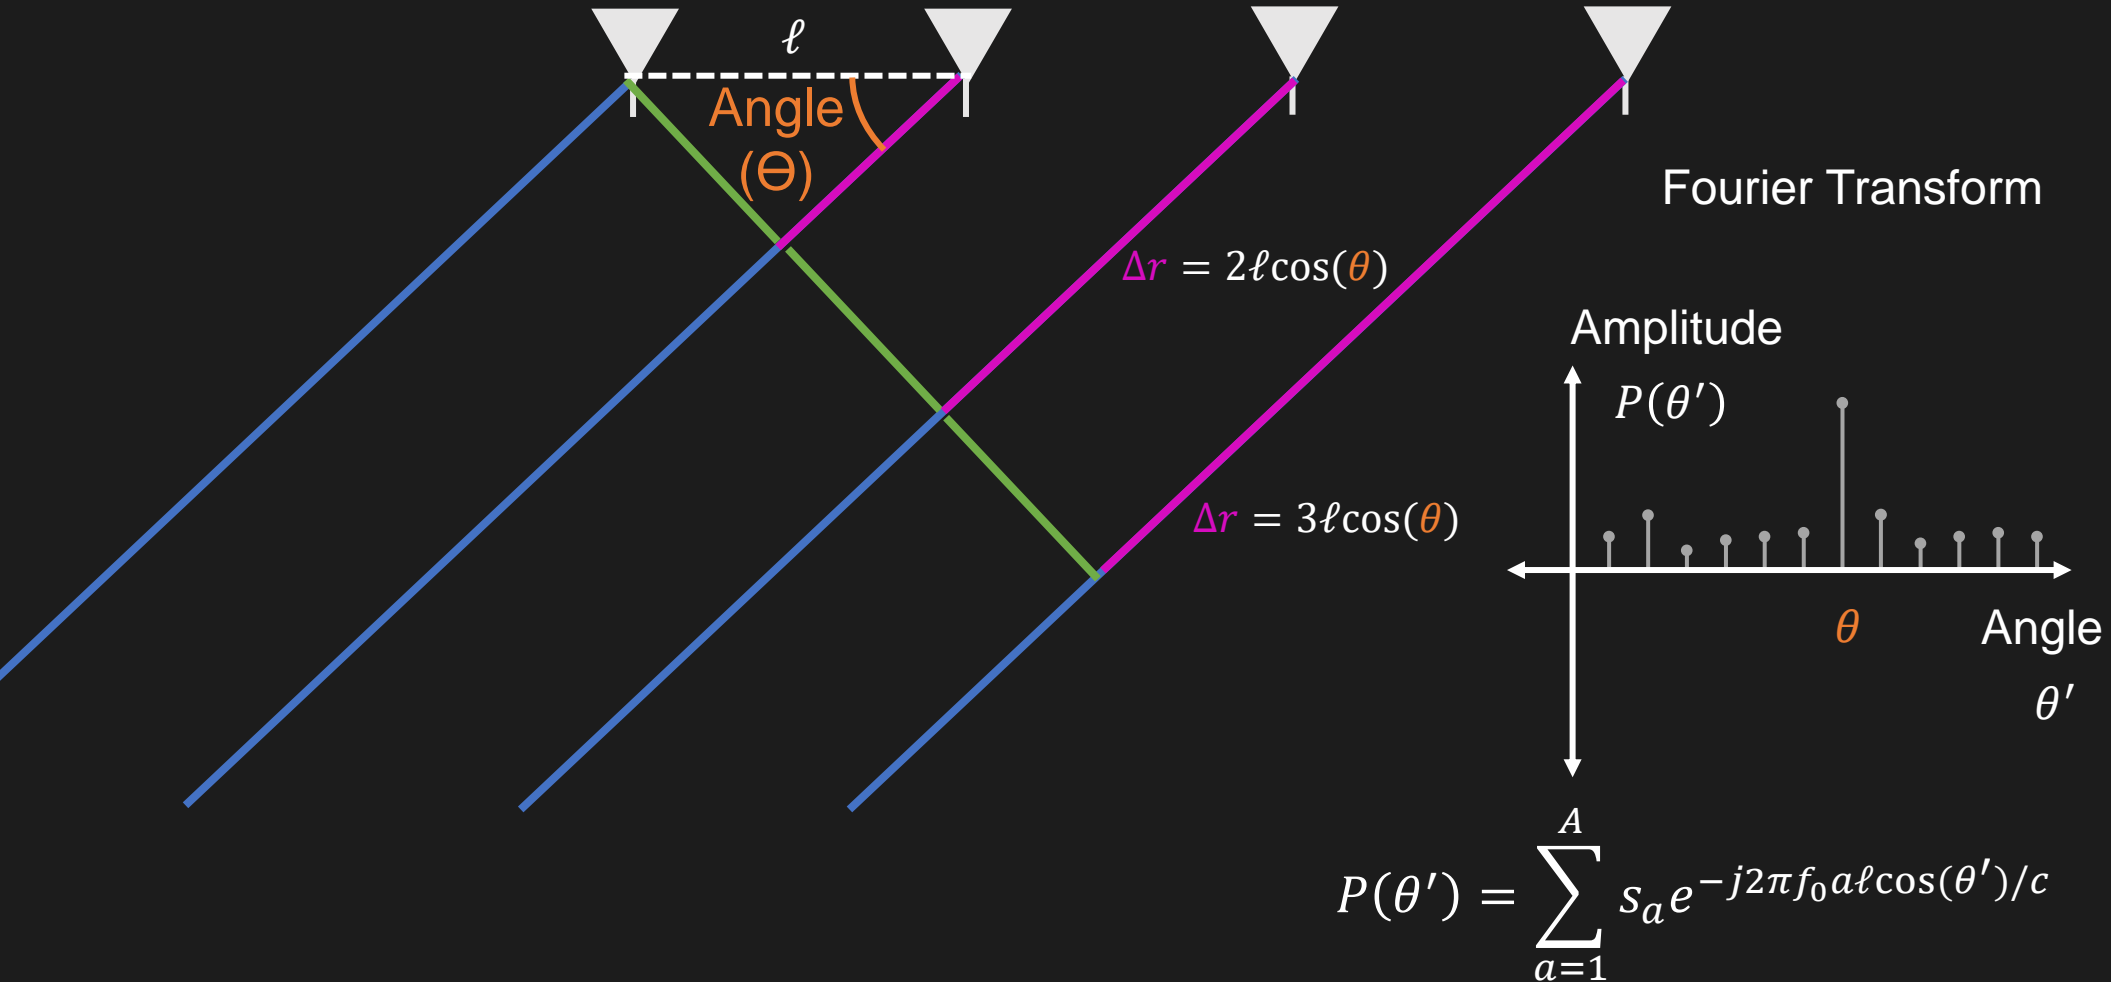
$$\angle s_1 - \angle s_2 = 2\pi f_0 \frac{r_1 - r_2}{c}$$

$$\angle s_1 - \angle s_2 = 2\pi f_0 \frac{\Delta r}{c}$$

$$\Delta r = \frac{(\angle s_1 - \angle s_2) c}{2\pi f_0}$$

$$\theta = \arccos\left(\frac{(\angle s_1 - \angle s_2) c}{2\pi f_0 \ell}\right)$$

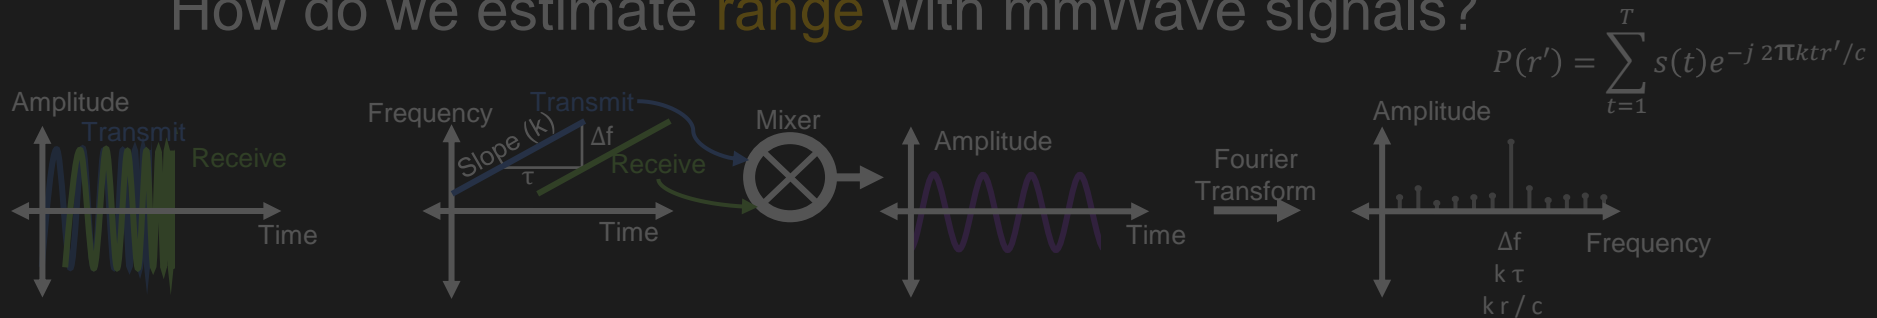
How do we estimate angle with mmWave signals?



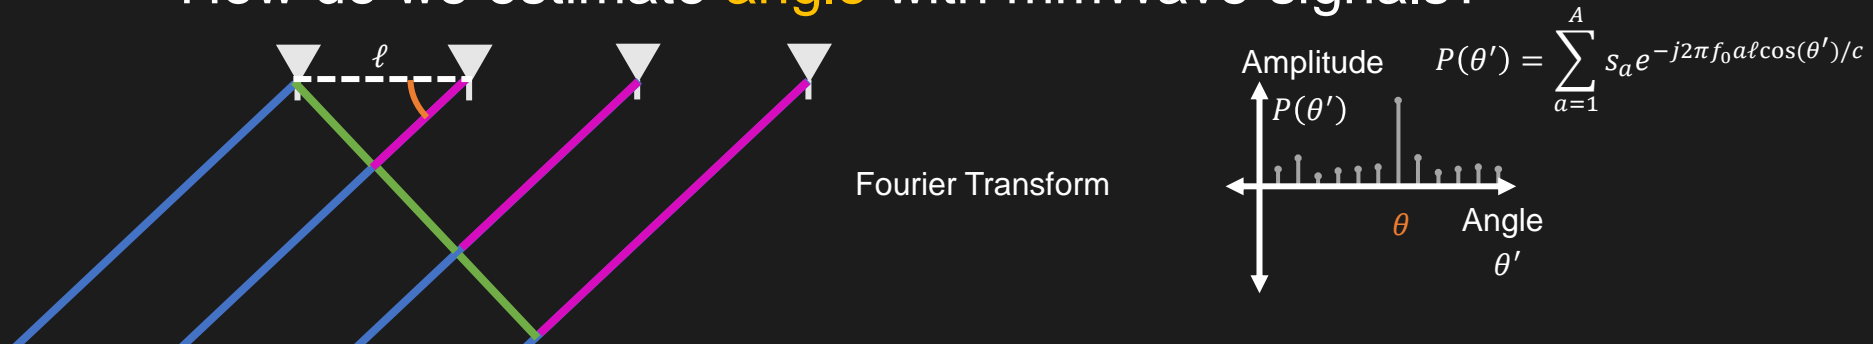
How do we measure a mmWave Image?

Estimate the **power** reflected from this **point (range, angle)** in space

How do we estimate **range** with mmWave signals?



How do we estimate **angle** with mmWave signals?



Estimate the **power** reflected from this **point (range, angle)** in space

range



$$P(r') = \sum_{t=1}^T s(t) e^{-j 2\pi k t r' / c}$$

angle

$$P(\theta') = \sum_{a=1}^A s_a e^{-j 2\pi f_0 a \ell \cos(\theta') / c}$$



$$P(r', \theta') = \sum_{a=1}^A \sum_{t=1}^T s_a(t) e^{-j 2\pi k t r' / c} e^{-j 2\pi f_0 a \ell \cos(\theta') / c}$$

Estimate the **power** reflected from this **point (range, angle)** in space





$$P(r', \theta') = \sum_{a=1}^A \sum_{t=1}^T s_a(t) e^{-j 2\pi k t r' / c} e^{-j 2\pi f_0 a \ell \cos(\theta') / c}$$

Extending from 2D to 3D




$$P(r', \theta') = \sum_{a=1}^A \sum_{t=1}^T s_a(t) e^{-j 2\pi k r' / c} e^{-j 2\pi f_0 a \ell \cos(\theta') / c}$$

Extending from 2D to 3D



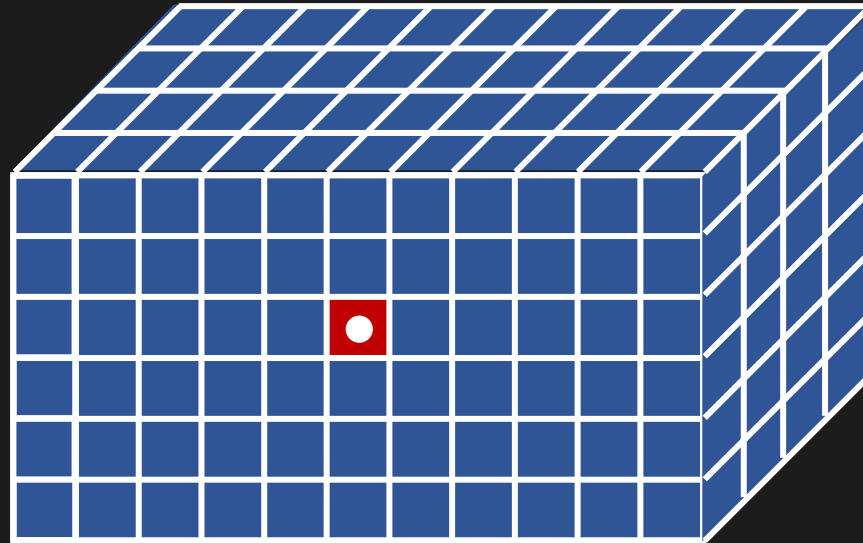
$$P(r', \theta', \phi') = \sum_{a_y=1}^{A_y} \sum_{a_x=1}^{A_x} \sum_{t=1}^T s_{a_x, a_y}(t) e^{-j 2\pi k r' / c} e^{-j 2\pi f_0 \sin(\phi') (a_x \ell \cos(\theta') + a_y \ell \sin(\theta')) / c}$$

What are the challenges of mmWave?




What are the challenges of mmWave?

Resolution of mmWave image is low compared to cameras

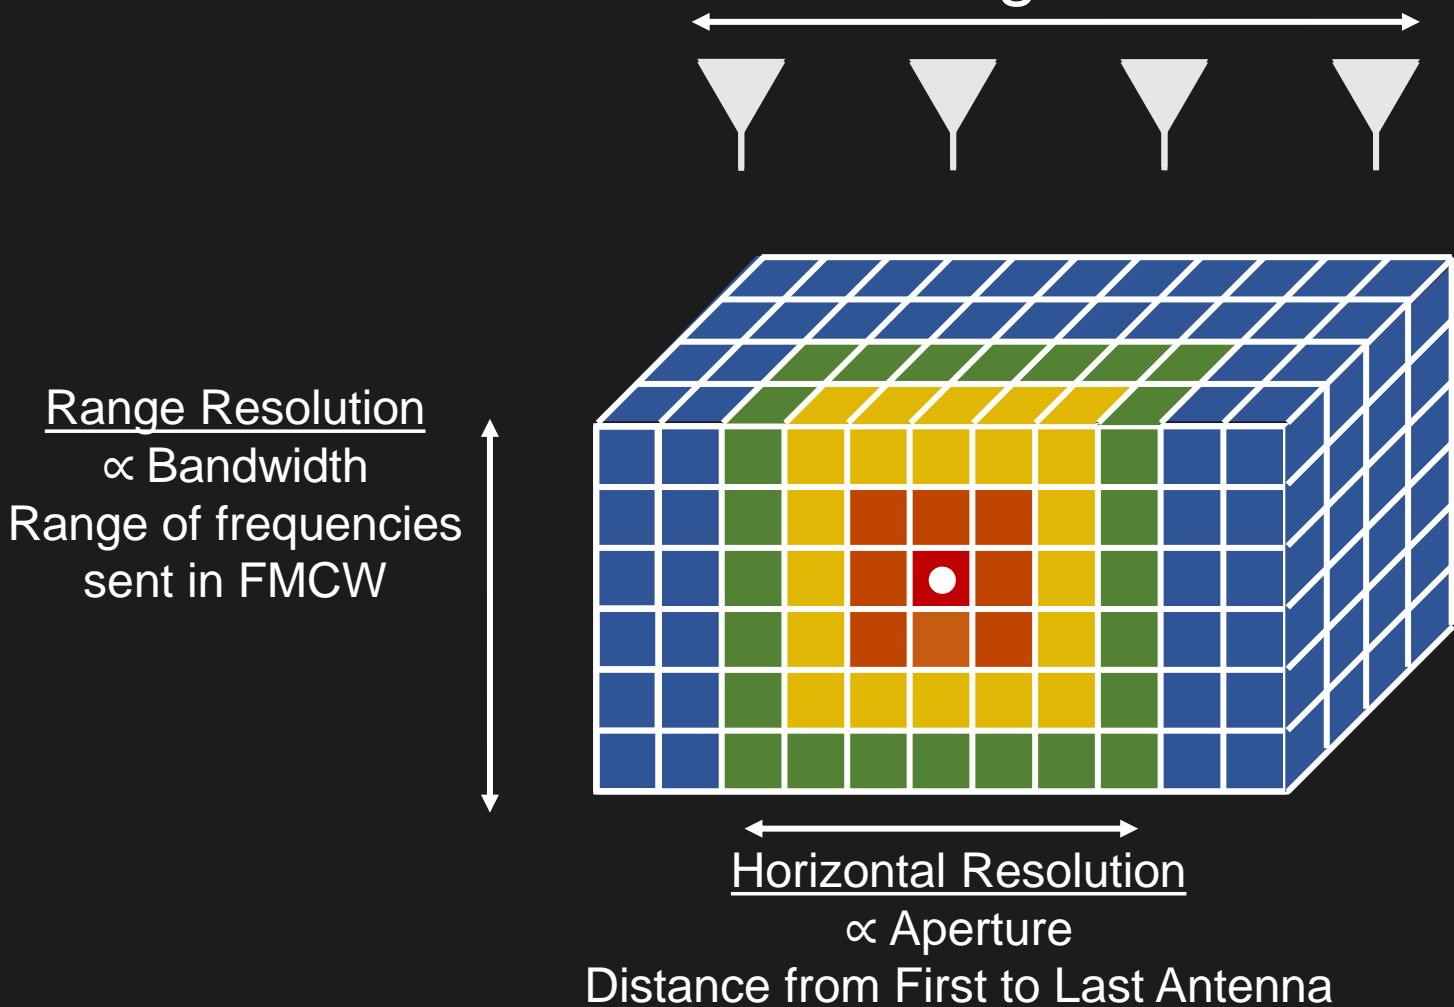


What are the challenges of mmWave?

Resolution of mmWave image is low compared to cameras



Resolution of mmWave image is low compared to cameras




Resolution of mmWave image is low compared to cameras



Range Resolution

\propto Bandwidth

Range of frequencies
sent in FMCW



Horizontal Resolution

\propto Aperture

Distance from First to Last Antenna

Resolution of mmWave image is low compared to cameras




Synthetic Aperture Radar
Move radar to new locations

Horizontal resolution can be increased by taking more measurements
Range resolution cannot be increased

sent in FMCW

✗ Limited by government regulations



antennas on the radar

✓ Can increase horizontal resolution



Horizontal Resolution

$$\propto \text{Aperture}$$



Distance from First to Last Antenna

Resolution of mmWave image is low compared to cameras

Top View



Side View



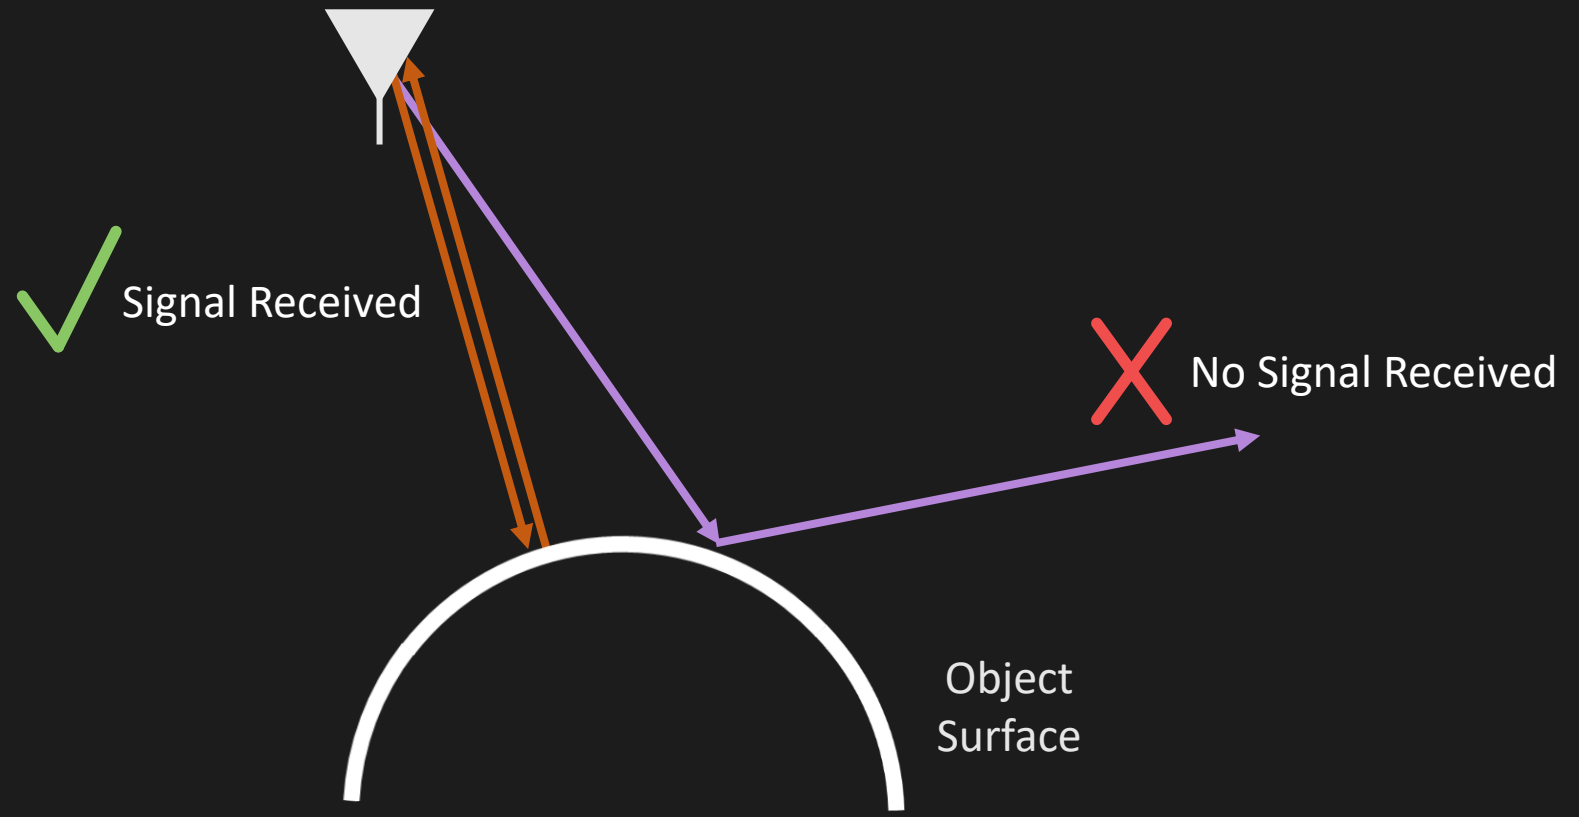
Large
Power
Spread

What are the challenges of mmWave?


Resolution of mmWave image is low compared to cameras

Coverage of mmWave image is limited by specularities

Coverage of mmWave image is limited by specularity

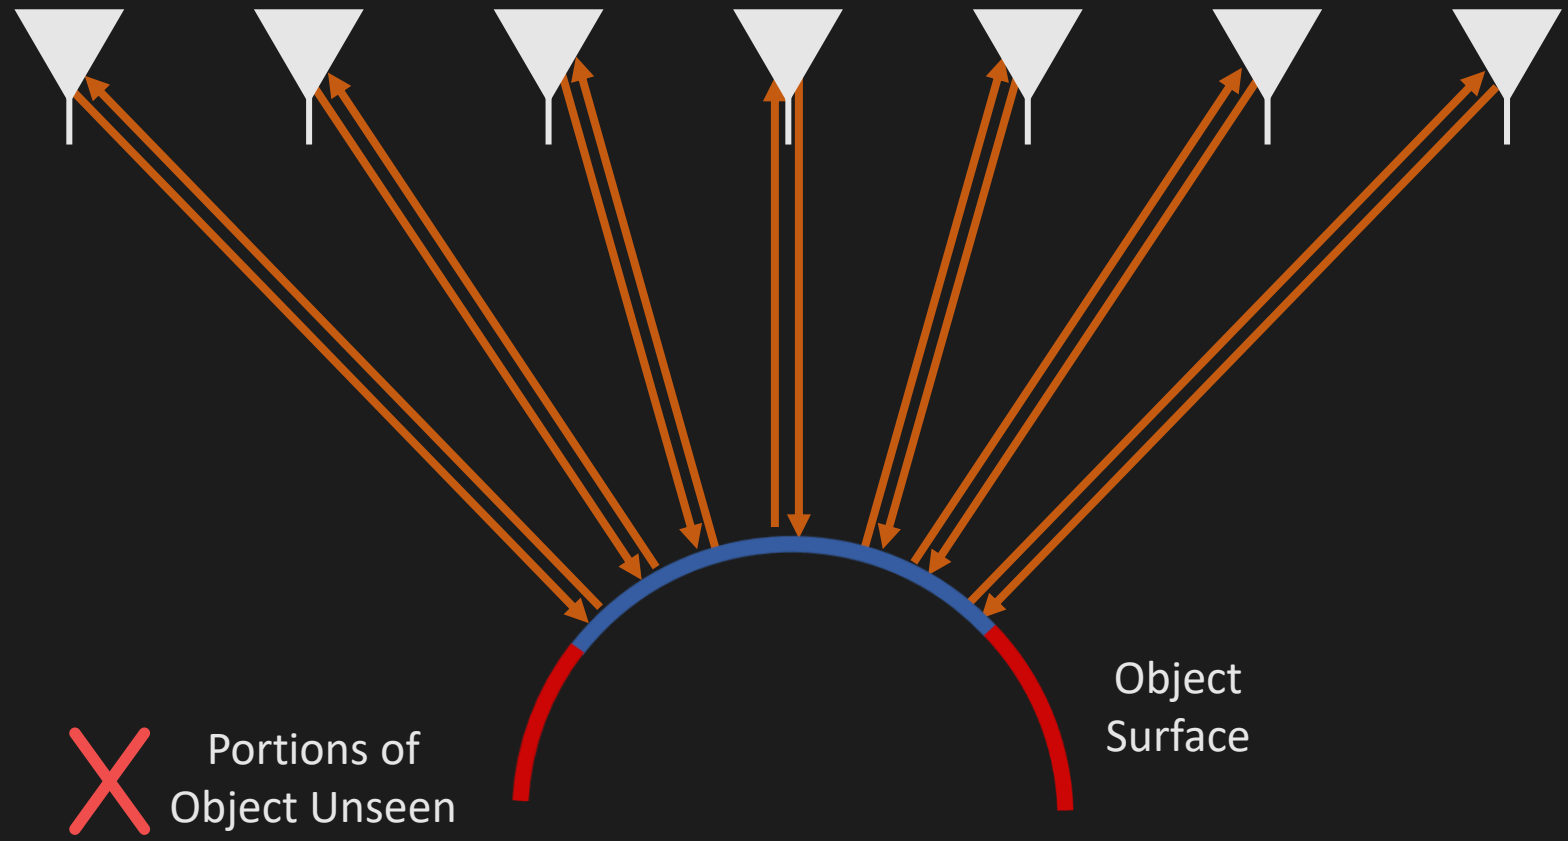


Coverage of mmWave image is limited by specularity



Object
Surface

Coverage of mmWave image is limited by specularity



What are the challenges of mmWave?

Resolution of mmWave image is low compared to cameras

Coverage of mmWave image is limited by specularities

mmWave signals suffer high noise & do not provide color information

We will be discussing several state-of-the-art methods to overcome these challenges

Extending Computer Vision to Hidden Objects

A Tutorial on Millimeter-Wave Imaging and Reconstruction of Occluded Scenes

mmWave 101: Basics of mmWave Imaging, Limitations

State-of-the-art Reconstruction Methods


mmWave Neural Radiance Fields

Coffee Break

Vision-quality mmWave Scene Understanding

Open Problems; Tour of Existing Datasets & Tools

Hands-On Tutorial: mmWave Imaging & Reconstruction



Extending Computer Vision to Hidden Objects

A Tutorial on Millimeter-Wave Imaging and Reconstruction of Occluded Scenes

mmWave 101: Basics of mmWave Imaging, Limitations

State-of-the-art Reconstruction Methods


mmWave Neural Radiance Fields

Coffee Break

Vision-quality mmWave Scene Understanding

Open Problems; Tour of Existing Datasets & Tools

Hands-On Tutorial: mmWave Imaging & Reconstruction



Non-Line-of-Sight 3D Object Reconstruction via mmWave Surface Normal Estimation

Laura Dodds, Tara Boroushaki, Kaichen Zhou, Fadel Adib



How can we go from 3D imaging to 3D reconstruction?

How can we go from 3D imaging to 3D reconstruction?

How can we go from 3D imaging to 3D reconstruction?

Naïve Solution: Select
every voxel above a
power threshold

How can we go from 3D imaging to 3D reconstruction?

Hidden Object



**Existing mmWave
Reconstruction***



mmNorm



High-Accuracy Non-Line-of-Sight 3D Reconstruction has Many Applications



Robotics



Shipping & Logistics




Augmented Reality


mmNorm

- A new & highly-accurate method for 3D object reconstruction leveraging an **entirely first-principles based approach**
- Introduces techniques for **mmWave Surface Normal Estimation** and **mmWave Structural Isosurface Optimization**
- Achieves a **qualitative leap** in mmWave reconstruction over state-of-the-art methods

How do existing methods perform 3D reconstruction?



Real-World
3D Point Cloud

 Reconstruction quality is limited by SAR resolution (i.e., bandwidth)

How do existing methods perform 3D reconstruction?

First-Principles Methods

Machine-Learning Methods




We need a fundamentally new approach for mmWave 3D object reconstruction

Our Idea: Directly estimate the curvature of the object


⚡ SAR resolution (i.e., bandwidth)

⚡ subset of objects


Our Idea: Directly estimate the curvature of the object by estimating surface normal vectors



Our Idea: Directly estimate the curvature of the object by estimating surface normal vectors




Our Idea: Directly estimate the curvature of the object
by estimating surface normal vectors




Our Idea: Directly estimate the curvature of the object by estimating surface normal vectors

1. Normal Field Estimation




Our Idea: Directly estimate the curvature of the object by estimating surface normal vectors

1. Normal Field Estimation



2. Multiple Valid Surfaces





Our Idea: Directly estimate the curvature of the object by estimating surface normal vectors

1. Normal Field Estimation


2. Multiple Valid Surfaces

3. Single Unifying Function

4. Isosurface Optimization




Relative Signed
Distance Function



Isosurface: Surface
of constant value

Our Idea: Directly estimate the curvature of the object by estimating surface normal vectors

1. Normal Field Estimation




2. Multiple Valid Surfaces




3. Single Unifying Function

Relative Signed Distance Function




Isosurface: Surface of constant value

4. Isosurface Optimization



Final Surface Estimate


Our Idea: Directly estimate the curvature of the object
by estimating surface normal vectors



Synthetic
Aperture



Our Idea: Directly estimate the curvature of the object by estimating surface normal vectors



Synthetic
Aperture



Our Idea: Directly estimate the curvature of the object
by estimating surface normal vectors




Synthetic
Aperture




Object
Surface

mmWave signals experience primarily specular reflections

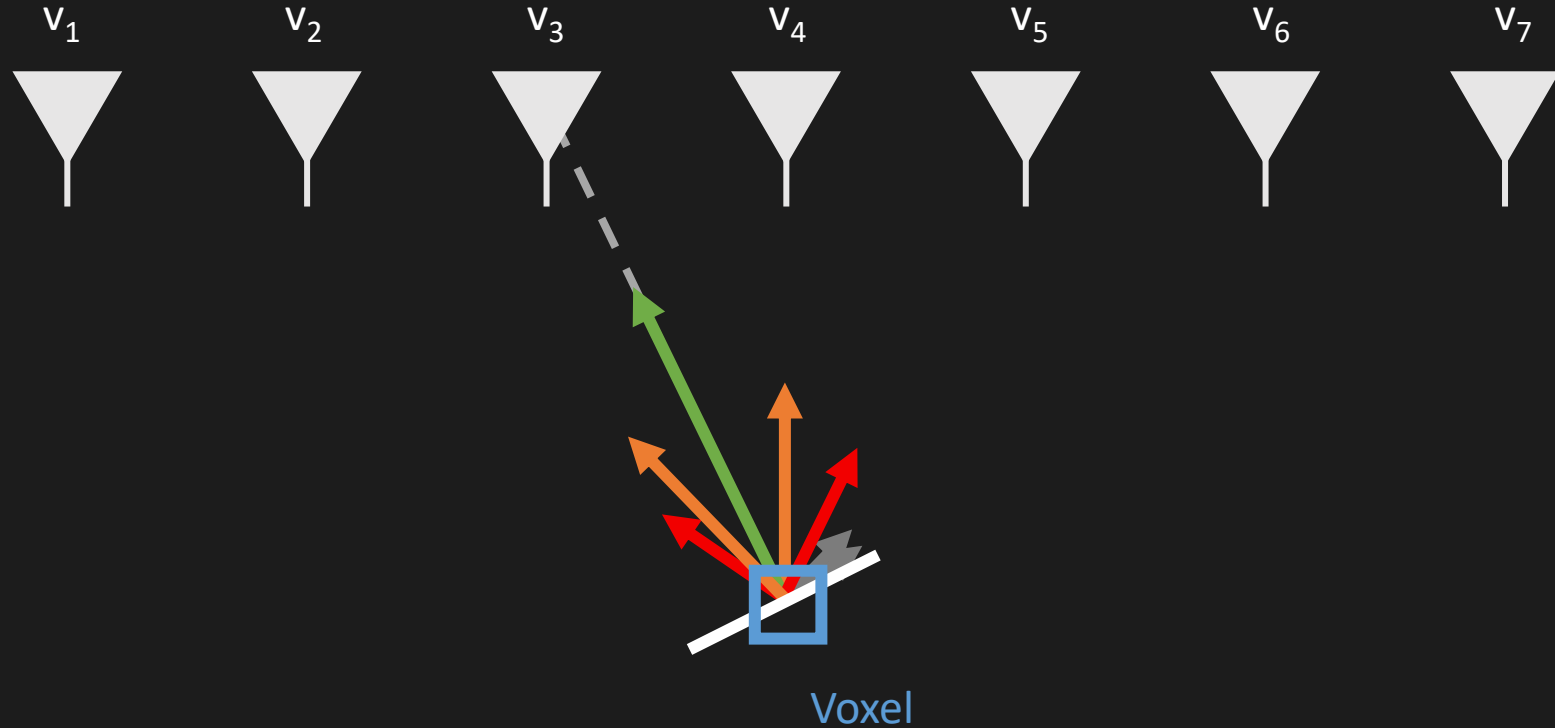


At a given voxel, what direction is the surface normal?

Lets say one of these antennas receive a reflection from this voxel

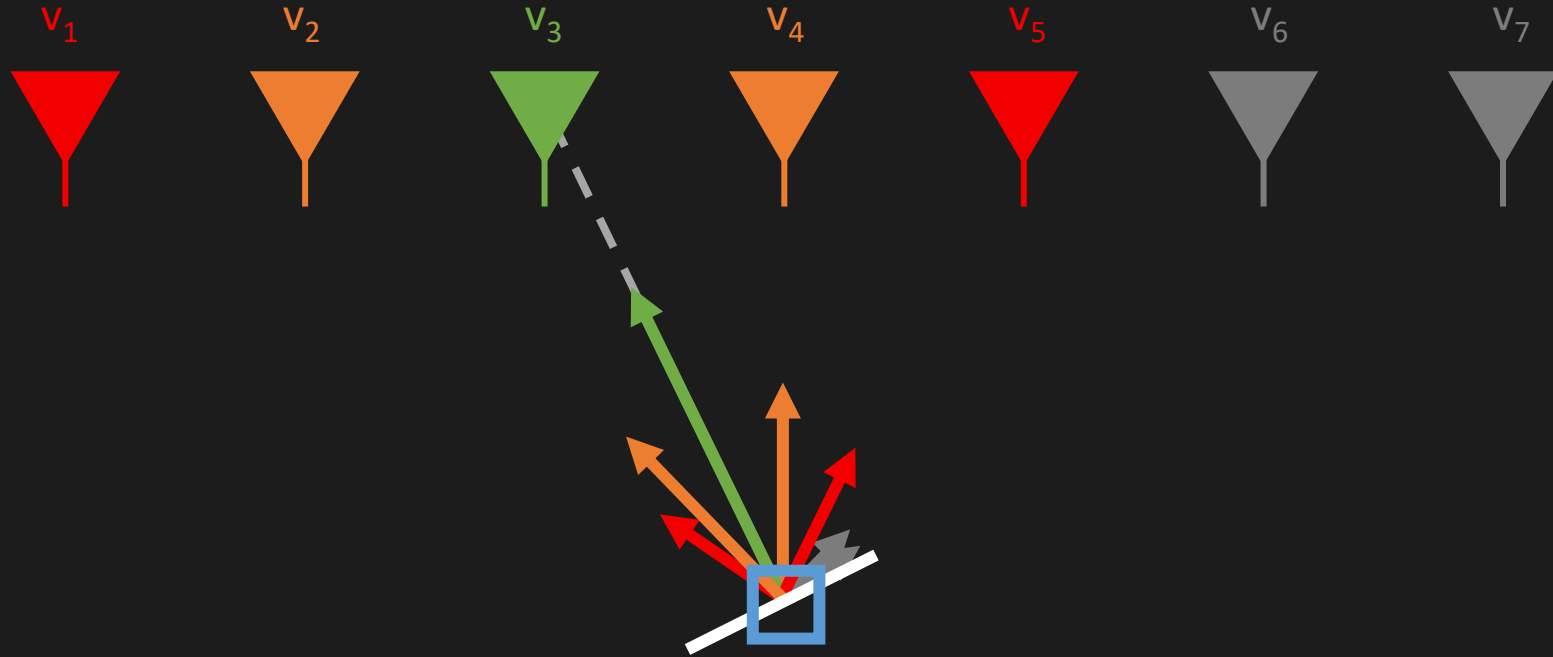


At a given voxel, what direction is the surface normal?



Reflection amplitude acts as a vote


At a given voxel, what direction is the surface normal?



Reflection amplitude acts as a vote

At a given voxel, what direction is the surface normal?

Example 1




Example 2


Reflection amplitude acts as a vote

At a given voxel, what direction is the surface normal?

Example 1




Example 2




Reflection amplitude acts as a vote

At a given voxel, what direction is the surface normal?

Example 1




Example 2




Reflection amplitude acts as a vote

At a given voxel, what direction is the surface normal?

Example 1




Example 2



Reflection amplitude acts as a vote

At a given voxel, what direction is the surface normal?




$$\vec{n}(u) = \sum_{k=0}^N v_k \frac{p_k - u}{\|p_k - u\|}$$

arriving from this voxel in SAR image


Reflection amplitude acts as a vote

Our Idea: Directly estimate the curvature of the object by estimating surface normal vectors

1. Normal Field Estimation




2. Multiple Valid Surfaces



3. Single Unifying Function

Relative Signed Distance Function



Isosurface: Surface of constant value


4. Isosurface Optimization




Final Surface Estimate

How can we capture all ambiguous surfaces in one function?

Computer Graphics

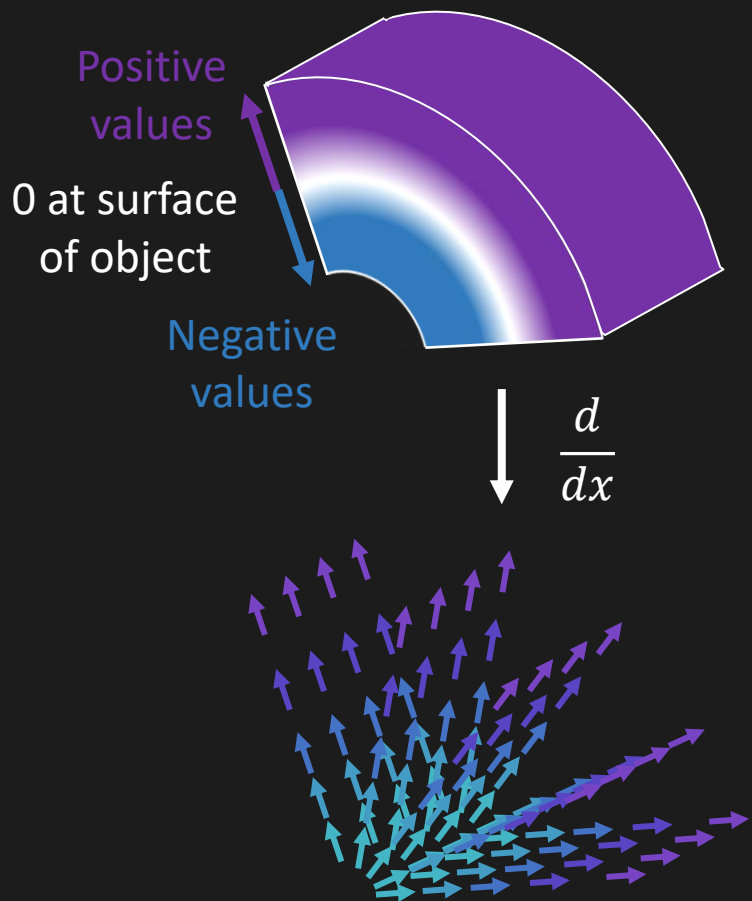


mmNorm

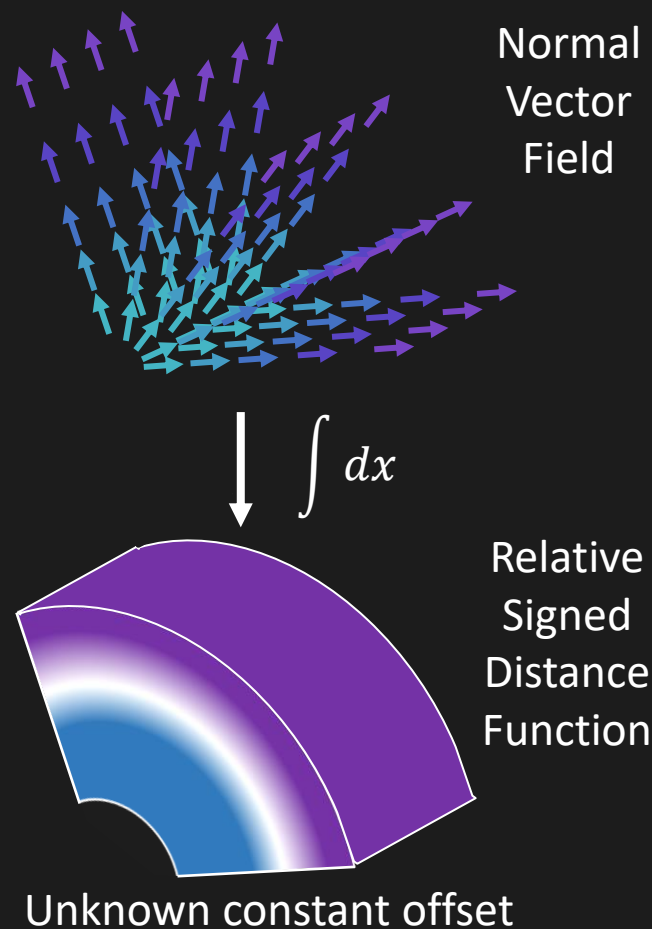


How can we capture all ambiguous surfaces in one function?

Computer Graphics




mmNorm




How can we capture all ambiguous surfaces in one function?

Computer Graphics




$$\downarrow \frac{d}{dx}$$




Signed
Distance
Function

Normal
Vector
Field

mmNorm




$$\downarrow \int dx$$




How can we capture all ambiguous surfaces in one function?

Computer Graphics




$$\downarrow \frac{d}{dx}$$




Signed
Distance
Function

Normal
Vector
Field

mmNorm



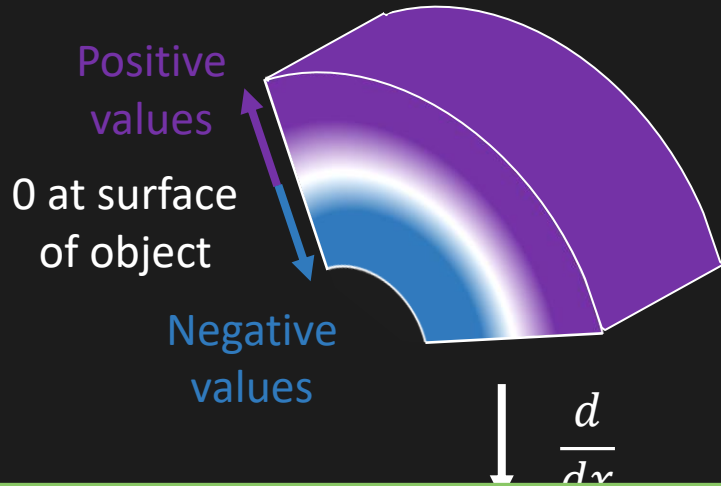
$$\downarrow \int dx$$




Unknown constant offset

How can we capture all ambiguous surfaces in one function?



Computer Graphics



mmNorm




We can determine the surface by finding this unknown constant offset




Our Idea: Directly estimate the curvature of the object by estimating surface normal vectors

1. Normal Field Estimation




2. Multiple Valid Surfaces



3. Single Unifying Function


Relative Signed Distance Function



Isosurface: Surface of constant value


4. Isosurface Optimization

$$L(C) = \sum_{k=1}^N \left\| \left| \frac{FFT(h(i,k))}{|FFT(h(i,k))|} \right| - C \right\|$$



Final Surface Estimate

Implementation



- UR5e Robot Arm
- TI IWR1443Boost 77GHz Radar
- Intel Realsense RGB-D camera for ground truth
- Measured experiments in line-of-sight and non-line-of-sight (cardboard occlusion)
- Collected experiments across 61 different everyday objects

Qualitative Results

Select every voxel in SAR
image above power
threshold

Qualitative Results

Full Object




RGB Surface



Baseline



mmNorm



Qualitative Results

Full Object



RGB Surface



Baseline



mmNorm




Measured 61 diverse everyday objects (YCB dataset)

Measure the distance from each point in the reconstruction to its nearest point in the ground truth

1.4 cm

0.25 cm

mmNorm: A paradigm shift in mmWave reconstruction



Novel Downstream Tasks

Large Scale Scenes


mmNorm Code



Robotics



Augmented Reality



Indoor
Scenes



Autonomous Driving



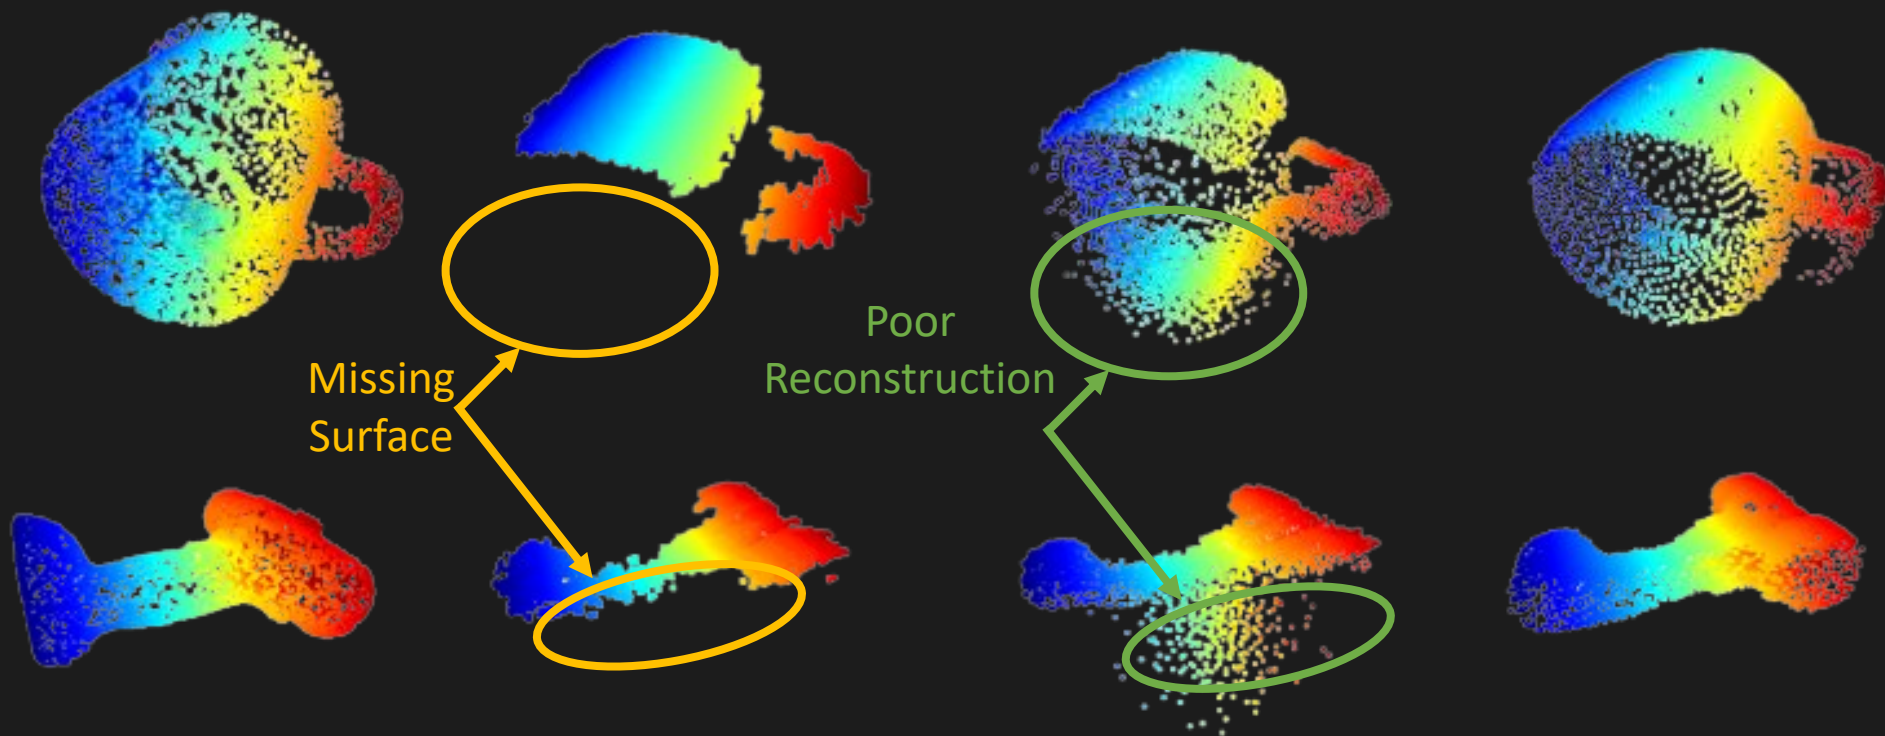
Limitation of mmNorm: Coverage

Ground-Truth


mmNorm

Applying Vision-Based
Shape-Completion


Wave-Former:
Physics-Aware
Shape Completion



Wave-Former: Physics-Aware Shape Completion



Incorporate :



Allows for training entirely on synthetic data and testing on real world data

Specularity-Aware

Inductive Bias:

Predict visible portions
based on mmWave Physics

Extending Computer Vision to Hidden Objects

A Tutorial on Millimeter-Wave Imaging and Reconstruction of Occluded Scenes

mmWave 101: Basics of mmWave Imaging, Limitations

State-of-the-art Reconstruction Methods


mmWave Neural Radiance Fields

Coffee Break

Vision-quality mmWave Scene Understanding

Open Problems; Tour of Existing Datasets & Tools

Hands-On Tutorial: mmWave Imaging & Reconstruction



Extending Computer Vision to Hidden Objects

A Tutorial on Millimeter-Wave Imaging and Reconstruction of Occluded Scenes

mmWave 101: Basics of mmWave Imaging, Limitations

State-of-the-art Reconstruction Methods


mmWave Neural Radiance Fields

Coffee Break

Vision-quality mmWave Scene Understanding

Open Problems; Tour of Existing Datasets & Tools

Hands-On Tutorial: mmWave Imaging & Reconstruction



Extending Computer Vision to Hidden Objects

A Tutorial on Millimeter-Wave Imaging and Reconstruction of Occluded Scenes

mmWave 101: Basics of mmWave Imaging, Limitations

State-of-the-art Reconstruction Methods


mmWave Neural Radiance Fields

Coffee Break

Vision-quality mmWave Scene Understanding

Open Problems; Tour of Existing Datasets & Tools

Hands-On Tutorial: mmWave Imaging & Reconstruction



Extending Computer Vision to Hidden Objects

A Tutorial on Millimeter-Wave Imaging and Reconstruction of Occluded Scenes

mmWave 101: Basics of mmWave Imaging, Limitations

State-of-the-art Reconstruction Methods


mmWave Neural Radiance Fields

Coffee Break

Vision-quality mmWave Scene Understanding

Open Problems; Tour of Existing Datasets & Tools

Hands-On Tutorial: mmWave Imaging & Reconstruction



Existing Tools

Objects

Humans

Self-Driving

Scenes

Real-World
Datasets

[MITO – 500 images](#)

[SquiggleMilli –
Handheld imaging –
1500 samples](#)

[mmBody – 200k
frames](#)

[Milipoint – 500k
frames](#)

[Mmvr – 300k
frames](#)

[Radatron – 500k
frames](#)

[Seeing through fog
without seeing fog
– 100k labeled pts](#)

[nuScenes - 1k
scenes](#)

[K-Radar -](#)

[PanoRadar – 12
buildings](#)

[CartoRadar – 6k
Frames](#)

[GRT – 1M](#)

Open-Source
Simulations

[MITO – Ray casting
based simulator](#)

[GeRaF – mmWave
NeRF](#)

[NVIDIA Sionna RT –
Designed for
mmWave
Communication](#)

Open Problems

Can we enable generalizable camera-quality mmWave perception?

Perception Tasks

- Achieving camera-quality:
 - Segmentation
 - Classification
 - 6 DoF Pose Estimation
 - Material Segmentation / Estimation
- Uncertainty Aware Reconstruction
- Zero-shot perception

Reconstruction Under Practical Constraints

- Smaller Apertures
 - Train Diffusion Models on reducing aperture sizes
- Sparse Apertures
 - Diffusion Models as Reconstruction Priors
- Real-time Performance
- Low Signal-to-Noise Ratio (SNR) environments
 - Thicker occlusions, fewer measurements, etc

mmWave Foundation Models

- Generalizing physical insights between different scales of mmWave imaging & other domain shifts
- Reducing Sim2Real Gap
- mmWave Self-Supervised Learning Techniques
- Multi-modal Fusion
- **Datasets & Benchmarking:**
 - **Make way towards large-scale, cross device, diverse, real-world datasets**

Extending Computer Vision to Hidden Objects

A Tutorial on Millimeter-Wave Imaging and Reconstruction of Occluded Scenes

mmWave 101: Basics of mmWave Imaging, Limitations

State-of-the-art Reconstruction Methods


mmWave Neural Radiance Fields

Coffee Break

Vision-quality mmWave Scene Understanding

Open Problems; Tour of Existing Datasets & Tools

Hands-On Tutorial: mmWave Imaging & Reconstruction



Extending Computer Vision to Hidden Objects

A Tutorial on Millimeter-Wave Imaging and Reconstruction of Occluded Scenes

mmWave 101: Basics of mmWave Imaging, Limitations

State-of-the-art Reconstruction Methods


mmWave Neural Radiance Fields

Coffee Break

Vision-quality mmWave Scene Understanding

Open Problems; Tour of Existing Datasets & Tools

Hands-On Tutorial: mmWave Imaging & Reconstruction



We would appreciate your feedback!



<https://tinyurl.com/cvpr-mmwave>

Hands-On Tutorial



<https://tinyurl.com/mmwave-tutorial>