

Wireless sensing

A person. A reflection. A measurement. Lecture 5 located a phone because the phone transmitted; this lecture locates and reads a person who carries no radio at all — from active localization to device-free sensing: reflection geometry, Doppler, Wi-Fi channel patterns a model can learn, and FMCW radar.

Yaxiong Xie




a person, a reflection, a measurement

nothing to carry, nothing to transmit

● — Active → passive

Last time, the phone sent the signal

LECTURE 5 · ACTIVE LOCALIZATION




We already know how to locate a signal source.

● — Active → passive

Now the person carries no transmitter

THIS LECTURE · DEVICE-FREE SENSING




Locate the reflector instead of the transmitter.

● — Active → passive

The target is passive; the sensor need not be

WHO TRANSMITS, WHO REFLECTS, WHO RECEIVES



the transmitter can still be active — the target is what is passive


Device-free sensing does not mean that no one transmits.

Three questions organize the lecture

01

Locate a reflector

TOF · AOA · Doppler



where is it?

02

Recognize an activity

CSI patterns · learning



what is the person doing?

03

Build the sensor

FMCW · structured data



how do we measure?

Geometry, recognition, and hardware answer different questions.

Can TOF and AOA locate the reflector of a signal?


01

Locate a reflector

Start from the geometry of lecture 5, with one change: the thing we locate reflects a signal instead of sending one. Echo delay draws a range circle, an array gives the echo a bearing, and a separate transmitter turns the circle into an ellipse. 15 min.

Radar sends a signal and measures its return

ONE RADAR, ONE TARGET: OUT AND BACK



Monostatic geometry; propagation greatly slowed.

The delay contains the outward and return journeys.

1 Geometry

2 Human echo

3 Doppler


4 Wi-Fi + learning

5 FMCW

6 Discussion

Echo delay gives a distance constraint

TOP VIEW · THE RADAR AT THE CENTRE



30 ns round-trip delay

4.5 m target range

range = $c \times \text{delay} / 2$: there and back

one delay, many possible positions

every point of it is 4.5 m from the radar

2-D slice; calibrated timing and a direct target echo assumed.

For a direct monostatic echo: range = light speed $\times$ delay / 2.

1 Geometry

2 Human echo

3 Doppler


4 Wi-Fi + learning

5 FMCW


6 Discussion

The array gives the echo an arrival direction

AN INCOMING WAVEFRONT



RELATIVE TIMING ACROSS ANTENNAS



Far-field plane-wave illustration; geometry not to scale.

AOA now describes the direction of an echo.

1 Geometry

2 Human echo

3 Doppler


4 Wi-Fi + learning

5 FMCW

6 Discussion

Range and bearing meet at the reflector

TOP VIEW · THE RADAR AT THE CENTRE



TOF → a range circle

AOA → an arrival bearing

their intersection → a location

one circle, one bearing, one reflector

A resolved direct echo can support geometric localization.

1 Geometry

2 Human echo

3 Doppler


4 Wi-Fi + learning

5 FMCW

6 Discussion

Separate TX and RX: the geometry changes

TOP VIEW · TRANSMITTER AND RECEIVER APART



2-D geometry; absolute delay or calibrated excess delay.

two distances add

the delay measures $d_1 + d_2$, not a range

same sum → an ellipse

TX and RX are its foci

AOA narrows the set

a bearing from RX picks one point of it

A bistatic delay constrains the sum of two distances.

1 Geometry

2 Human echo

3 Doppler


4 Wi-Fi + learning

5 FMCW

6 Discussion

A reflected arrival is not always the direct echo

TOP VIEW · THE RADAR ONLY KNOWS: HOW LATE, AND FROM WHERE



direct echo

straight out and back: its delay and its direction both point at the body ✓

the same pulse, via the wall

radar → wall → body → wall → radar: longer, and it returns from the wall

the wall acts as a mirror

wall → body = wall → image, so the echo looks like a straight trip to the image: a ghost

Real echo and imagined straight echo leave and return together: the radar cannot tell them apart.

Path association matters as much as delay and angle.

1 Geometry

2 Human echo

3 Doppler

4 Wi-Fi + learning

5 FMCW

6 Discussion

Which part of the mixture belongs to the person?


02

Find the human echo

A room returns many echoes: the direct path, the walls, the furniture — and the person. Delay, angle and Doppler are the three axes along which echoes can be told apart, background subtraction exposes what changed, and a track ties detections together. 20 min.

Which of these returns belongs to the person?

THE ROOM, ONE PERSON IN IT



direct path

fixed: the strongest arrival

wall reflection

fixed: the room does not move

body reflection

the one that moves with the person

the receiver gets all three at once — one mixed signal

Selected paths in an illustrative room.

Which contribution carries the human motion?

1 Geometry

2 Human echo

3 Doppler

4 Wi-Fi + learning

5 FMCW

6 Discussion

Separate echoes along measurable dimensions

01

Delay


different total path lengths



02

Angle


different arrival directions



03

Doppler

different path-change rates



Paths are resolvable only when measurements distinguish them.

1 Geometry

2 Human echo

3 Doppler


4 Wi-Fi + learning

5 FMCW

6 Discussion

Match delayed copies of a known probe

A KNOWN PROBE, ITS DELAYED COPIES, AND A SLIDING COMPARISON



Ideal known-code matching; not single-tone ranging.

Correlation asks how much of the probe arrived at each delay.

1 Geometry

2 Human echo

3 Doppler


4 Wi-Fi + learning

5 FMCW

6 Discussion

A delay profile contains candidate paths

ONE RECEIVER, ONE PROBE: HOW MUCH ARRIVED AT EACH DELAY



Synthetic delay profile; object labels are teaching ground truth.

Peak delay and peak height do not identify a person.

1 Geometry

2 Human echo

3 Doppler


4 Wi-Fi + learning

5 FMCW

6 Discussion

Finite bandwidth blurs nearby delays

THE SAME TWO PATHS, MEASURED WITH TWO BANDWIDTHS



Ideal response-width model; same two paths in every frame.

More measured bandwidth can separate closer path delays.

1 Geometry

2 Human echo

3 Doppler


4 Wi-Fi + learning

5 FMCW

6 Discussion

An array can separate echoes by direction


SAME DELAY, DIFFERENT ANGLES



test one arrival direction at a time

Illustrative receive-beam scan; finite-width beams are not rays.

SCAN THE RECEIVE BEAM



two angular peaks: two echoes

Spatial filtering distinguishes returns with different AOA.

1 Geometry

2 Human echo

3 Doppler

4 Wi-Fi + learning

5 FMCW


6 Discussion

Combine delay and angle before choosing a target


RANGE ALONE



RANGE × ARRIVAL ANGLE



THE 4 m SLICE



two peaks
along angle

two returns,
one cell each

Synthetic map; one strong cell is not necessarily one object.

A joint map can resolve overlap in one dimension.

1 Geometry

2 Human echo

3 Doppler

4 Wi-Fi + learning

5 FMCW


6 Discussion

Suppress stable clutter to expose change

01

Background

the empty room, measured once



02

Current scene


the room with the person in it



03

After comparison

subtract: only the change remains



what changed — human or not

Ideal coherent subtraction; calibration and stability assumed.

Background removal highlights changes, including nonhuman motion.

1 Geometry

2 Human echo

3 Doppler


4 Wi-Fi + learning

5 FMCW

6 Discussion

Associate detections across successive frames

TOP VIEW · WHAT THE RADAR DETECTS



every frame: a few detections

reflections, not yet people


match plausible motion

the person moves a little per frame

a person track

with an uncertainty margin,
and a guess for the next frame

THE SAME FLOOR, IN PERSPECTIVE



Illustrative association of reflection detections.

A human track is inferred from repeated measurements.

1 Geometry

2 Human echo

3 Doppler





4 Wi-Fi + learning

5 FMCW

6 Discussion

Motion helps, but it does not name the source

- 01 Strongest $\neq$ human**
a wall may dominate
- 02 Still $\neq$ absent**
background removal can hide it
- 03 Moving $\neq$ human**
a fan can also move



Combine geometry, dynamics, and task-specific evidence.

1 Geometry

2 Human echo

3 Doppler

4 Wi-Fi + learning

5 FMCW

6 Discussion

Why does a moving reflector shift the frequency?

03

Measure motion with Doppler

A reflector that moves returns its echoes on a changing schedule. That schedule is a frequency shift, a phase ramp, and a peak away from zero in a spectrum — motion as its own separation dimension, down to the swing of a limb. 15 min.


Motion changes the echo return schedule

EQUALLY SPACED MARKERS, OUT AND BACK

ONE COMMON CLOCK



each later marker meets him closer: a shorter trip



Equal-phase markers; speeds and time scales greatly exaggerated.

Approaching motion makes successive return paths shorter.

1 Geometry

2 Human echo

3 Doppler

4 Wi-Fi + learning

5 FMCW

6 Discussion

Closer arrival spacing means higher frequency

01


Stationary

the path length stays the same

sent, equally spaced



received



same frequency

02


Approaching

each path is shorter than the last

sent, equally spaced



received



higher frequency

03


Receding

each path is longer than the last

sent, equally spaced



received



lower frequency

Monostatic radial motion; differences exaggerated.

Doppler is a frequency shift from a changing propagation path.

1 Geometry

2 Human echo

3 Doppler


4 Wi-Fi + learning

5 FMCW

6 Discussion


Phase slope is another view of Doppler

SENT AND ECHO, COMPARED AT FIVE MOMENTS



twice as fast: +72° every moment
the echo's frequency is higher still

THE PHASE GAP OVER TIME



slope = how fast the gap grows = the Doppler frequency

Unwrapped relative phase; sign follows the receiver convention.

Doppler measures phase change per unit time.

1 Geometry

2 Human echo

3 Doppler


4 Wi-Fi + learning

5 FMCW

6 Discussion

A Doppler spectrum separates changing rates

HOW MUCH ECHO ARRIVES AT EACH RATE OF CHANGE



Synthetic spectrum; positive Doppler means approaching here.

Nonzero Doppler is motion evidence, not a human identity.

1 Geometry

2 Human echo

3 Doppler

4 Wi-Fi + learning

5 FMCW


6 Discussion

Same range, different motion: separable returns


RANGE ALONE



RANGE × RADIAL VELOCITY



THE 4 m SLICE



Synthetic range-velocity map; approaching velocity is positive.

Range and Doppler together can resolve an ambiguity.

1 Geometry

2 Human echo

3 Doppler


4 Wi-Fi + learning

5 FMCW

6 Discussion

Body parts produce different Doppler components

A WALKING PERSON




torso and limbs

torso: one steady Doppler

limbs: swinging to and fro

DOPPLER OVER TIME: A SPECTROGRAM



a time-varying Doppler distribution: a motion signature

Synthetic walking-like pattern; color encodes signal energy.

Micro-Doppler can provide a motion signature for recognition.

1 Geometry

2 Human echo

3 Doppler

4 Wi-Fi + learning

5 FMCW

6 Discussion

Can we recognize activity without resolving paths?

04

Recognize the mixed pattern

Wi-Fi packets already cross the room. Their channel state keeps amplitude and phase on many frequencies, and a moving body stirs that mixture in ways a model can learn — if the features, the training examples and the test split are chosen with care. 25 min.


Must we reconstruct every reflection?

SO FAR · PARTS 1–3

model the physics

- 1 **split the paths**
by delay, angle and Doppler
- 2 **keep the body's echo**
its range and its angle
- 3 **locate and track it**
a position, frame after frame
- 4 **read the activity**
from how that position moves

ONE MIXED OBSERVATION



the receiver gets one mixed signal:
direct + walls + the body, added up

NOW · PART 4

learn from the change

- 1 **the whole signal over time**
nothing separated, no path model
- 2 **an ML model, or DSP**
trained on labelled examples
- 3 **read the activity**
walk · sit down · fall



the input: how the signal changes in time

Activity recognition can work without a complete path reconstruction.

1 Geometry

2 Human echo

3 Doppler


4 **Wi-Fi + learning**

5 FMCW

6 Discussion

Wi-Fi packets provide repeated channel observations

AN ACCESS POINT, A LAPTOP, AND THE PERSON BETWEEN THEM



every packet carries known training symbols: a measurement for free



Packet timing shown schematically.

Communication signals can also probe human-induced changes.

1 Geometry

2 Human echo

3 Doppler


4 Wi-Fi + learning

5 FMCW


6 Discussion

CSI keeps amplitude and phase across components

ONE PACKET, ONE FREQUENCY AT A TIME



Blue: what was sent · coral: what arrived. Each frequency gives one amplitude ratio and one phase shift.



CSI is a channel measurement, not a body map.

1 Geometry

2 Human echo

3 Doppler


4 Wi-Fi + learning

5 FMCW

6 Discussion

A changing path changes the mixed CSI

THE ROOM: EVERY PATH ARRIVES




fixed: the direct path and the wall

moving: the person's path

6 cm more path = one full turn

EACH PATH IS AN ARROW




static part: direct + wall, added once

dynamic part: the body's arrow turns


total = static + dynamic: what we measure

WHAT THE RECEIVER SEES

amplitude: the total's length



phase: the total's angle



both wobble as the body arrow turns

One dynamic contribution plus a static reference; synthetic CSI.

Useful temporal structure can survive unresolved multipath.

1 Geometry

2 Human echo

3 Doppler

4 Wi-Fi + learning

5 FMCW


6 Discussion

A recognizer learns from labeled signal windows

01

CSI window


a few seconds of CSI,
one curve per subcarrier



02

Representation

features a model can compare;
here a spectrogram
how fast the CSI changes




colour = energy; fast change
≈ fast-moving body parts

03

Trained model


learns which pattern goes
with which label



04

Activity

the label with the
highest score



Illustrative recognition pipeline; no performance claim.

The model connects a measured pattern to a task label.

1 Geometry

2 Human echo

3 Doppler


4 Wi-Fi + learning

5 FMCW

6 Discussion

Training requires examples of variation

FOUR RECORDINGS OF EACH ACTIVITY, MADE AT DIFFERENT TIMES



Split by whole recording: windows of one recording on both sides is data leakage.

Learn the activity across examples, not one recording.

1 Geometry

2 Human echo

3 Doppler

4 Wi-Fi + learning

5 FMCW


6 Discussion

Choose features that retain useful motion

01

Statistics


amplitude distributions



02

Time-frequency


motion-energy patterns



03

Learned sequence

multi-component input



The representation determines what the classifier can use.

1 Geometry

2 Human echo

3 Doppler




4 Wi-Fi + learning

5 FMCW

6 Discussion

E-eyes: compare location–activity profiles

LOCATION–ACTIVITY PROFILES · ONE HISTOGRAM OF CSI AMPLITUDE EACH



a new sample matches this profile

the location context is part of the model: a profile belongs to a place

Method redraw; illustrative profiles, not paper measurements.

A profile can identify activity without resolving each ray.

1 Geometry

2 Human echo

3 Doppler

4 Wi-Fi + learning

5 FMCW


6 Discussion

CARM: preserve the sequence of motion speeds

01

Denoise CSI

PCA across subcarriers




the correlated part survives
the noise does not

02

Motion-energy sequence

wavelet features over time

motion speed




time

how fast the body parts move, over time

03

Activity dynamics

HMM activity models



an activity is a sequence of states
each with its own speed profile

Conceptual pipeline; synthetic feature sequence.

Physical structure can guide features used for learning.

1 Geometry

2 Human echo

3 Doppler


4 Wi-Fi + learning

5 FMCW

6 Discussion


The same gesture can look different elsewhere

SAME GESTURE · SETUP A




move the receiver

SAME GESTURE · SETUP B



the room is
in the data too



its signal pattern



a different pattern

Synthetic example of domain-dependent signals.

A classifier can learn room cues along with activity.

1 Geometry

2 Human echo

3 Doppler

4 Wi-Fi + learning

5 FMCW


6 Discussion

Widar3.0: use physics before the neural network

01

Several Wi-Fi links

Doppler on each link




02

Body-coordinate velocity

a BVP sequence

forward velocity




needs the person's position and orientation

BVP is a velocity distribution, not a body image.

03

Spatial-temporal network

a gesture label out



a physical representation first,
so the network learns the gesture, not the room

A stable physical representation can help transfer recognition.

1 Geometry

2 Human echo

3 Doppler

4 Wi-Fi + learning

5 FMCW


6 Discussion

Test the change expected after deployment

01

Leave-one-session-out

test set: a new recording




one whole group is the test set

02

Leave-one-subject-out

cross-person: a new participant




one whole group is the test set

03

Leave-one-room-out

cross-environment: a new room



one whole group is the test set

blue = training set · coral = test set · rotate the coral group: cross-validation

... and put unknown and poor-quality inputs in the test set too

A good split tests the intended generalization claim.

1 Geometry

2 Human echo

3 Doppler

4 Wi-Fi + learning

5 FMCW


6 Discussion

Wi-Fi recognition still needs a quality decision

THREE INPUTS, THREE DECISIONS

Supported input


regular, quiet observations



✓ use for inference

Unseen interference

large unrelated motion



✗ not reliable

Too few observations

packets missing



✗ not reliable

Illustrative quality decisions; no measured performance claim.

"Unknown" or "not reliable" can be a valid output.

1 Geometry

2 Human echo

3 Doppler

4 Wi-Fi + learning

5 FMCW

6 Discussion

How does one radar measure range, motion and angle?

05

Build the measurement with FMCW

A radar built for sensing: a chirp turns a few nanoseconds of delay into a slow beat, repeated chirps expose motion, and several antennas add angle — the joint range–velocity–angle measurement the earlier parts asked for. 30 min.

Tiny echo delays need a practical measurement

A REFLECTOR 3 METRES AWAY



20 ns

round-trip delay

far too short to time directly on a cheap chip



measure a slow beat instead



FMCW converts propagation delay into a frequency difference.

1 Geometry

2 Human echo

3 Doppler


4 Wi-Fi + learning

5 FMCW

6 Discussion

mmWave is the band; FMCW is the method

THE RADAR OF THIS PART



a 77 GHz FMCW radar chip
one transmitter, four receive antennas

millimeter wave

a short radio wavelength: $\lambda \approx 3.9$ mm at 77 GHz




a 2 mm move of the chest turns the phase by a full cycle

FMCW

frequency-modulated continuous wave: a designed sweep

frequency



time

77 → 81 GHz in 40 μ s,
then again, chirp after chirp

This example uses millimeter-wave FMCW radar.

1 Geometry

2 Human echo

3 Doppler

4 Wi-Fi + learning

5 FMCW

6 Discussion

A chirp sweeps through frequency

ONE CHIRP: FREQUENCY AGAINST TIME



transmitted frequency

follows a known schedule:
at any instant we know
exactly what was sent



A linear sweep; carrier oscillations greatly slowed.

The transmitted frequency follows a known schedule.

1 Geometry

2 Human echo

3 Doppler


4 Wi-Fi + learning

5 FMCW

6 Discussion

The echo repeats an earlier part of the sweep

ONE CHIRP: FREQUENCY AGAINST TIME



the frequency gap

is the same at every instant,
and it grows with the delay:
farther target → larger gap

Stationary target; echo delay exaggerated.

For a stationary target, longer delay creates a larger frequency gap.

1 Geometry

2 Human echo

3 Doppler


4 Wi-Fi + learning

5 FMCW

6 Discussion

Mixing exposes the slow difference frequency

TRANSMIT REFERENCE $\times$ ECHO $\rightarrow$ THE DIFFERENCE FREQUENCY



Illustrative input frequencies and difference component.

The beat signal makes the delay measurable.

1 Geometry

2 Human echo

3 Doppler


4 Wi-Fi + learning

5 FMCW


6 Discussion

Move the reflector; watch the beat change


TARGET DISTANCE



ECHO DELAY



BEAT SIGNAL



Quasi-static range examples; motion between measurements is schematic.

Farther target → longer delay → larger gap → faster beat.

1 Geometry

2 Human echo

3 Doppler



4 Wi-Fi + learning

5 FMCW

6 Discussion

Several reflectors produce several beat components

TWO REFLECTORS: A PERSON AT 3 m, A WALL AT 7.2 m



beat frequency $\propto$ delay $\propto$ distance: the FFT is a range profile

Two ideal targets; schematic range responses.

The range FFT organizes the beat signal by distance.

1 Geometry

2 Human echo

3 Doppler


4 Wi-Fi + learning

5 FMCW


6 Discussion

Bandwidth controls how close two ranges can be

TWO REFLECTORS 20 cm APART



THE RANGE RESPONSE AROUND THEM



resolution = $c / 2B$ — but a 1 mm shift still moves the phase: sensitivity is another thing
Ideal power-response model, not measured radar data.

Range resolution and displacement sensitivity are different.

1 Geometry

2 Human echo

3 Doppler

4 Wi-Fi + learning

5 FMCW


6 Discussion

Repeat chirps to measure phase evolution

SELECT ONE DISTANCE REGION



REPEAT THE MEASUREMENT



Slow-time phase progression reveals radial motion.

1 Geometry

2 Human echo

3 Doppler

4 Wi-Fi + learning

5 FMCW

6 Discussion


Three sampling axes expose three dimensions

- 01

Within one chirp
fast time
- 02

Across chirps
slow time
- 03

Across antennas
space



samples of the beat



Range



the same bin, chirp after chirp



Doppler



the same bin on each antenna



Angle

Conceptual sampling axes; implementation order can vary.

The data organization follows the measurement physics.

1 Geometry

2 Human echo

3 Doppler

4 Wi-Fi + learning

5 FMCW


6 Discussion

Build a joint range–velocity map


RANGE ALONE



RANGE × RADIAL VELOCITY



THE 4 m SLICE



two peaks
along velocity

two returns,
one cell each

Synthetic range–velocity map; approaching velocity is positive.

FMCW supplies the joint evidence introduced earlier.

1 Geometry

2 Human echo

3 Doppler

4 Wi-Fi + learning

5 FMCW


6 Discussion

Estimate angle at the selected range-Doppler cell

01

Select a range-Doppler cell


one candidate return
velocity



02

Read this cell on each RX


four antennas, four phases



03

Estimate its AOA

the angle of that return
response



Associate range, Doppler, and angle from the same return.

1 Geometry

2 Human echo

3 Doppler


4 Wi-Fi + learning

5 FMCW

6 Discussion

Track a person through successive detections

TOP VIEW · WHAT THE RADAR DETECTS



every frame: a few detections

reflections, not yet people


match plausible motion

the person moves a little per frame

a person track

with an uncertainty margin,
and a guess for the next frame

THE SAME FLOOR, IN PERSPECTIVE



Illustrative person-level tracking, not a body reconstruction.

Applications need detection, association, and interpretation.

1 Geometry

2 Human echo

3 Doppler


4 Wi-Fi + learning

5 FMCW


6 Discussion

Select a person before measuring small motion

TWO PEOPLE IN FRONT OF THE RADAR



THE PHASE OF THE SELECTED ECHO



15 breaths a minute — and the ripple is the heartbeat

Synthetic scene and breathing trace; no accuracy claim.

Spatial separation helps isolate the motion of interest.

1 Geometry

2 Human echo

3 Doppler

4 Wi-Fi + learning

5 FMCW

6 Discussion

What must be measured before making the claim?

06

Choose evidence for the task

Two inference routes from one physical process. Match the measurement and the representation to the claim the system is supposed to make. 5 min.

Match the approach to the output you need

WHAT COMES OUT, AND WHAT HAS TO GO IN

Estimate physical state

path parameters → geometry: delay, angle, Doppler of a resolved return

position · radial velocity

Recognize a task label

a temporal representation of the mixture → a trained model

walk · sit · gesture

Combine the two

physical features first, learned inference on top

track, then classify

Wi-Fi and radar can both combine physics with learning.

1 Geometry

2 Human echo

3 Doppler

4 Wi-Fi + learning

5 FMCW

6 Discussion

Explain these three design choices

THREE QUESTIONS · CLICK FOR EACH ANSWER

1 Why is the strongest echo not enough?

✓ A wall may dominate; a peak has a delay and an angle, but no identity.

2 Why can ML work without resolving every ray?

✓ The mixture can keep useful activity patterns even when no path is resolved.

3 Why does FMCW repeat chirps?

✓ The phase of one range bin across chirps reveals motion: Doppler in slow time.

Connect each output to the measurements that support it.

1 Geometry

2 Human echo

3 Doppler


4 Wi-Fi + learning

5 FMCW

6 Discussion

One physical process, two inference routes

ONE PHYSICAL PROCESS



illumination → reflection → measurement

TWO INFERENCE ROUTES



Passive target

a device-free target changes propagation — nothing on it transmits



Geometry or patterns

resolve paths and measure geometry, or recognize a temporal pattern



Evidence for a task

validate the inference on the change the task will meet

Illumination → reflection → measurement → task-specific inference.

1 Geometry

2 Human echo

3 Doppler

4 Wi-Fi + learning

5 FMCW

6 Discussion

Research systems behind the Wi-Fi examples

THE SYSTEMS BEHIND THE EXAMPLES

01 E-eyes · MobiCom 2014

https://www.winlab.rutgers.edu/~gruteser/papers/E_eyes_paper.pdf

02 CARM · MobiCom 2015

<https://www.sigmobile.org/mobicom/2015/papers/p65-wangA.pdf>

03 Widar3.0 · MobiSys 2019

https://cswu.me/papers/mobisys19_widar3_paper.pdf

04 WiSee · Whole-home gesture recognition · MobiCom 2013

<https://ubicomplab.cs.washington.edu/pdfs/whole-home-gesture.pdf>

05 FarSense · Wi-Fi respiration sensing · UbiComp 2019

<https://arxiv.org/abs/1907.03994>

1 Geometry

2 Human echo

3 Doppler

4 Wi-Fi + learning

5 FMCW

6 Discussion

Further reading on radar measurements

FOLLOW-UP READING

01 **Texas Instruments · FMCW radar training**

<https://www.ti.com/video/series/mmwave-training-series.html>

02 **TI · Fundamentals of mmWave radar sensors**

<https://www.ti.com/lit/wp/spyy005a/spyy005a.pdf>

Both are free: the training videos for the pictures, the white paper for the equations.

1 Geometry

2 Human echo

3 Doppler

4 Wi-Fi + learning

5 FMCW

6 Discussion