COM-405: Mobile Networks

Lecture 9.1: Wireless Sensing & Imaging

Haitham Hassanieh

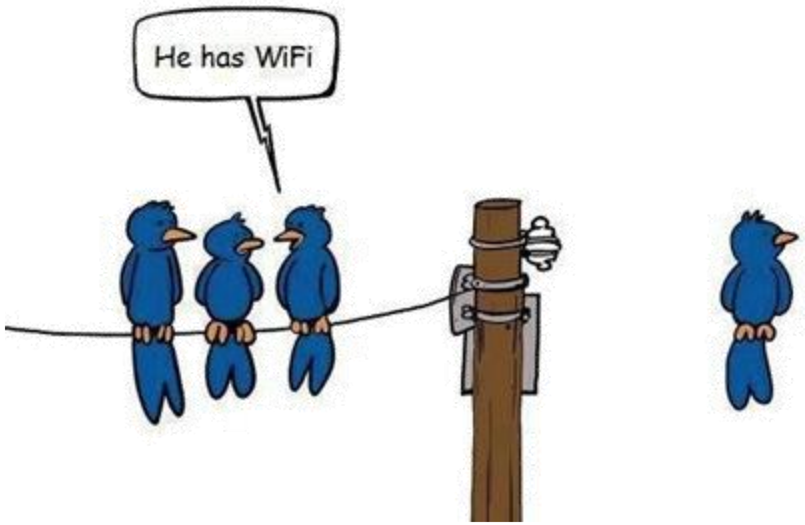

EPFL

Laboratory of SEnsing & Networking Systems

Can you tell people's emotions even if they don't show up on their faces?

Smart Homes that adapt to our mood

Did I get the Job? .... No

Does my advisor like my work?

Graduate student

Advisor

Combating Depression

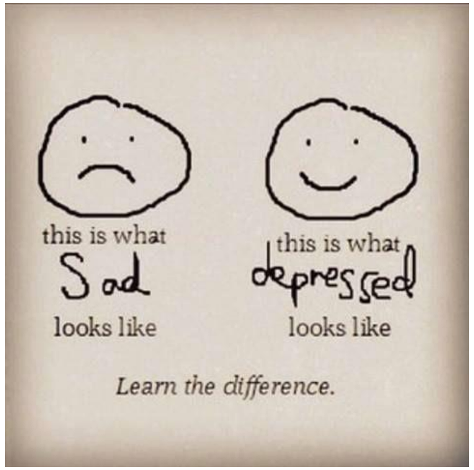

Is the date going well!

Existing approaches measure vital signs

- Use ECG to get very accurate heartbeats

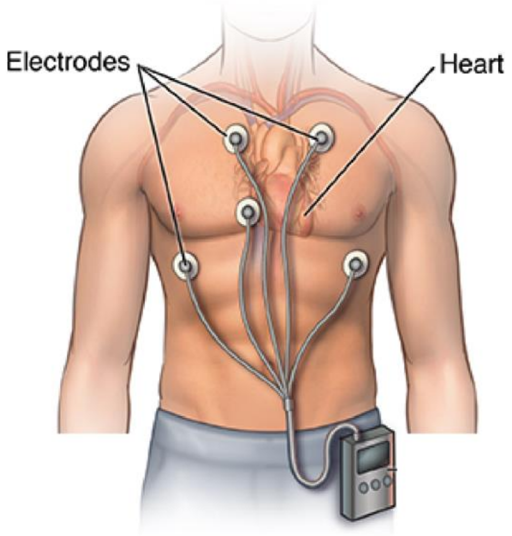

Input signal

Wireless reflection of the human body

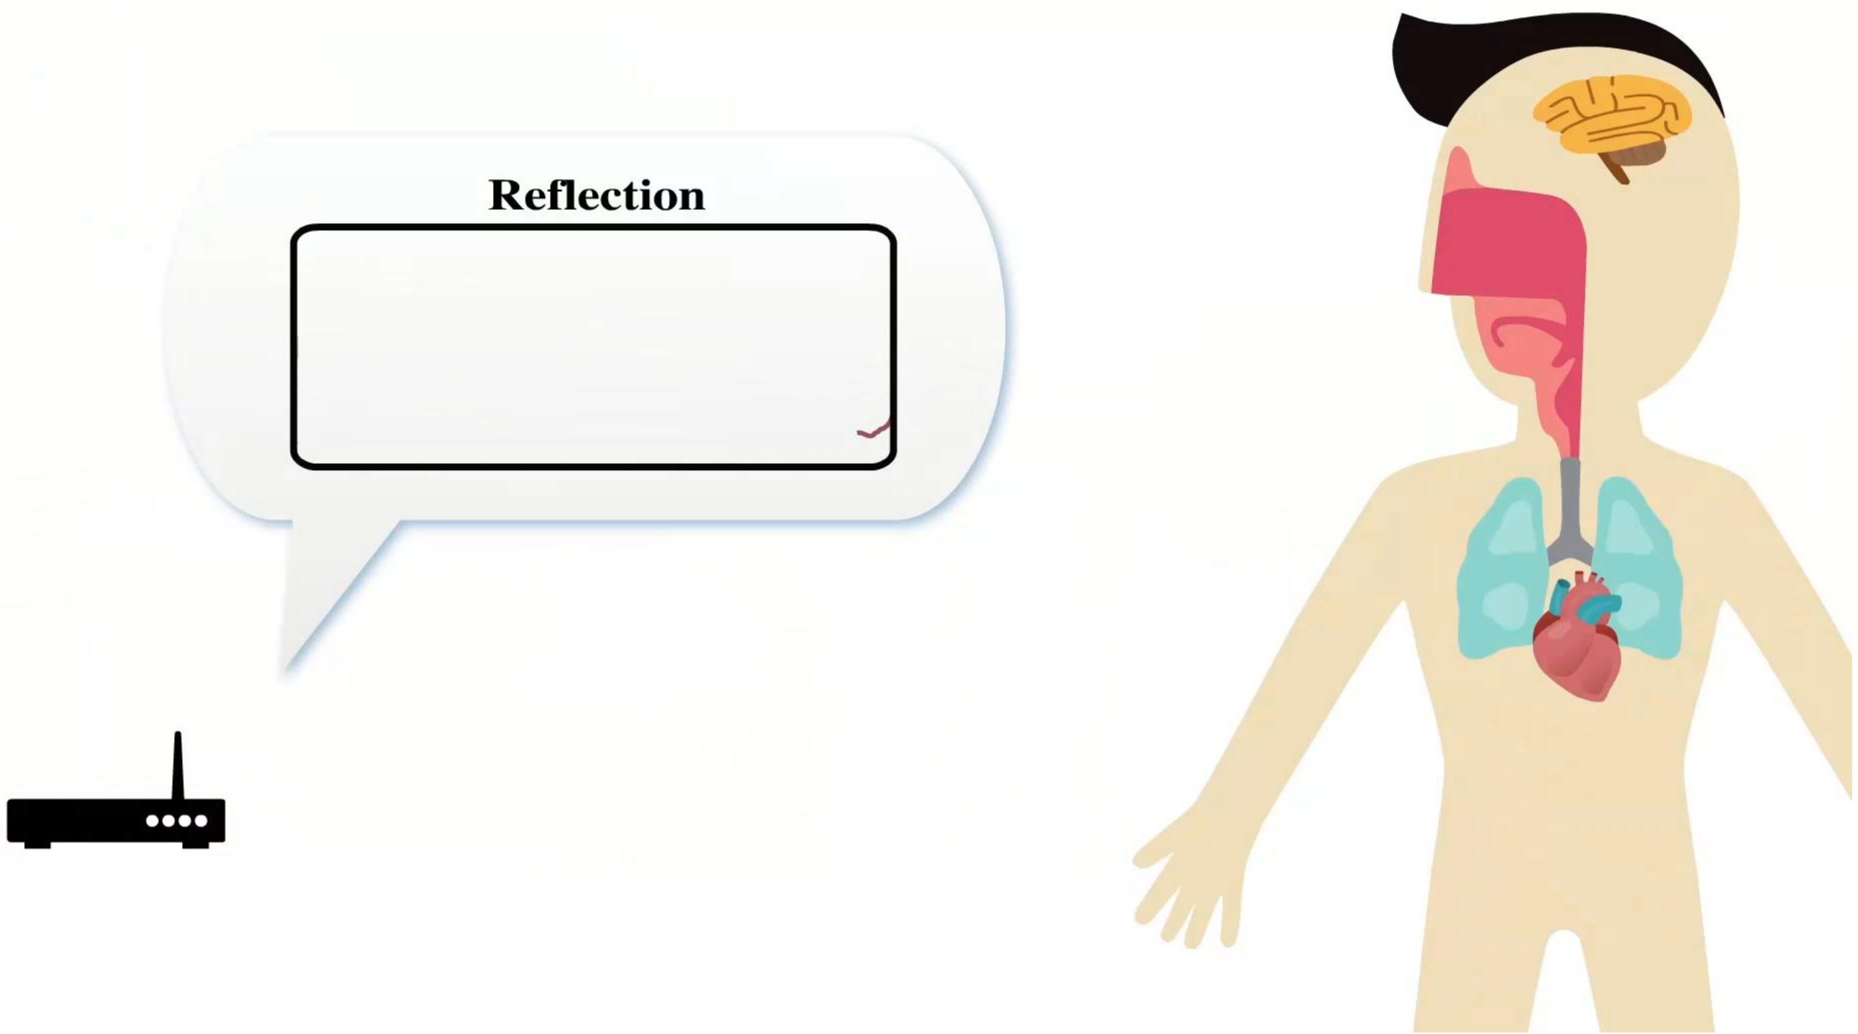

Emotion recognition using wireless signals

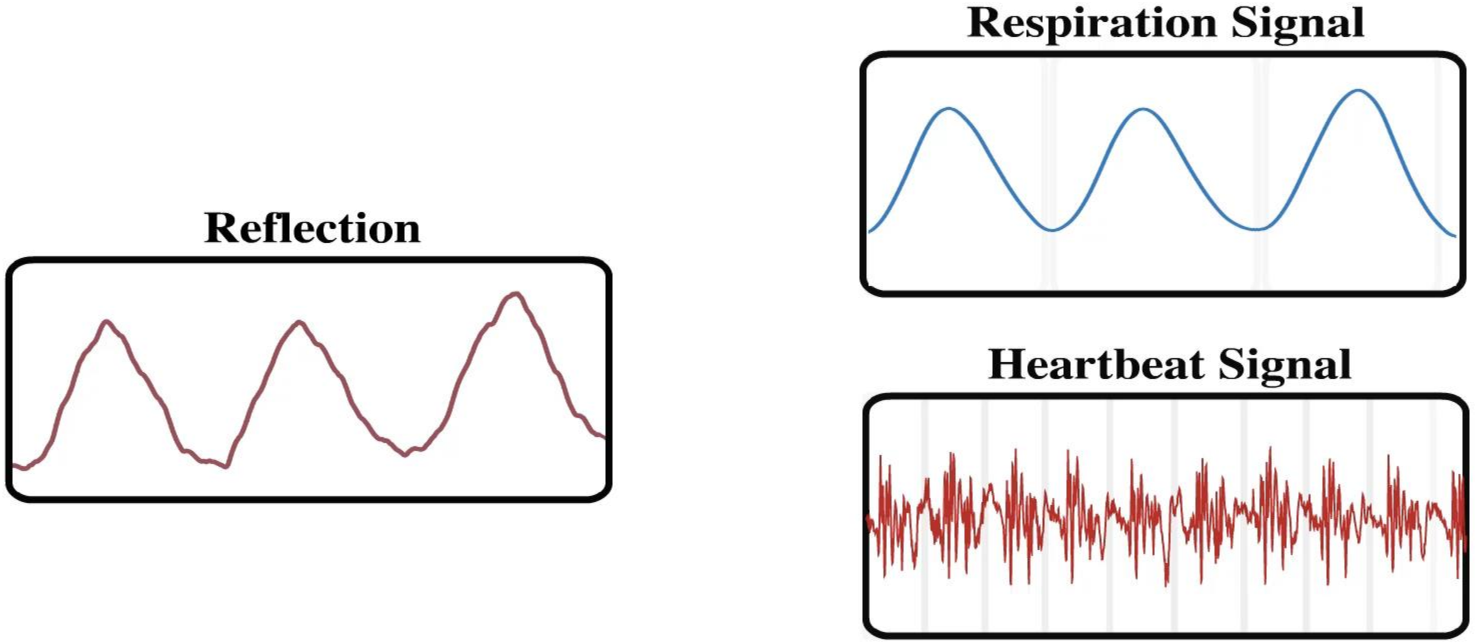

Key challenge: Inter-Beat Interval (IBI)

- Emotion recognition needs accurate measurements of the length of every single heartbeat

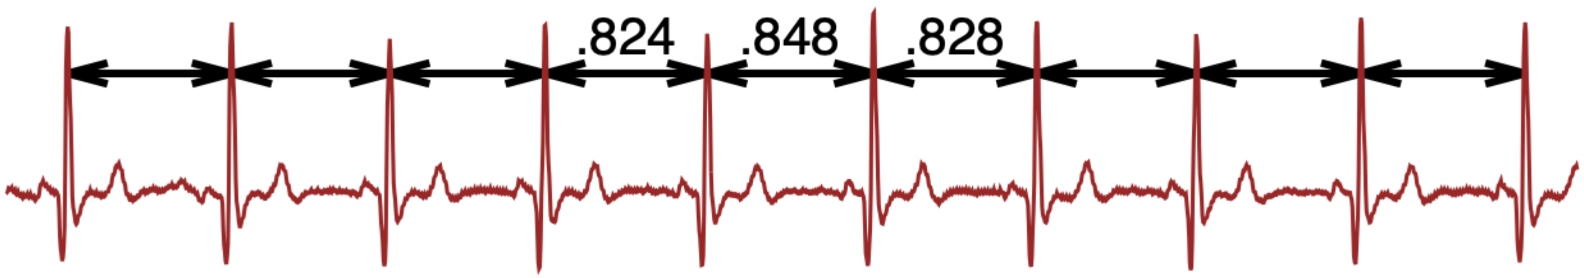

We need to extract IBI with accuracy over 99%

Emotion Model

- Standard 2D emotion model
- Classify into **anger**, **sadness**, **pleasure** and **joy**

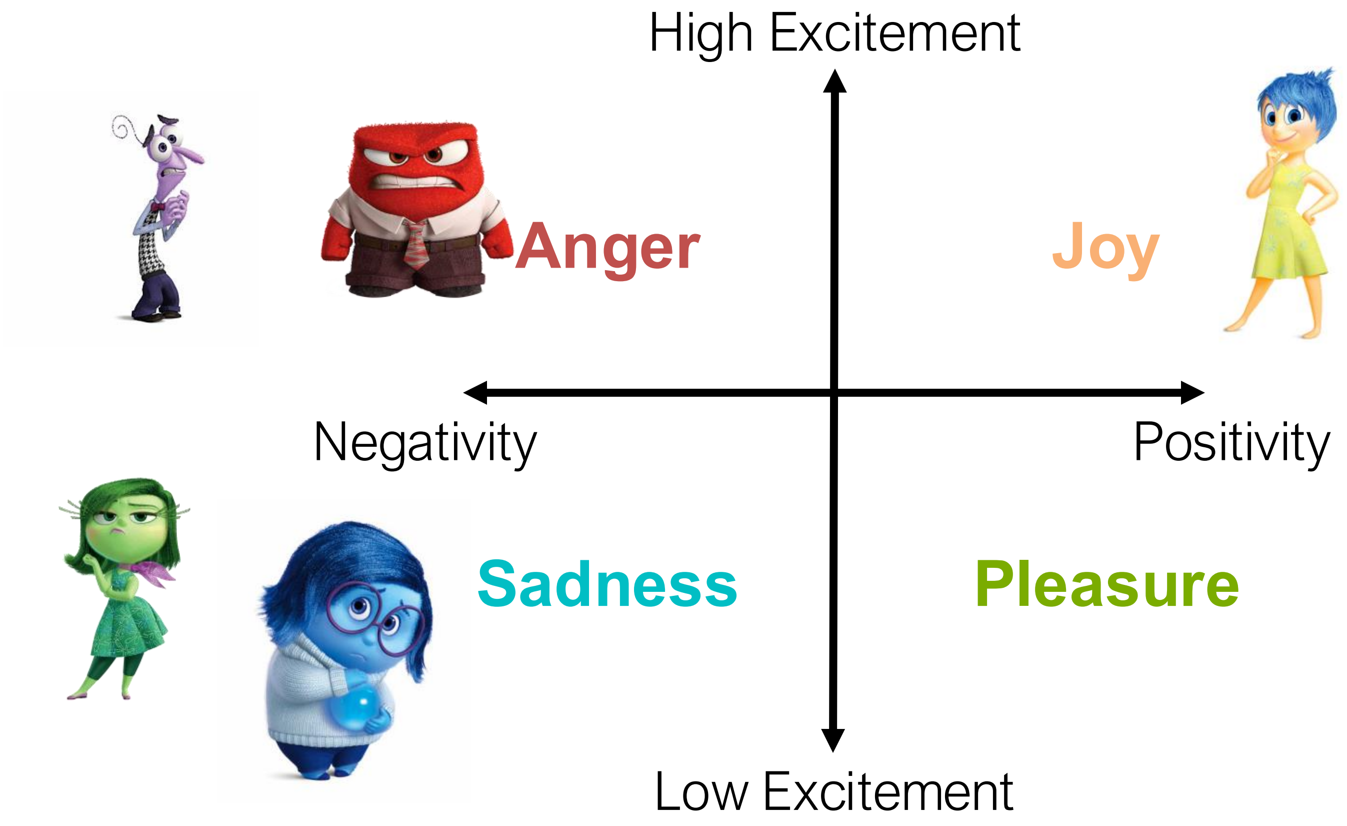

Person-dependent Classification

- Train and test on the same person

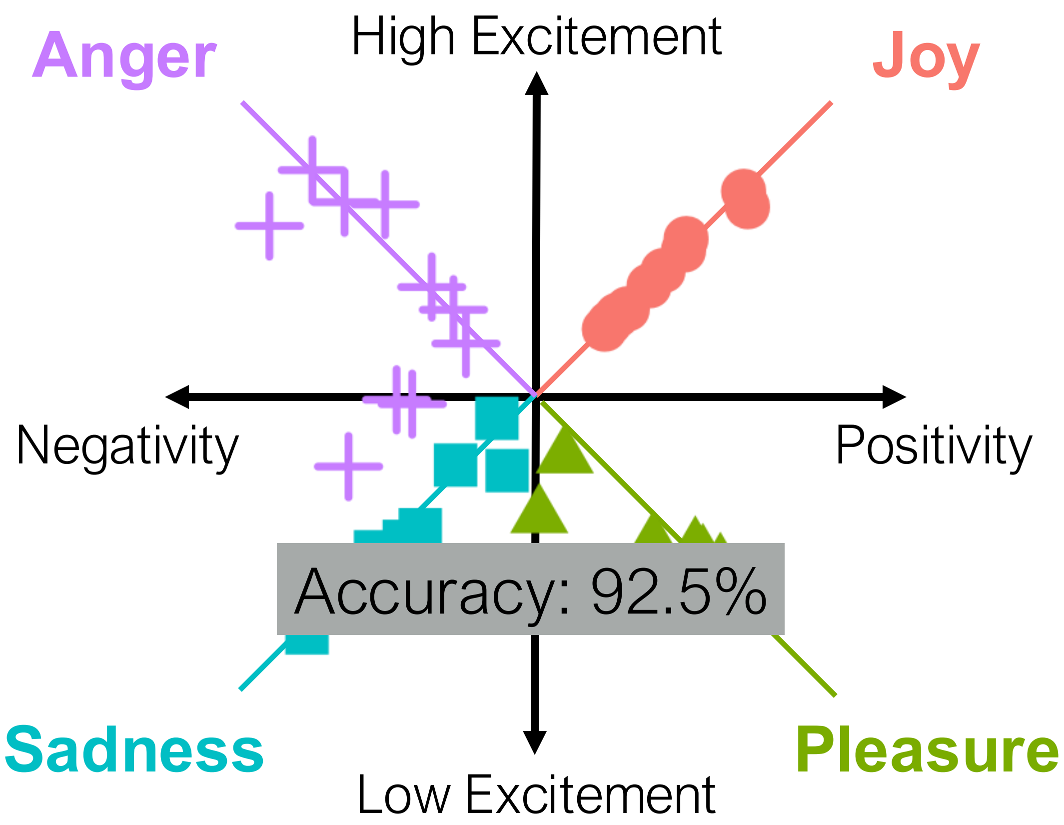

Understanding Diseases with Sleep Stages

But, monitoring sleep stages is difficult …

done in hospital with many electrodes on the body

Experience in Sleep Lab

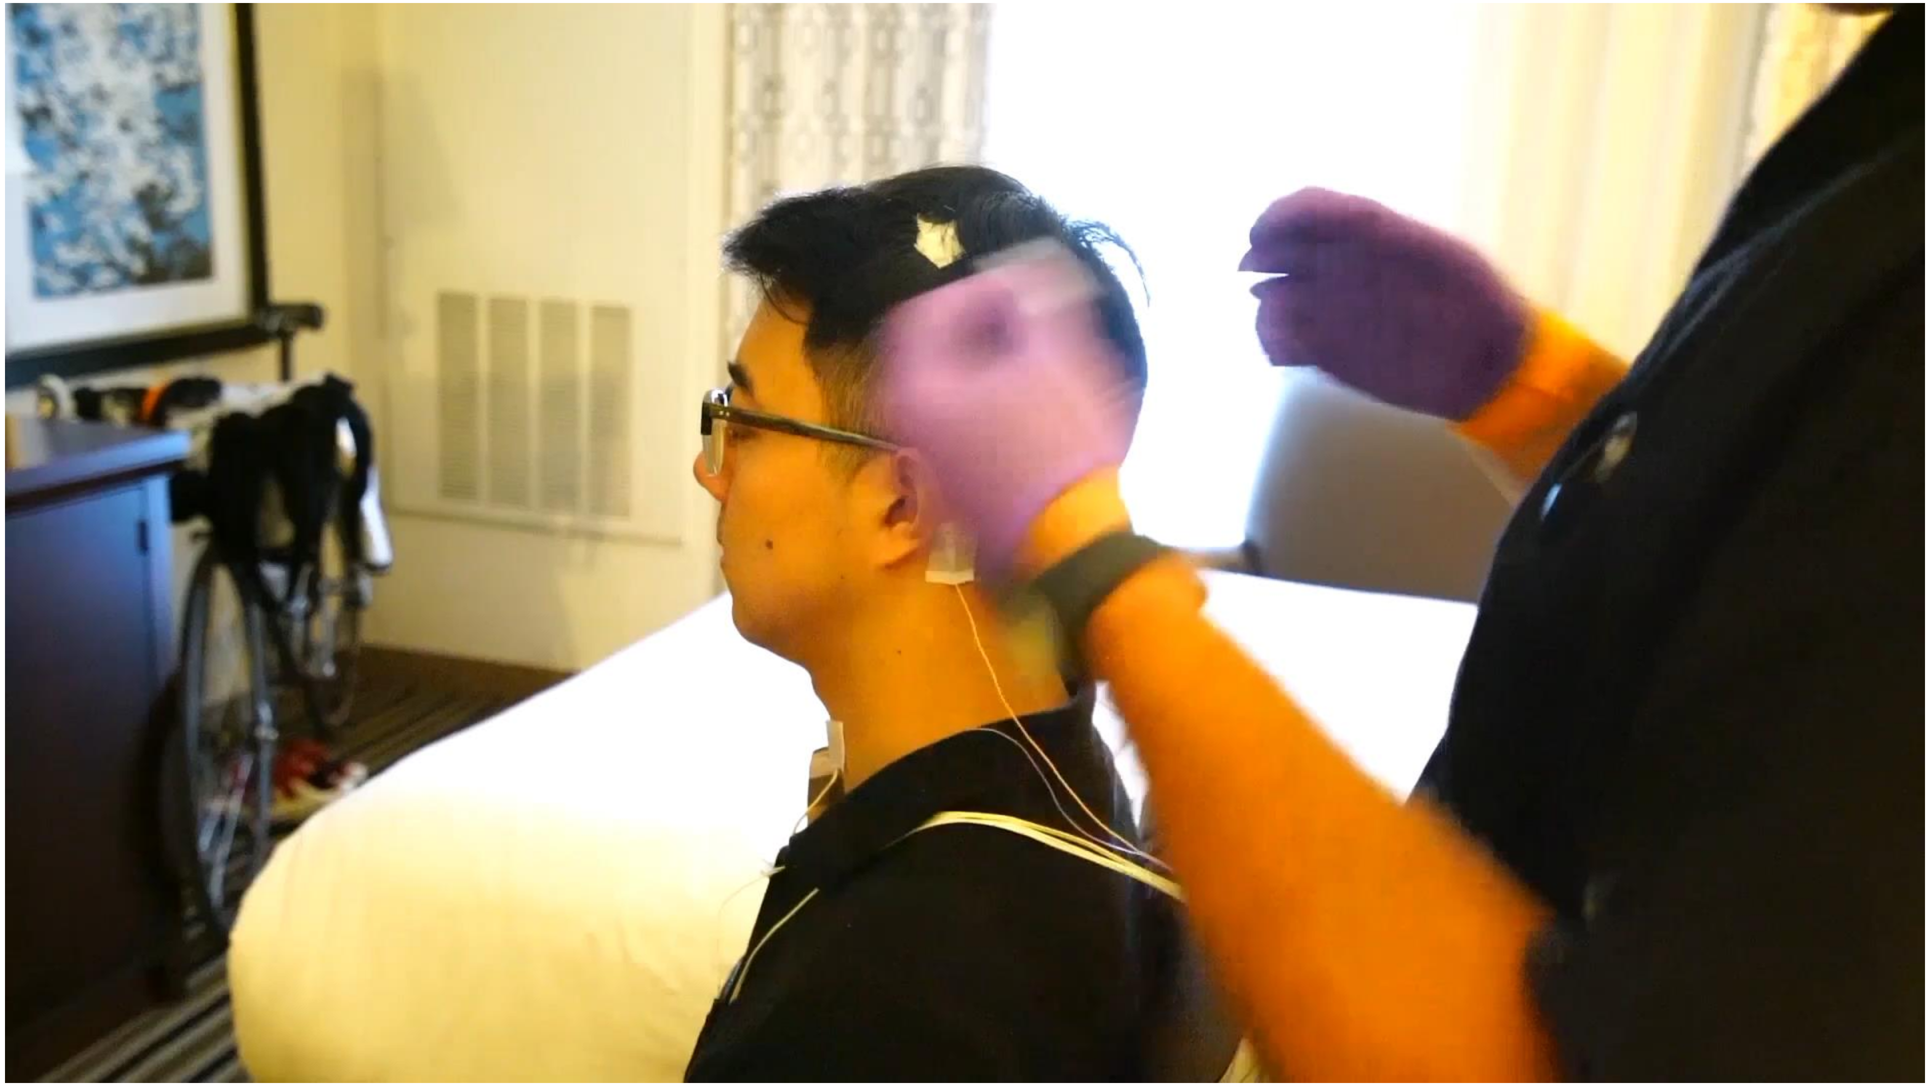

Experience in Sleep Lab

Wireless Sensing Sleep Stages

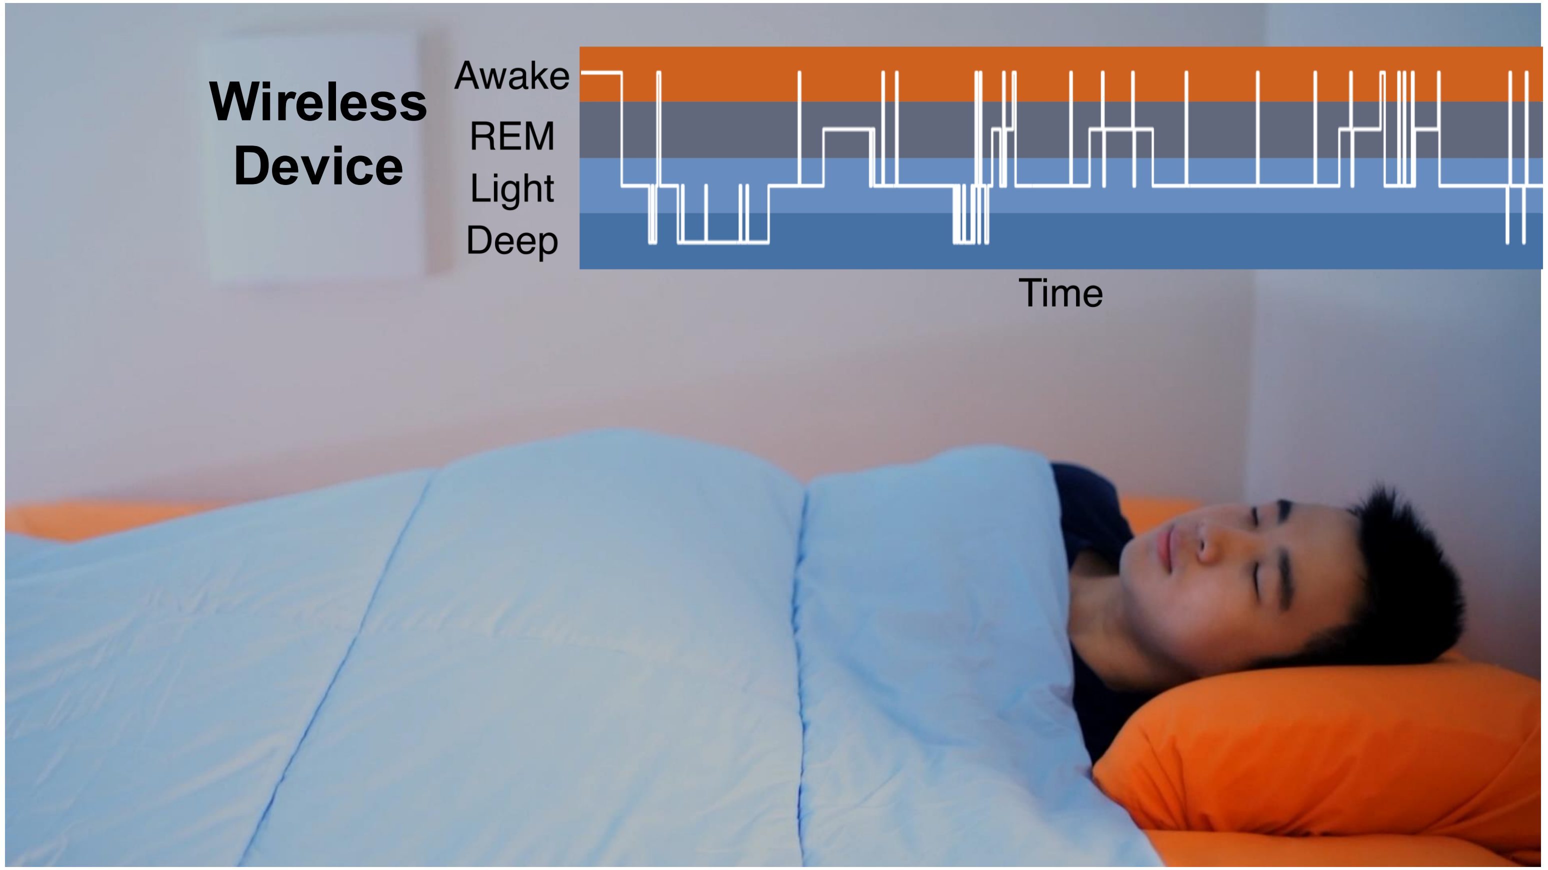

Key Challenge

RF reflections are highly dependent on the **measurement conditions** and the **individuals**.

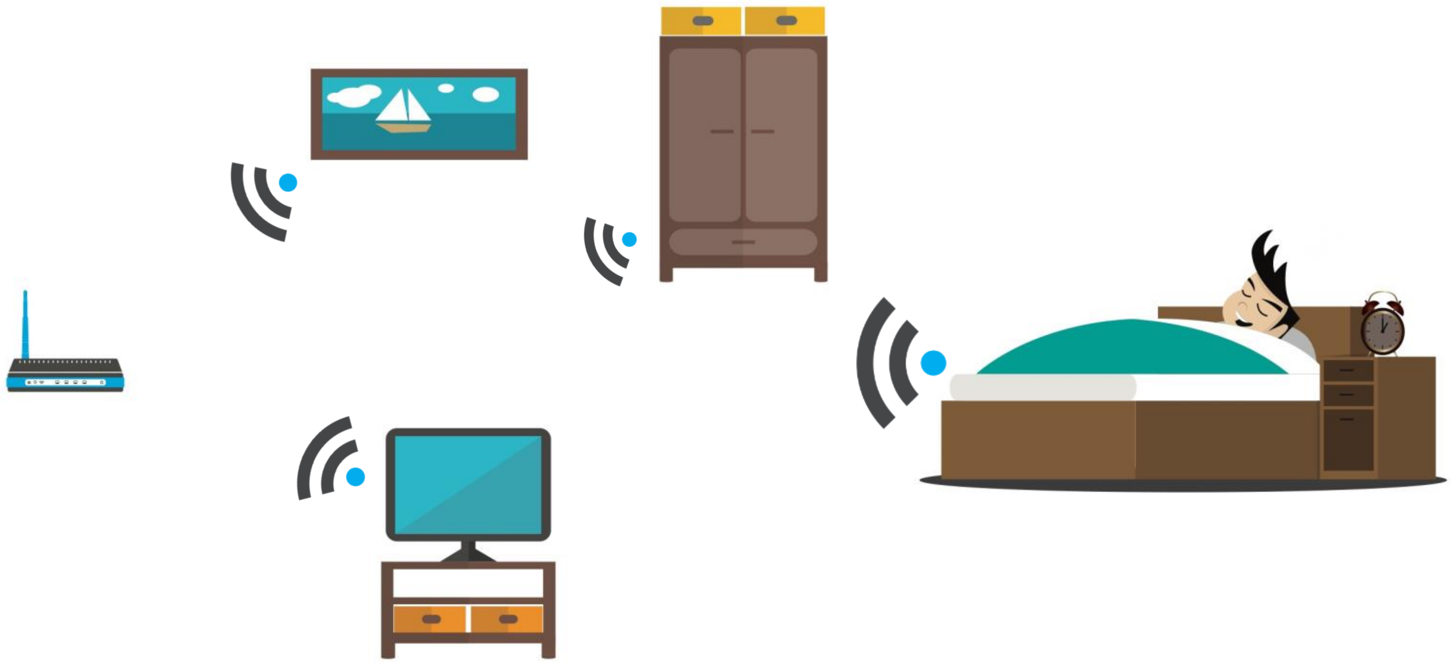

Need to remove such extraneous information!

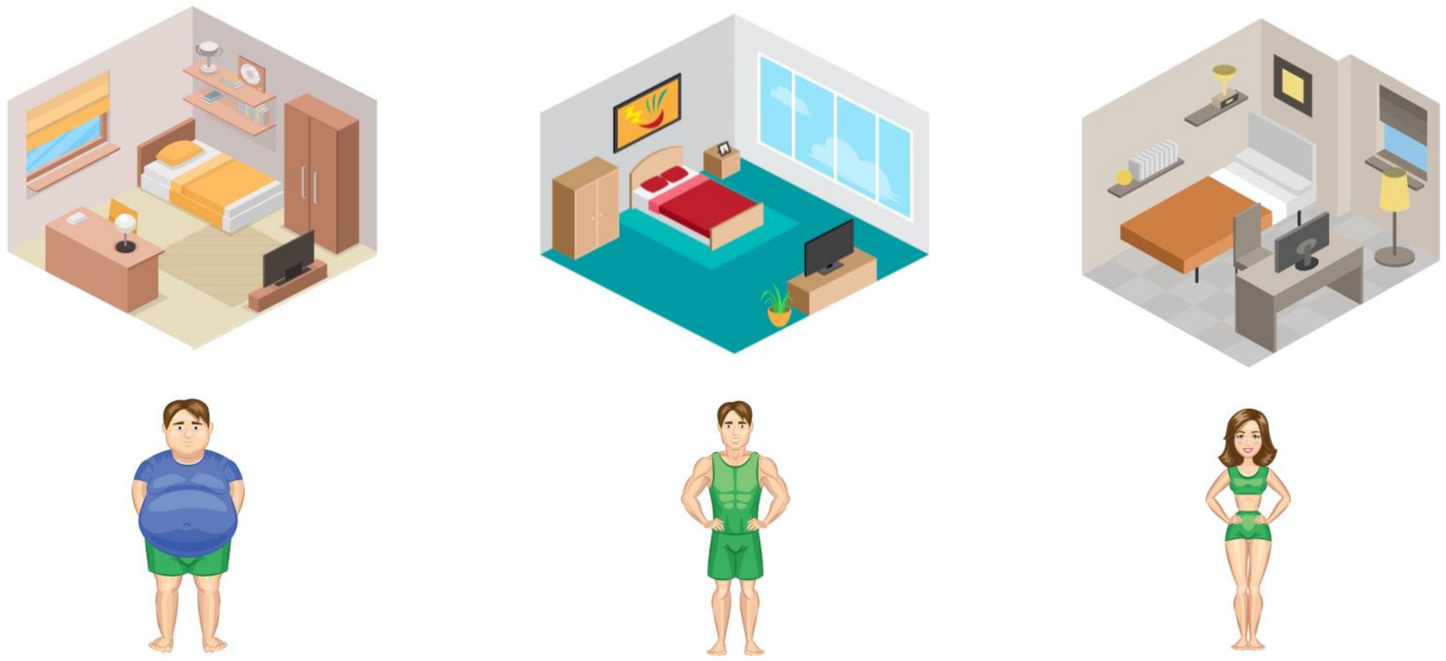

Multi-Source Domain Adaptation

domain = measurement condition + individual

Source domain A   Source domain B   Target domain C

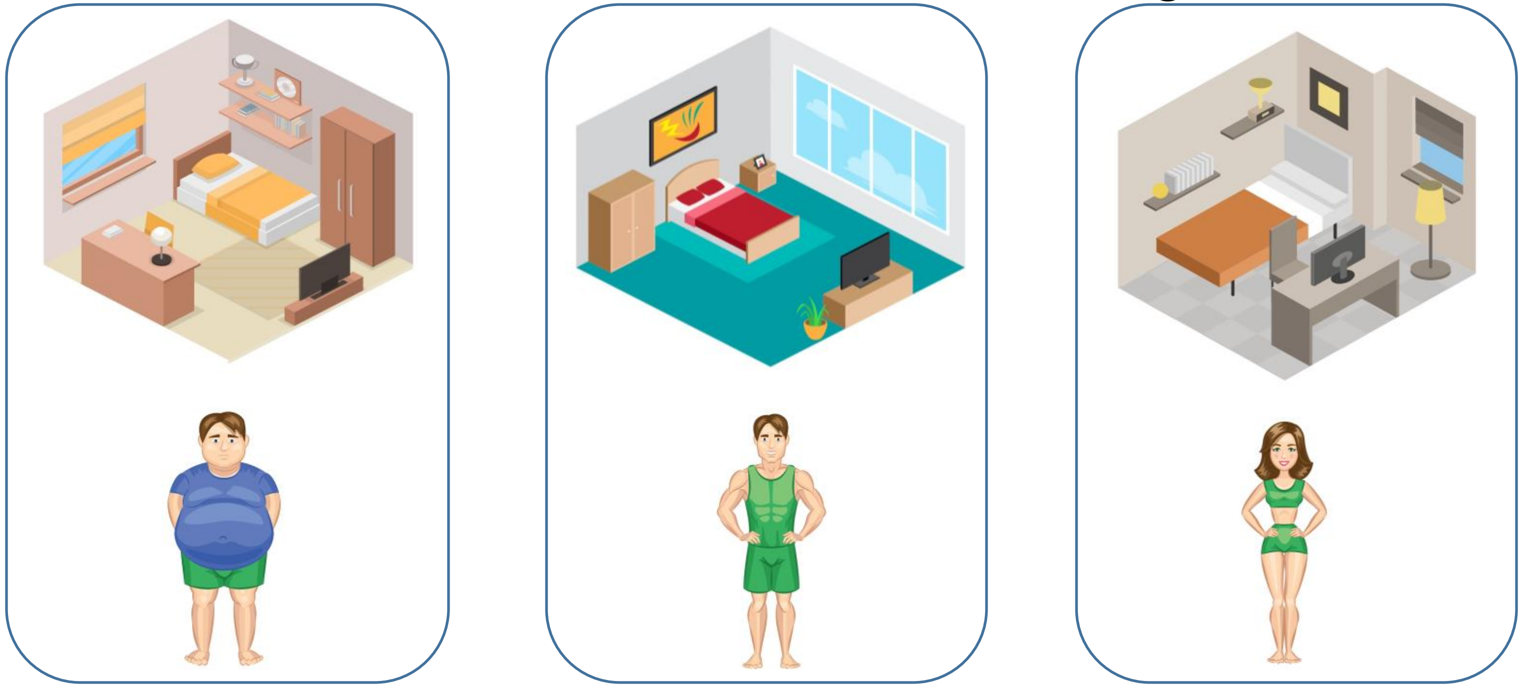

Multi-Source Domain Adaptation

domain = measurement condition + individual

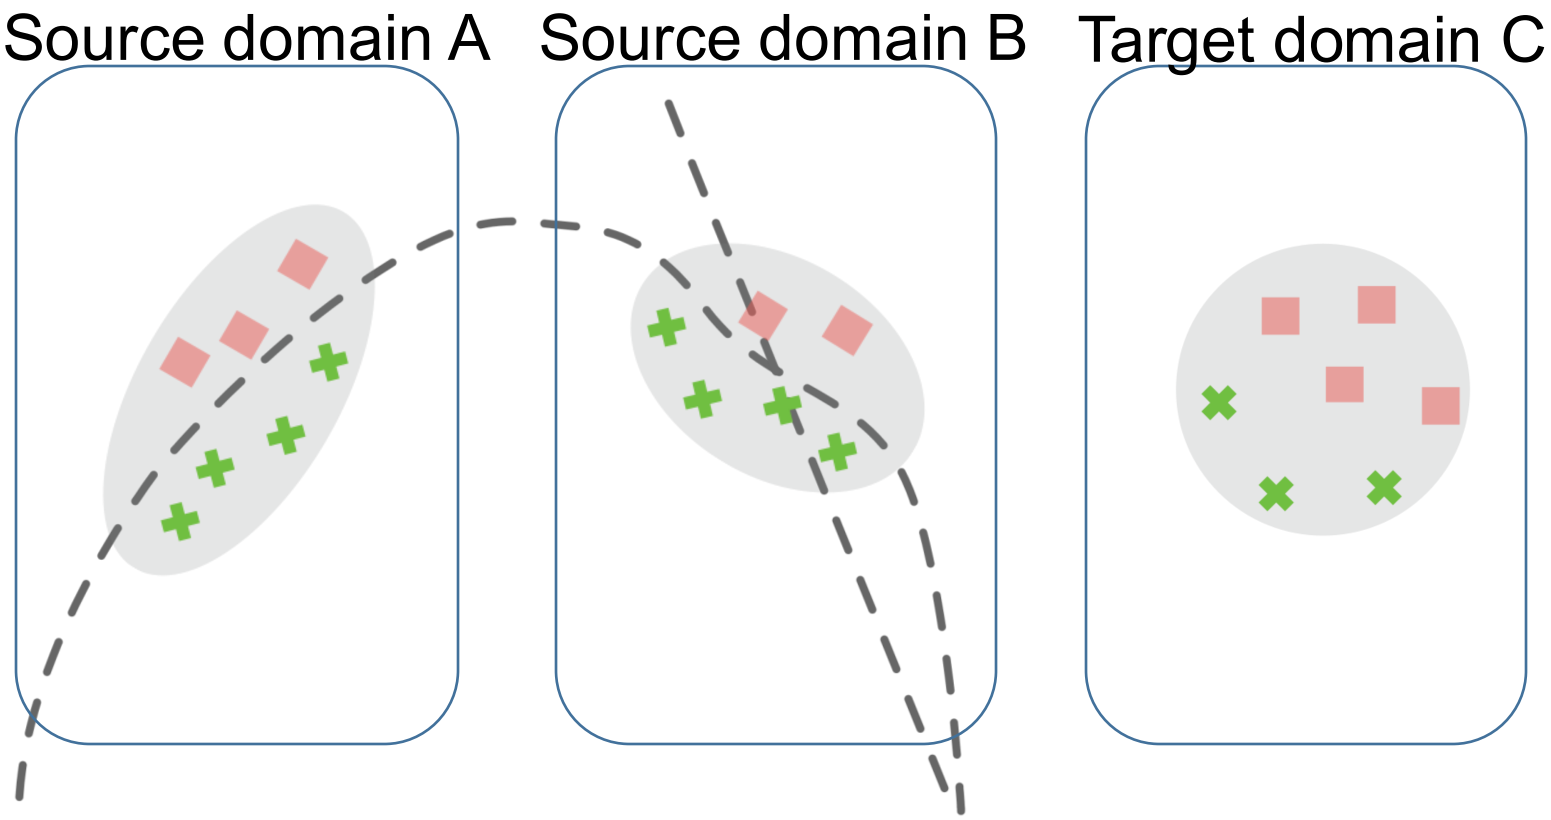

Initial Solution: Adversarial Domain Adaptation

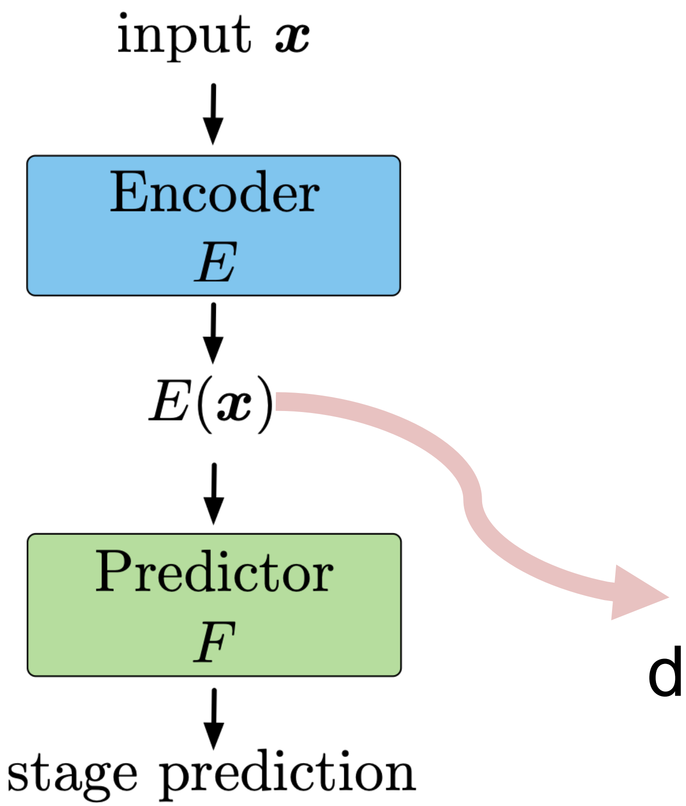

want extraneous domain information to be removed

Problem: Discriminator removes both extraneous and useful information

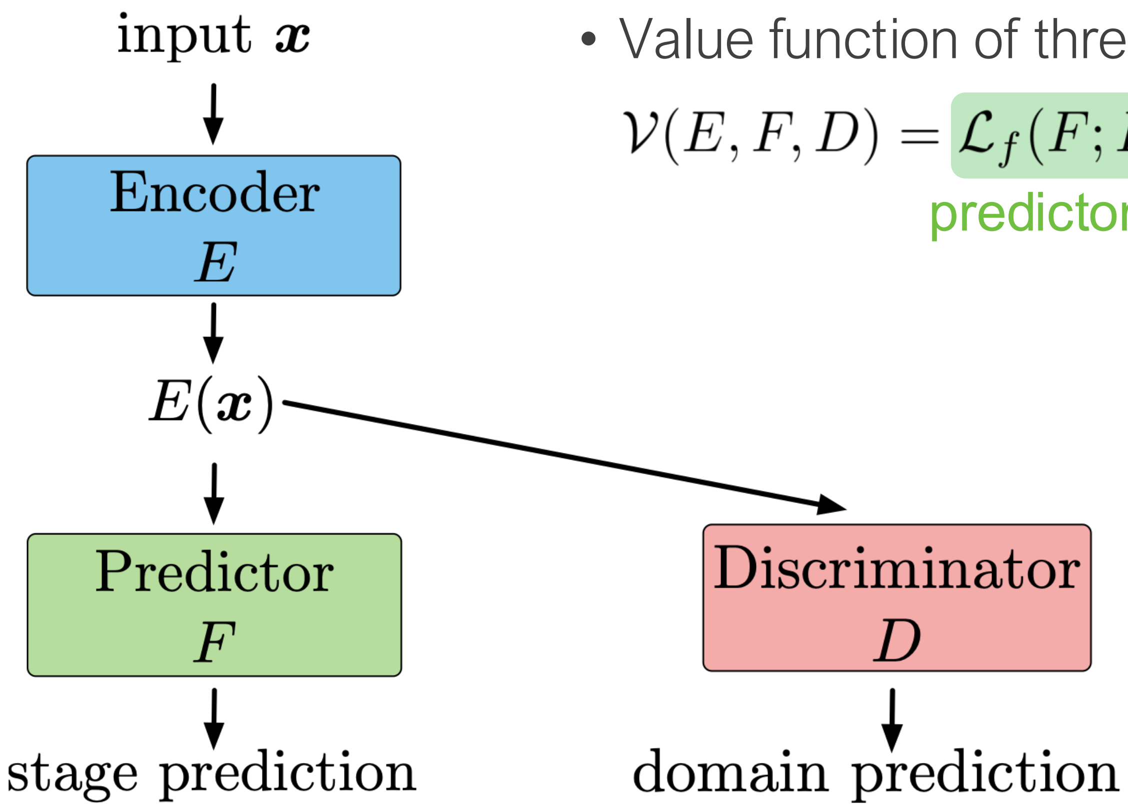

- Value function of three-player game:

$$\mathcal{V}(E, F, D) = \underbrace{\mathcal{L}_f(F; E)}_{\text{predictor loss}} - \lambda \cdot \underbrace{\mathcal{L}_d(D; E)}_{\text{discriminator loss}}$$

Conditional Adversary

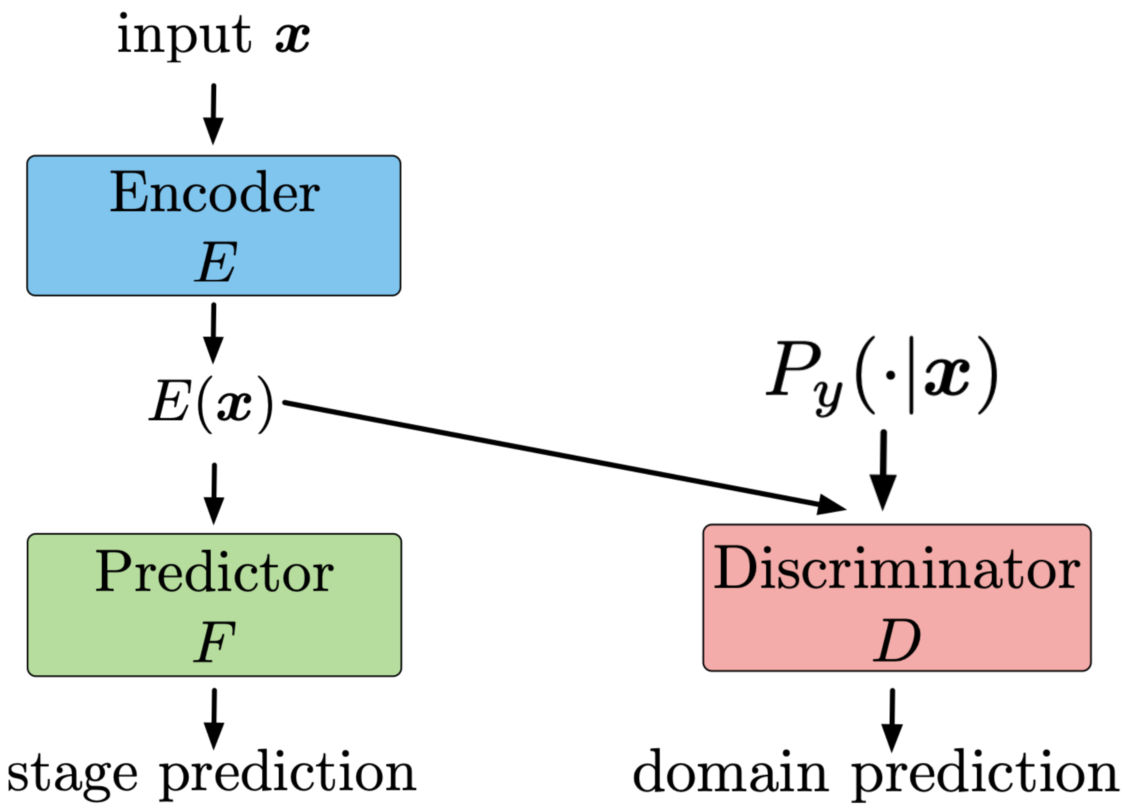

Role of Adversary

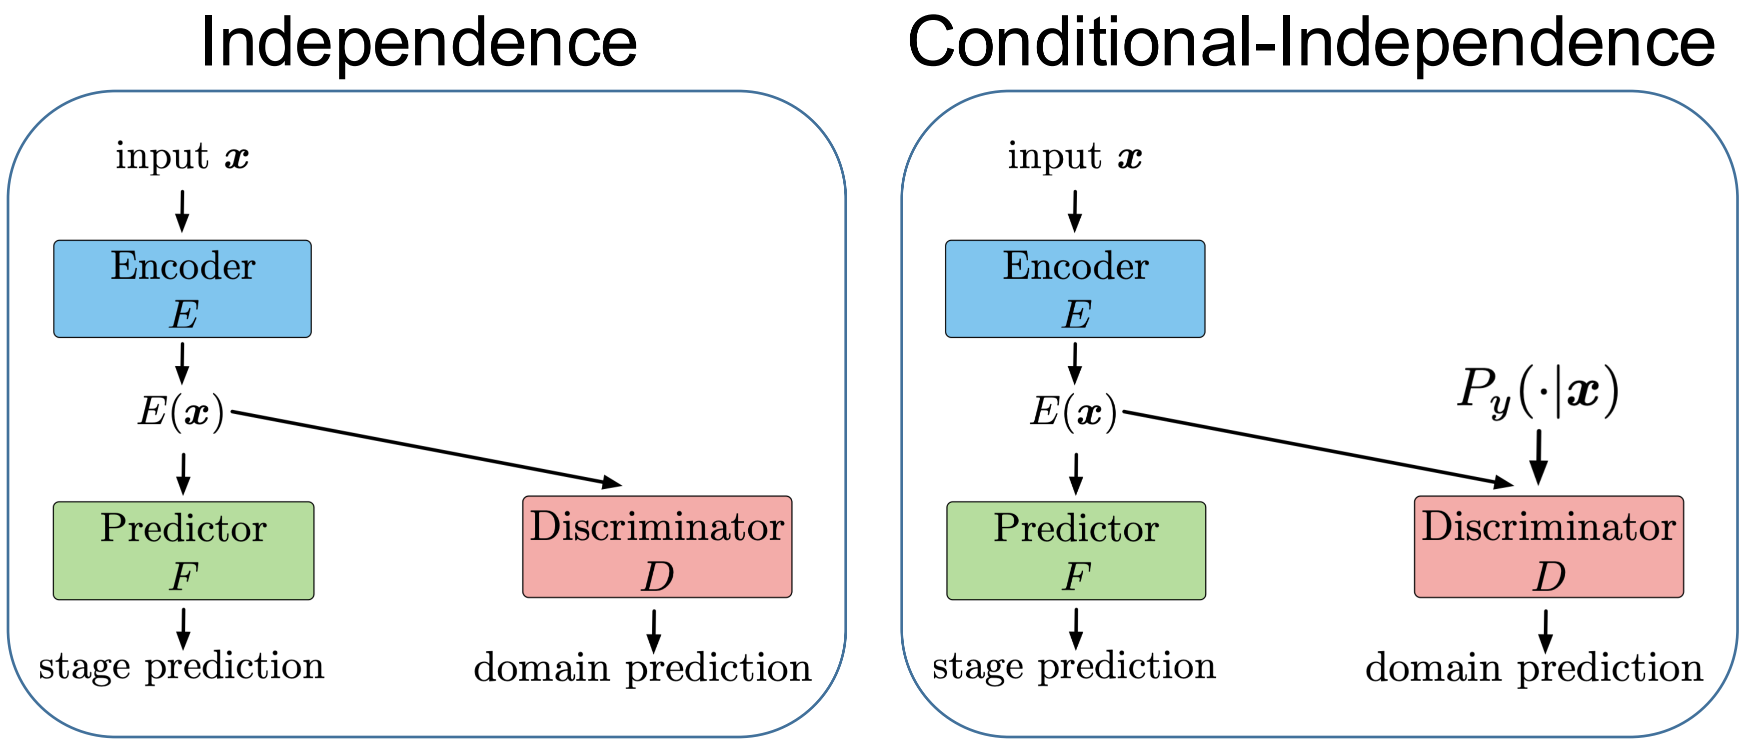

Representative Example

Accuracy = 80%

Ground-truth using EEG

Our Prediction

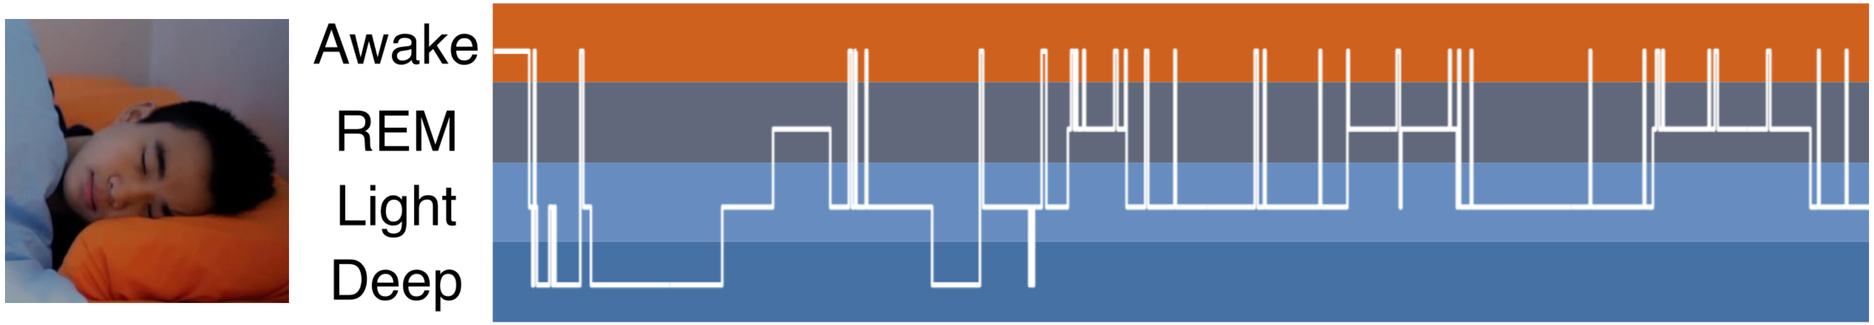

engadget

Reviews Buying Guides Gaming Gear Entertainment Tomorrow Deals News Video Podcasts

Amazon's Halo Rise is a $140 bedside sleep tracker that works by sensing you breathe

No cameras or wearables necessary.

Amazon

Cherlynn Low | @cherlynnlow | September 28, 2022 12:26 PM

Biotech & Health

Google's Soli radar returns to track sleep on the new Nest Hub

Brian Heater @bheater / 2:00 PM GMT+1 • March 16, 2021

Comment

Image Credits: Google

Talk about surprise comebacks. This morning Google announced the arrival of the next-gen Nest Hub. In spite of rebranding from Google Home Hub back in 2019, the smart screen hasn't seen many changes since its 2018 introduction. Today's arrival doesn't represent a huge upgrade from its predecessor, but it does support a familiar — and largely forgotten — face.

We haven't heard a peep from Project Soli since the technology was introduced with the Pixel in late-2019. The miniature, motion-sensing radar tech was positioned to be a major selling point, finally arriving on a device some four years after being announced. Applications were relatively few and far between — including gesture detection and a weird, one-off Pokémon app.

Sensing Humans in the Environment

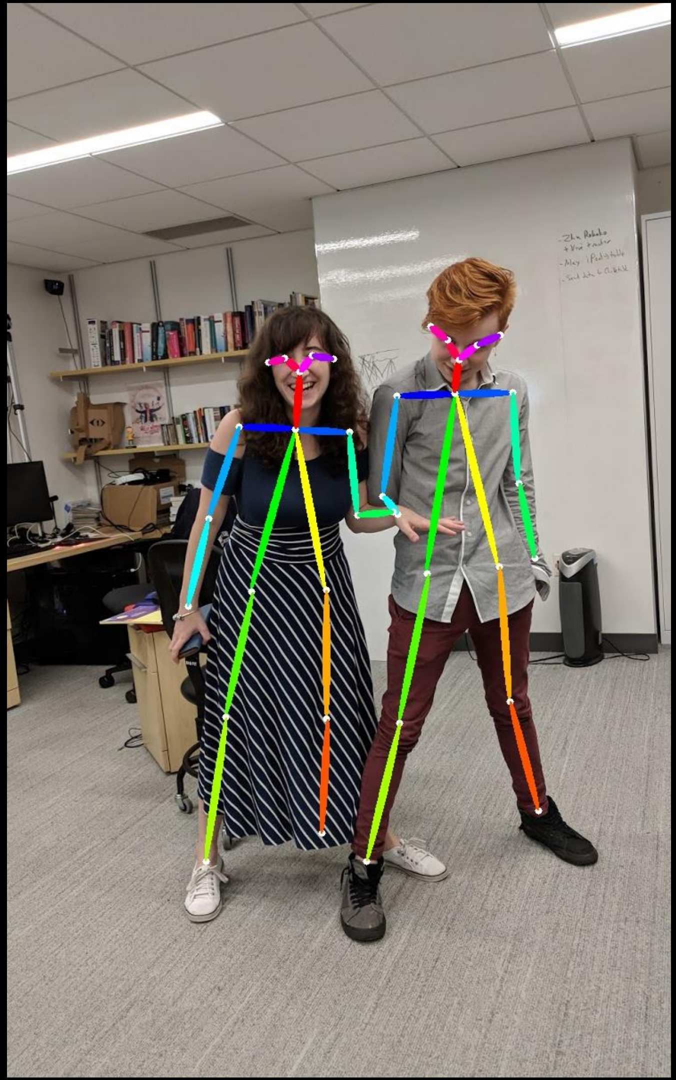

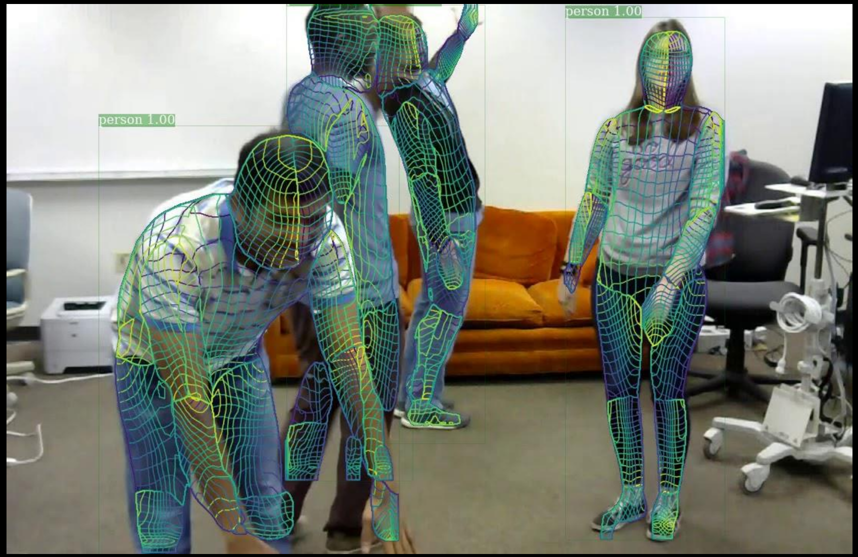

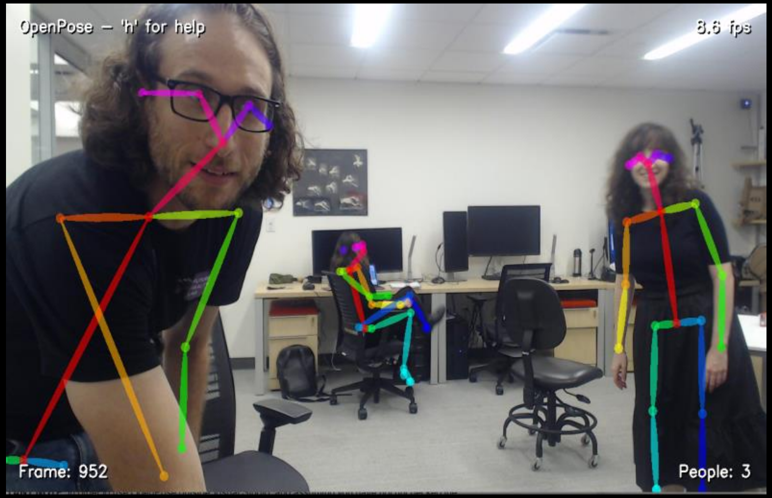

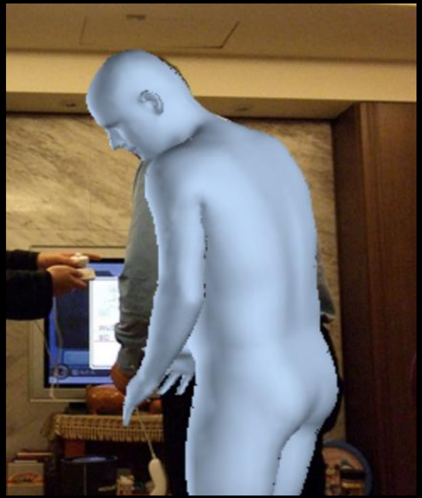

Occlusion is a fundamental challenge for vision

Vision also fails in bad lighting conditions

Want to see the human through walls & in the dark

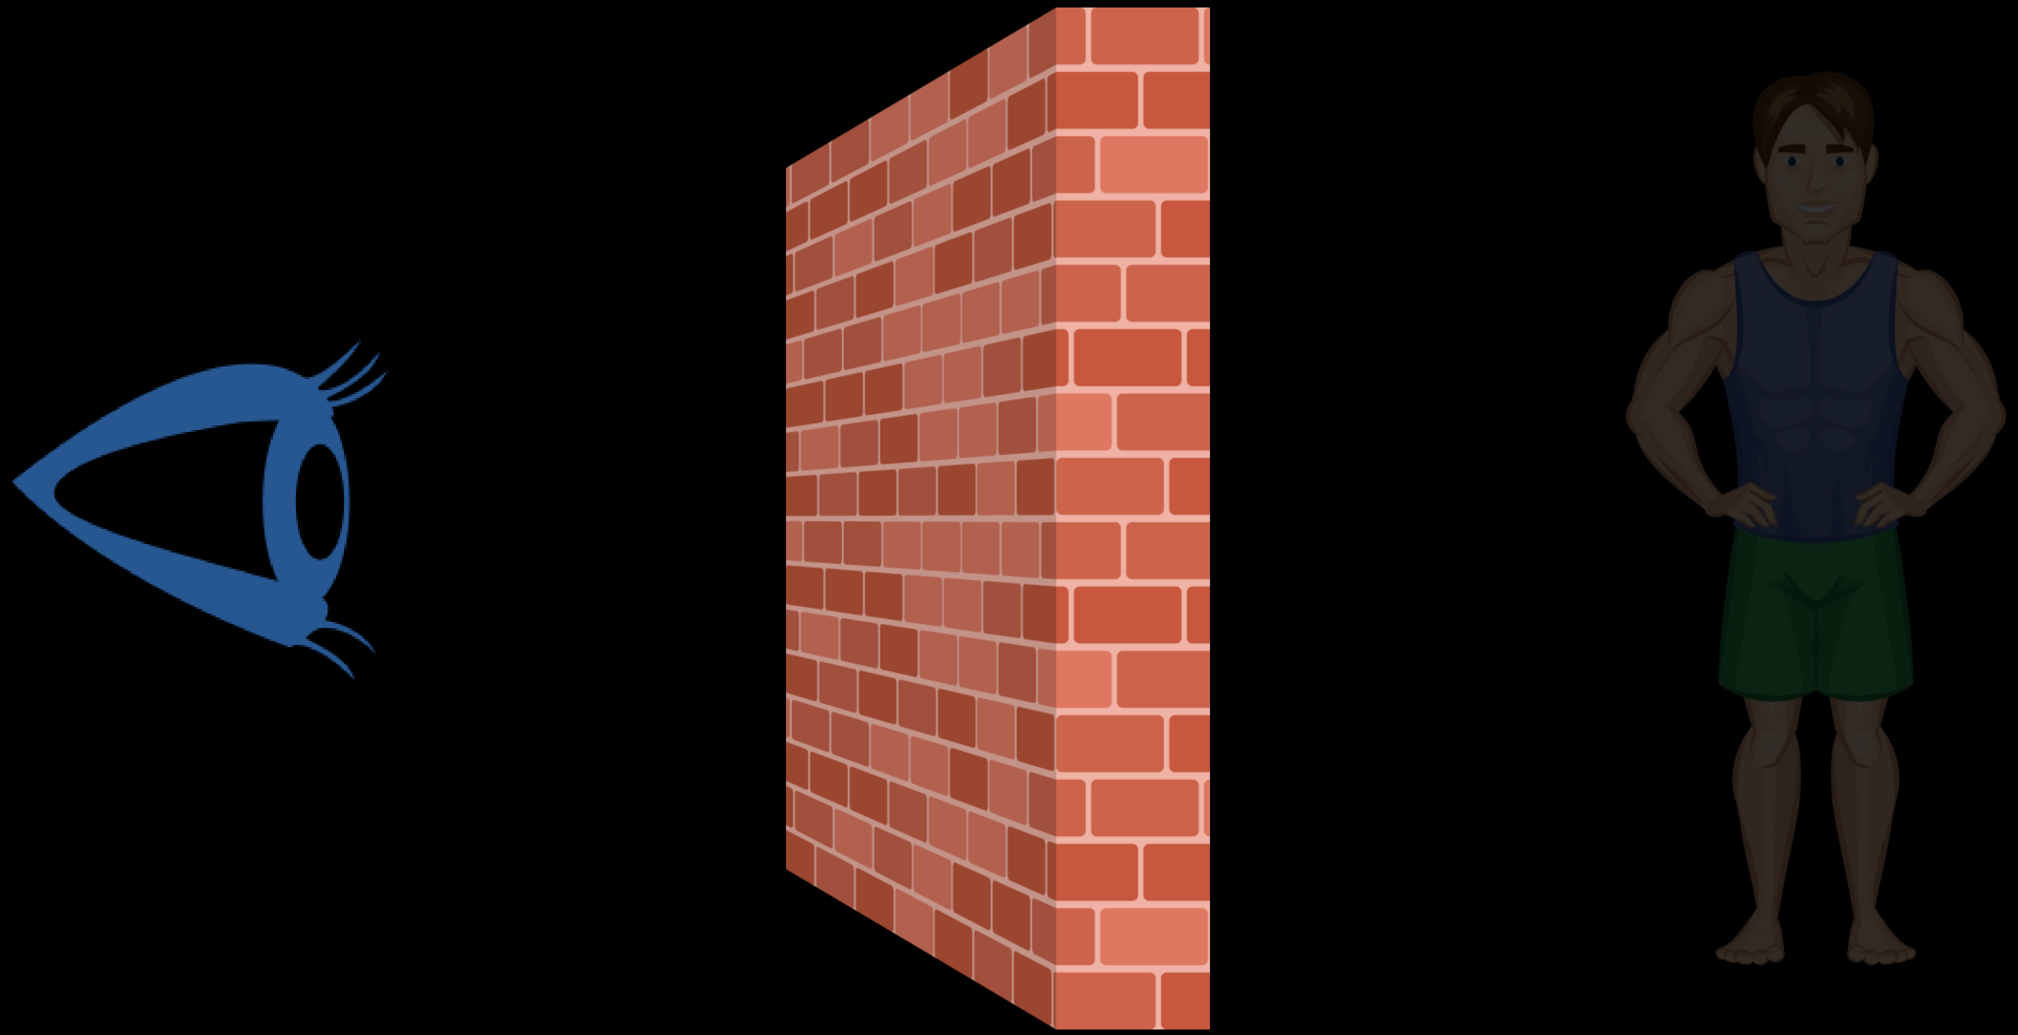

Want to see the human through walls & in the dark

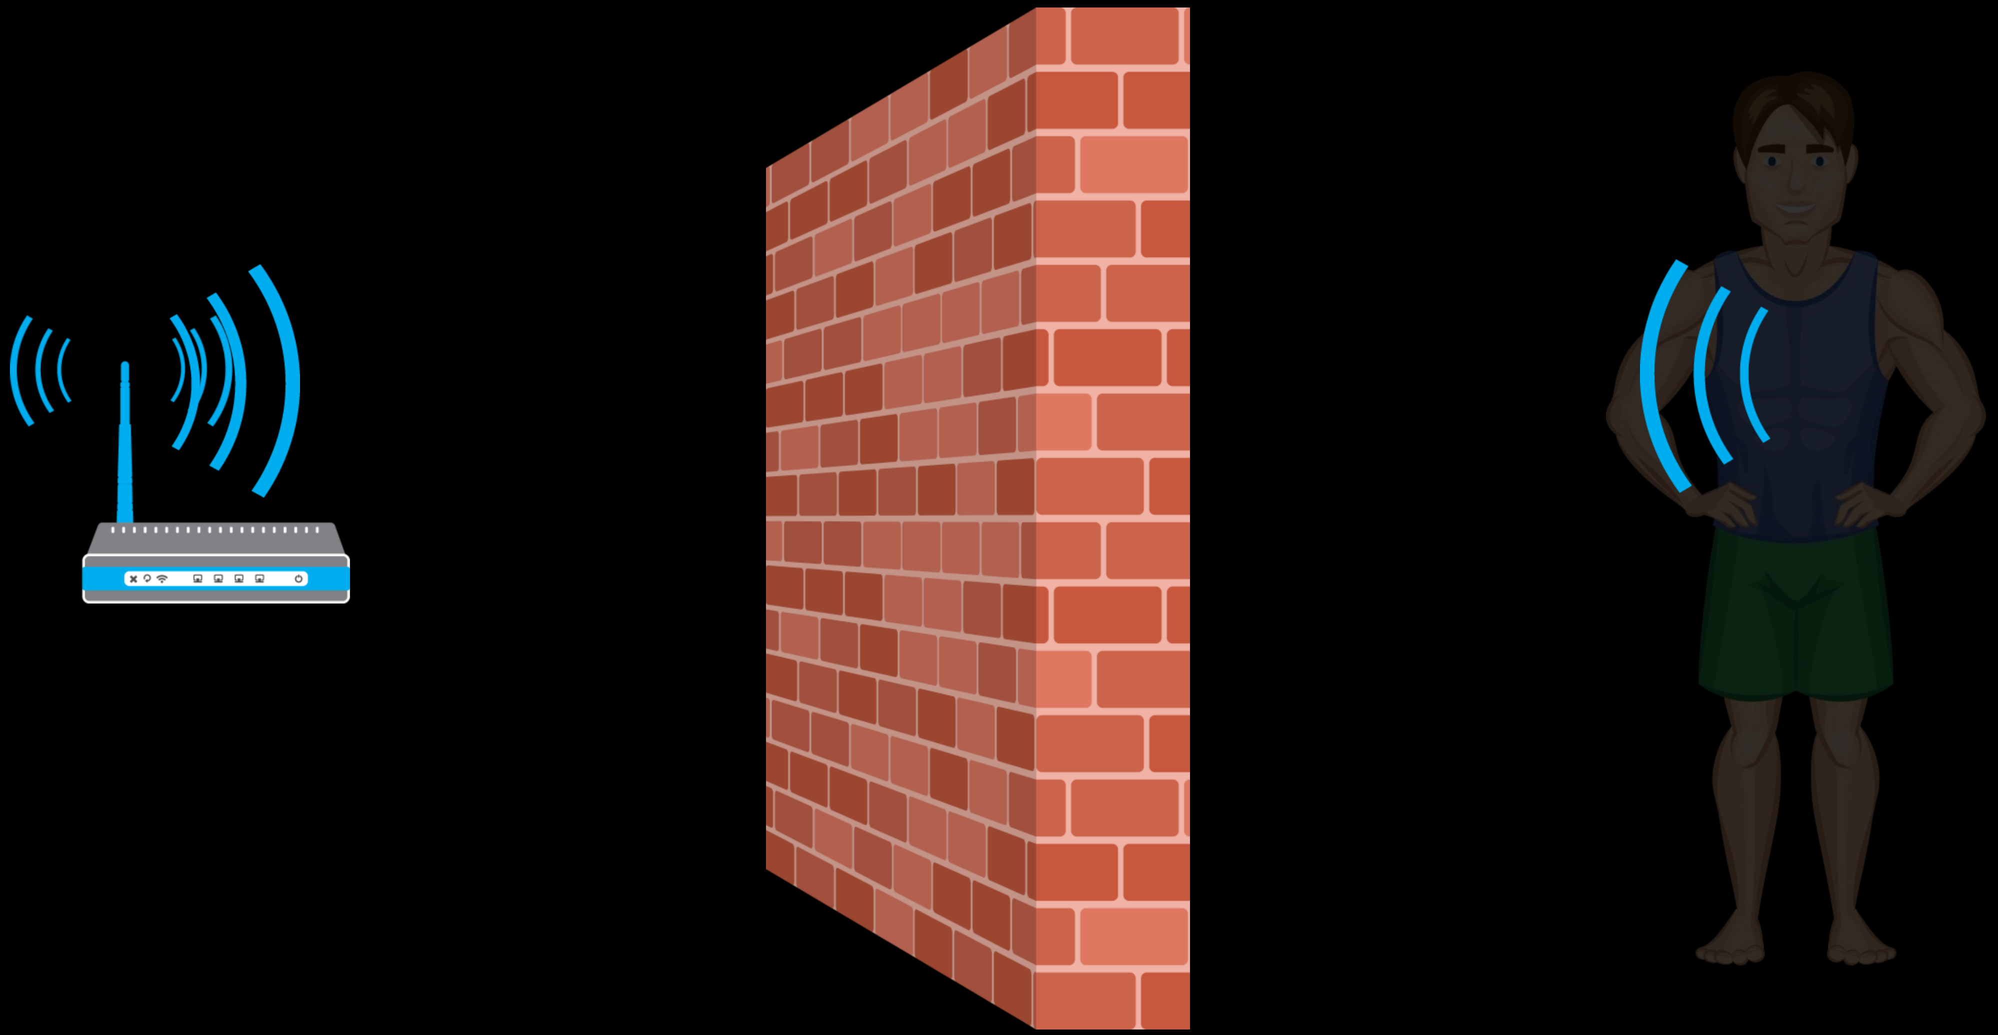

RF-based Approach

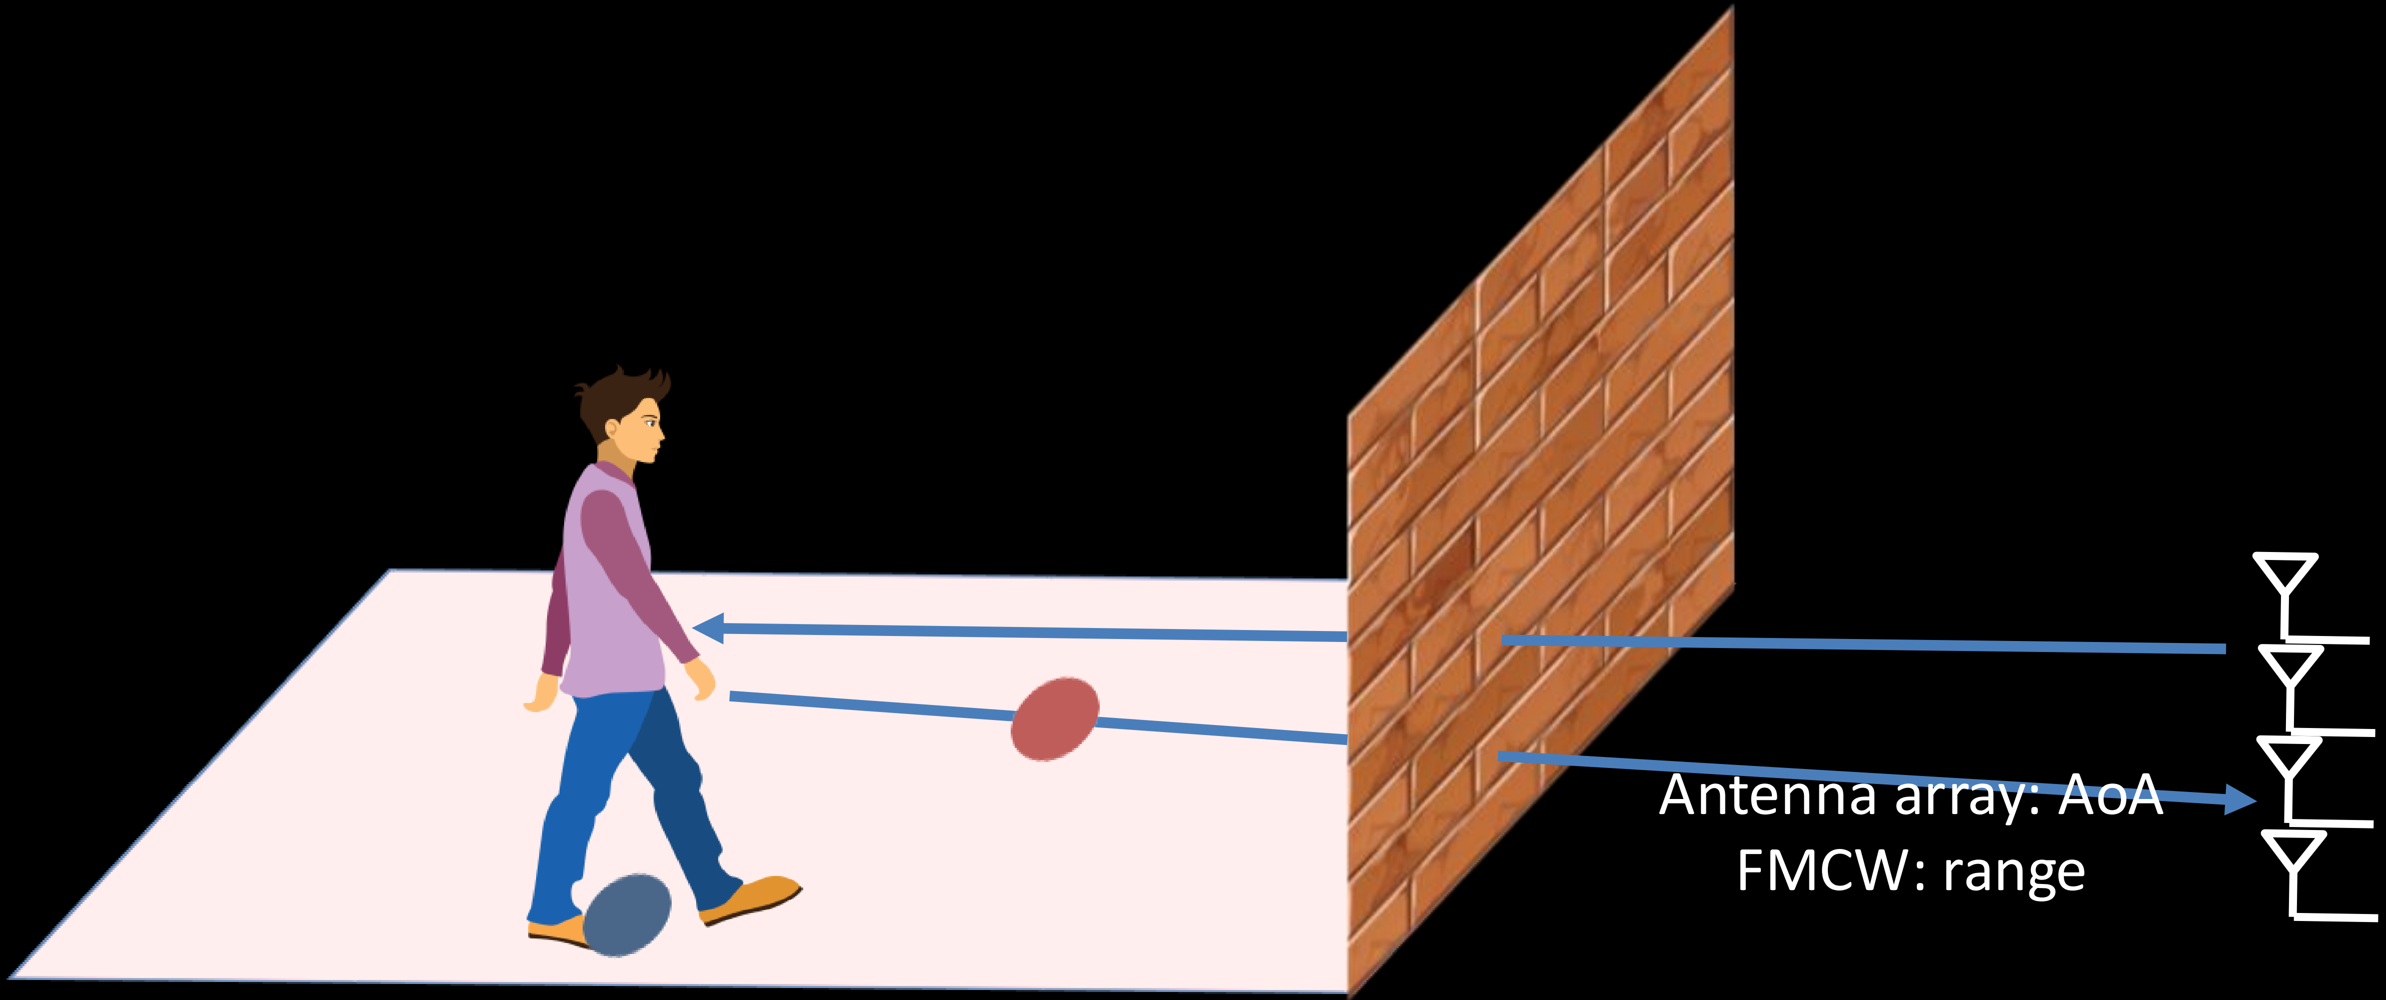

How to train a model to estimate pose from RF?

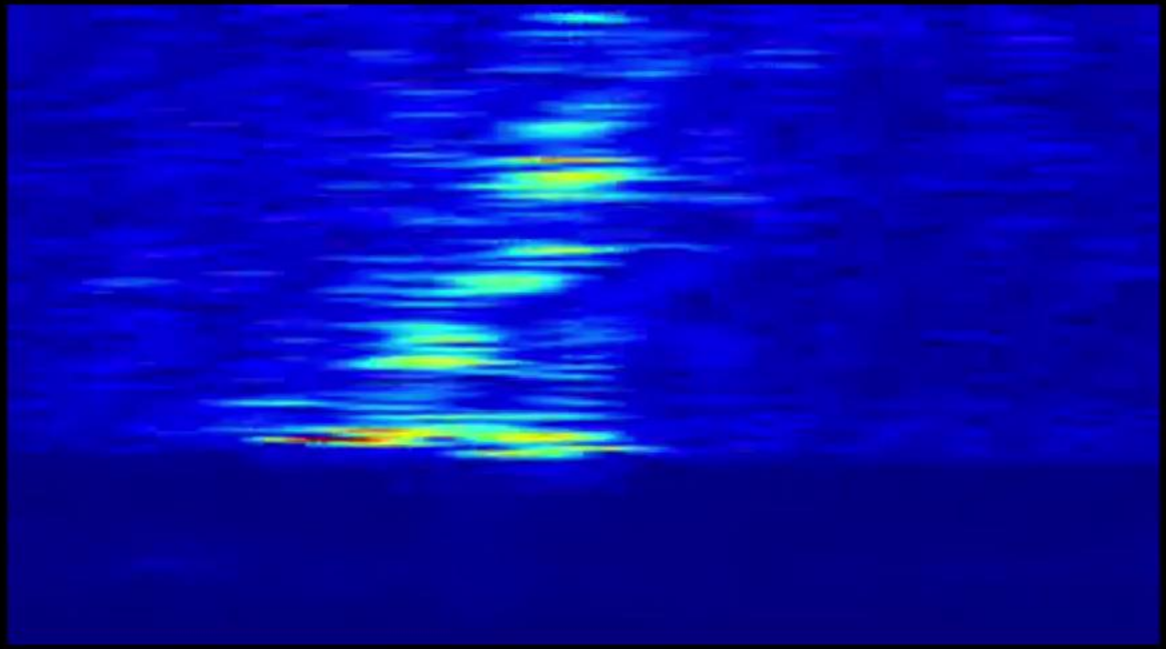

Vertical RF heatmaps

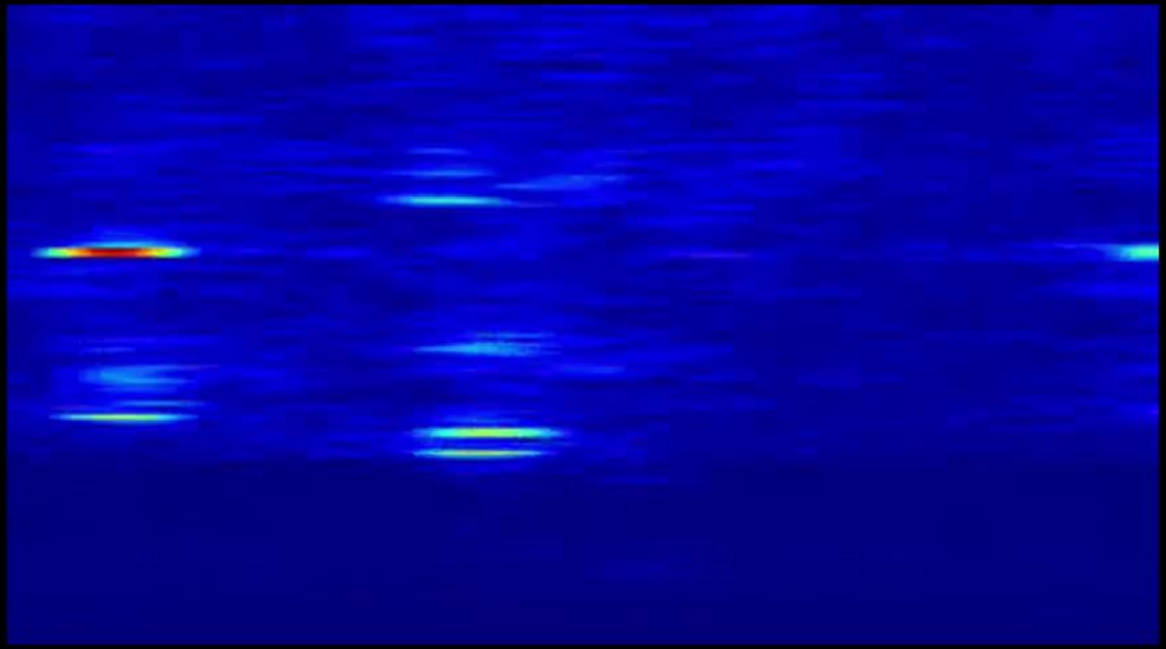

Horizontal RF heatmaps

Challenge: Specularity of Human Body

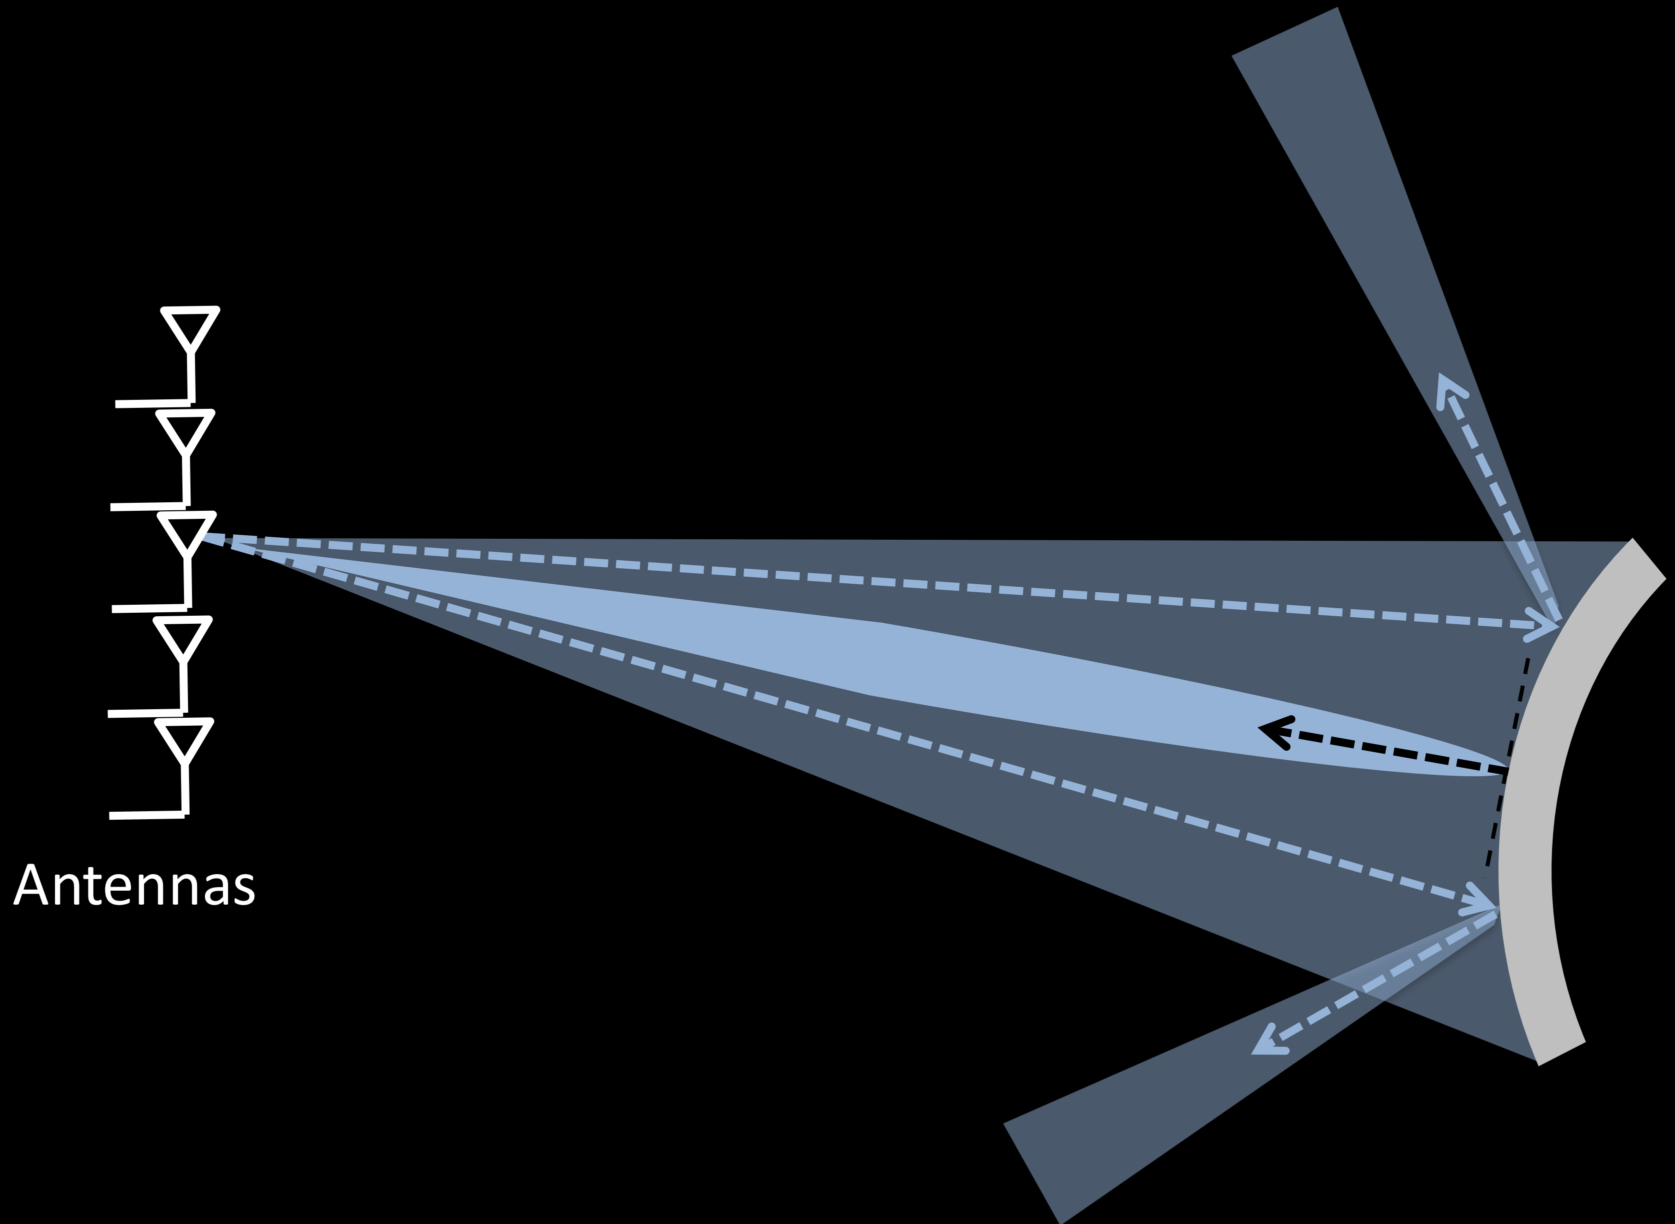

A Snapshot Doesn’t Have Skeleton

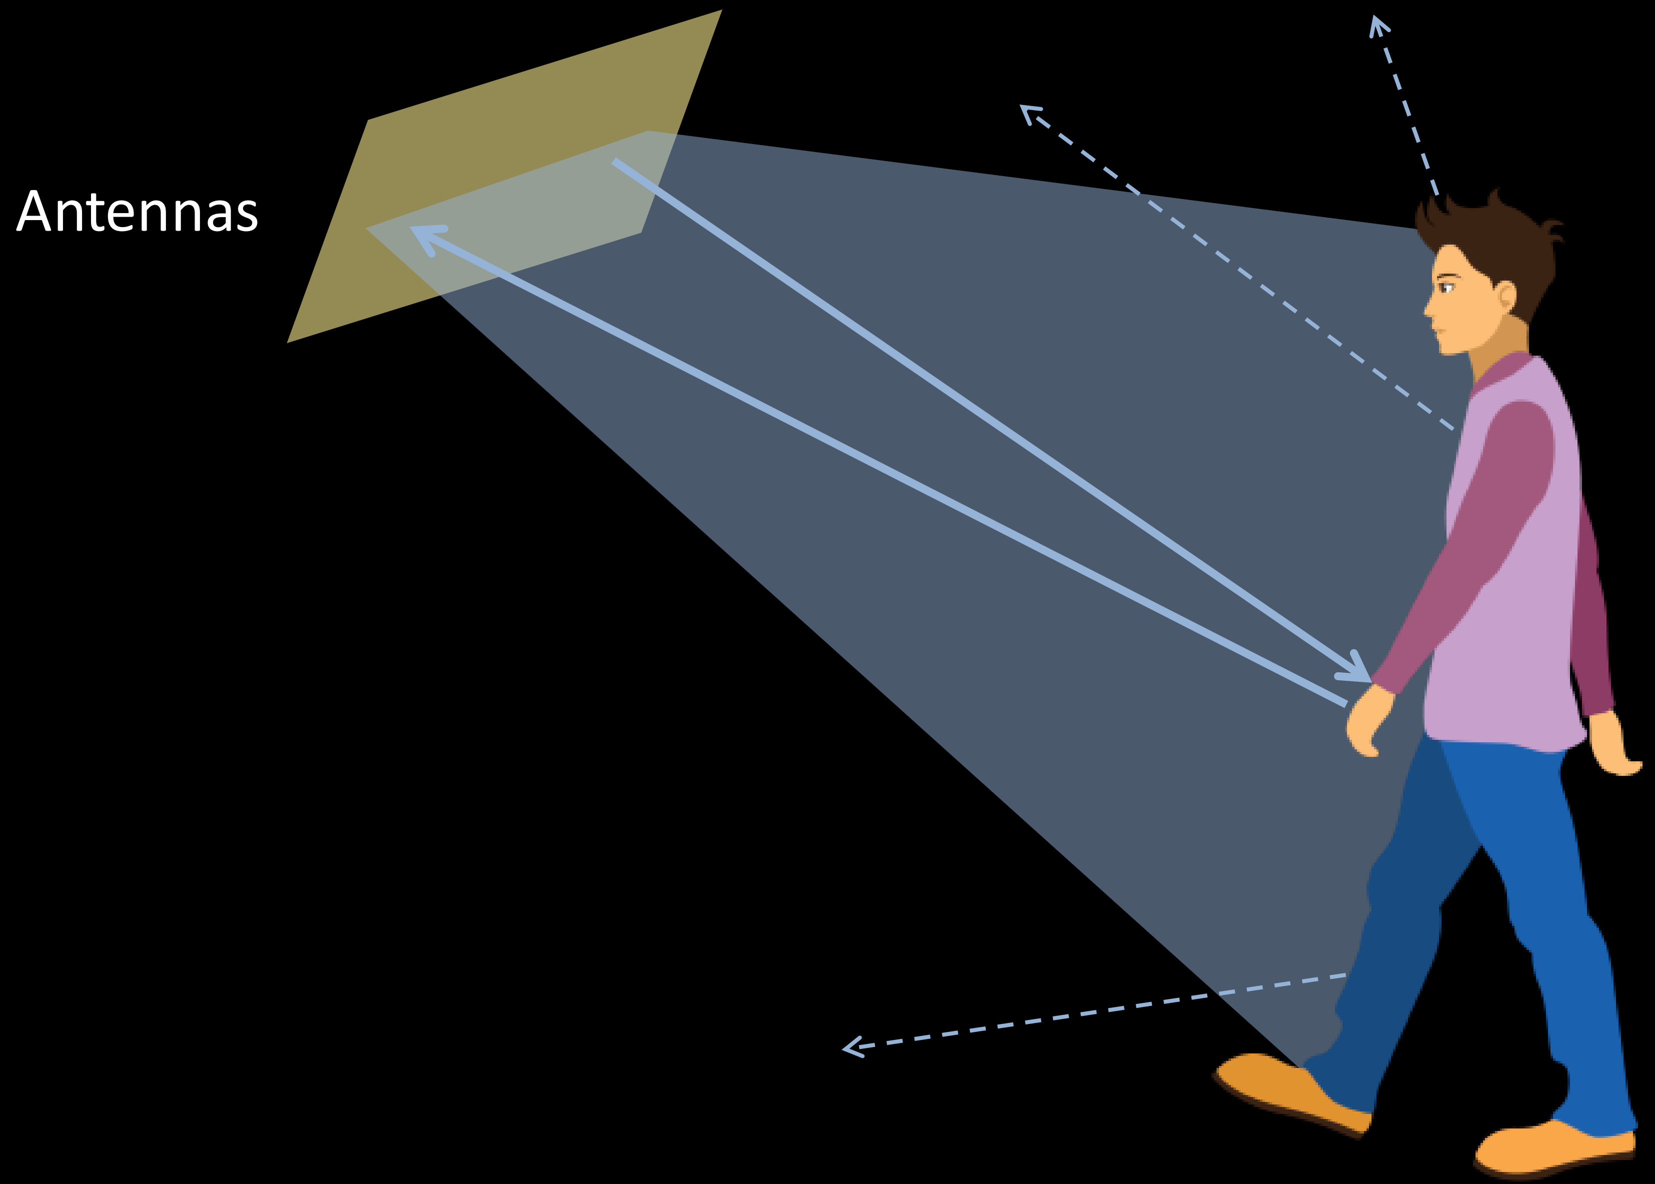

Solution: Use Human Motion Across Time

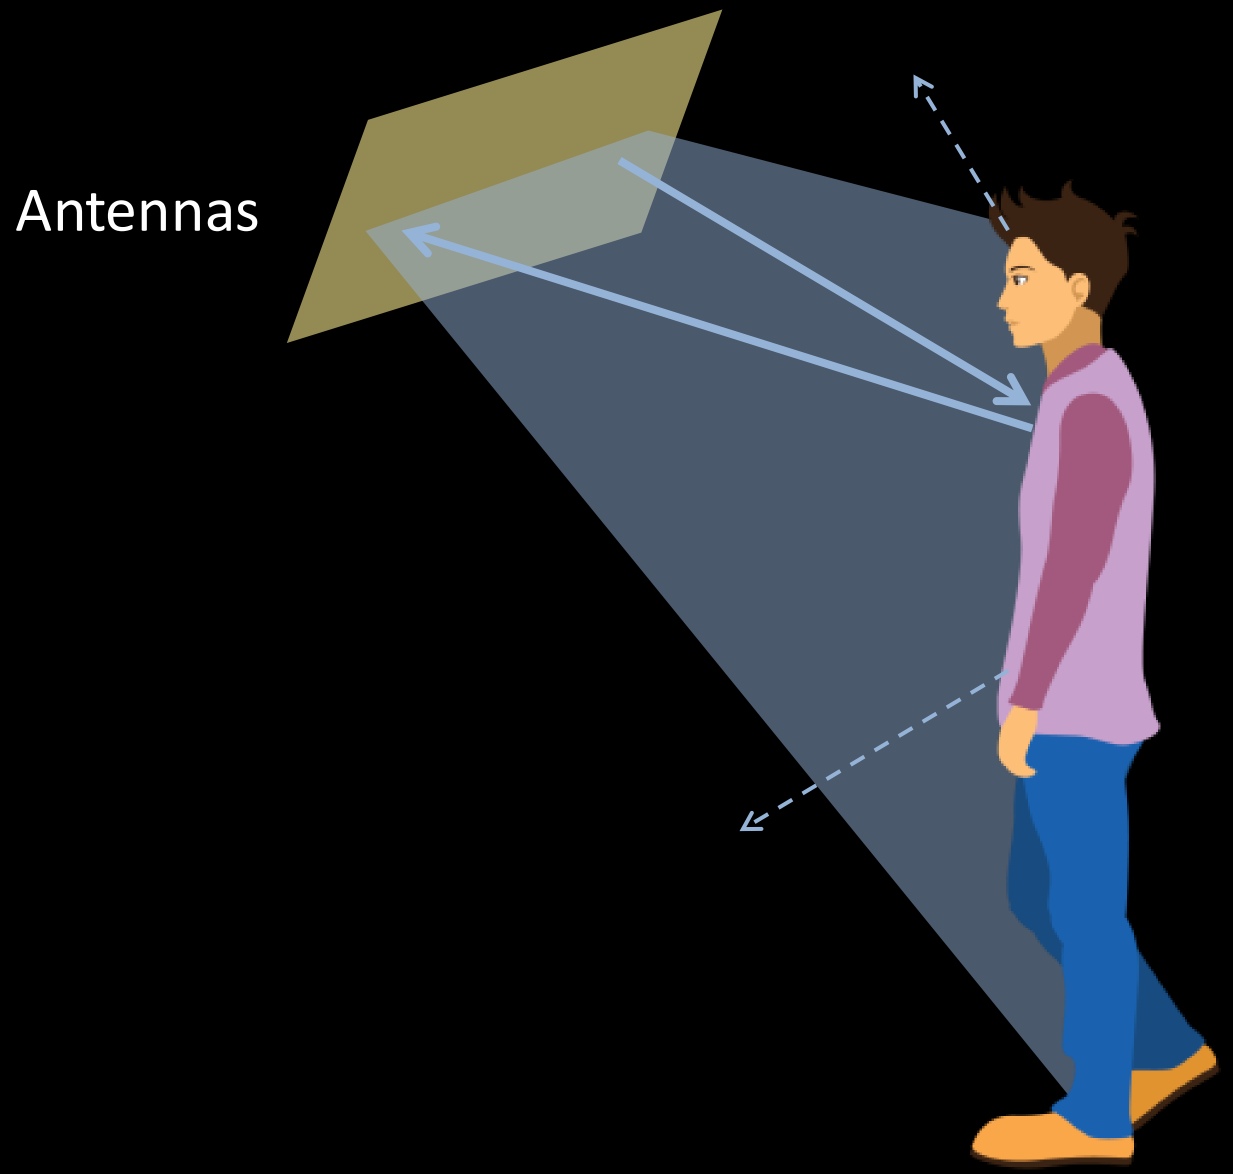

Solution: Use a series of RF snapshots

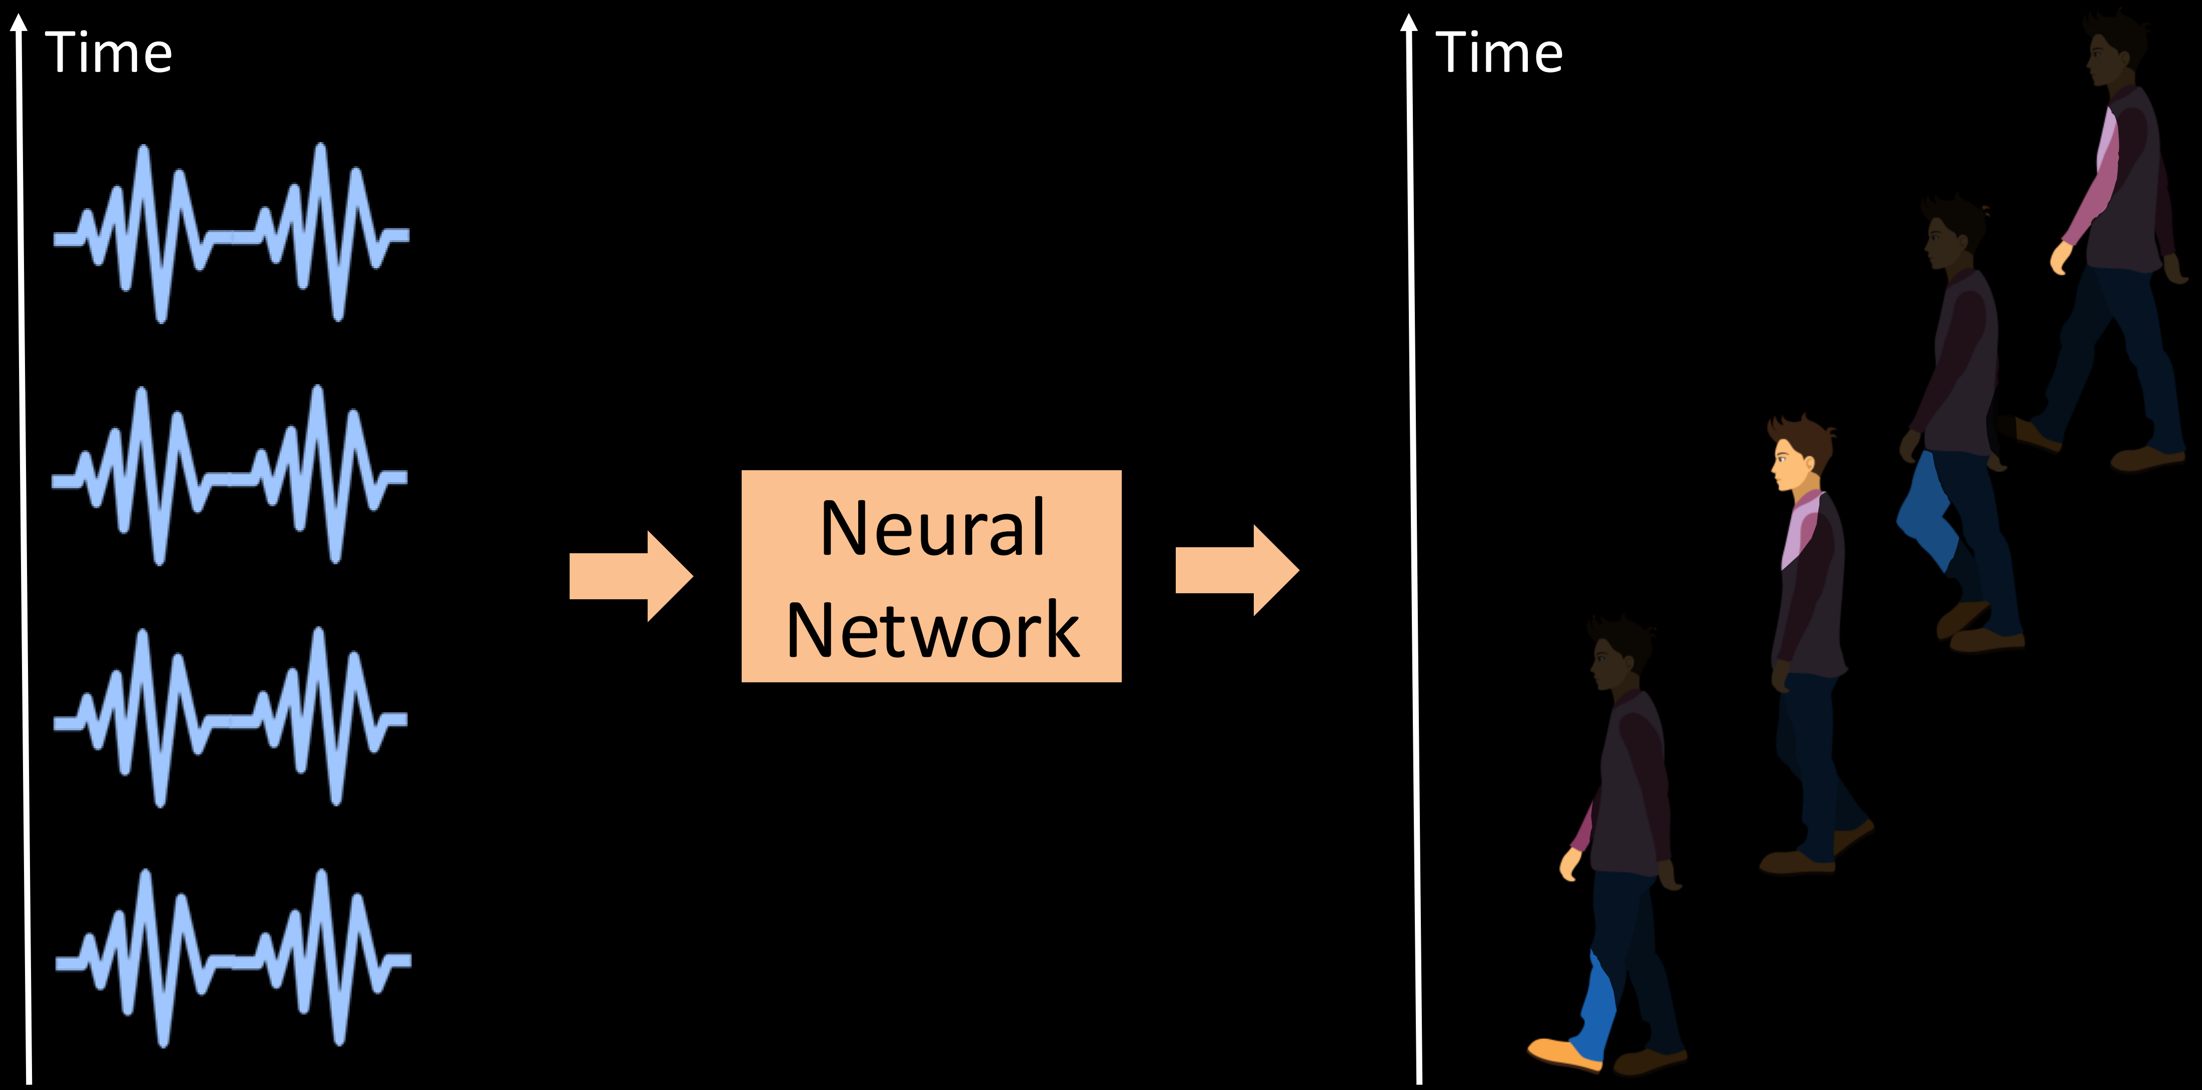

Solution: Extracts limbs and fills in missing parts

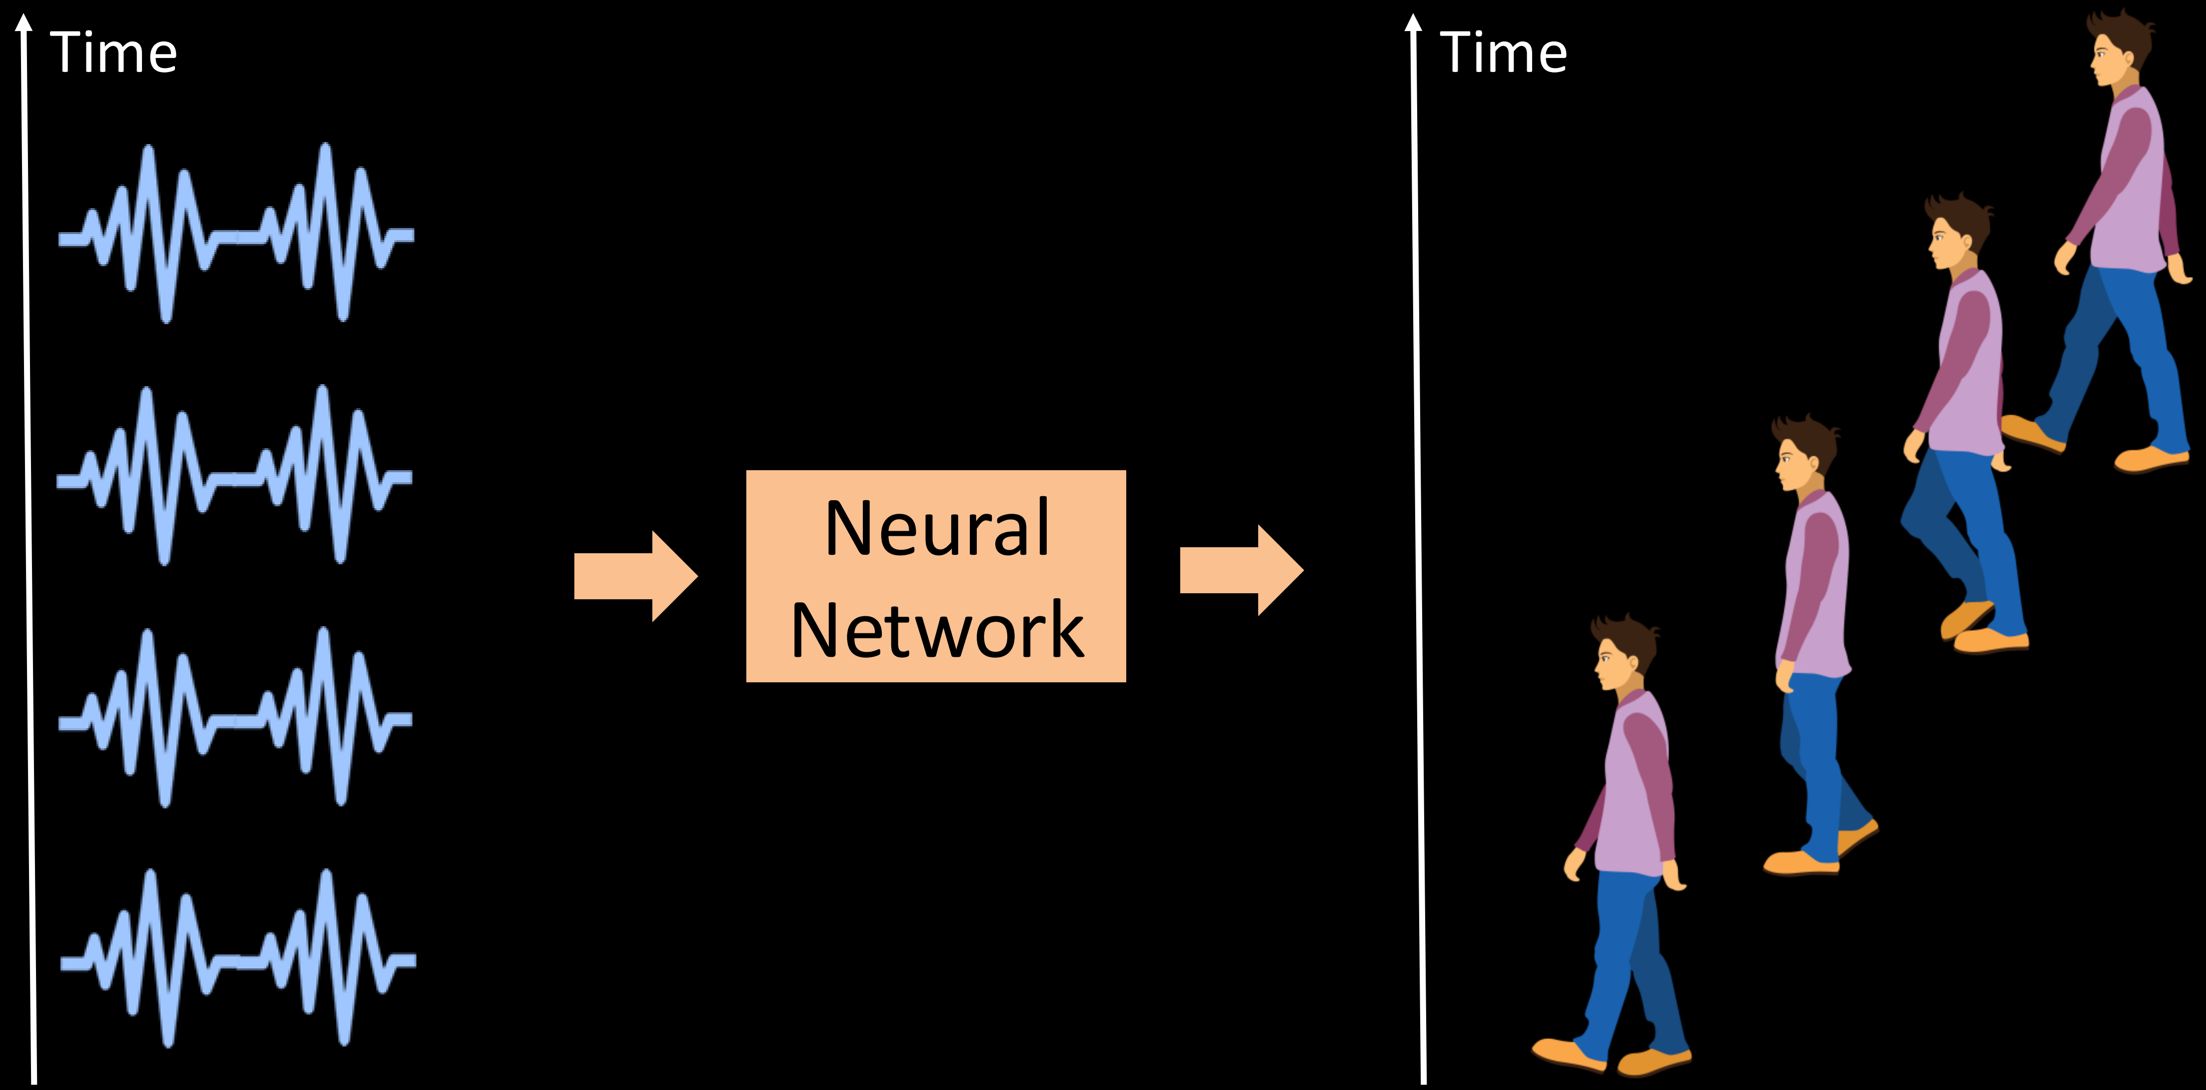

Challenge: 4D signals are too large for NN!

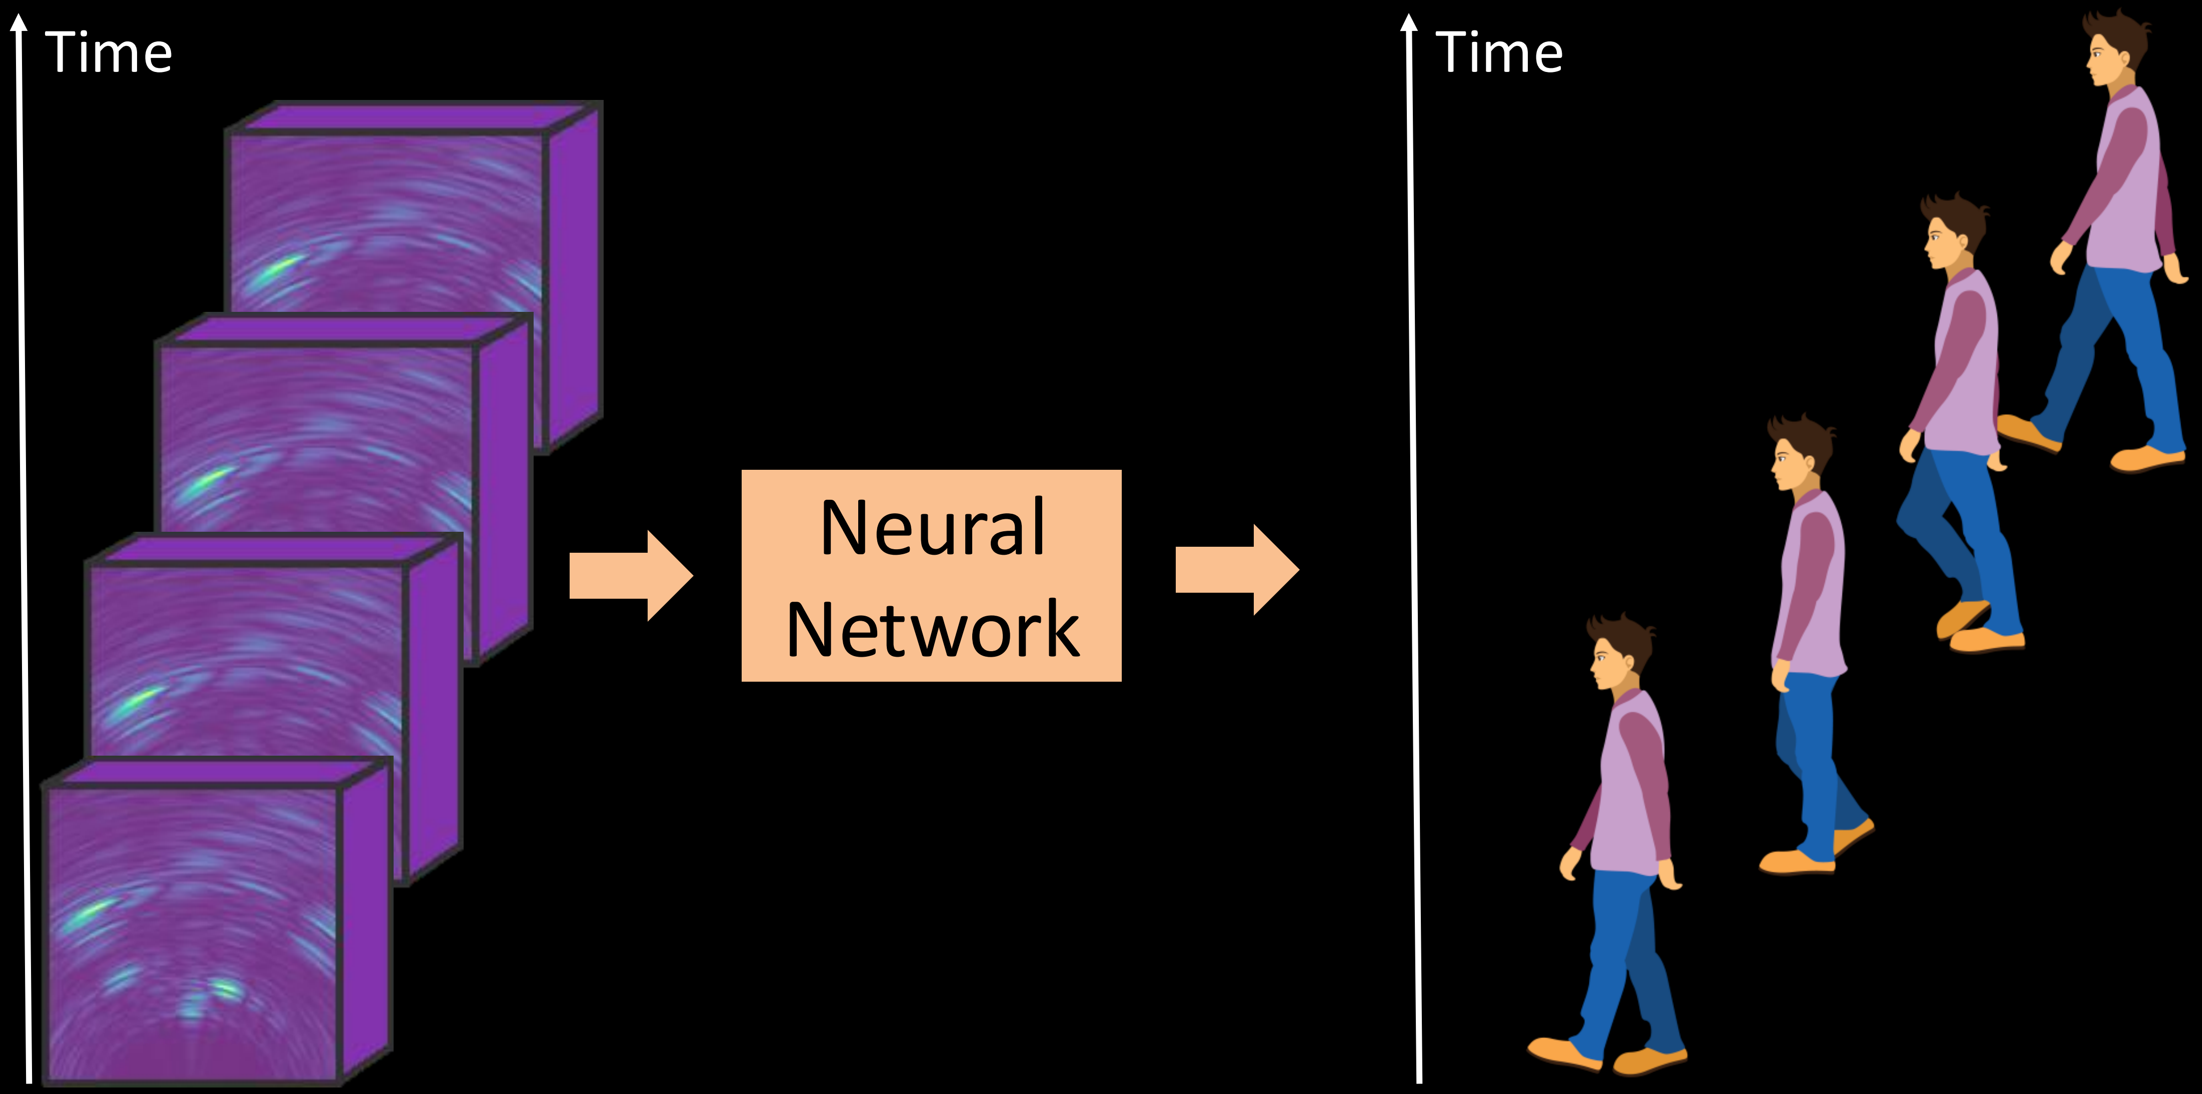

Challenge: How to obtain labeled data?

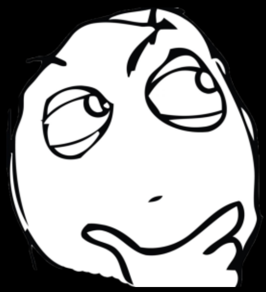

annotate skeleton?

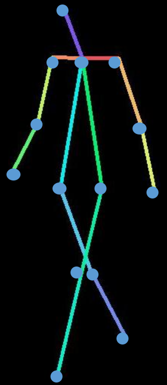

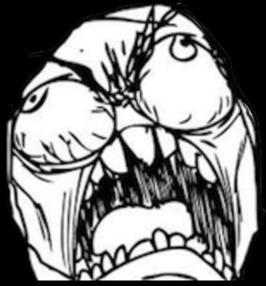

annotate skeleton?

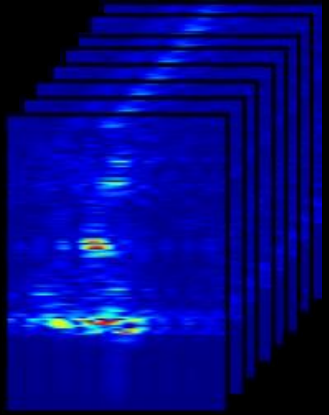

Idea: Cross-Modal Supervision

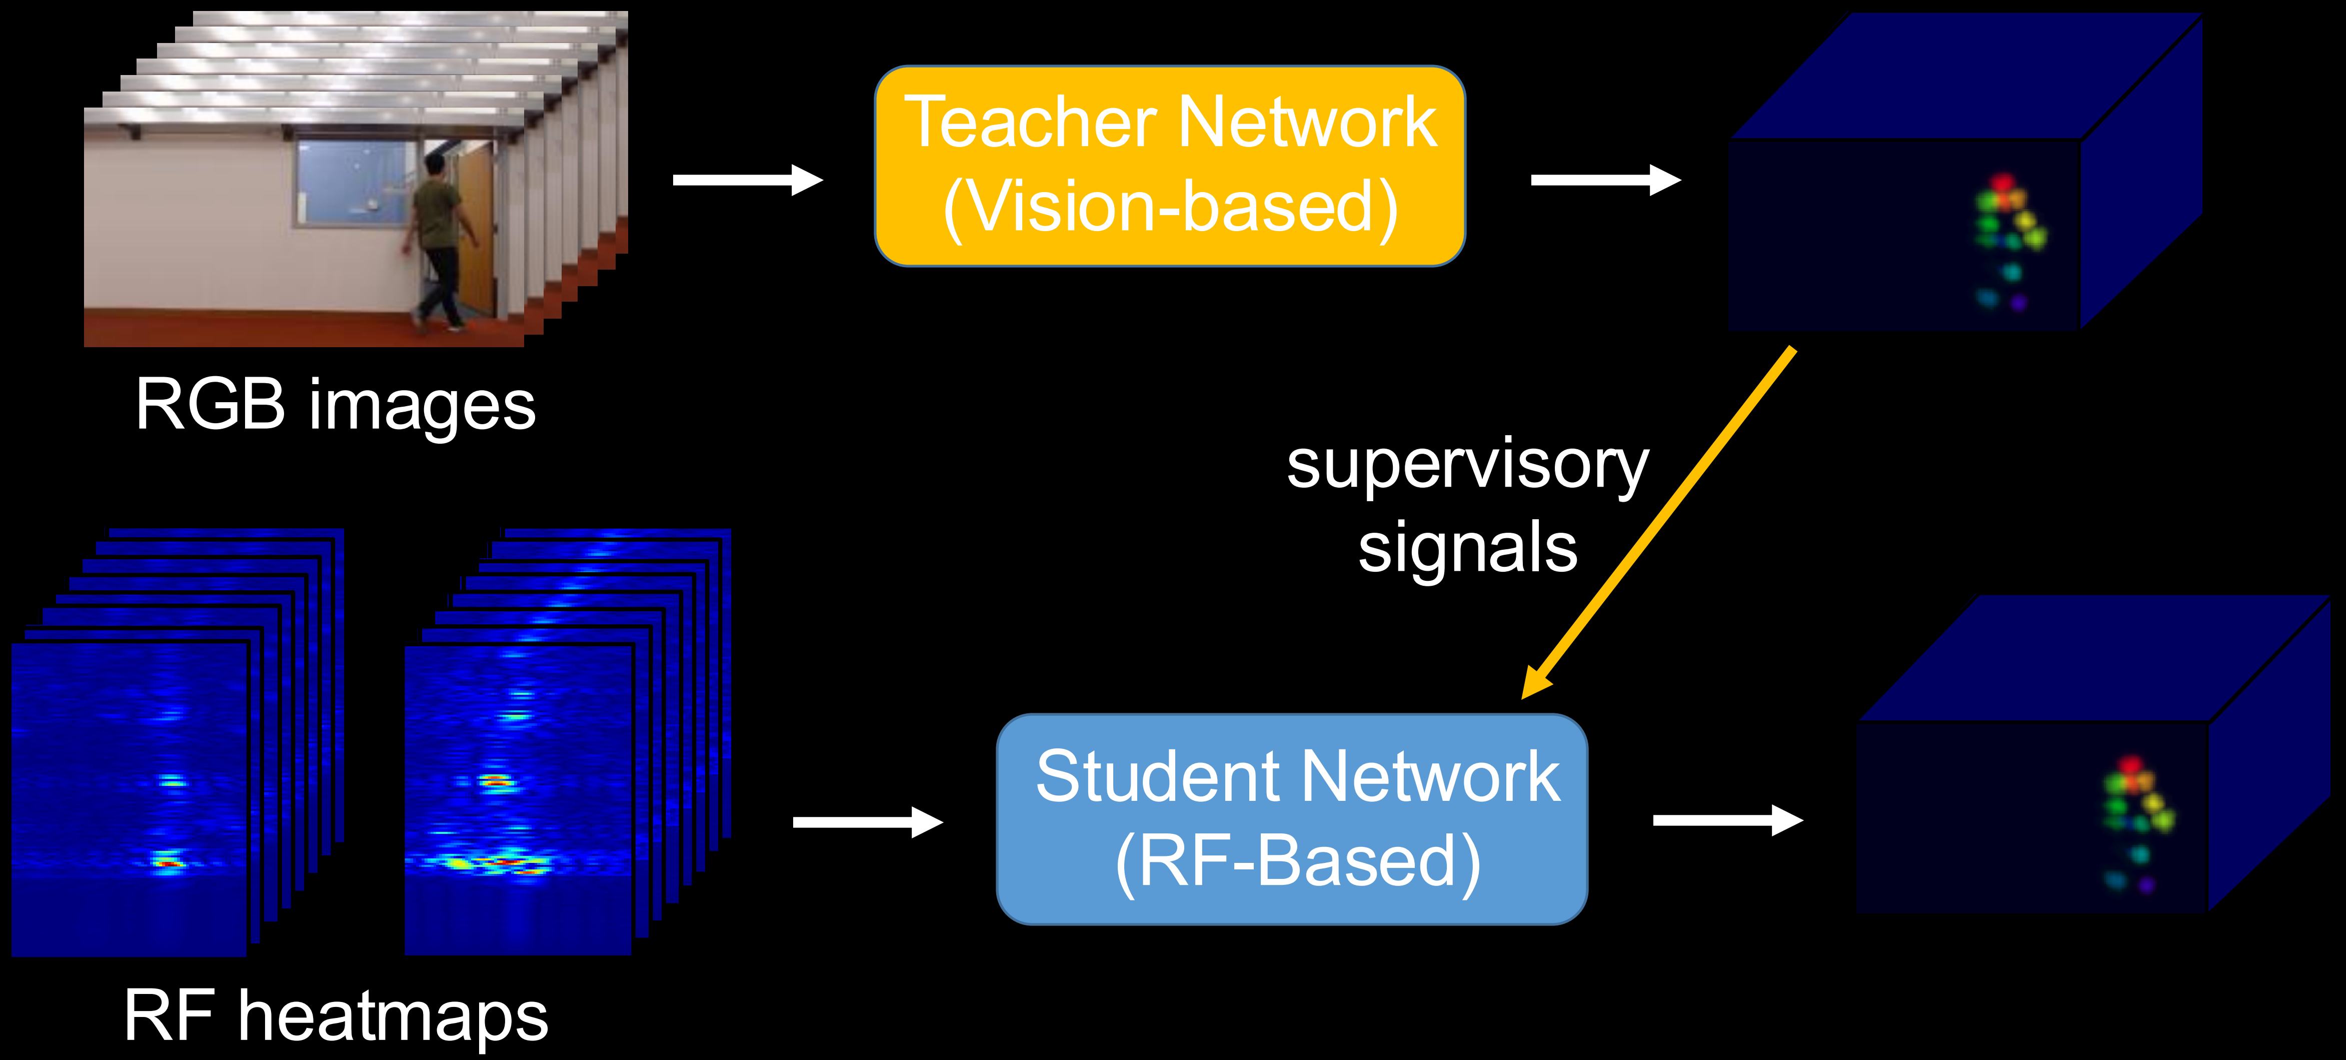

During inference

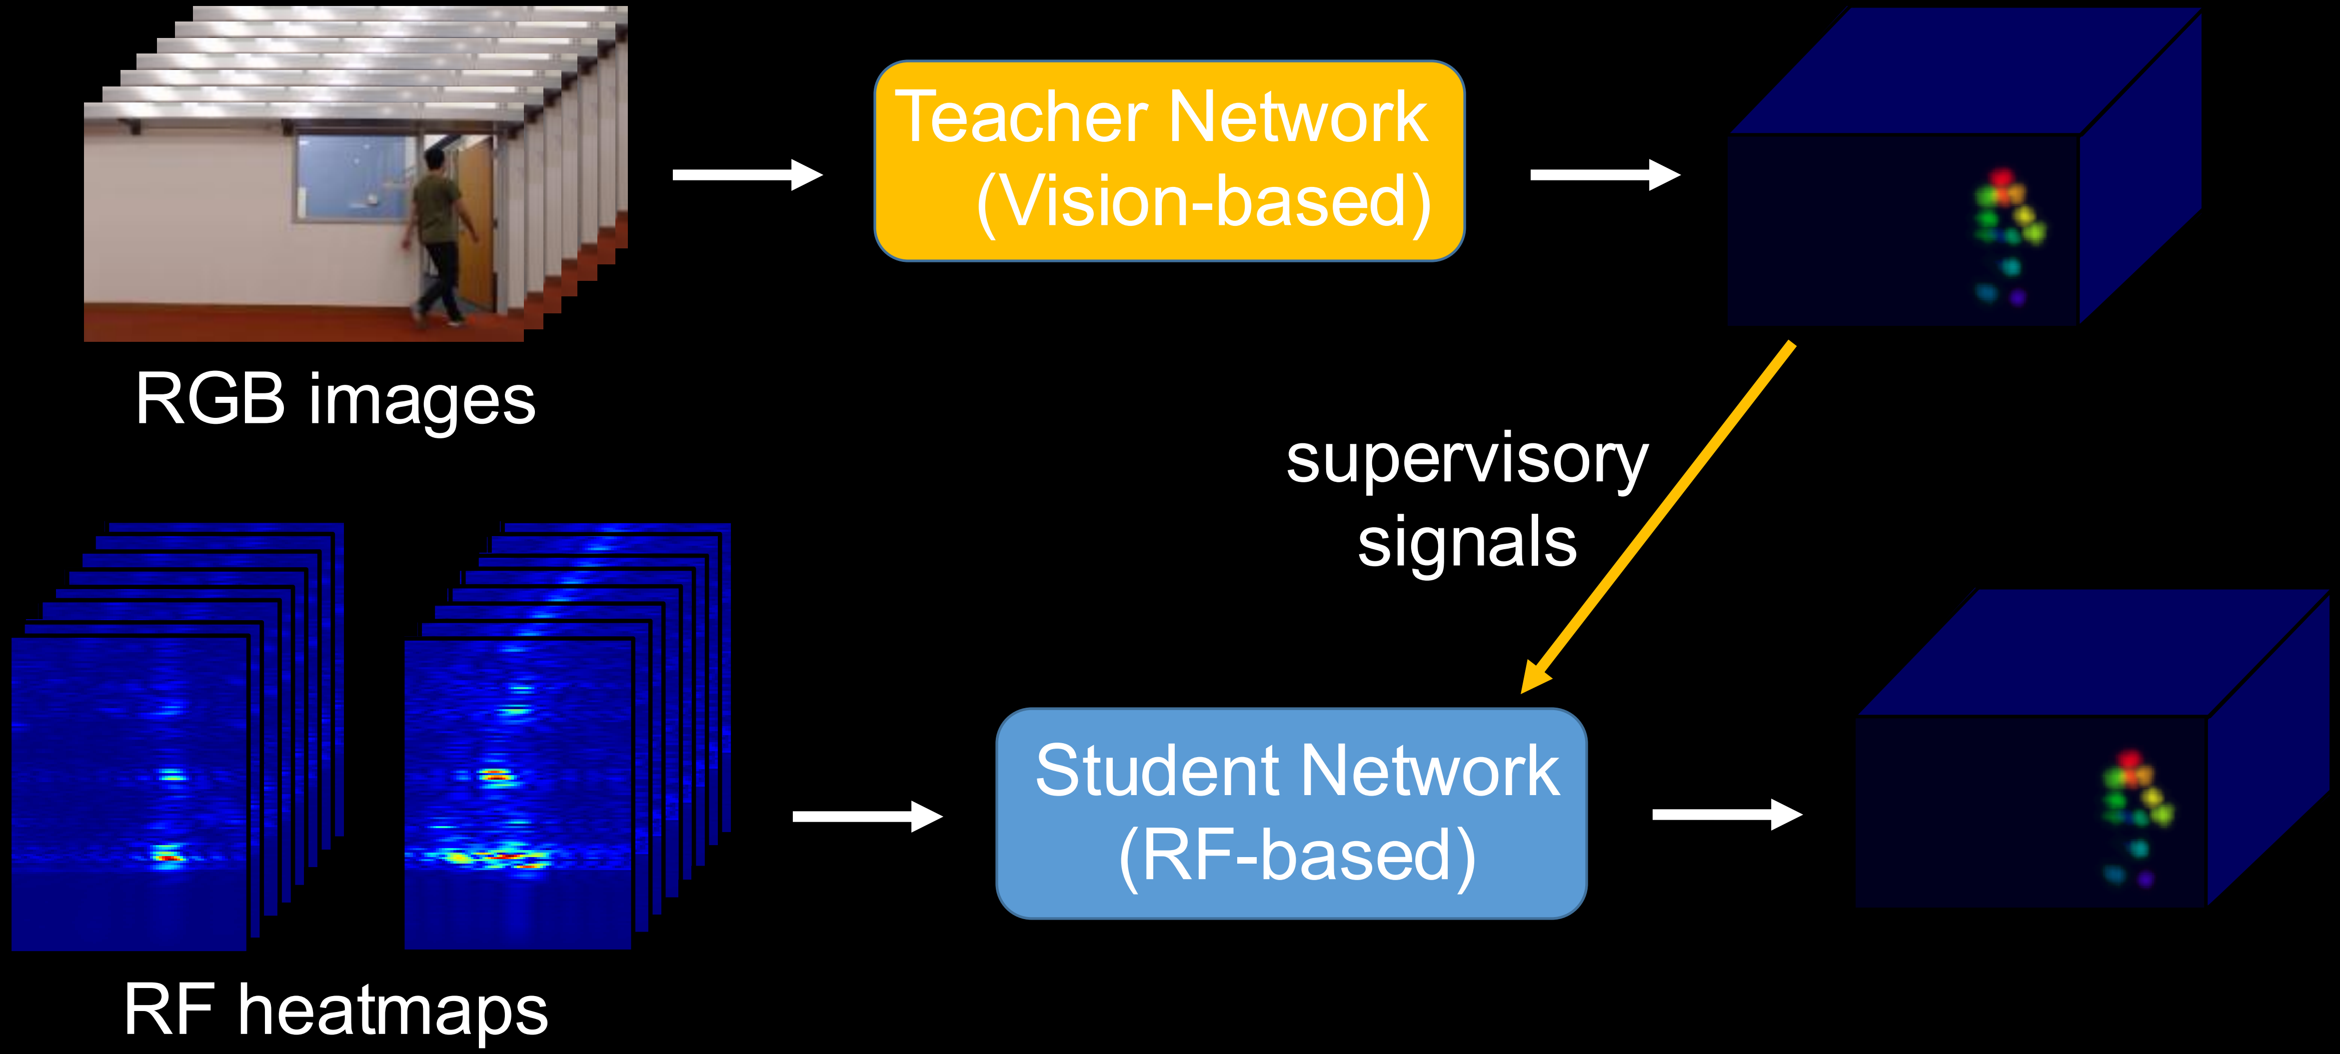

Through-wall poses using **only** RF

RGB (visualization only)

Skeletons

Confidence Maps

Works in bad lighting

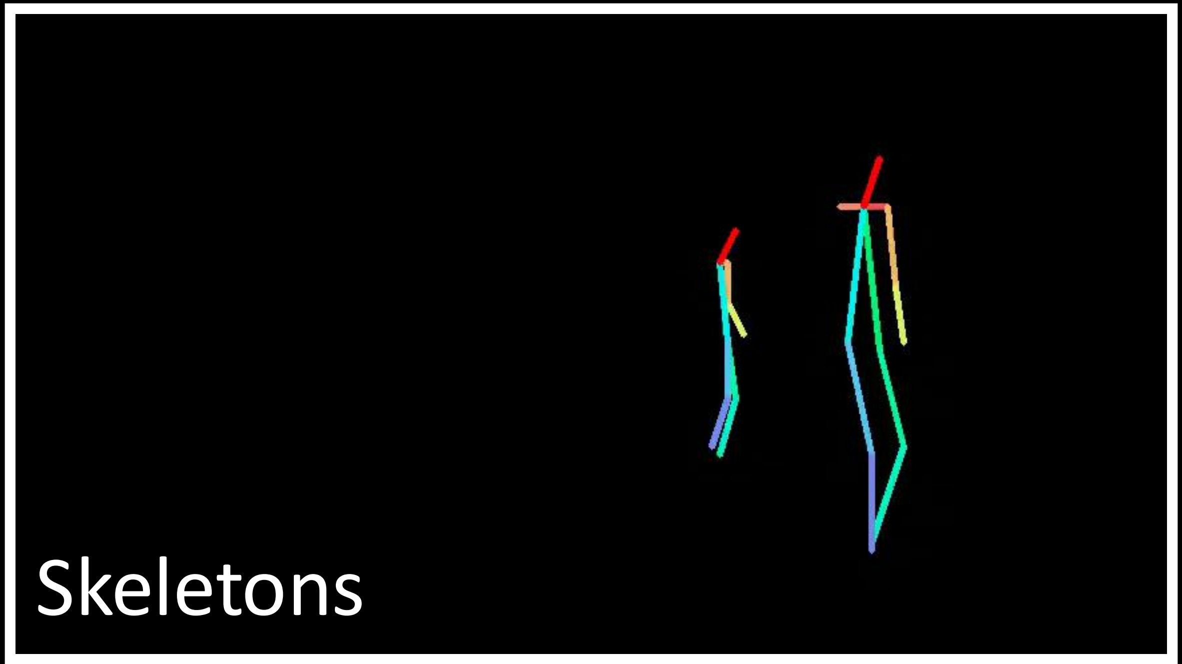

Works with different environment and daily activities

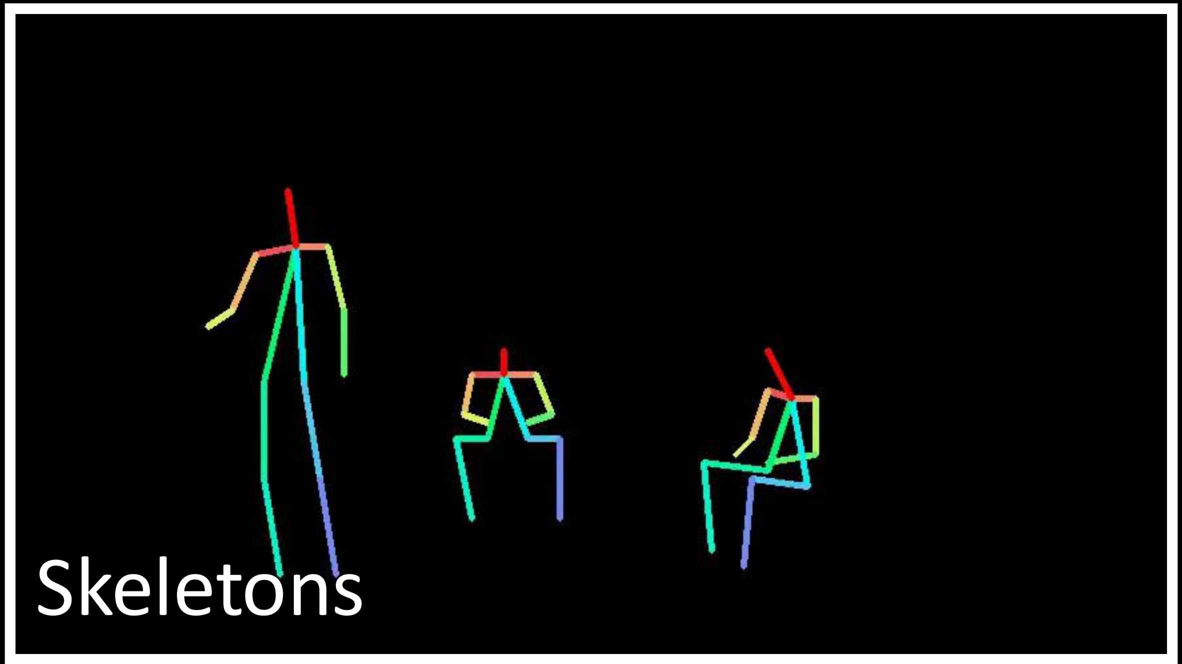

Model Design: Complexity

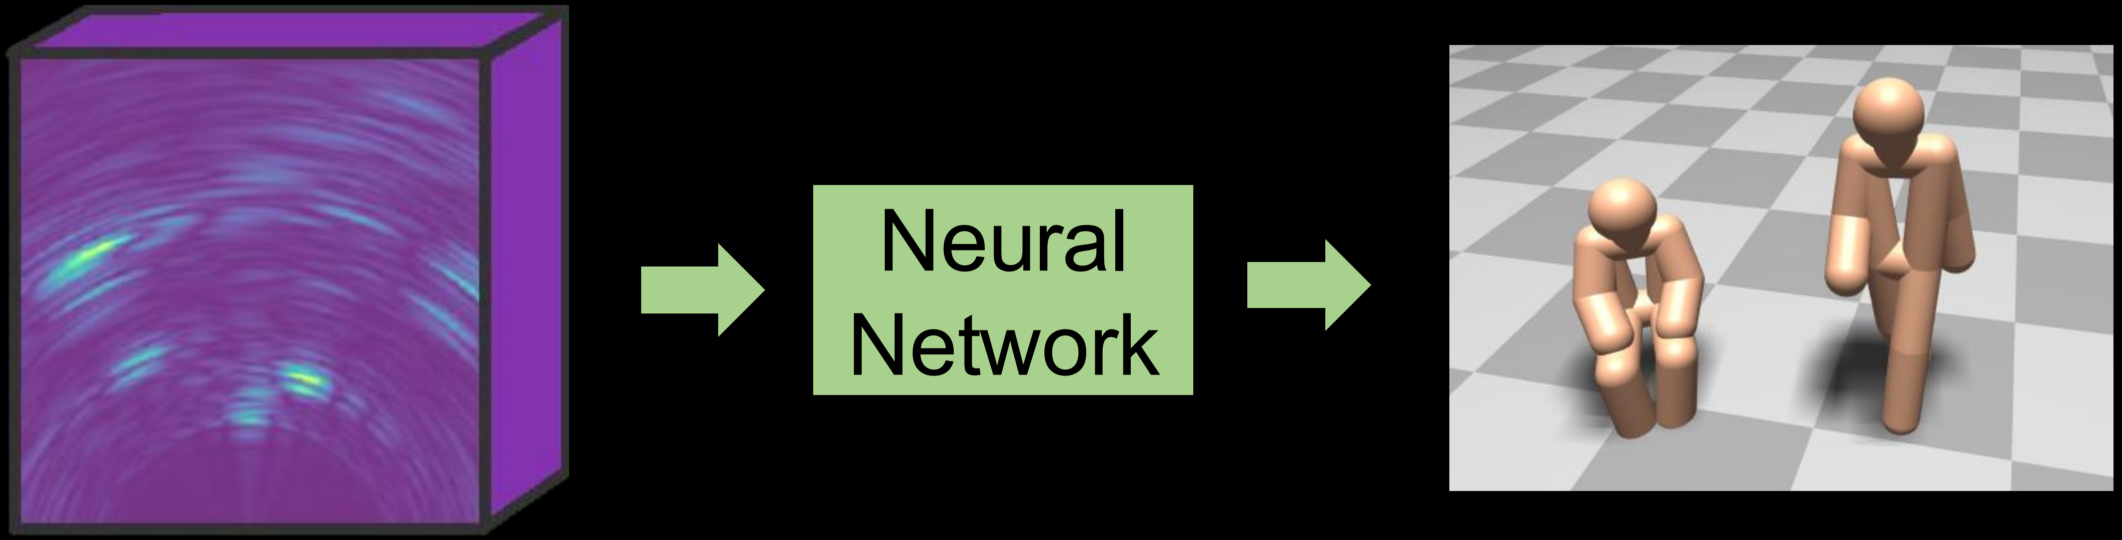

3D RF tensor

Model Design: Complexity

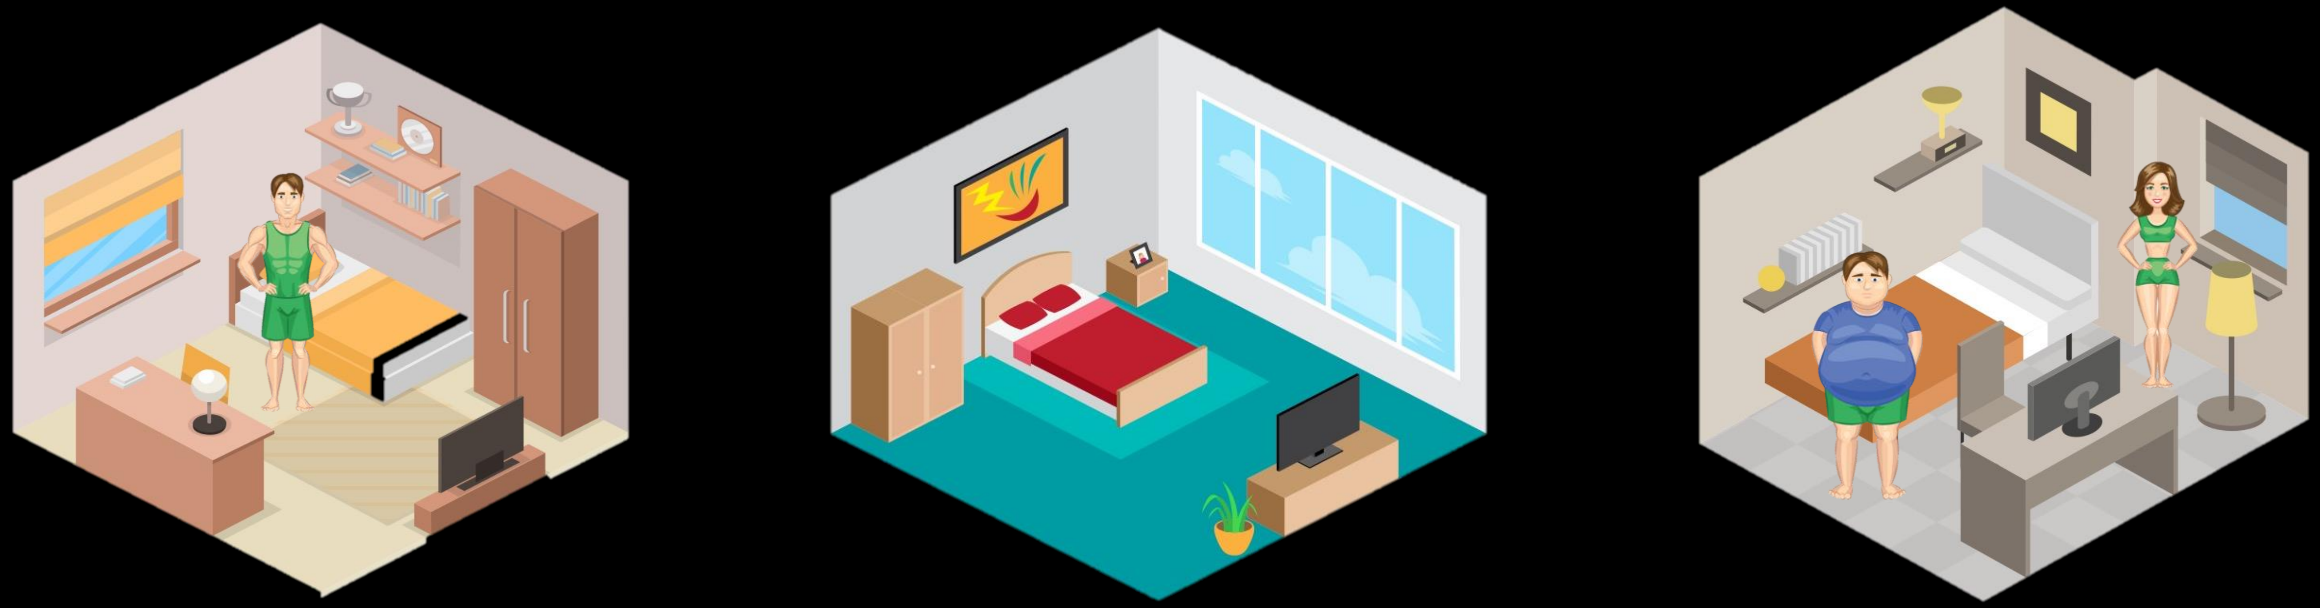

Implications:

- Large neural networks
- Huge amount of labeled examples
- Very hard to train

Solution: Two-Stage Model for Task Separation

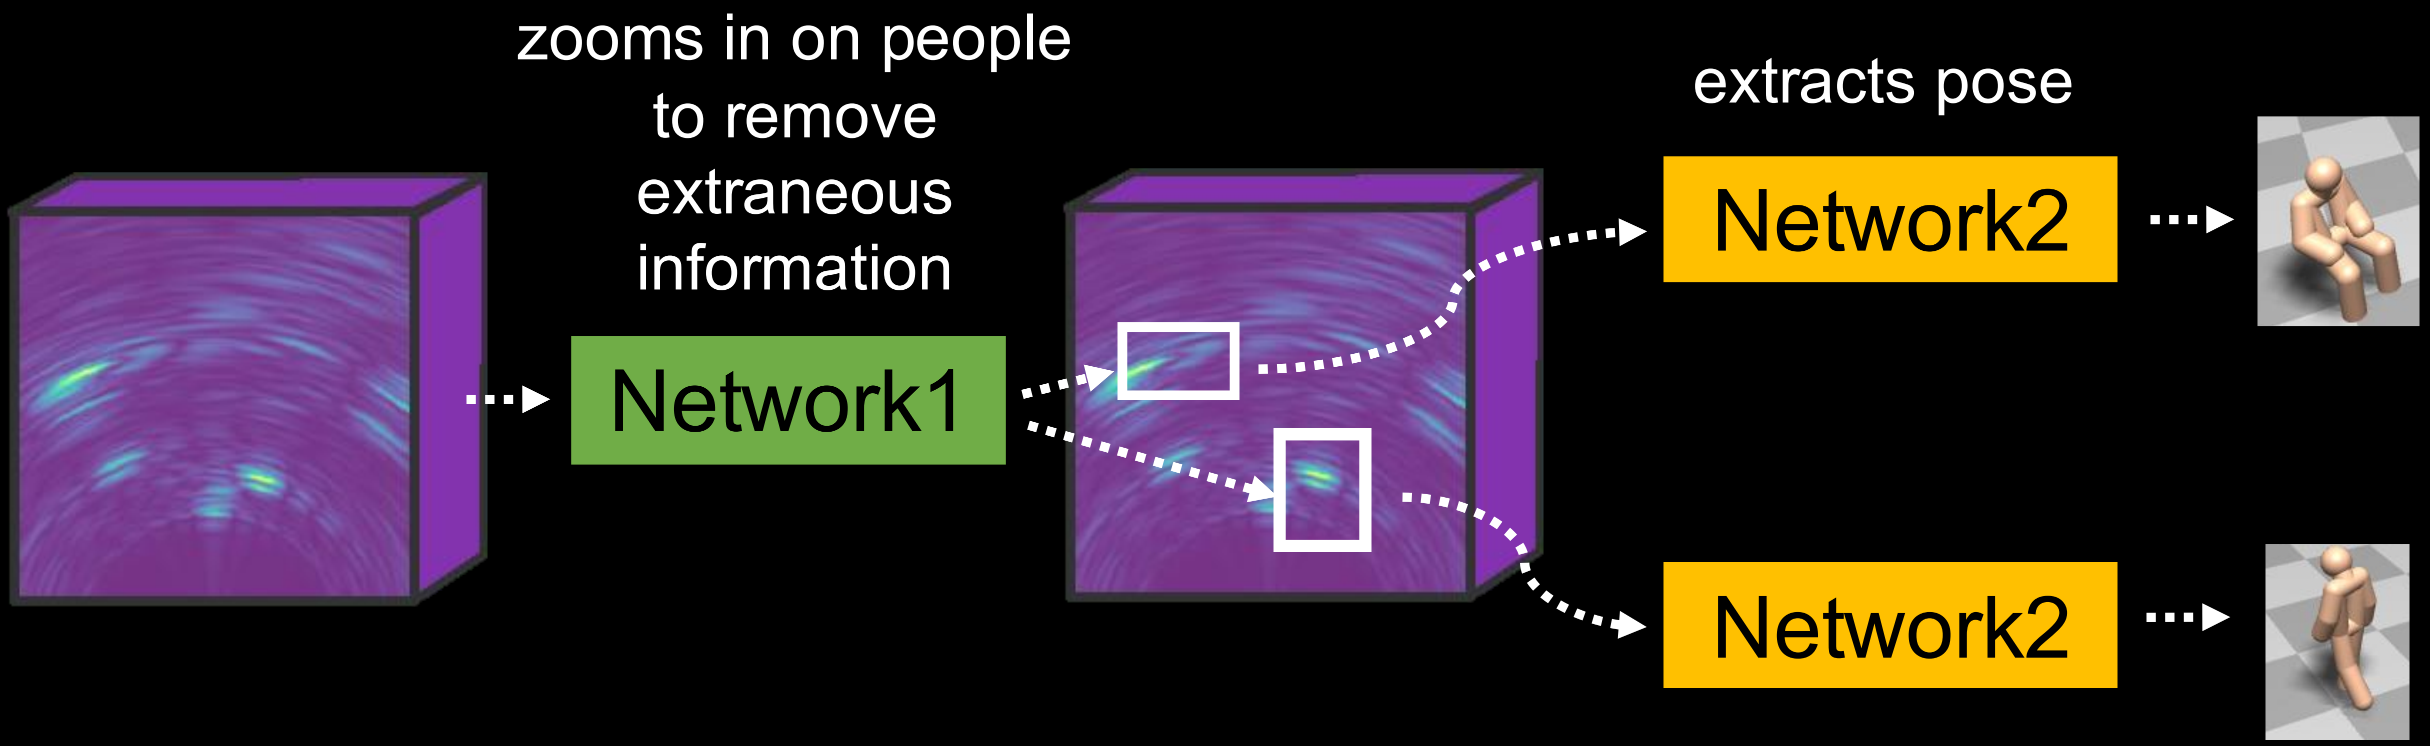

Automatic labeling of 3D poses with cameras

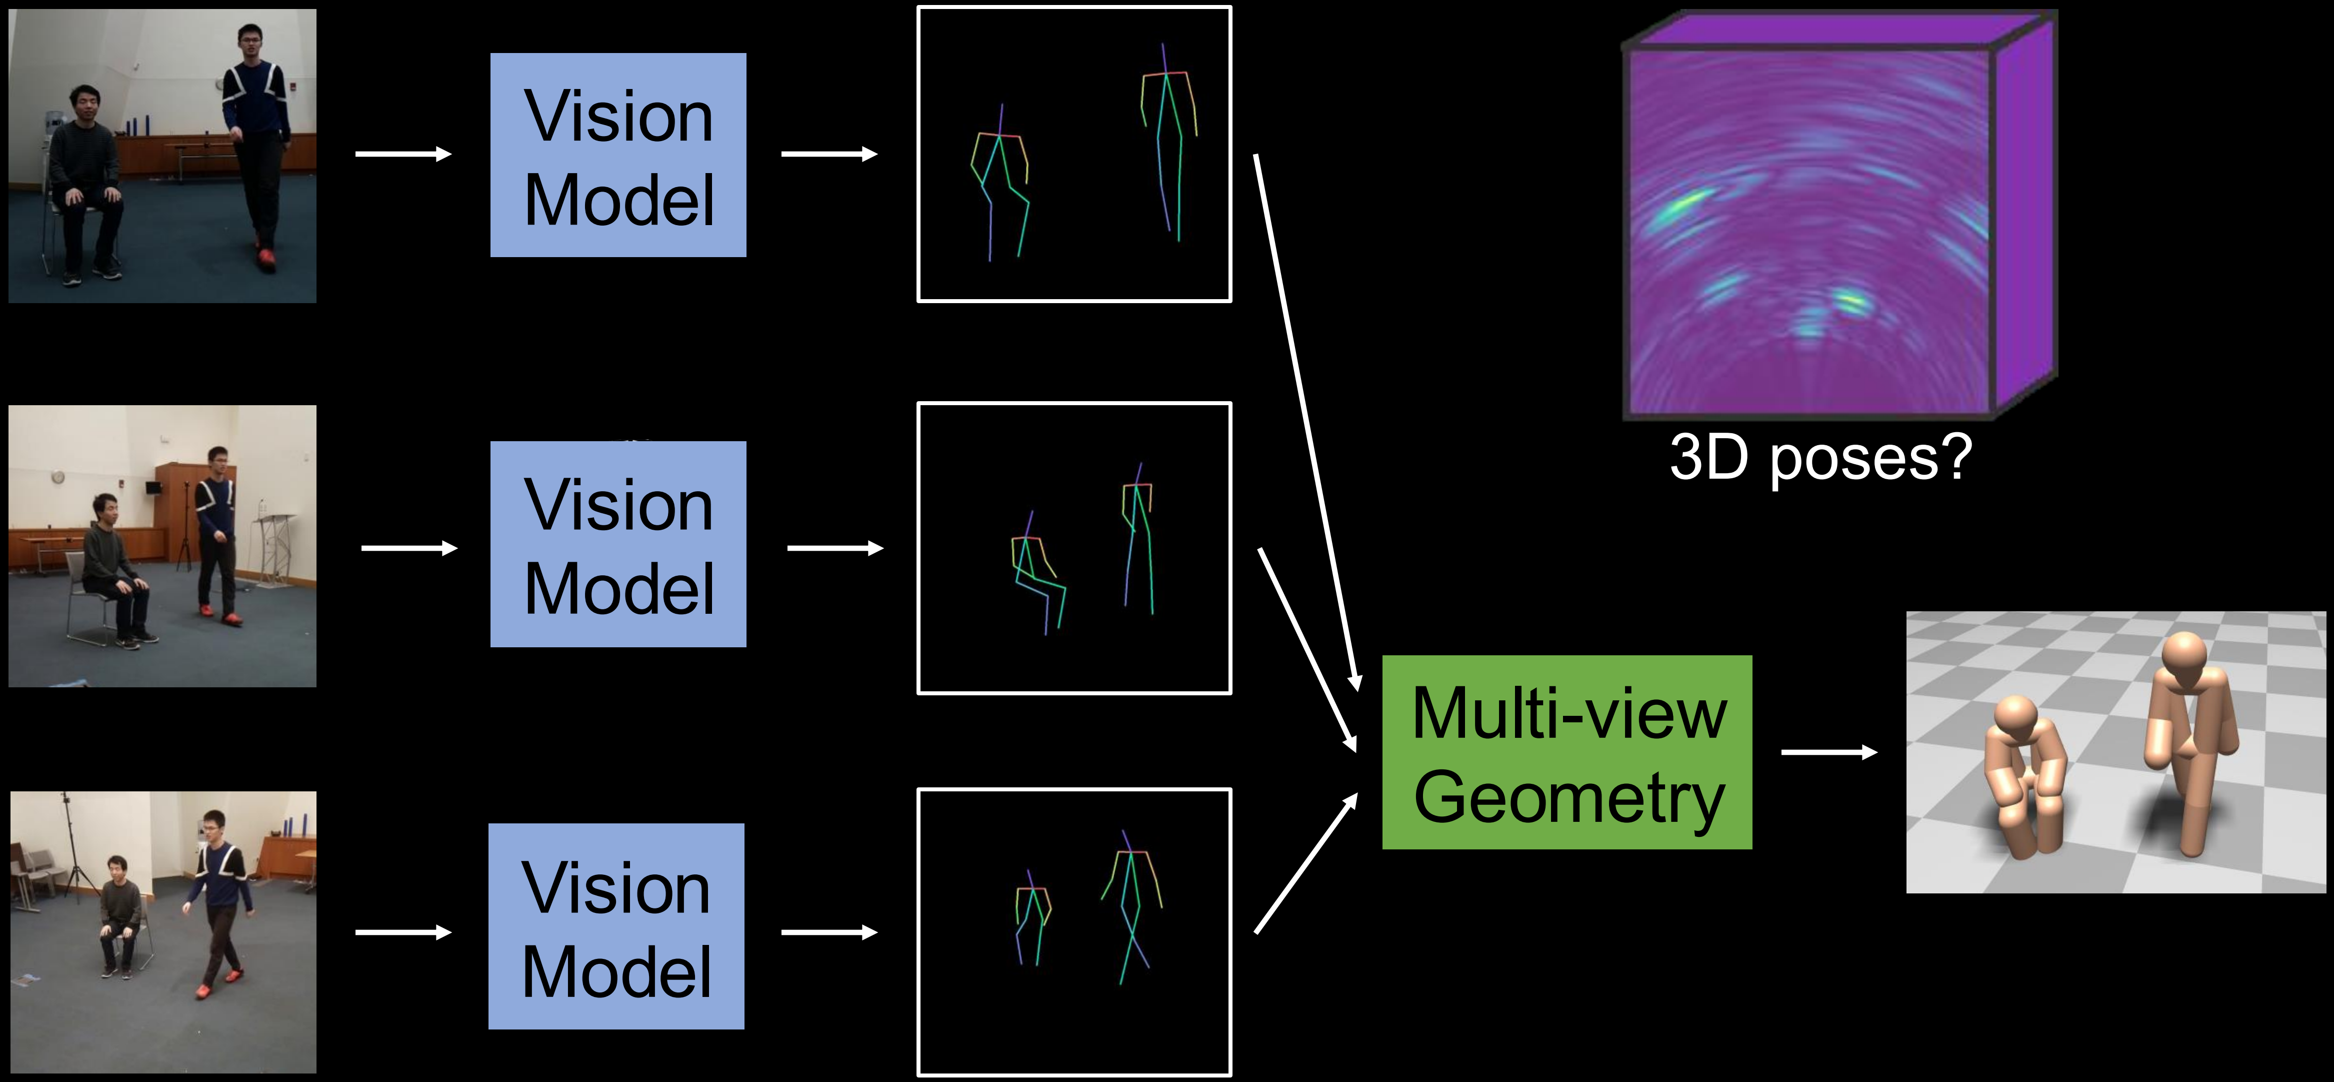

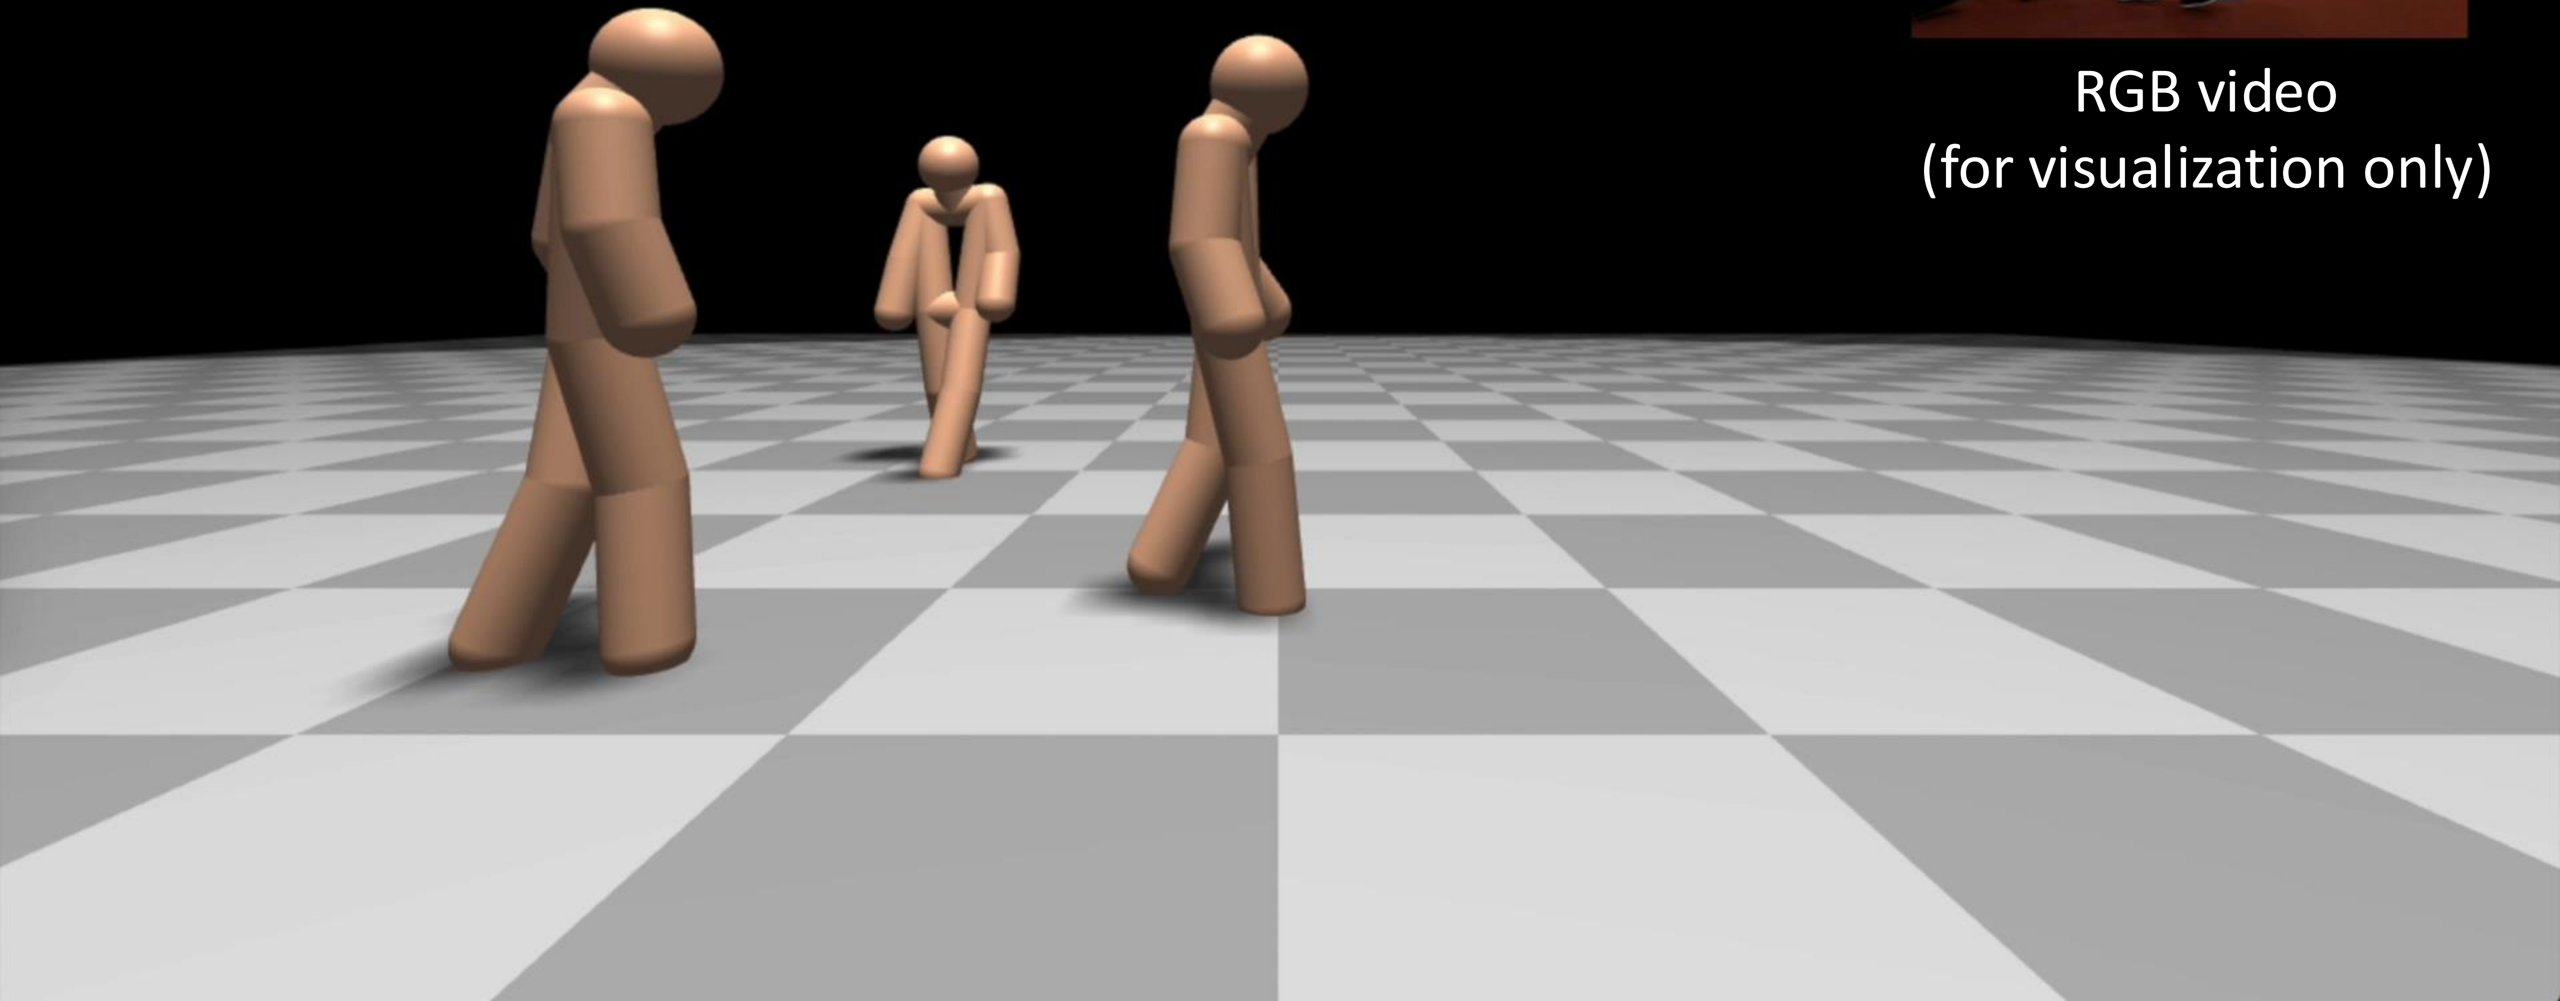

RGB video
(for visualization only)

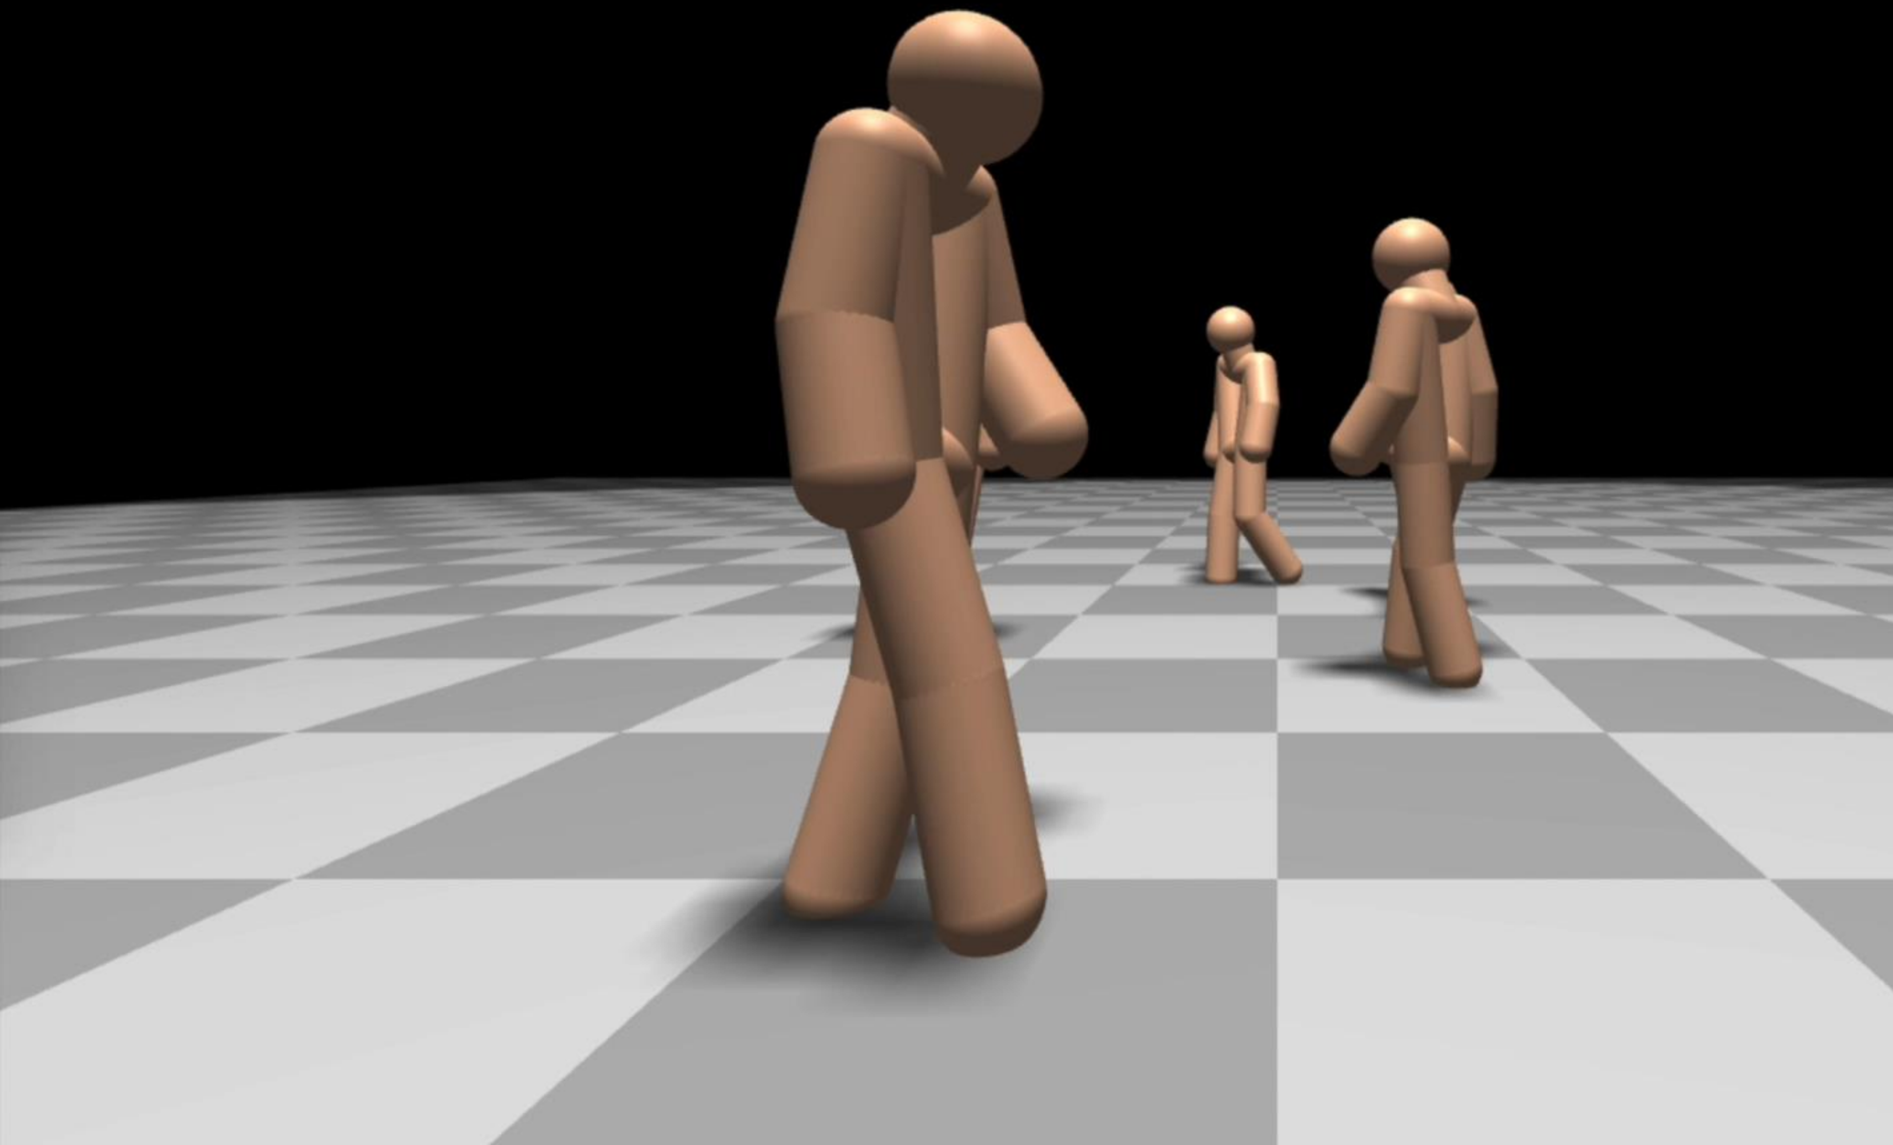

RGB video
(for visualization only)

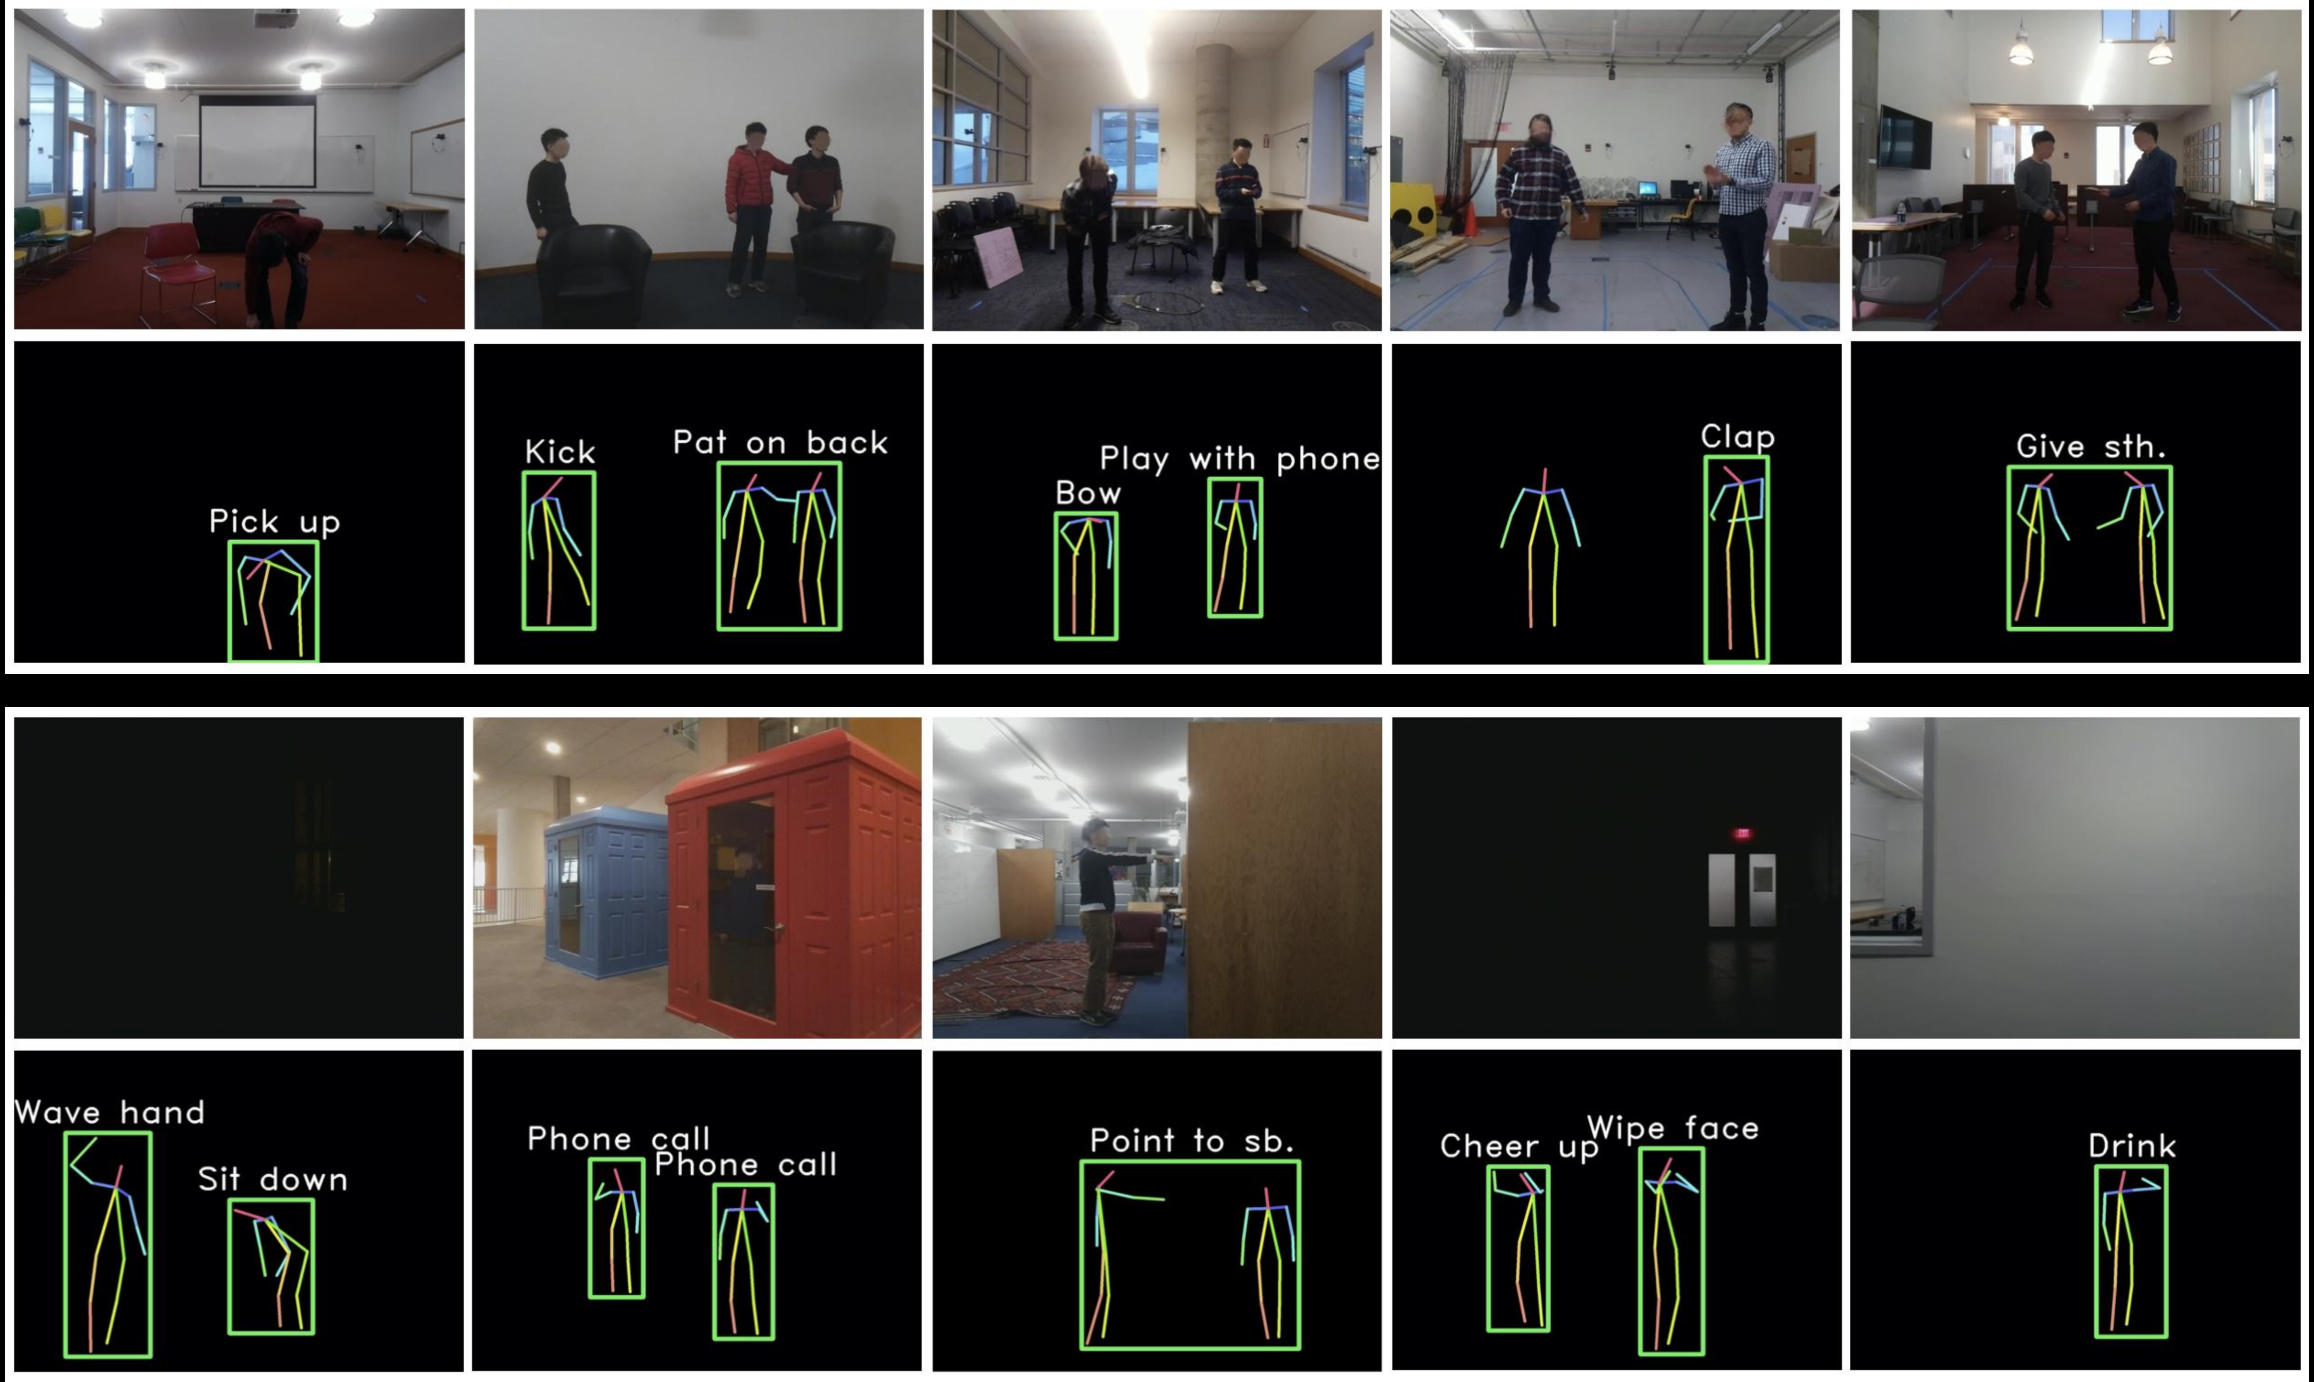
Pick up
Kick
Pat on back
Play with phone
Bow
Clap
Give sth.
Wave hand
Sit down
Phone call
Phone call
Point to sb.
Cheer up
Wipe face
Drink

RGB view

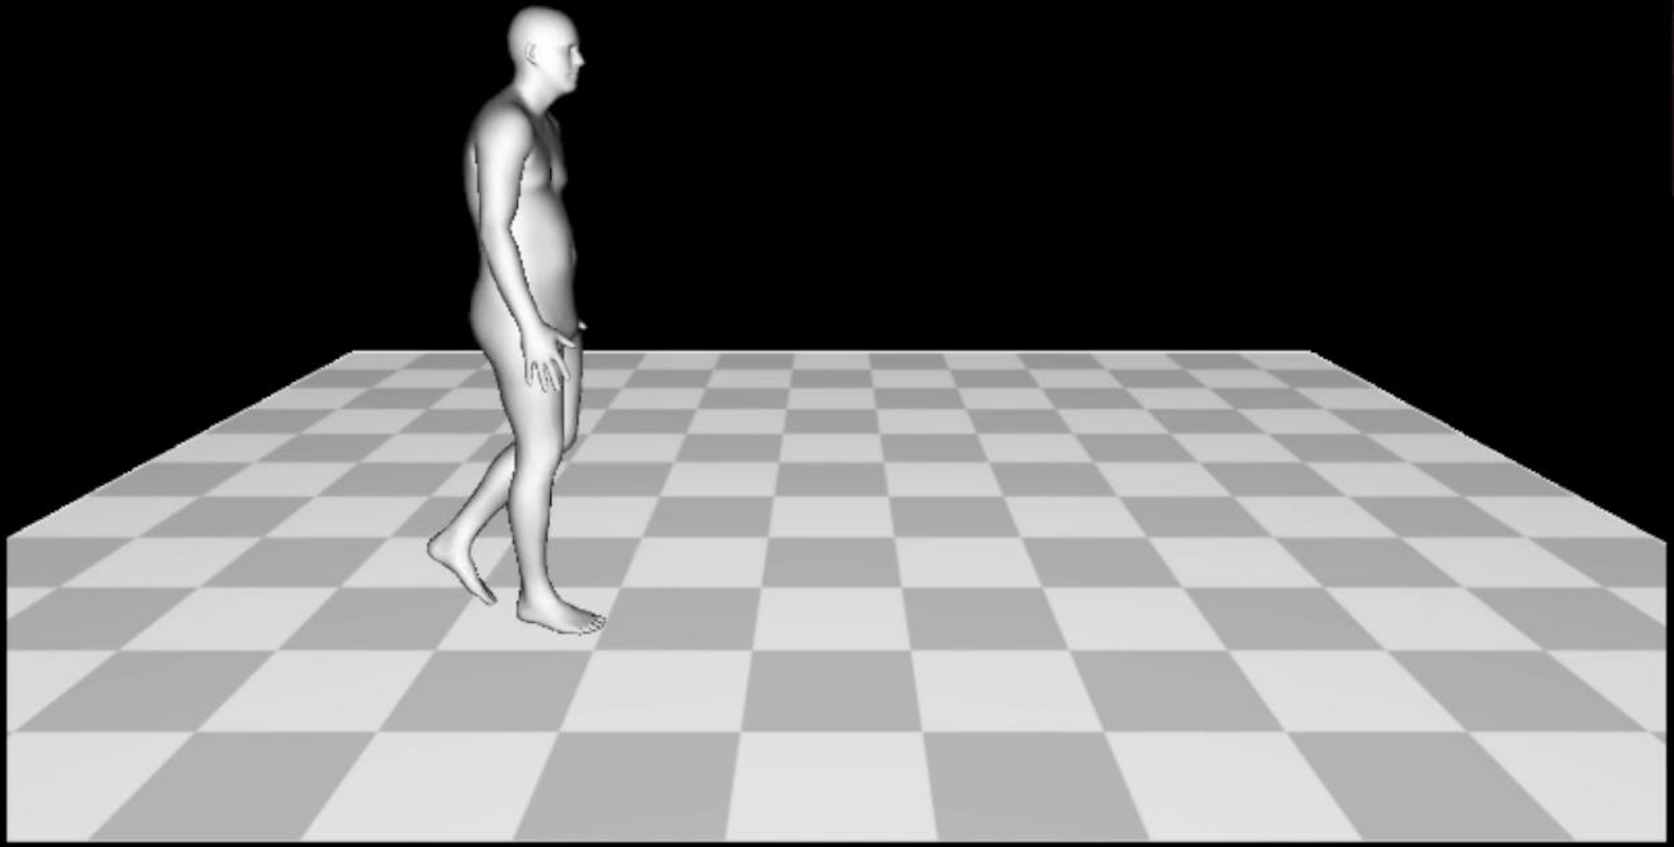

Dynamic human meshes from RF

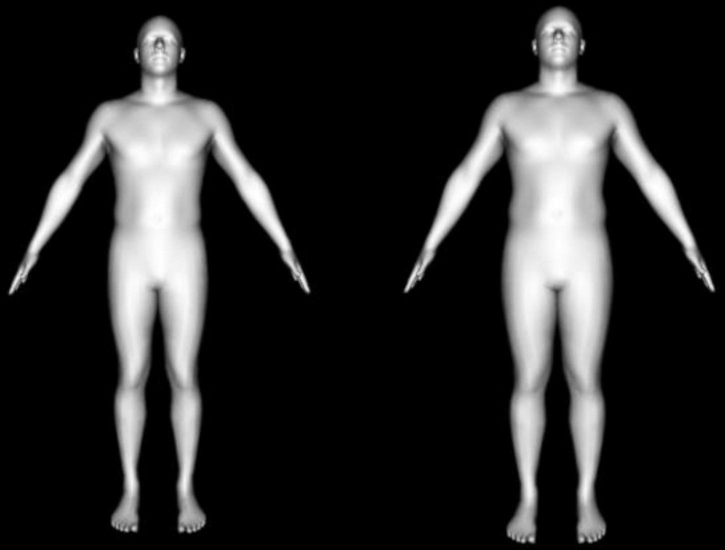

Body shapes from RF

Tesla in Fog

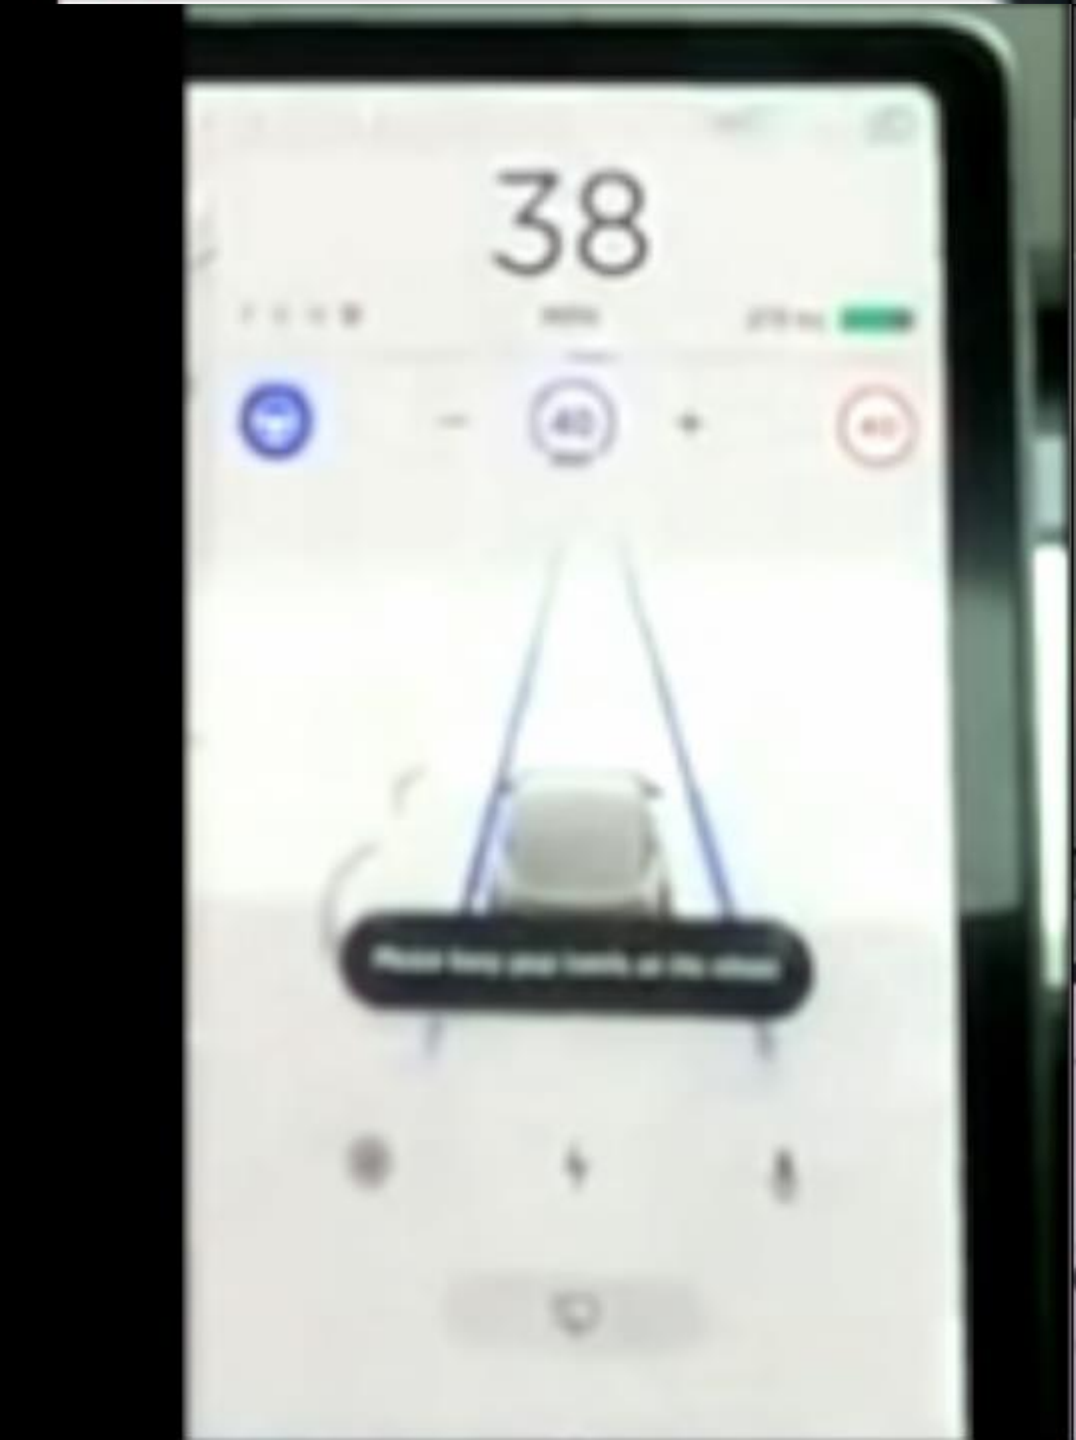

Tesla in Fog

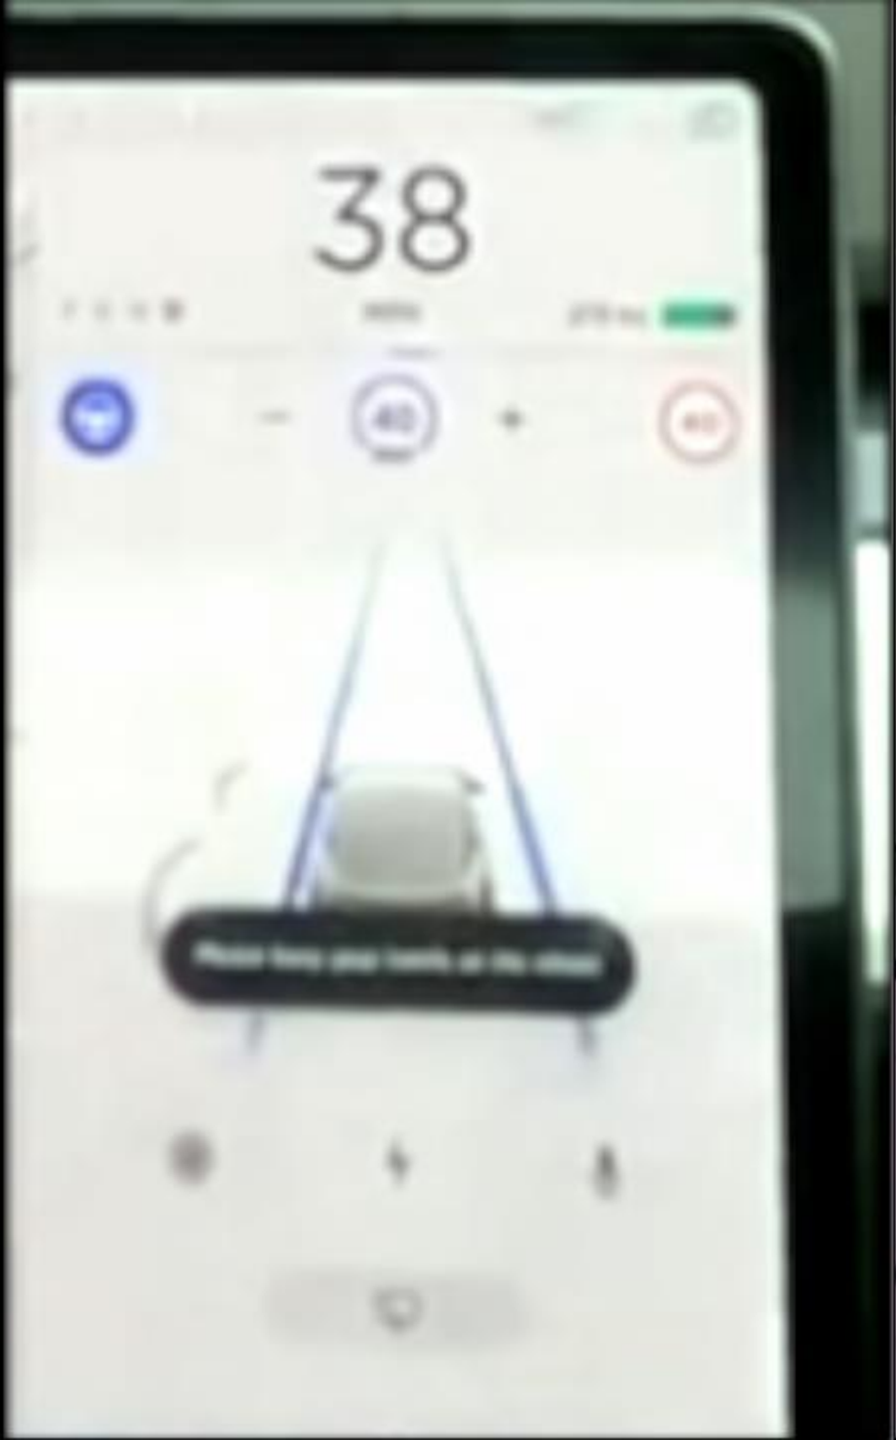

Tesla in Fog

Millimeter Wave Radar

Radar can function in adverse conditions

Millimeter Wave Radar

Radar can function in adverse conditions

Can we use millimeter wave radars in scenarios where LiDARs and cameras fail?


ZED

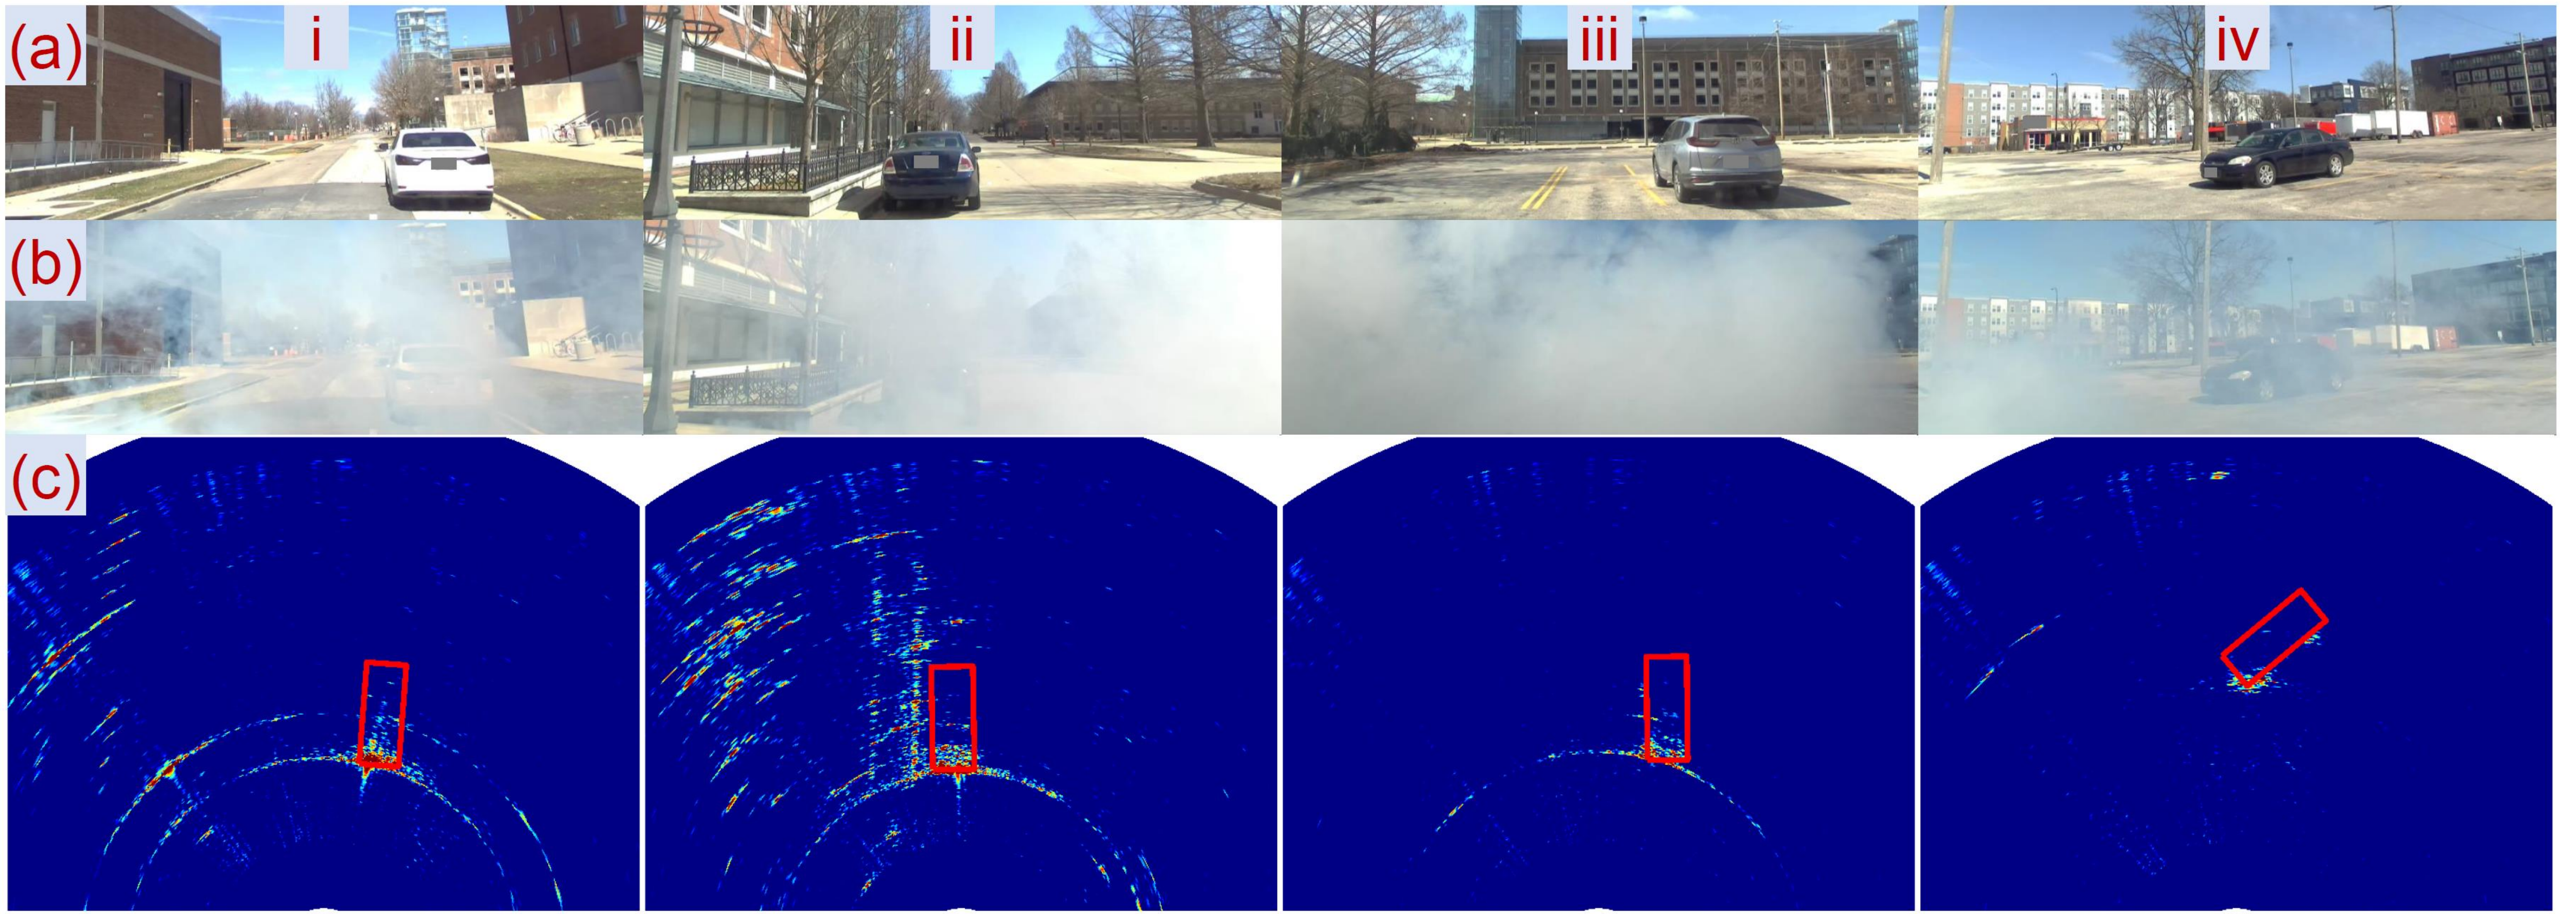

(a)
(b)
(c)
i
ii
iii
iv

Challenge: Labeled Dataset

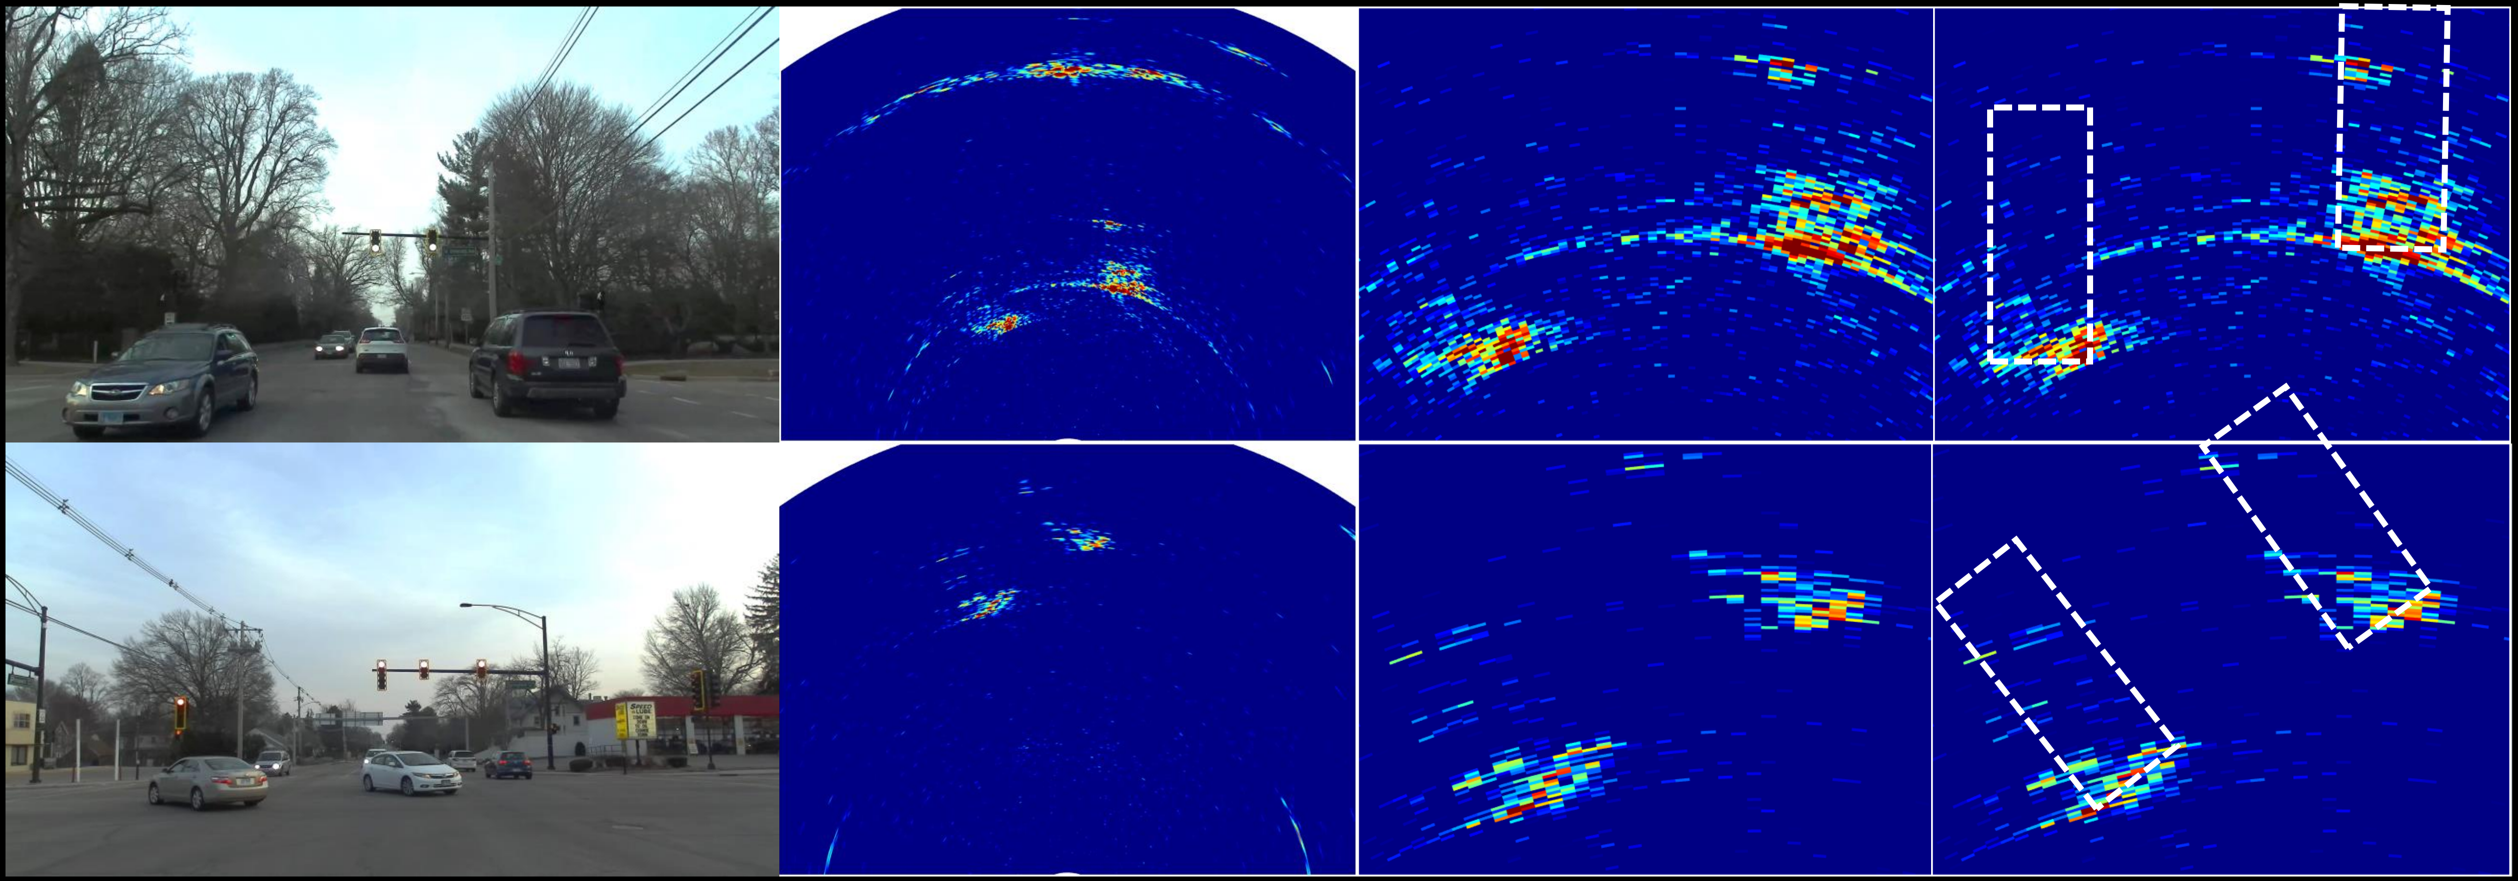

(a) Scene

(b) BEV Radar Heatmap

(c) Zoomed Into Car Area

(d) Annotation of (c)

Leveraging large-scale unlabeled radar data but bypassing the complexities of explicit annotation.

Self-Supervised Learning

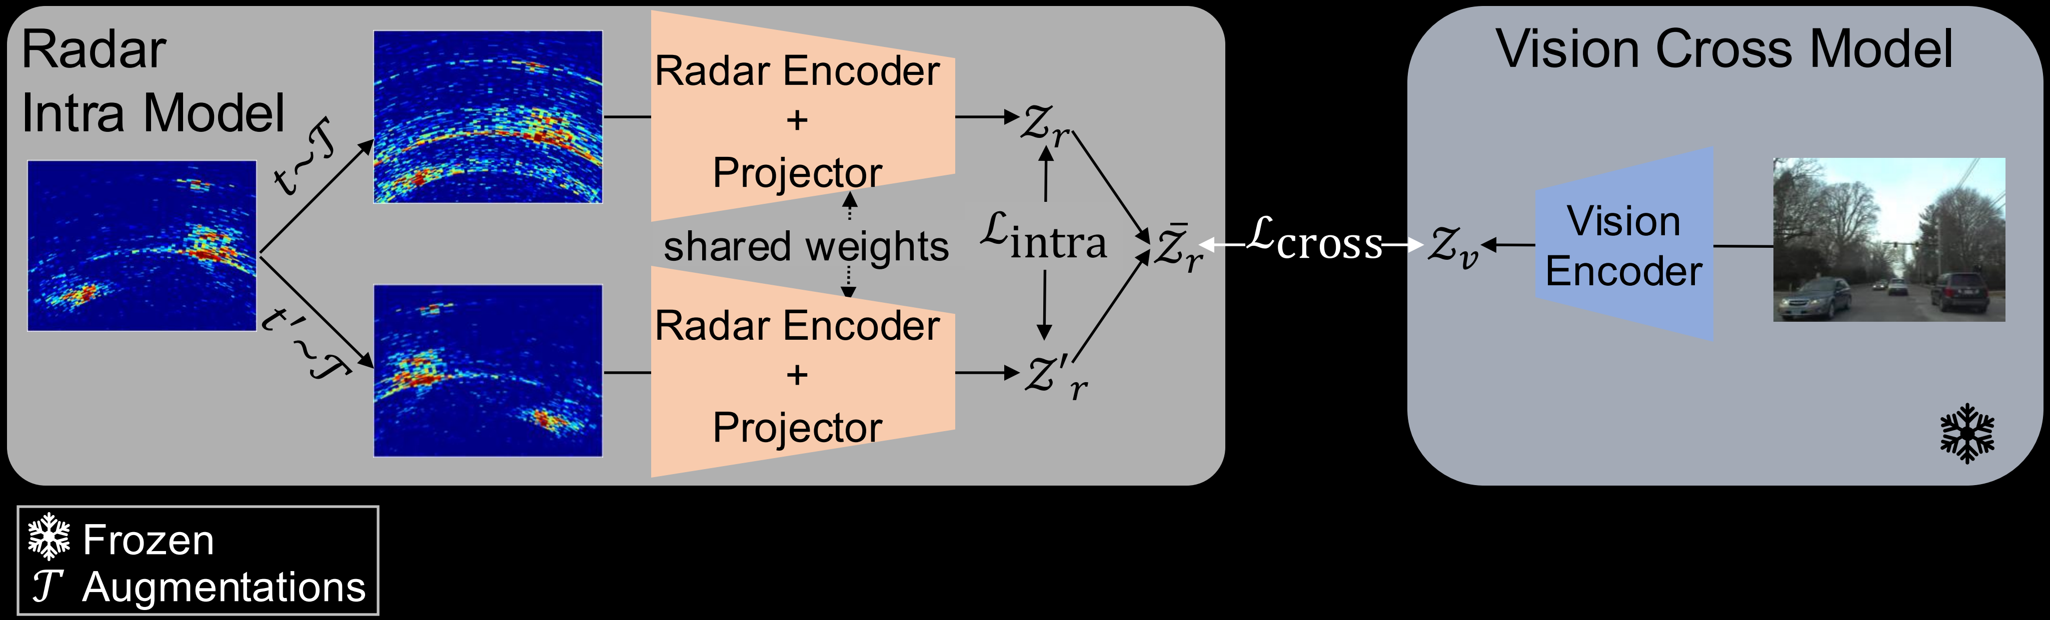

Leveraging large-scale unlabeled radar data but bypassing the complexities of explicit annotation.

Radar Specific Augmentations

Original Radar Heatmap | Zommed-in Region of Cars

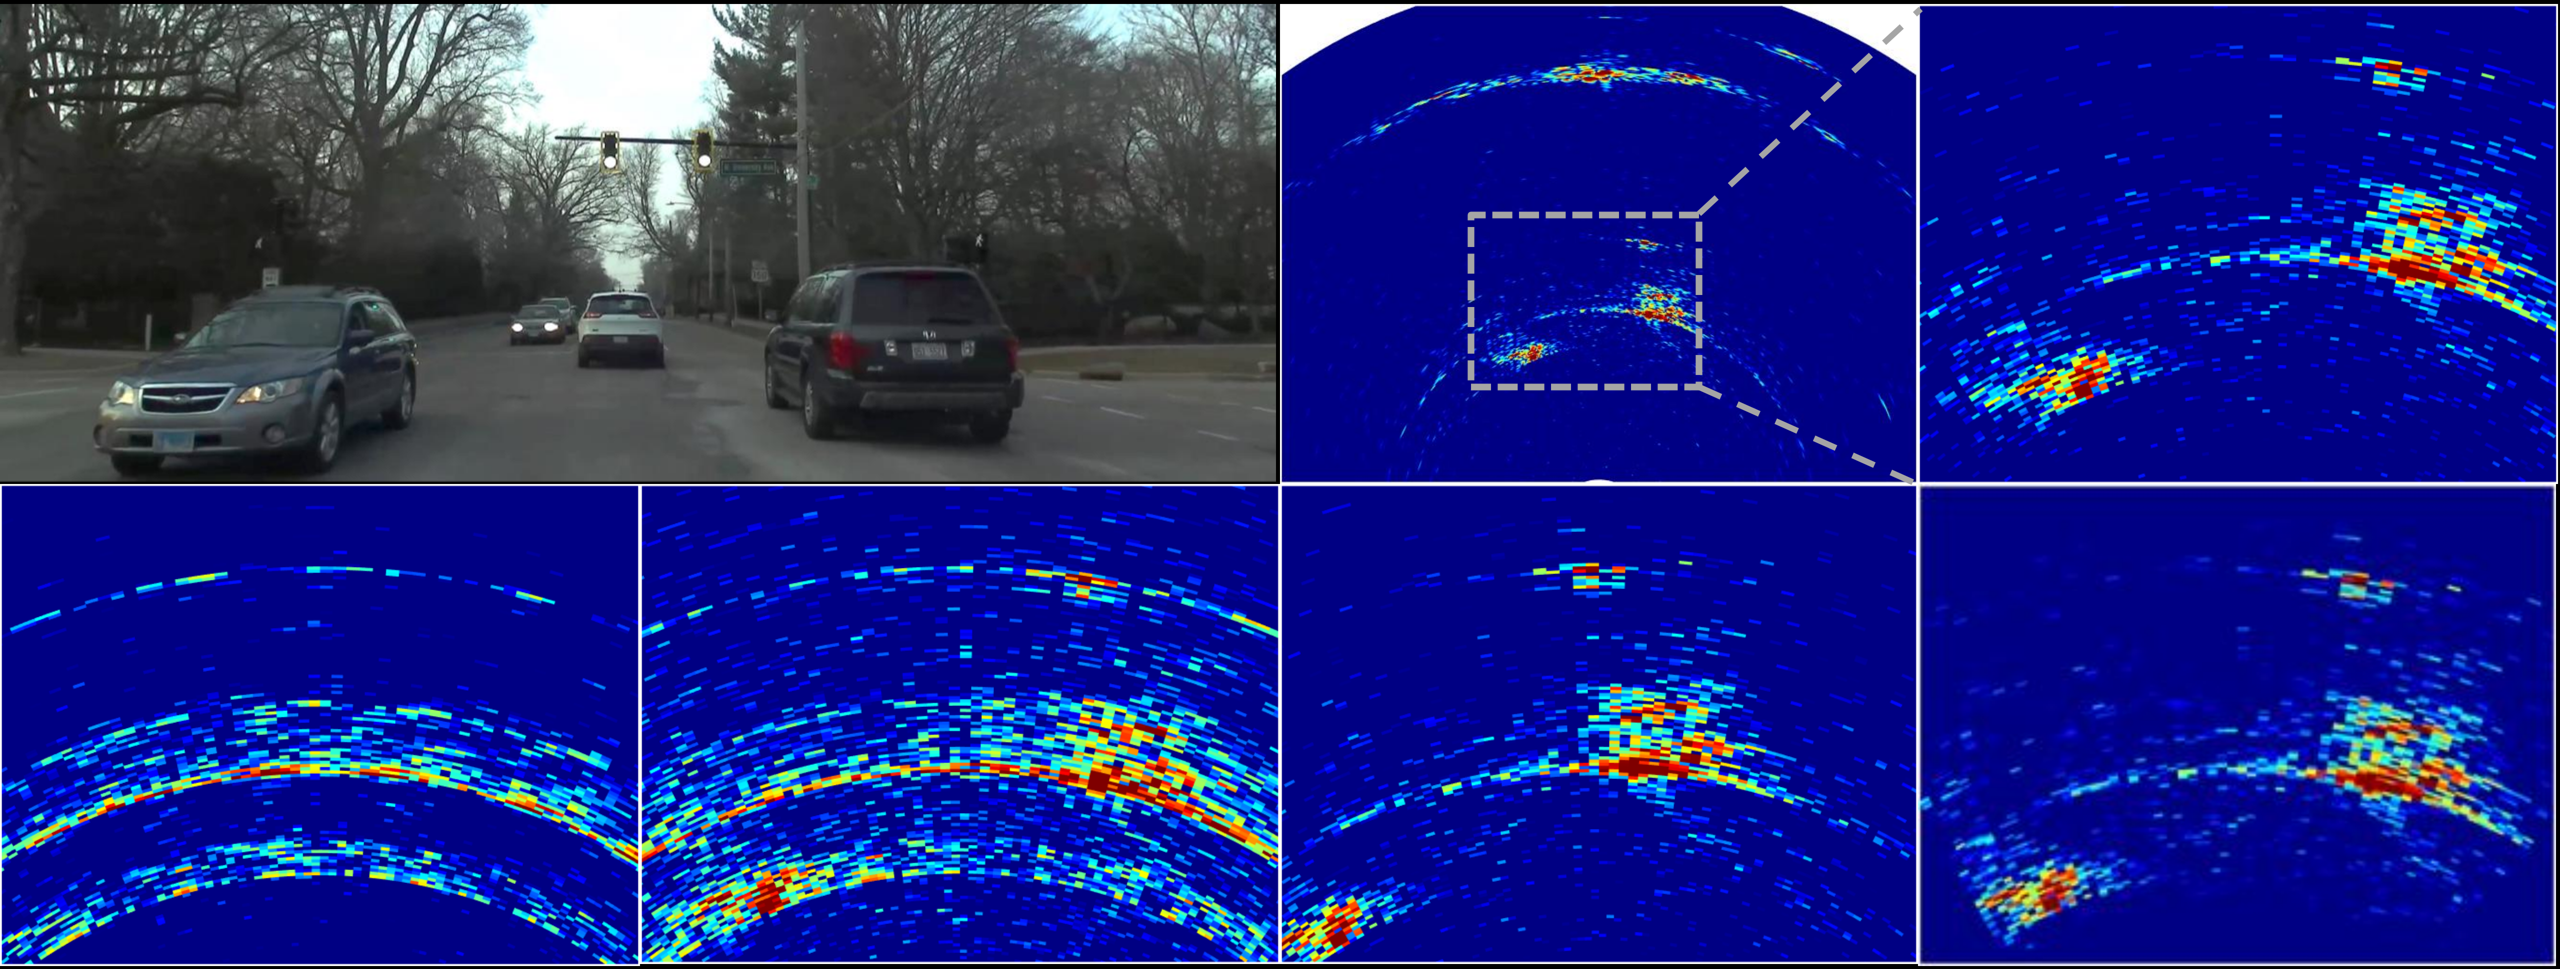

Random Phase | Antenna Dropout | Rotation (Polar) | Center Cropping (Polar)

Accurate Bounding Box Detection

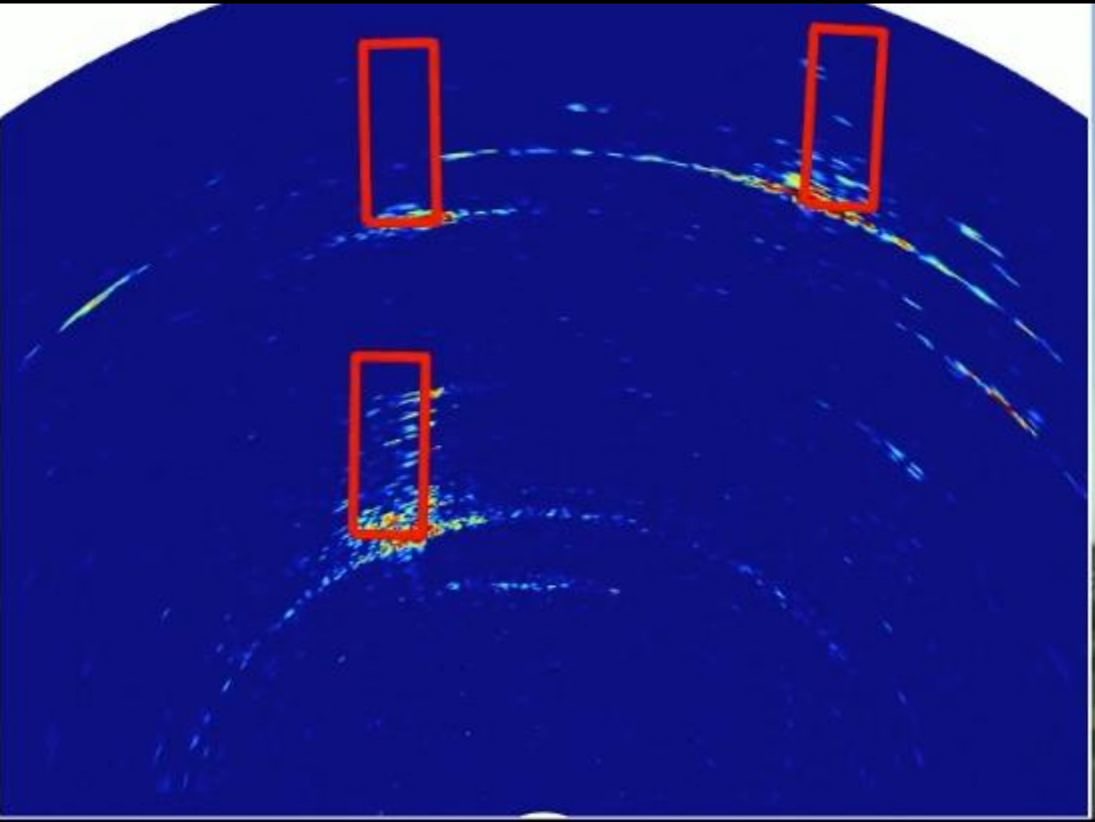

No Temporal Post Processing

Frame to frame detection in the range of 25 meters.

Challenges in Radar Perception

Camera Image

Point Cloud

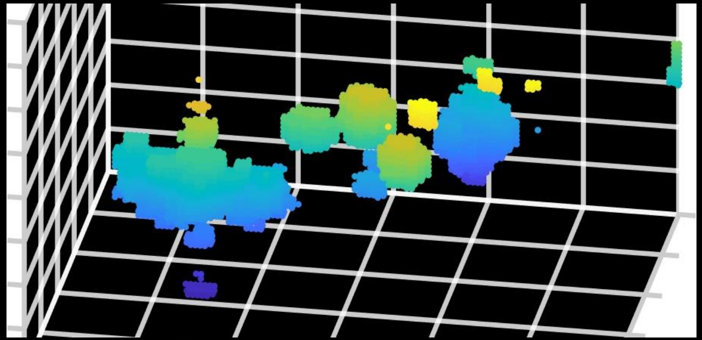

Top-View

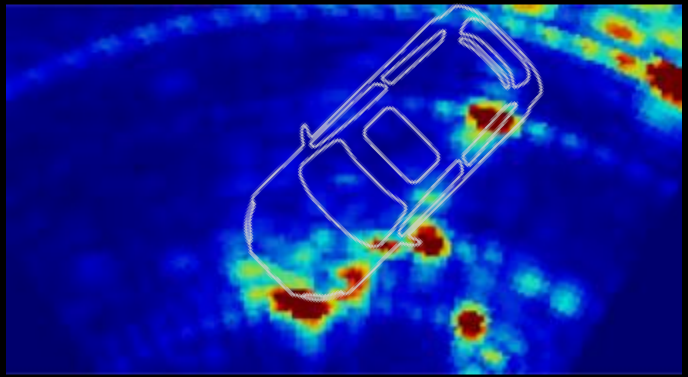

Front-View

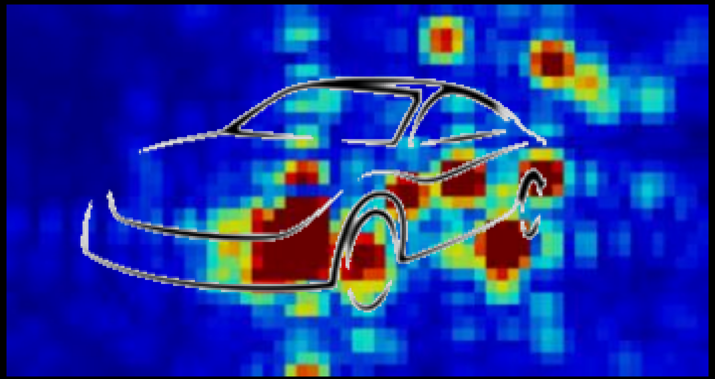

Huge Performance Gap Between Radar and Vision!

Challenges in Radar Perception

1. Low Angular Resolution
- Blobs of reflected power
- No sharp boundaries/shapes

2. Specularity
- Missing major parts of cars

3. Multipath
- Spurious Reflections

Camera Image

Point Cloud

Top-View

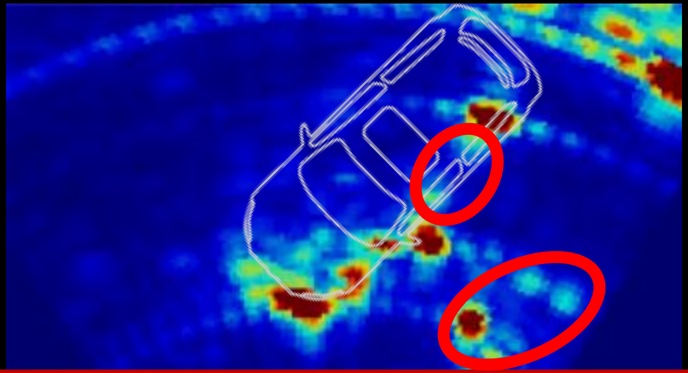

Front-View

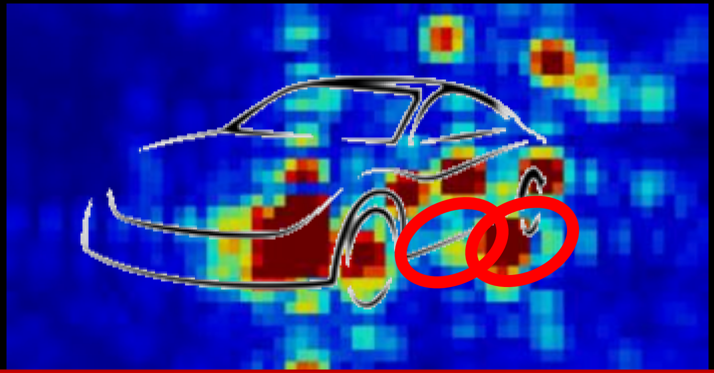

Cast mmWave Radar Perception as a Learning Problem

Our Solution

Learning *geometric priors* on structures of commonly found streetside objects.

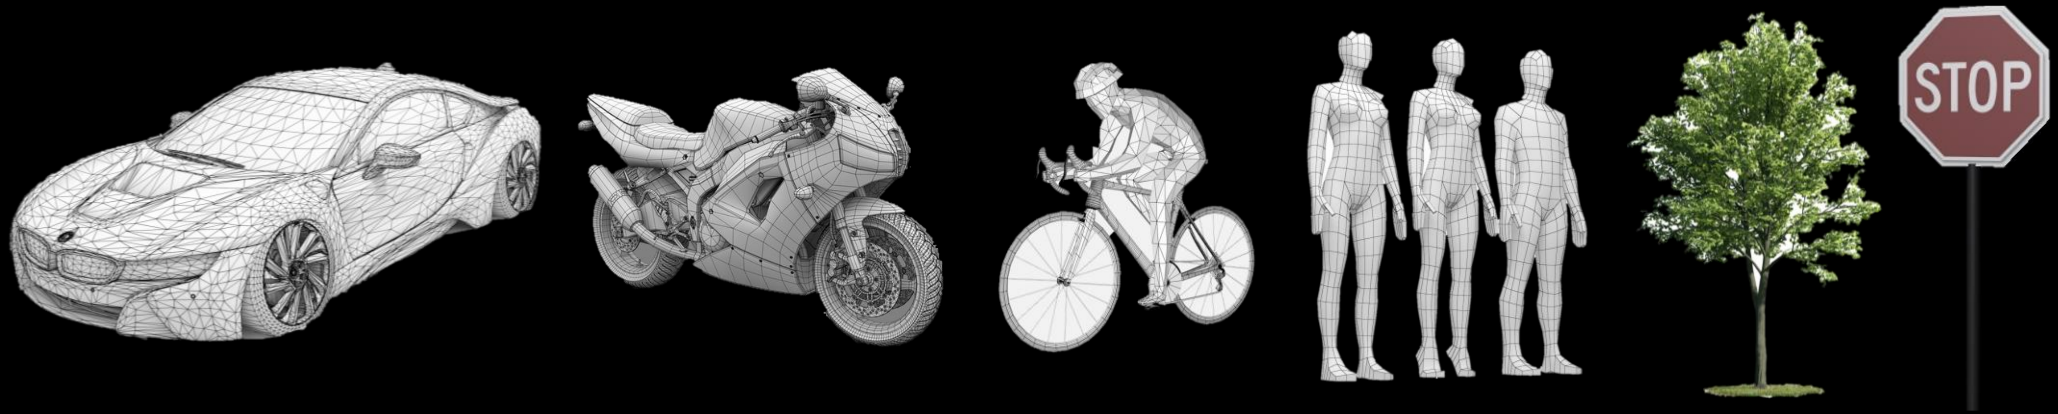

Deep Neural Networks are effective for various computer vision tasks: super-resolution, learning image prior, image style transformation, etc.

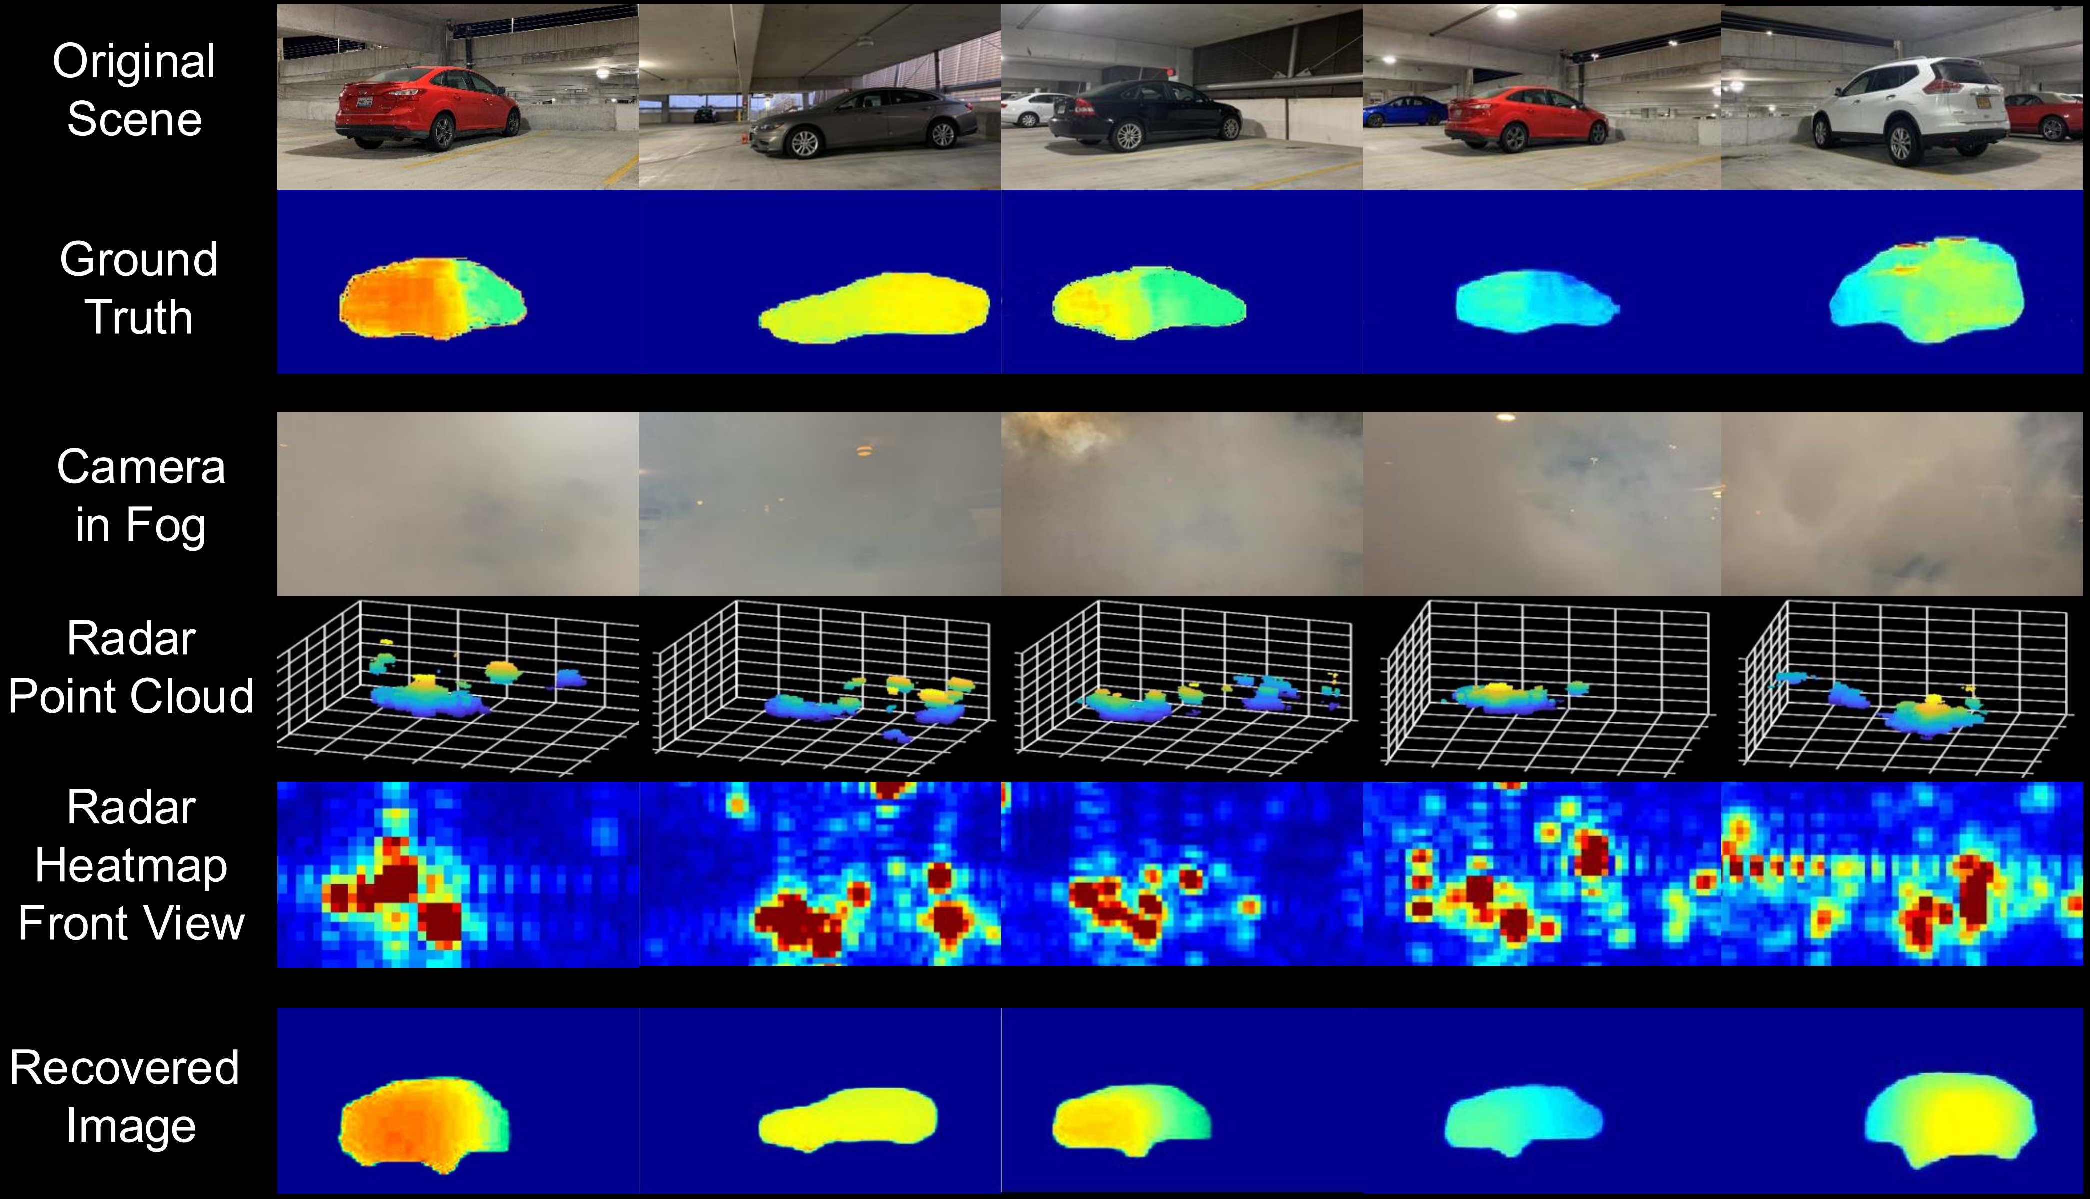

Trained using simulated data and tested using real data.

2D Arrays

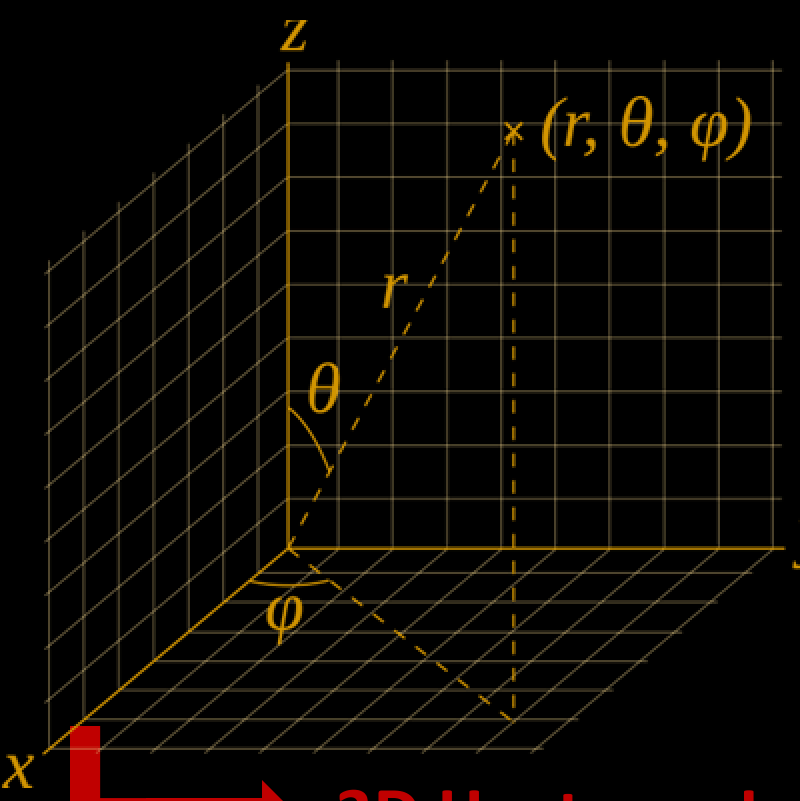

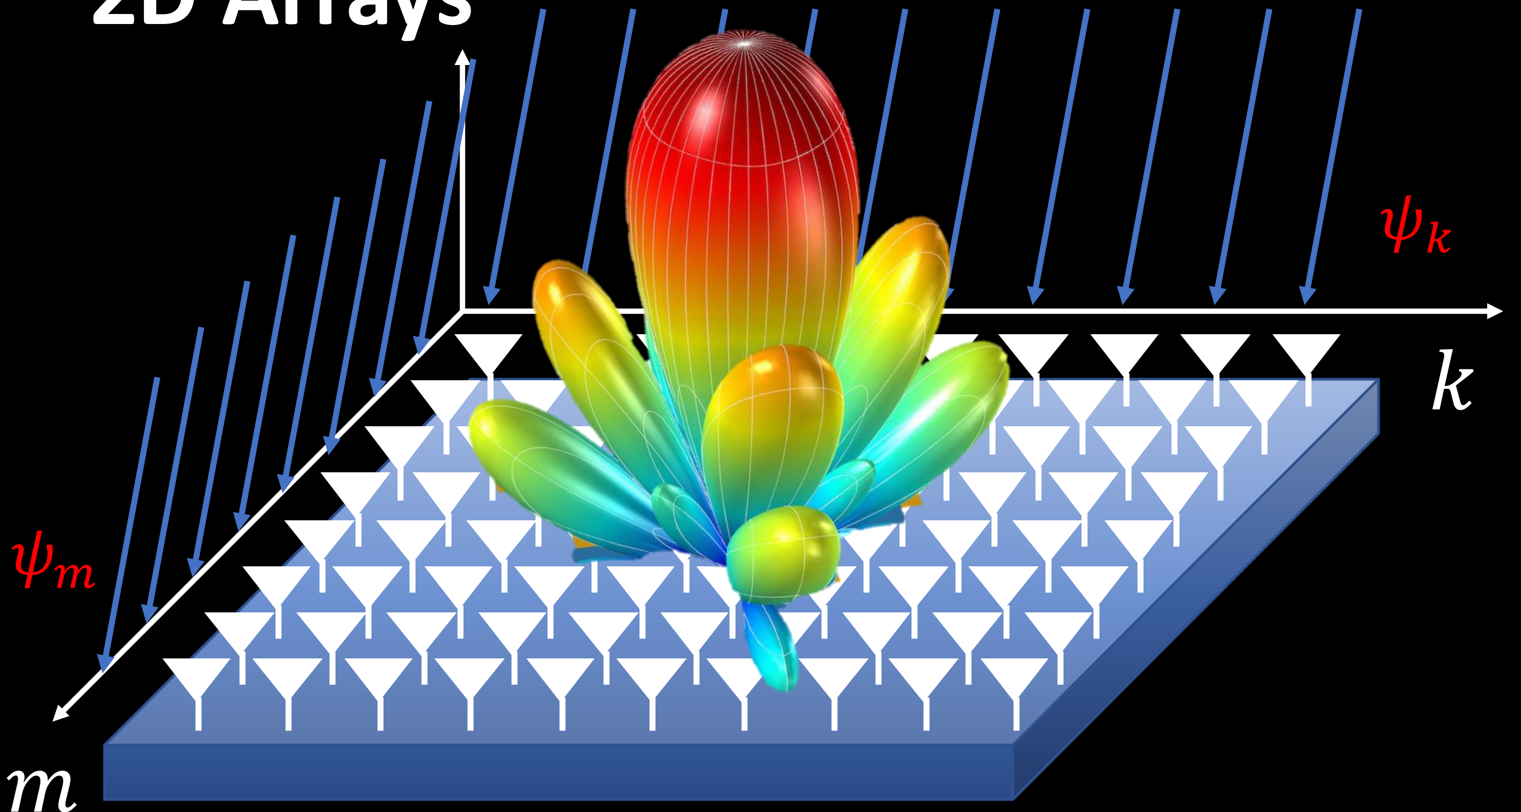

3D Heatmap Image

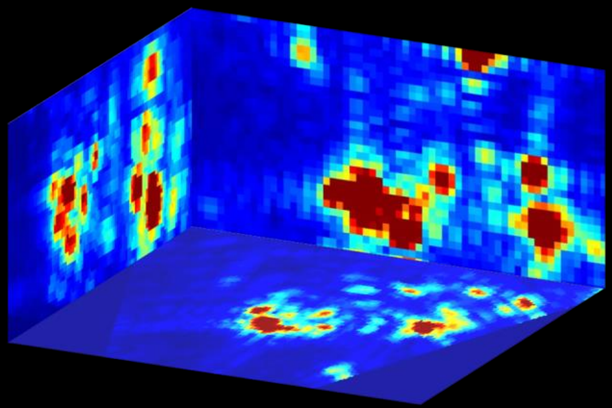

3D Point Cloud

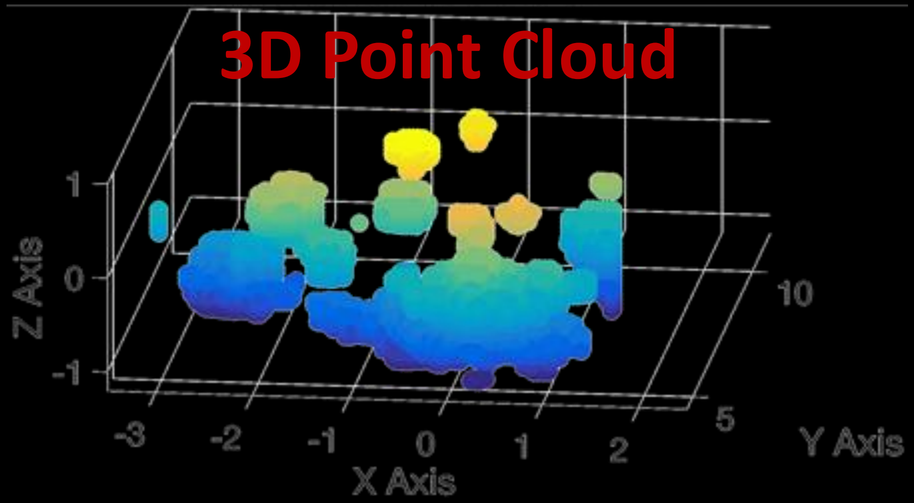

3D shape Reconstruction from Autonomous Driving Radars

Low 3D resolution

Combining two 1D antenna arrays!

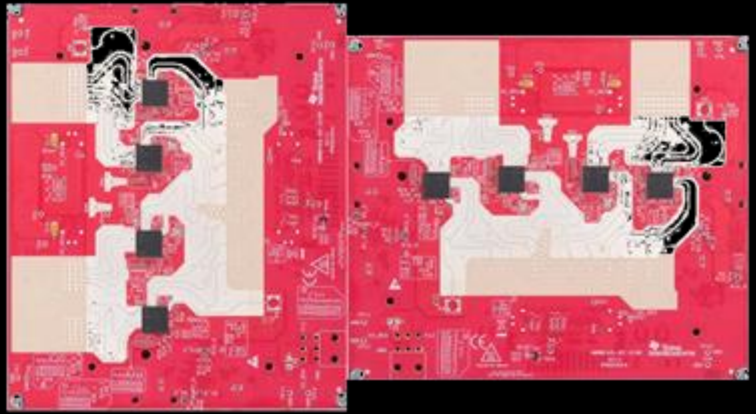

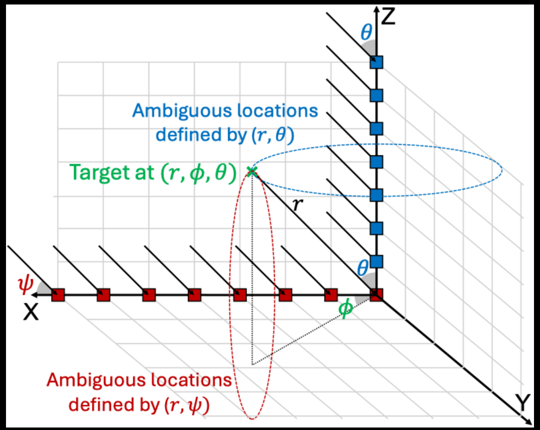

Specularity of reflections

Spatiotemporal fusion

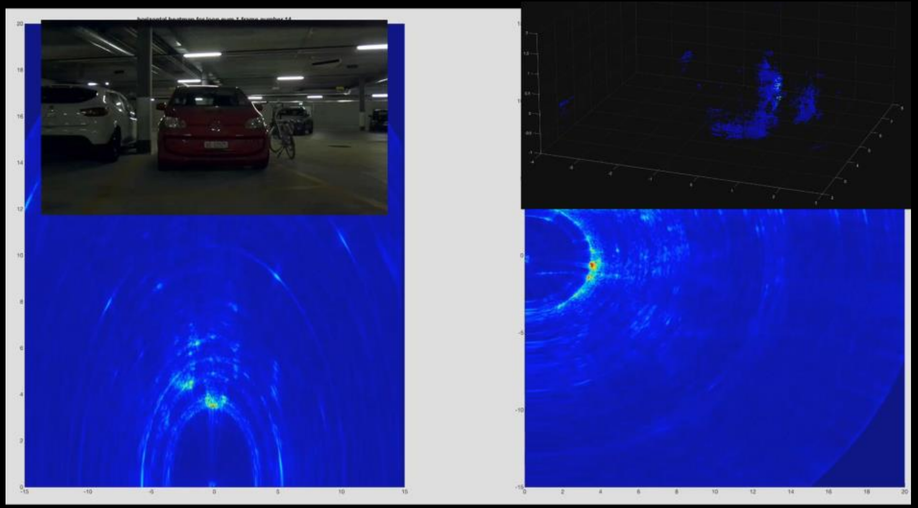

Point sparsity & practical scenarios

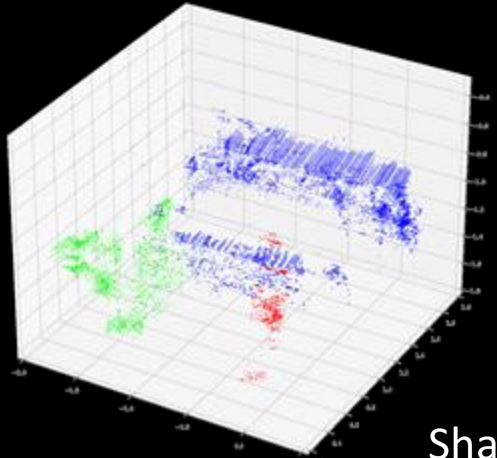

Shape Completion Network

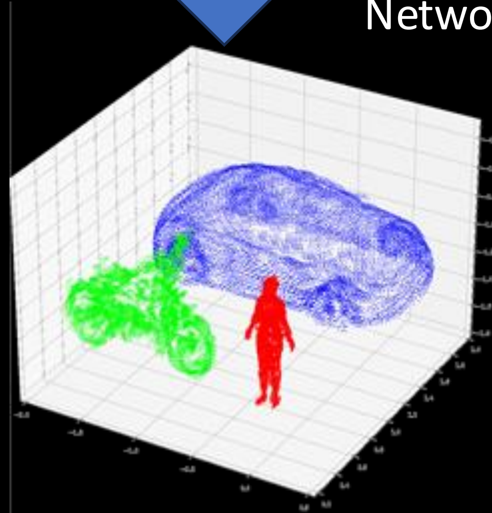

3D shape Reconstruction from Autonomous Driving Radars

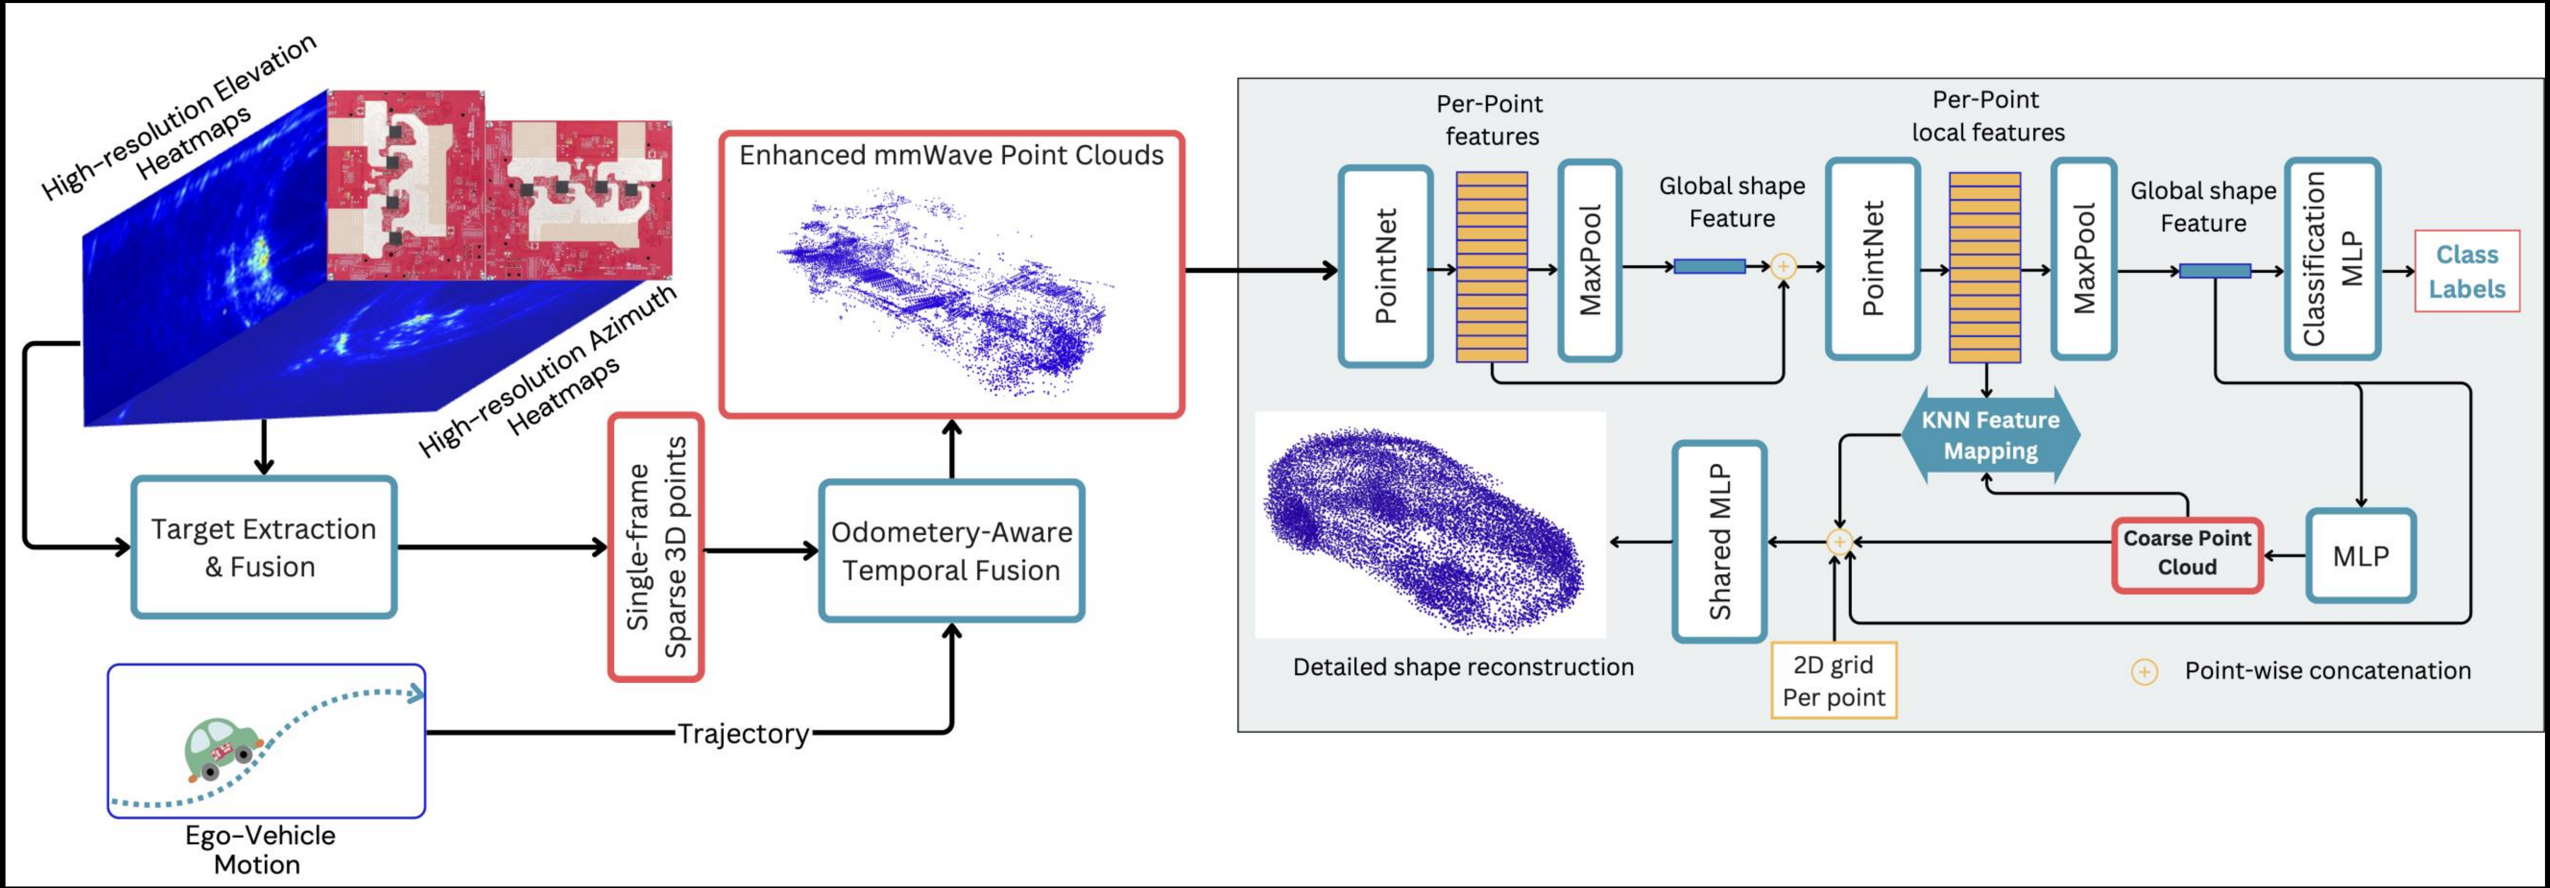

3D shape Reconstruction from Autonomous Driving Radars

Fused Radar Points

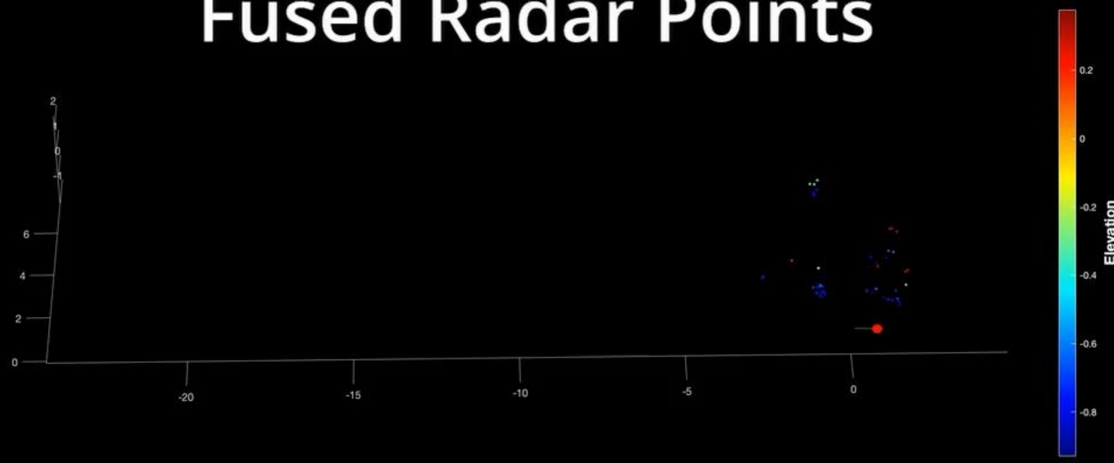

Radar setup

Our Reconstruction

Processed up to 8m range

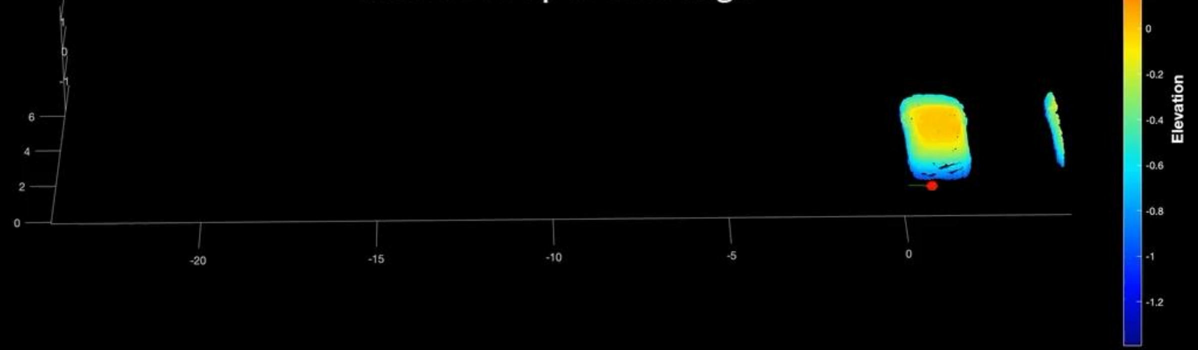

3D shape Reconstruction from Autonomous Driving Radars

Fused Radar Points

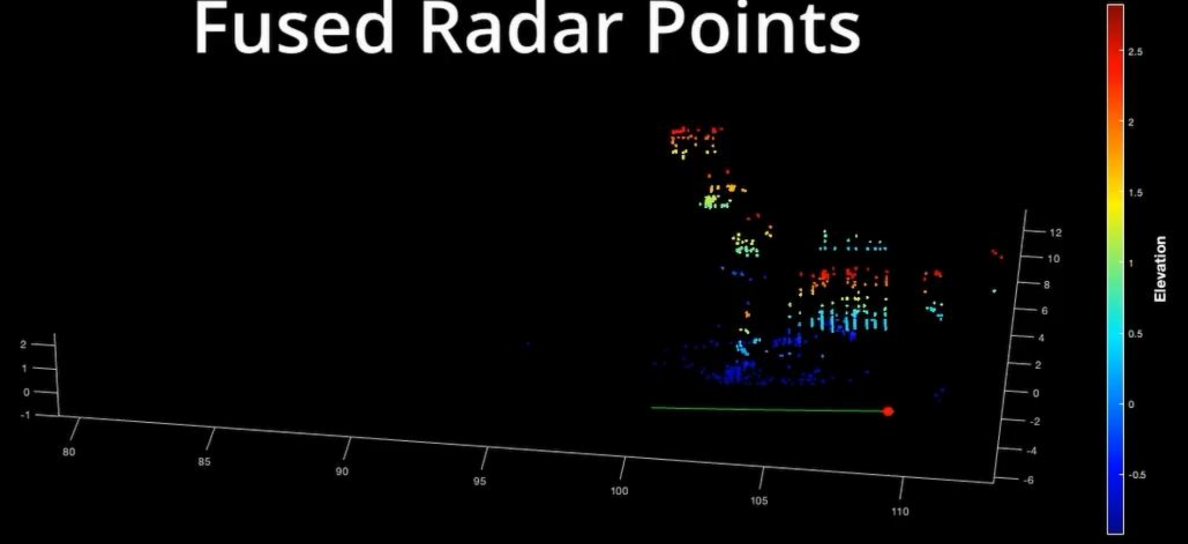

Radar setup

Our Reconstruction

Processed up to 8m range

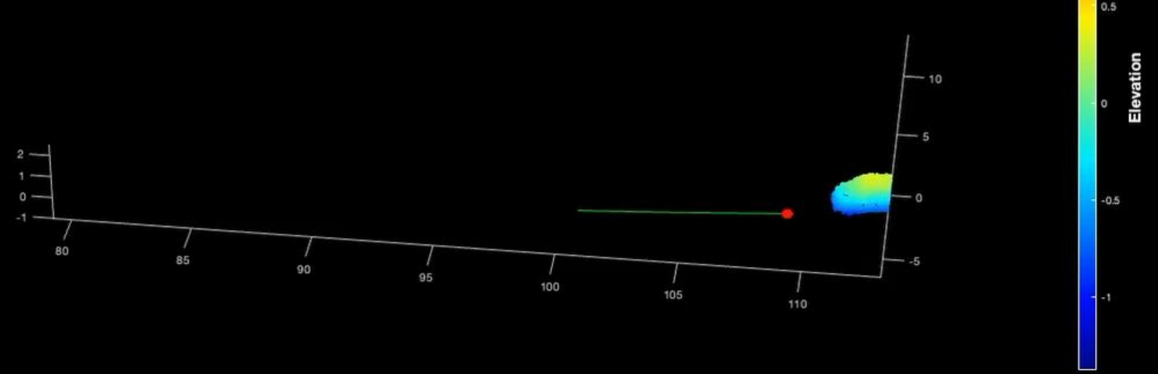

3D shape Reconstruction from Autonomous Driving Radars

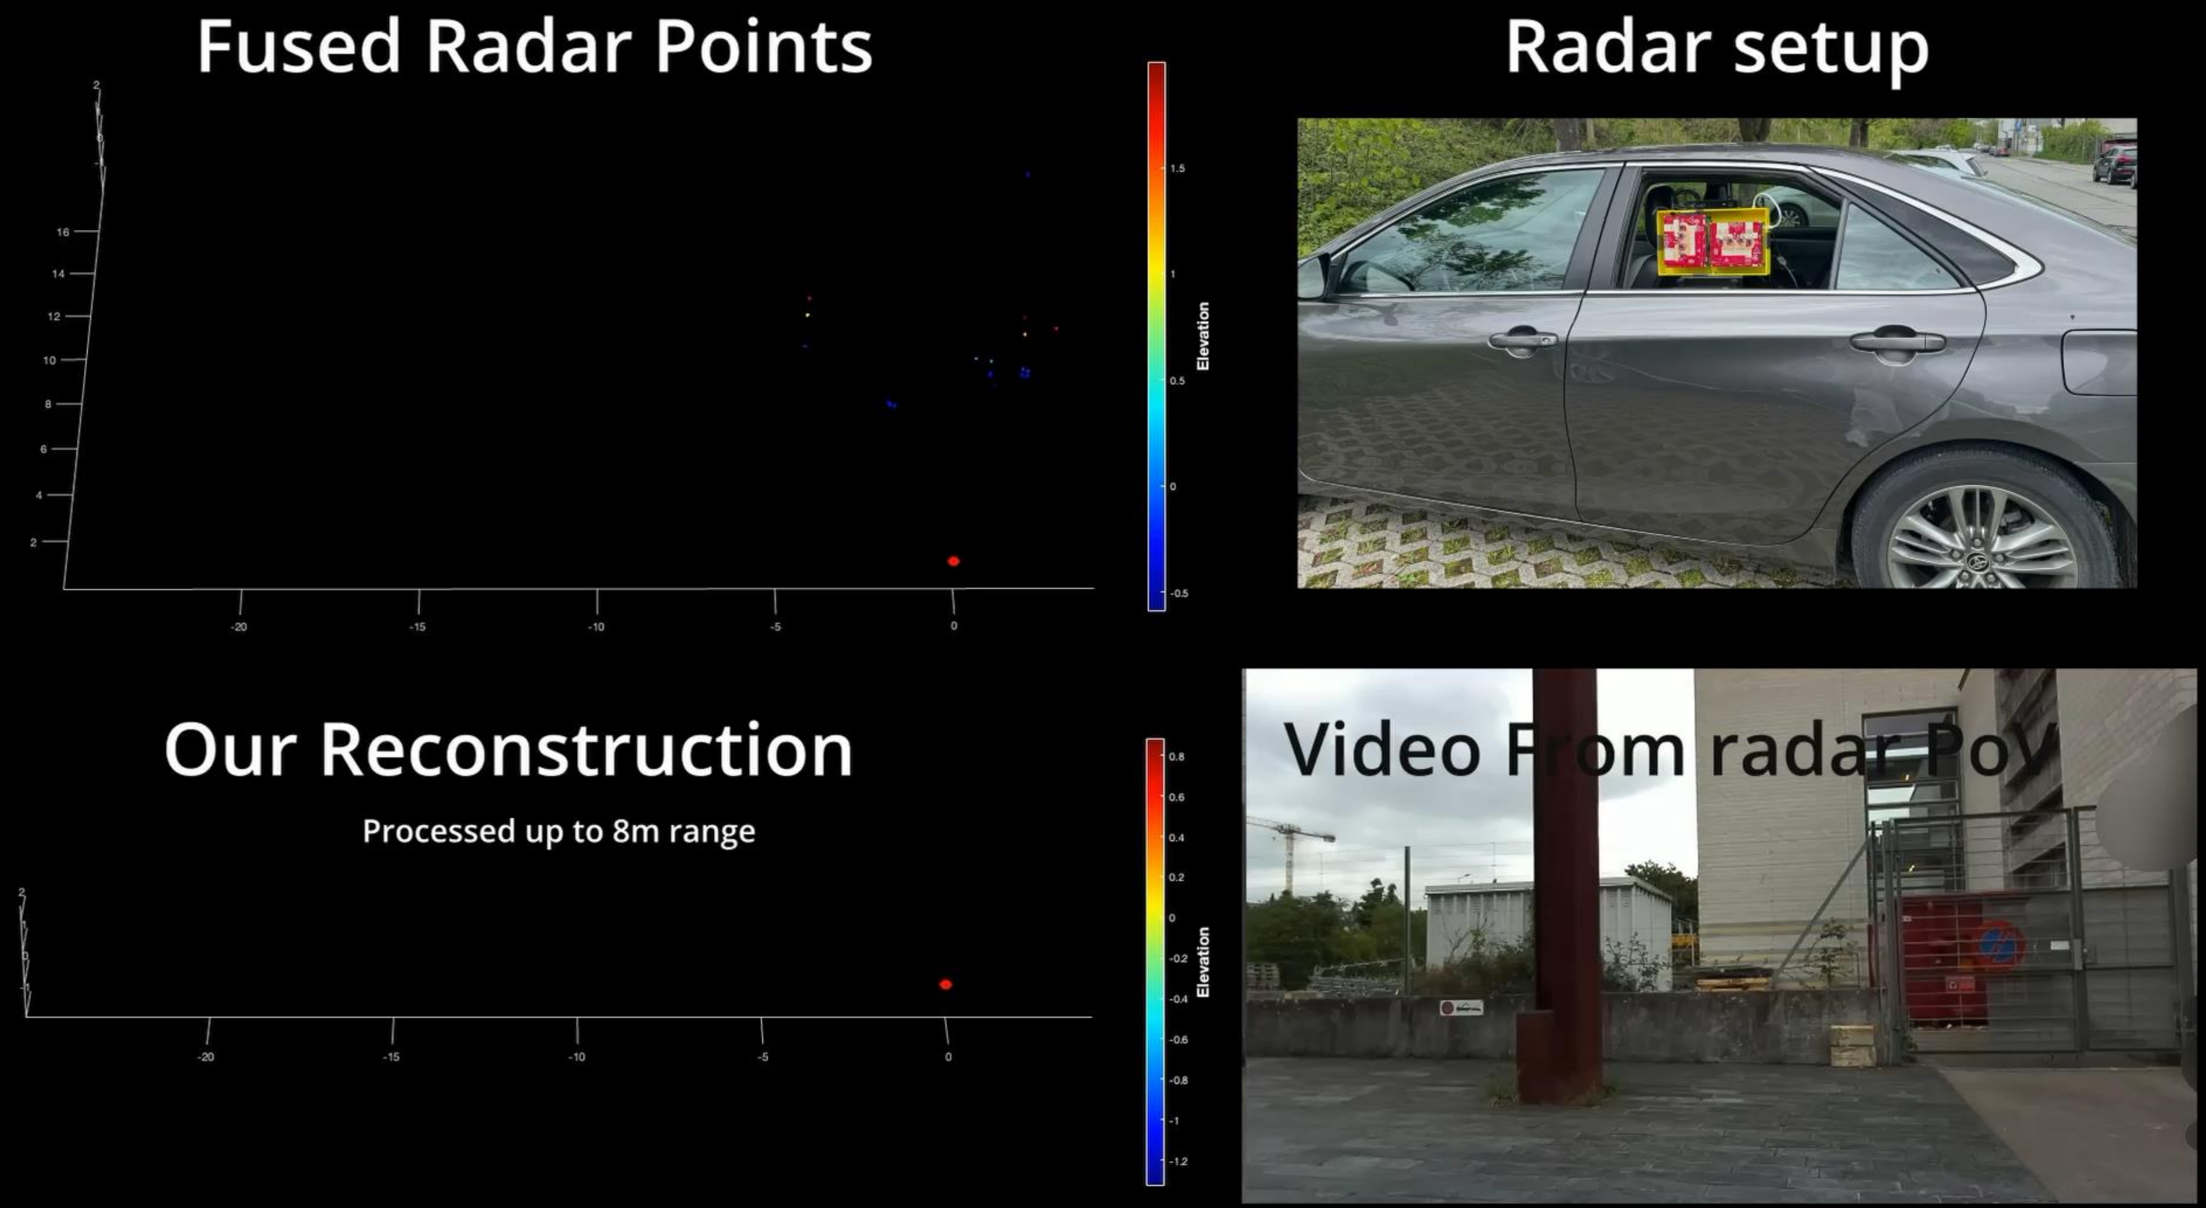

Classic Sensing Modalities are Limited to Line-of-Sight

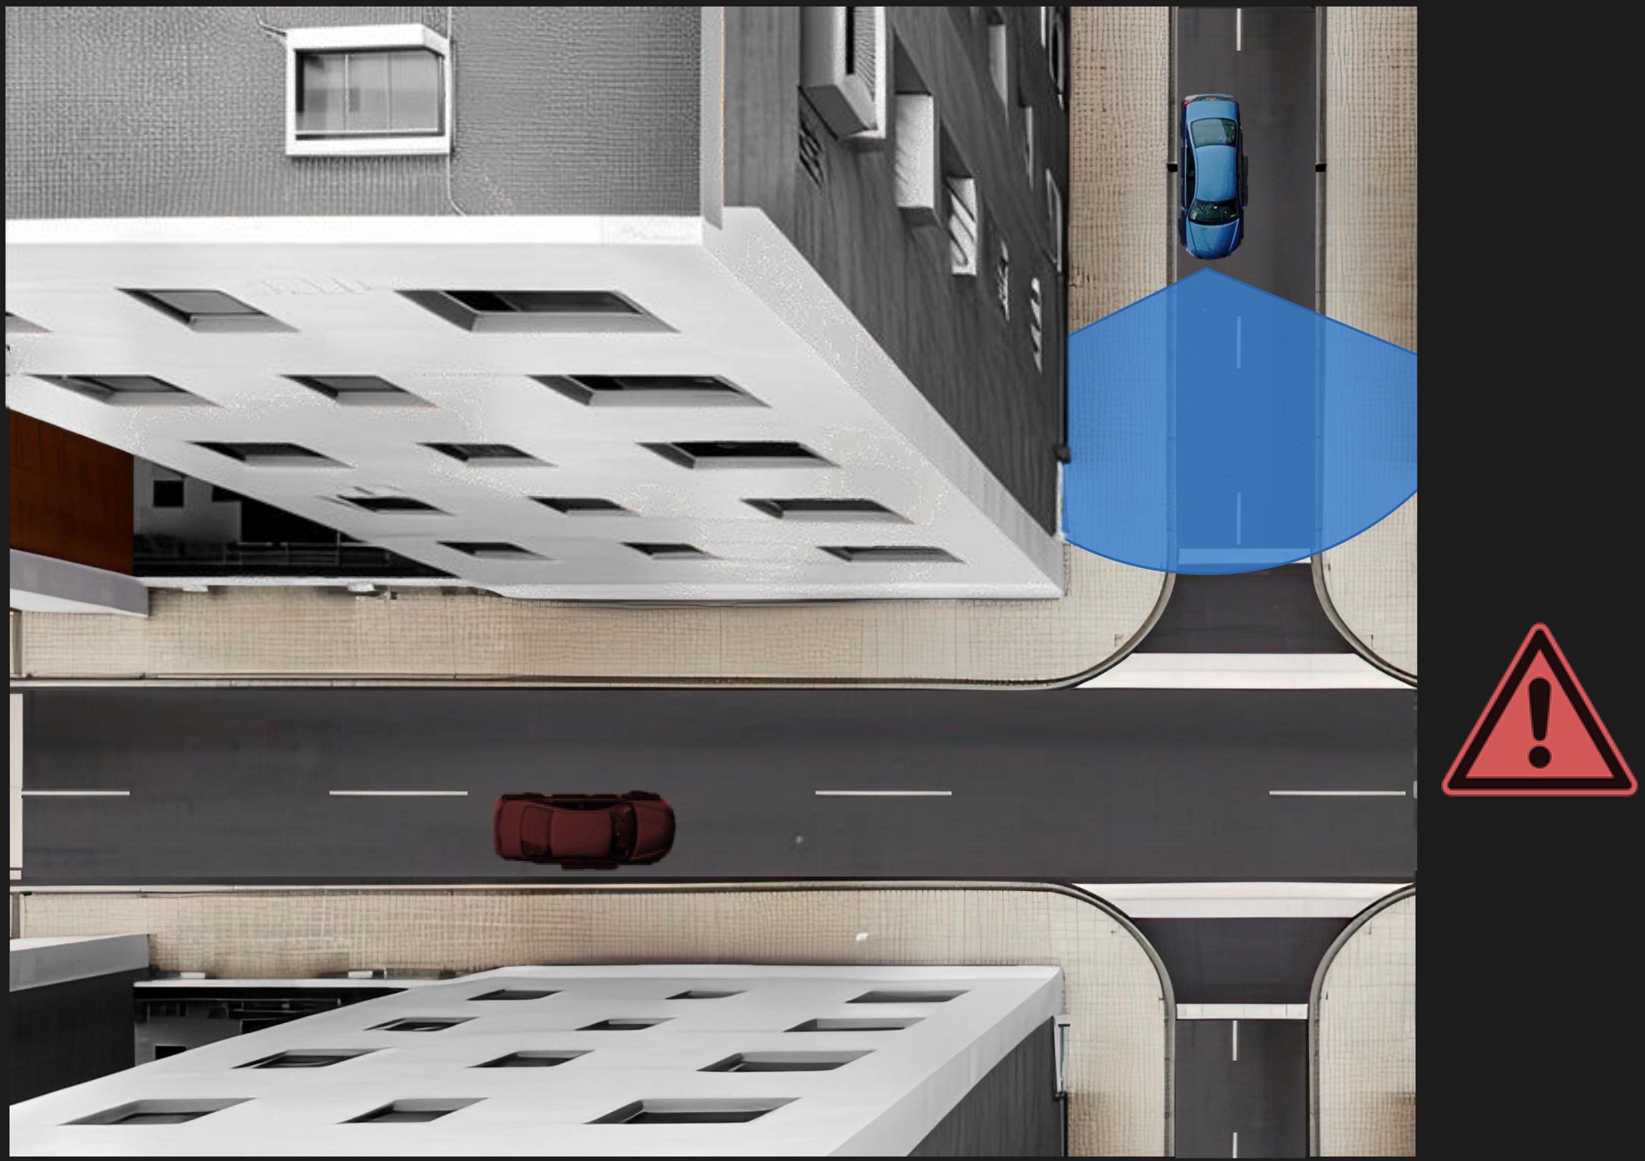

Can We Produce RF Images of Objects Around-the-Corner?

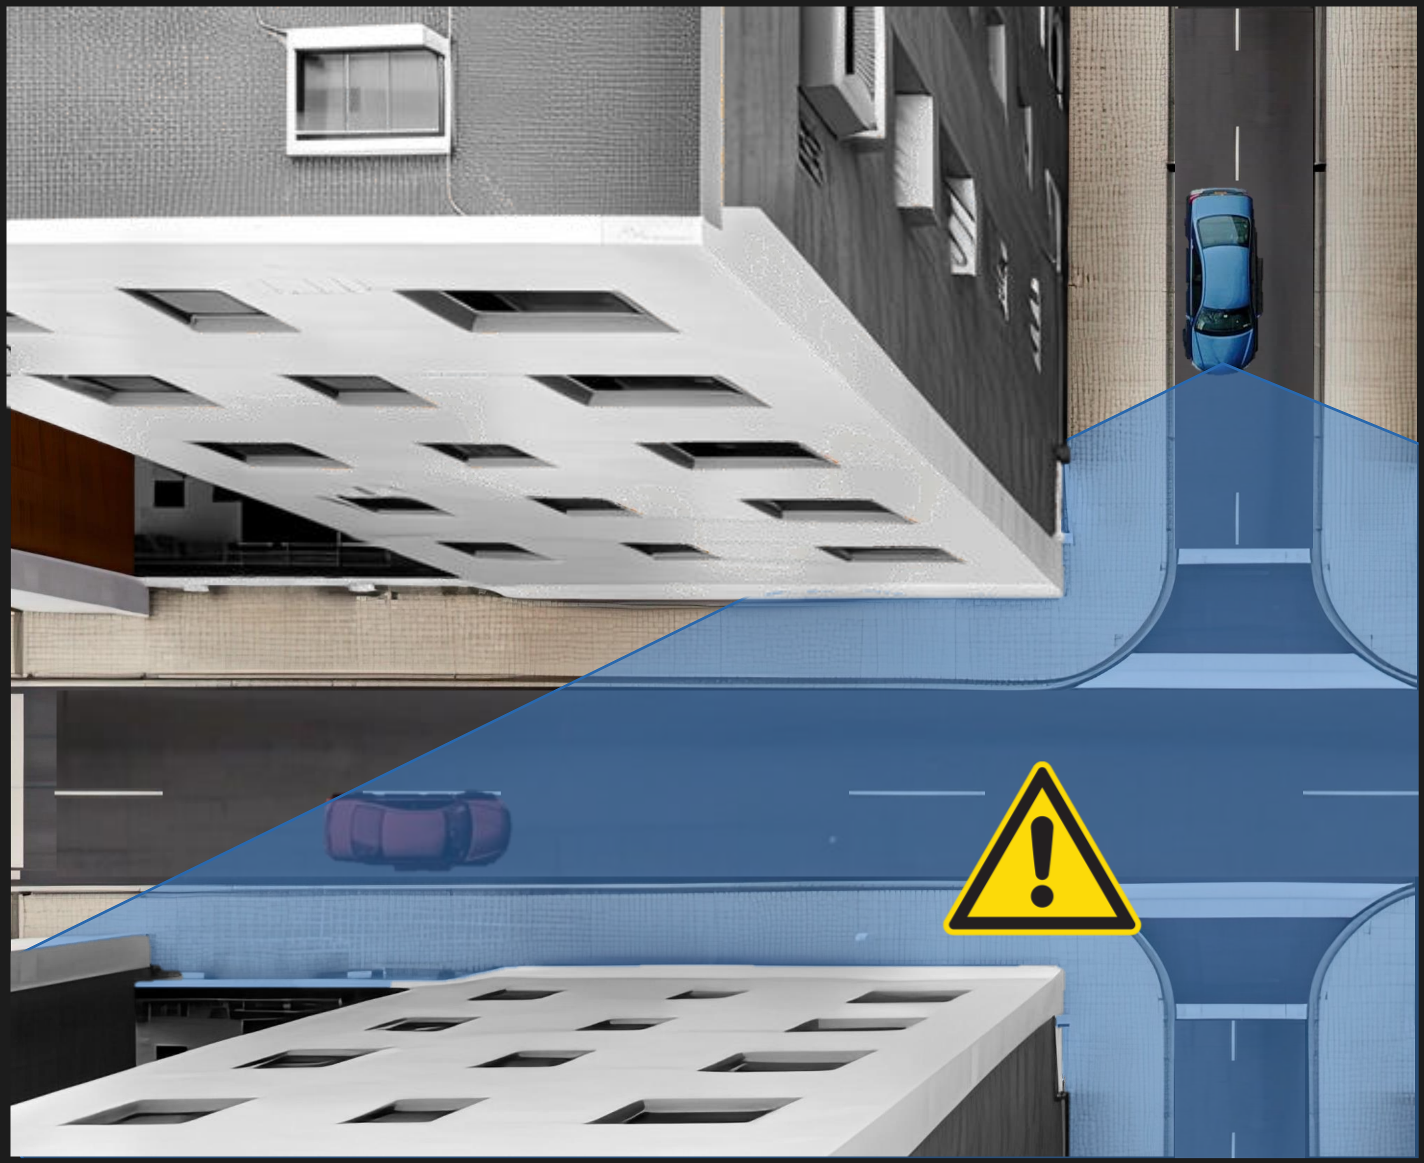

The Ability to Image Around-the-Corner has Many Applications

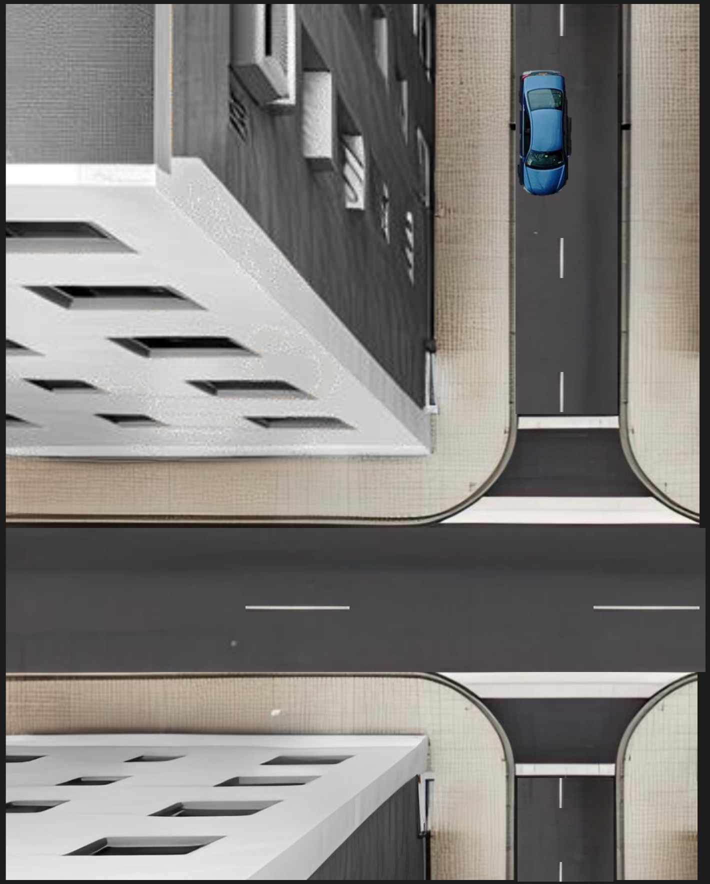

Autonomous Driving

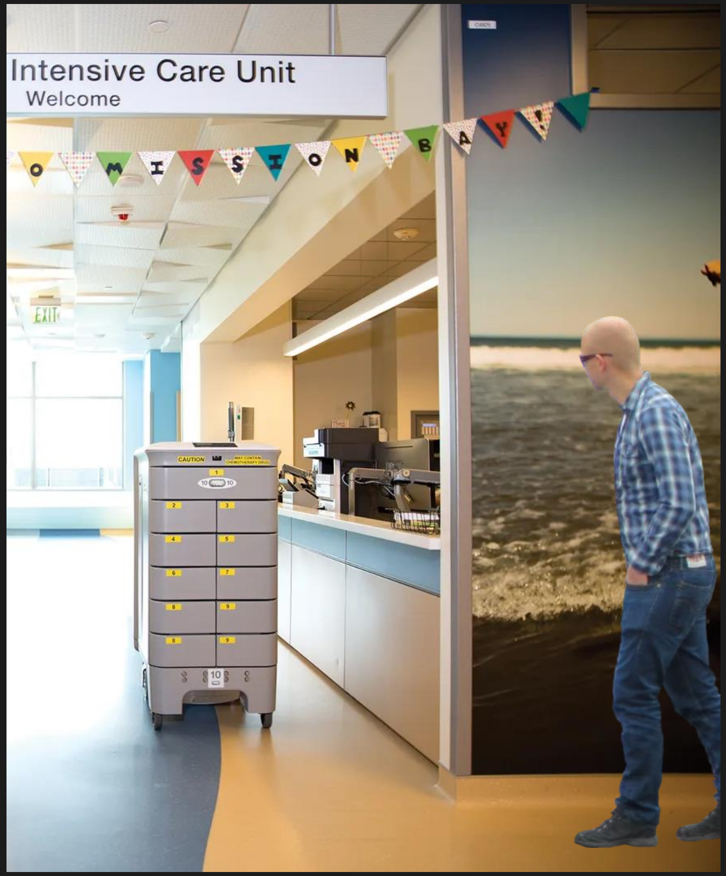

Robotic Navigation

Search and Rescue

Practical Environments Contain a Variety of Reflecting Surfaces

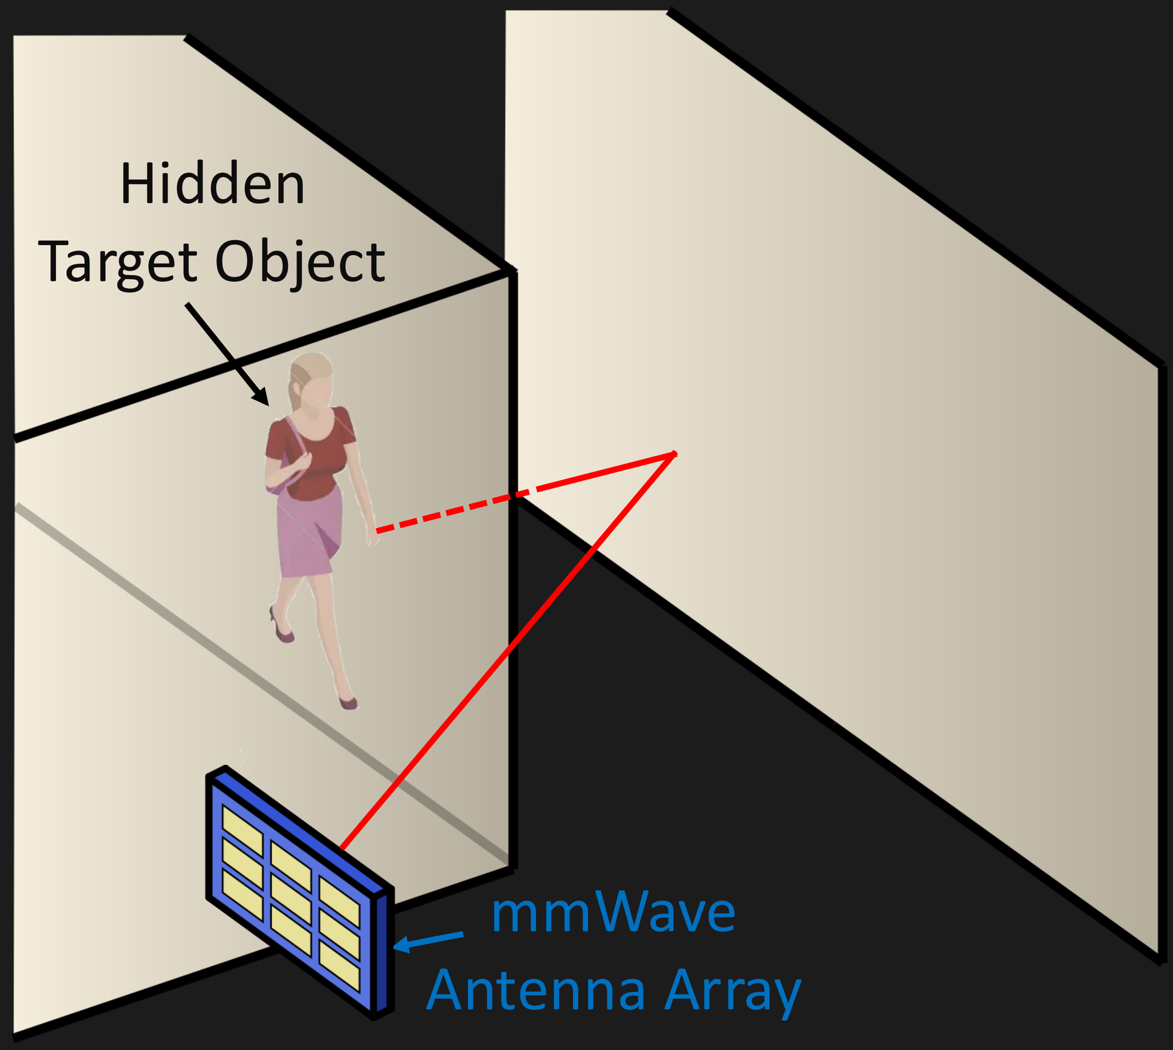

Planar Surfaces

Practical Environments Contain a Variety of Reflecting Surfaces

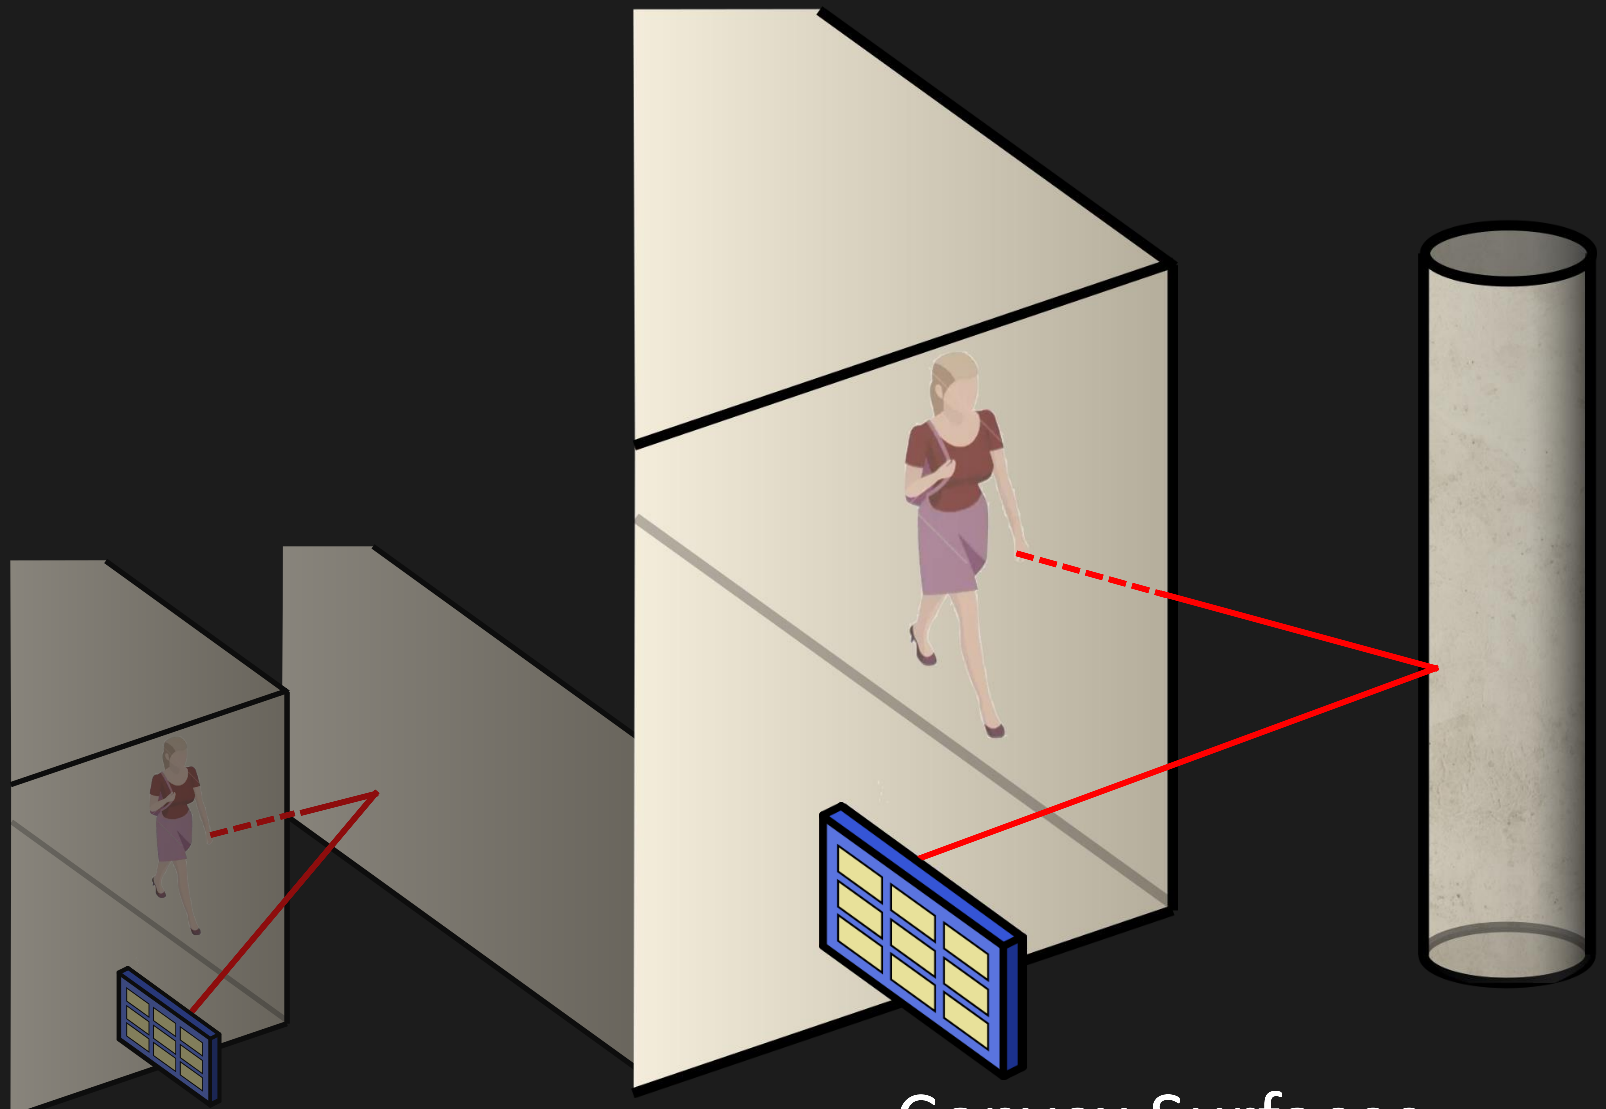

Convex Surfaces

Planar Surfaces

Practical Environments Contain a Variety of Reflecting Surfaces

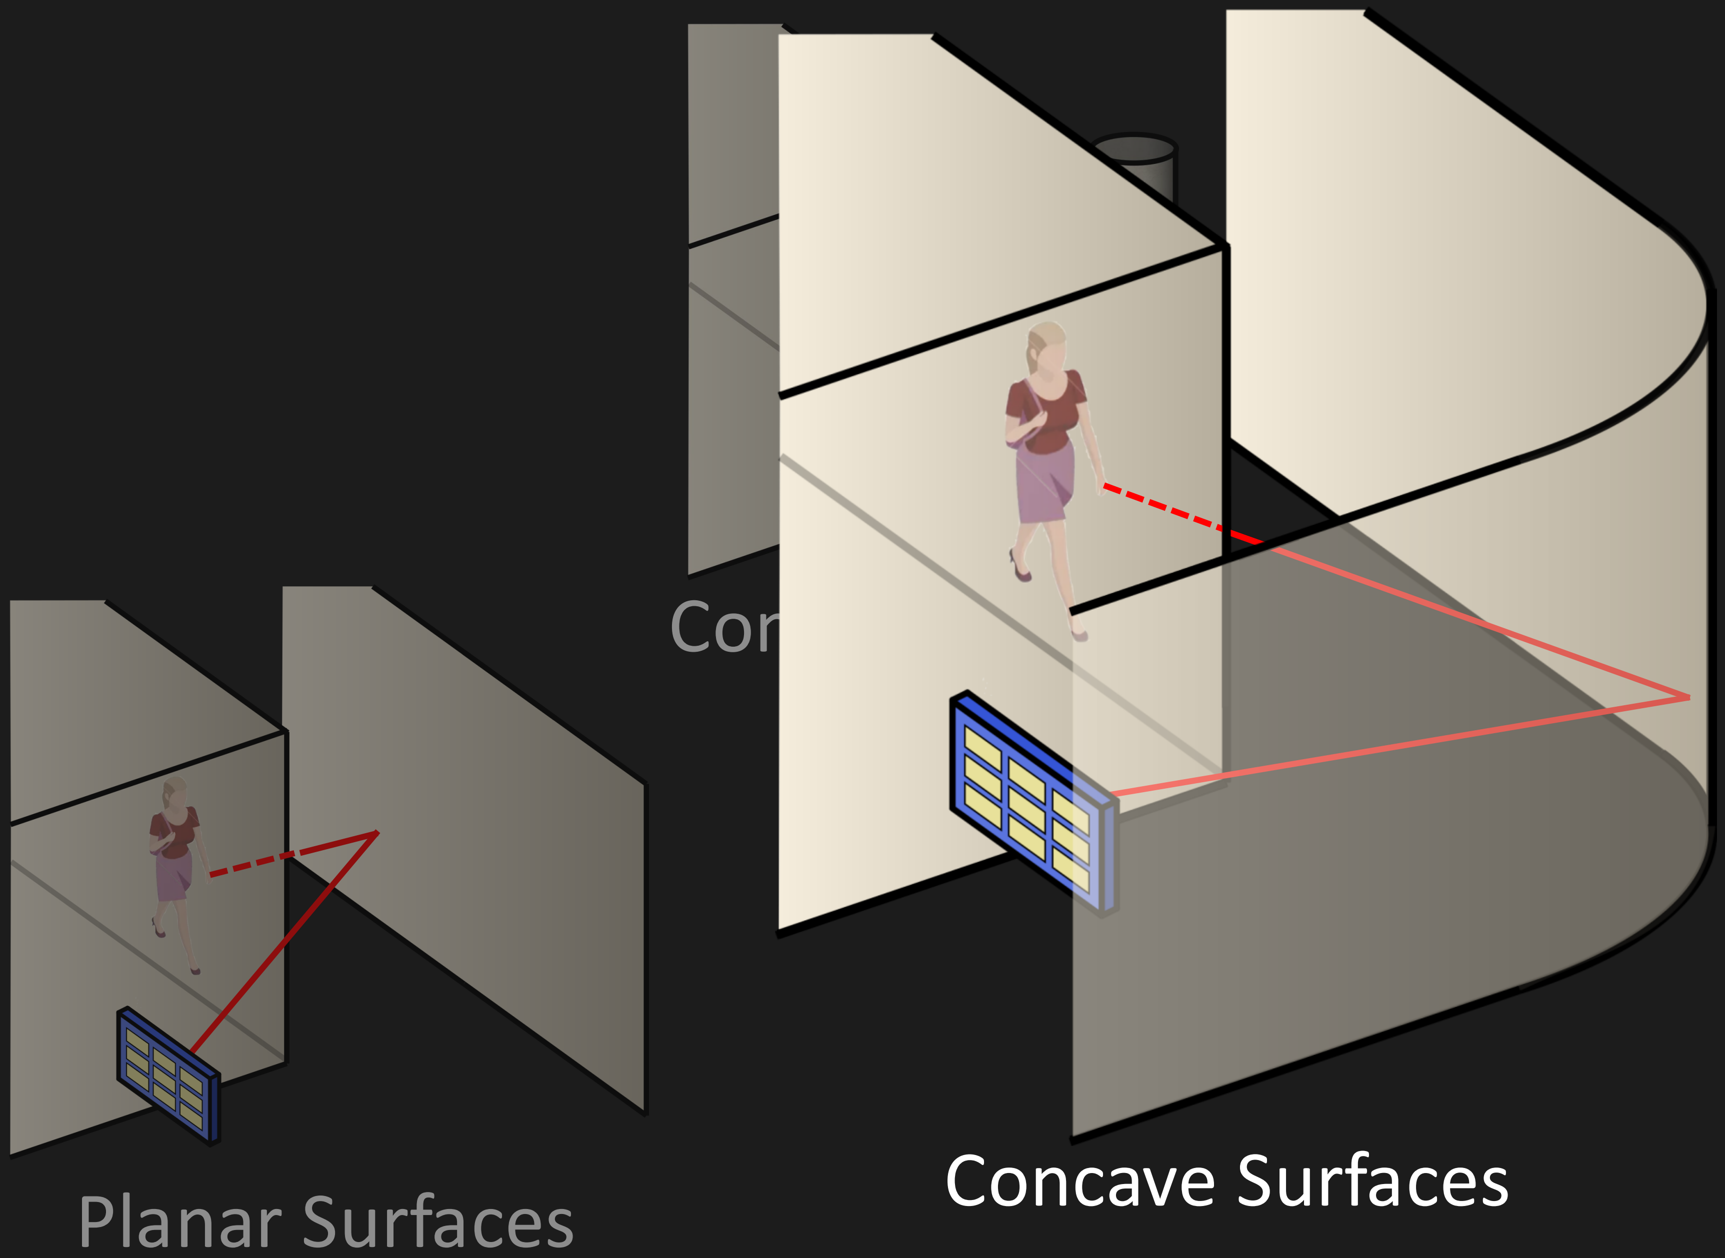

Practical Environments Contain a Variety of Reflecting Surfaces

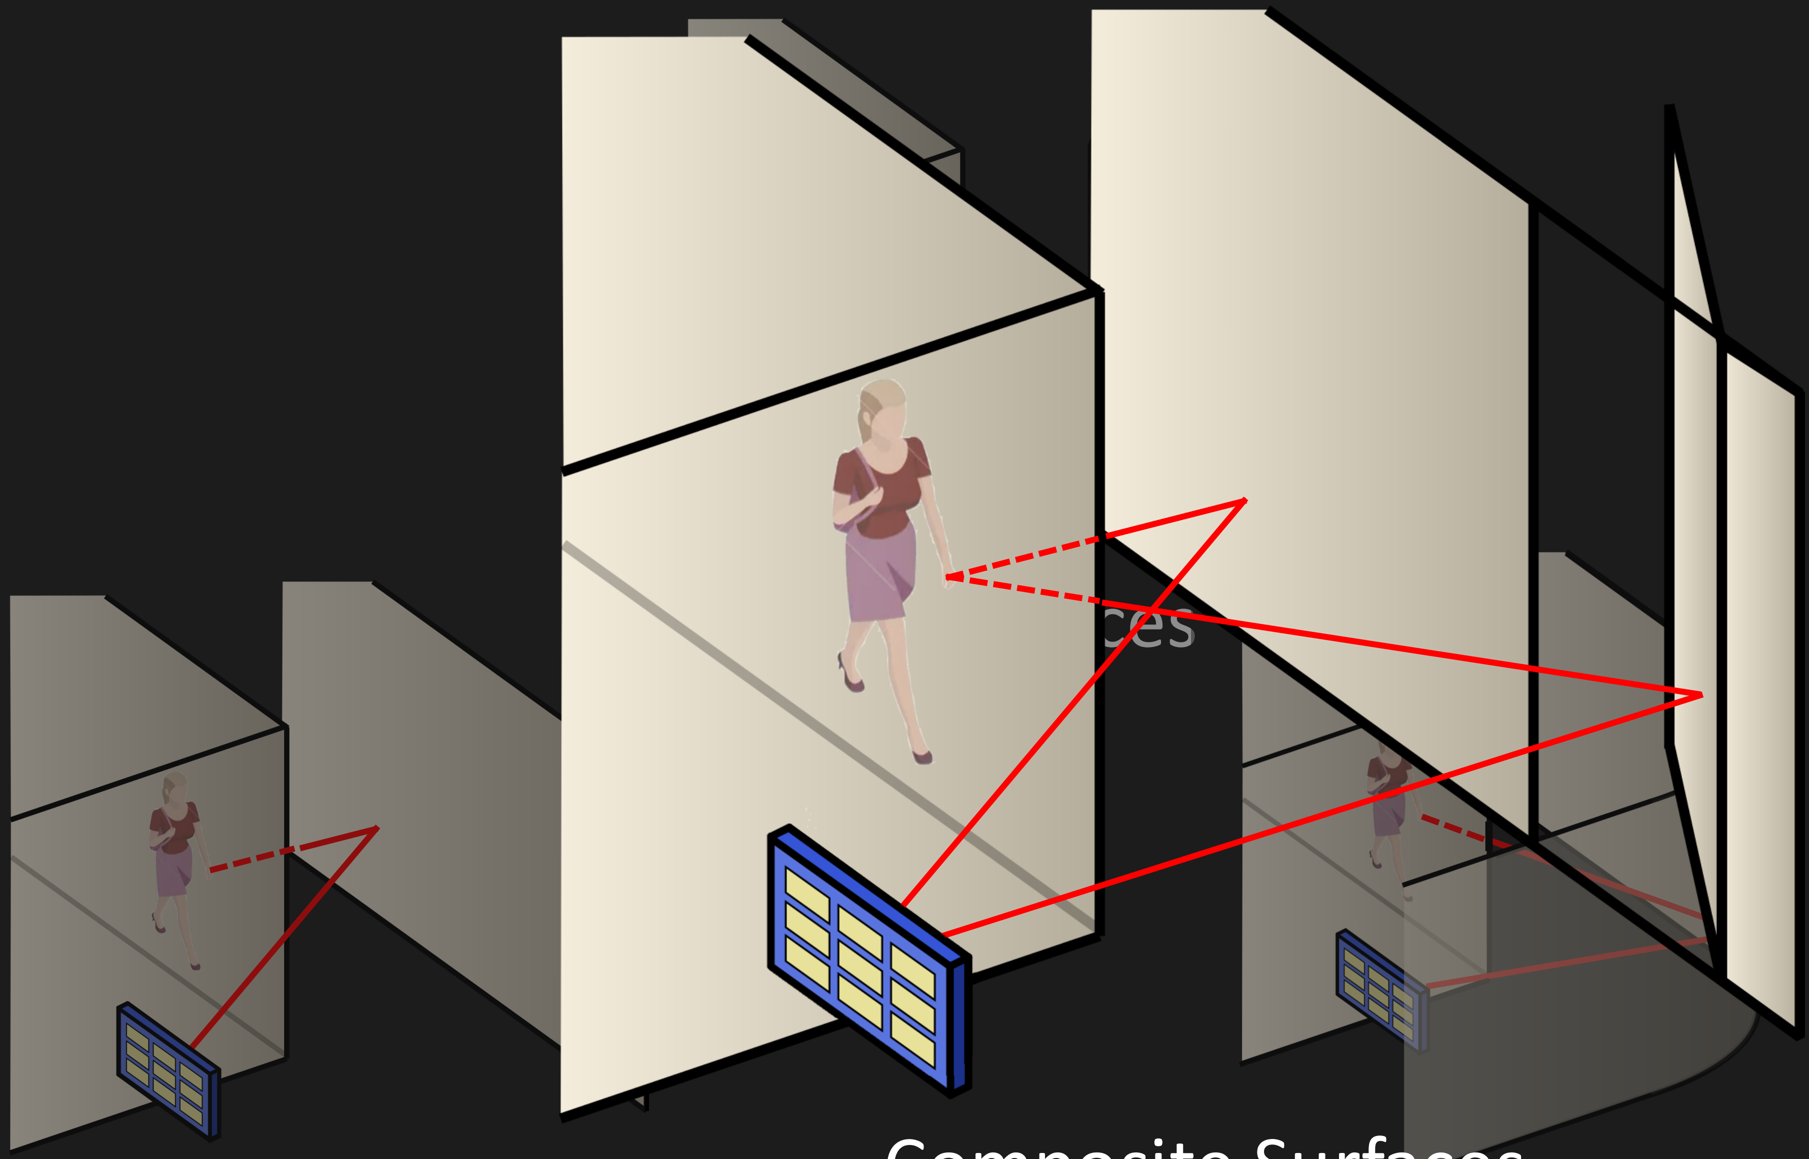

Practical Environments Contain a Variety of Reflecting Surfaces

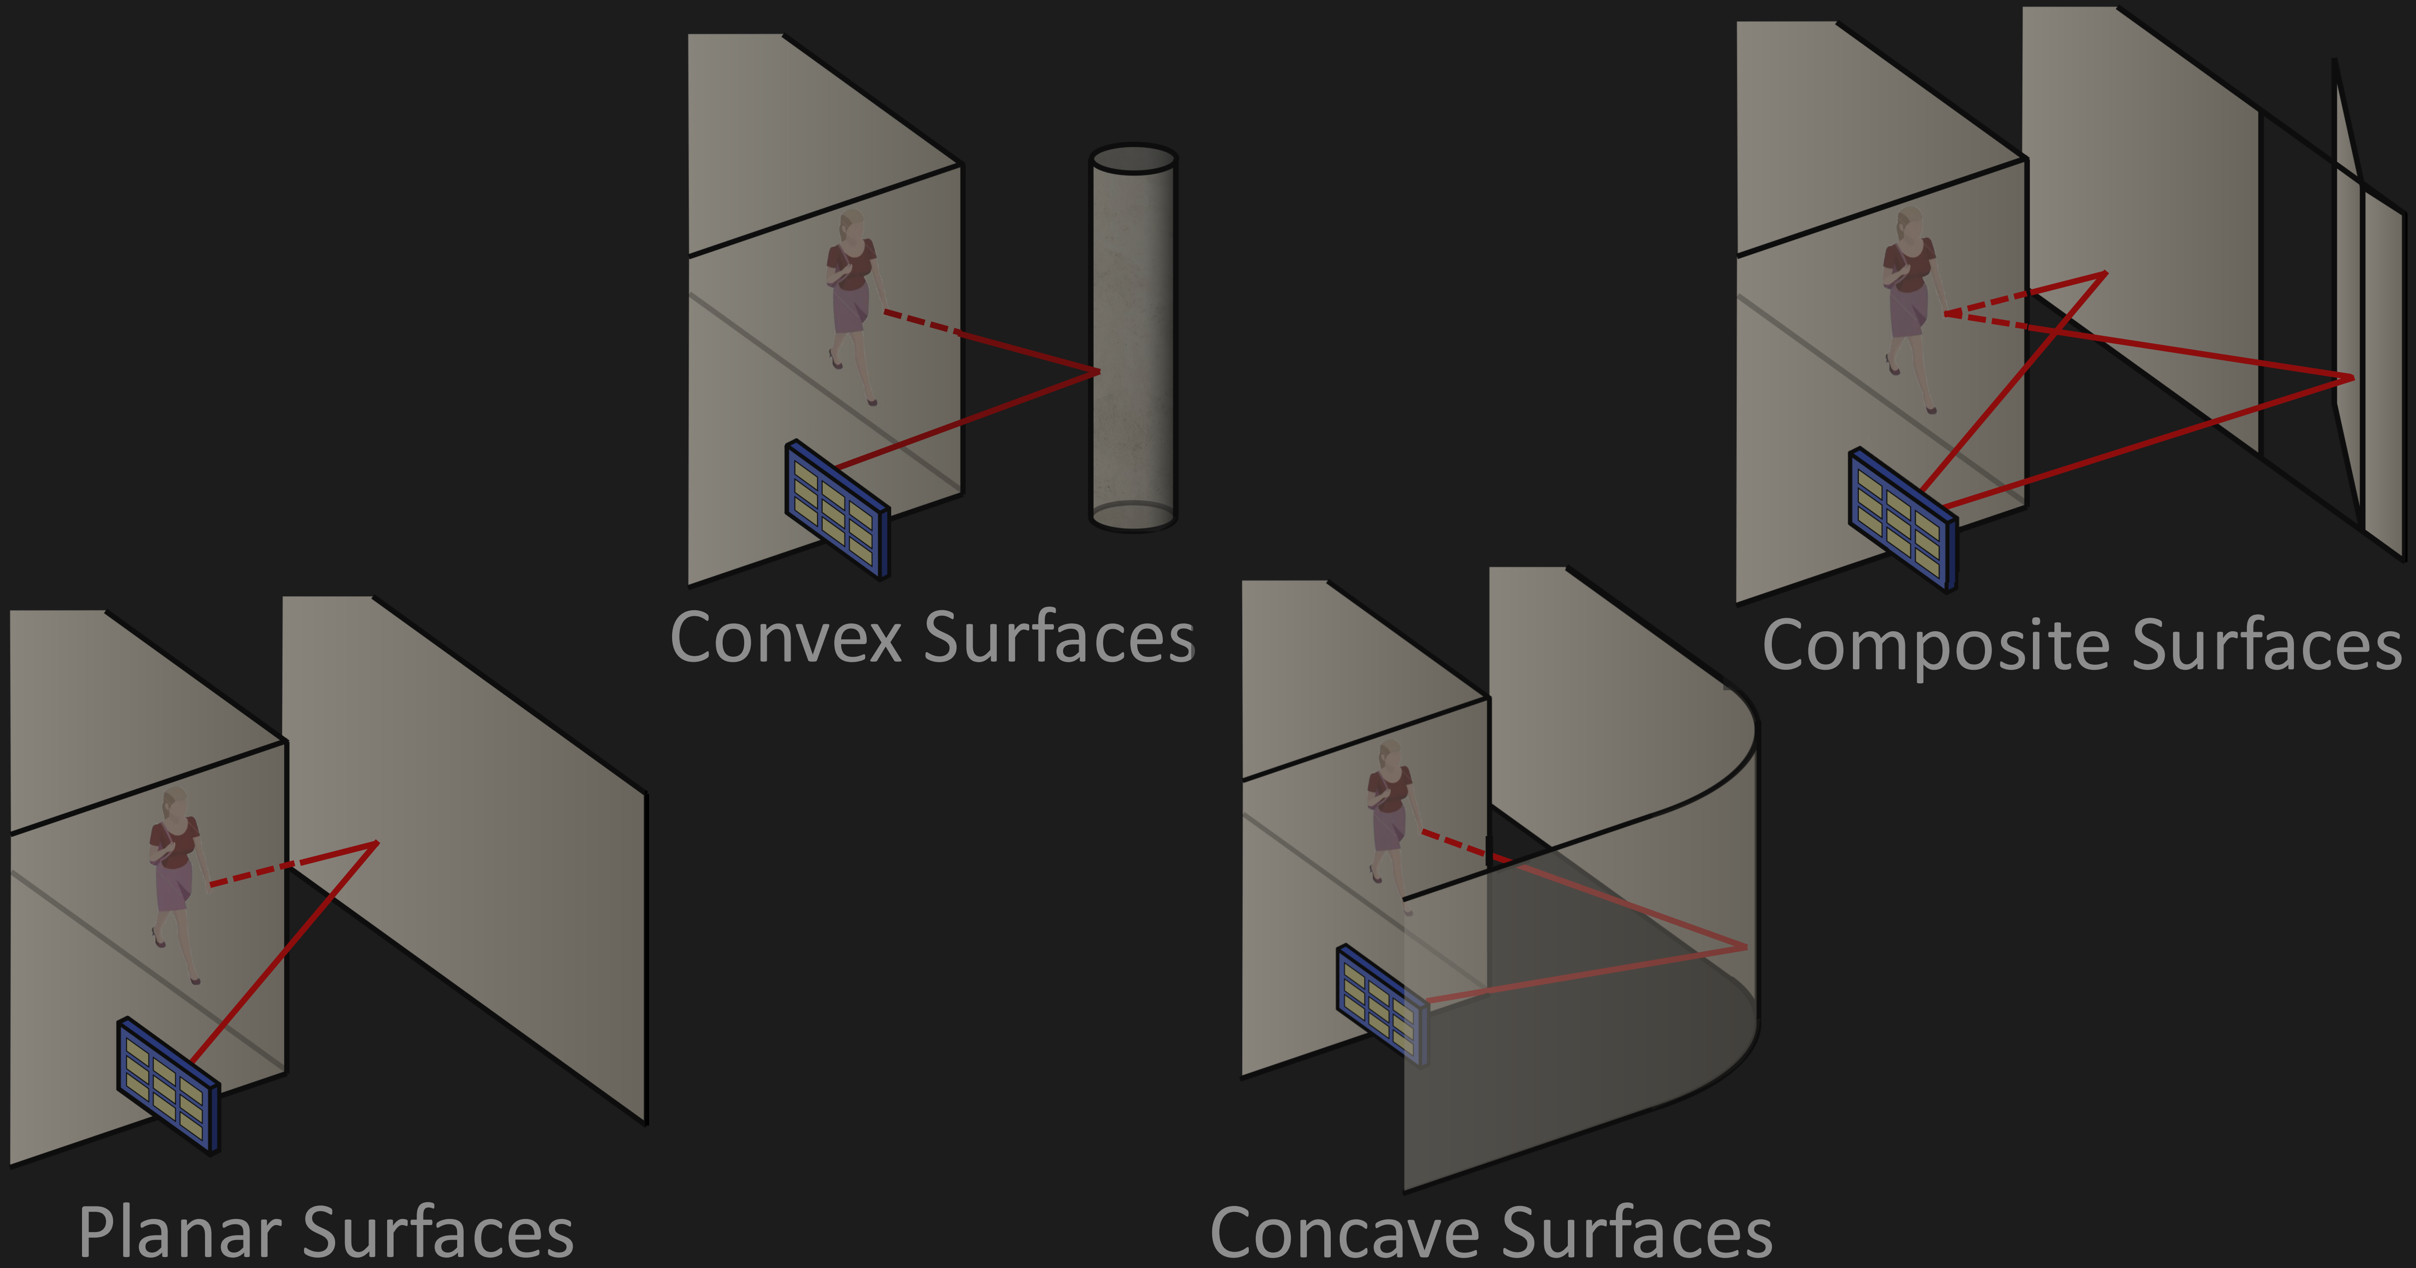

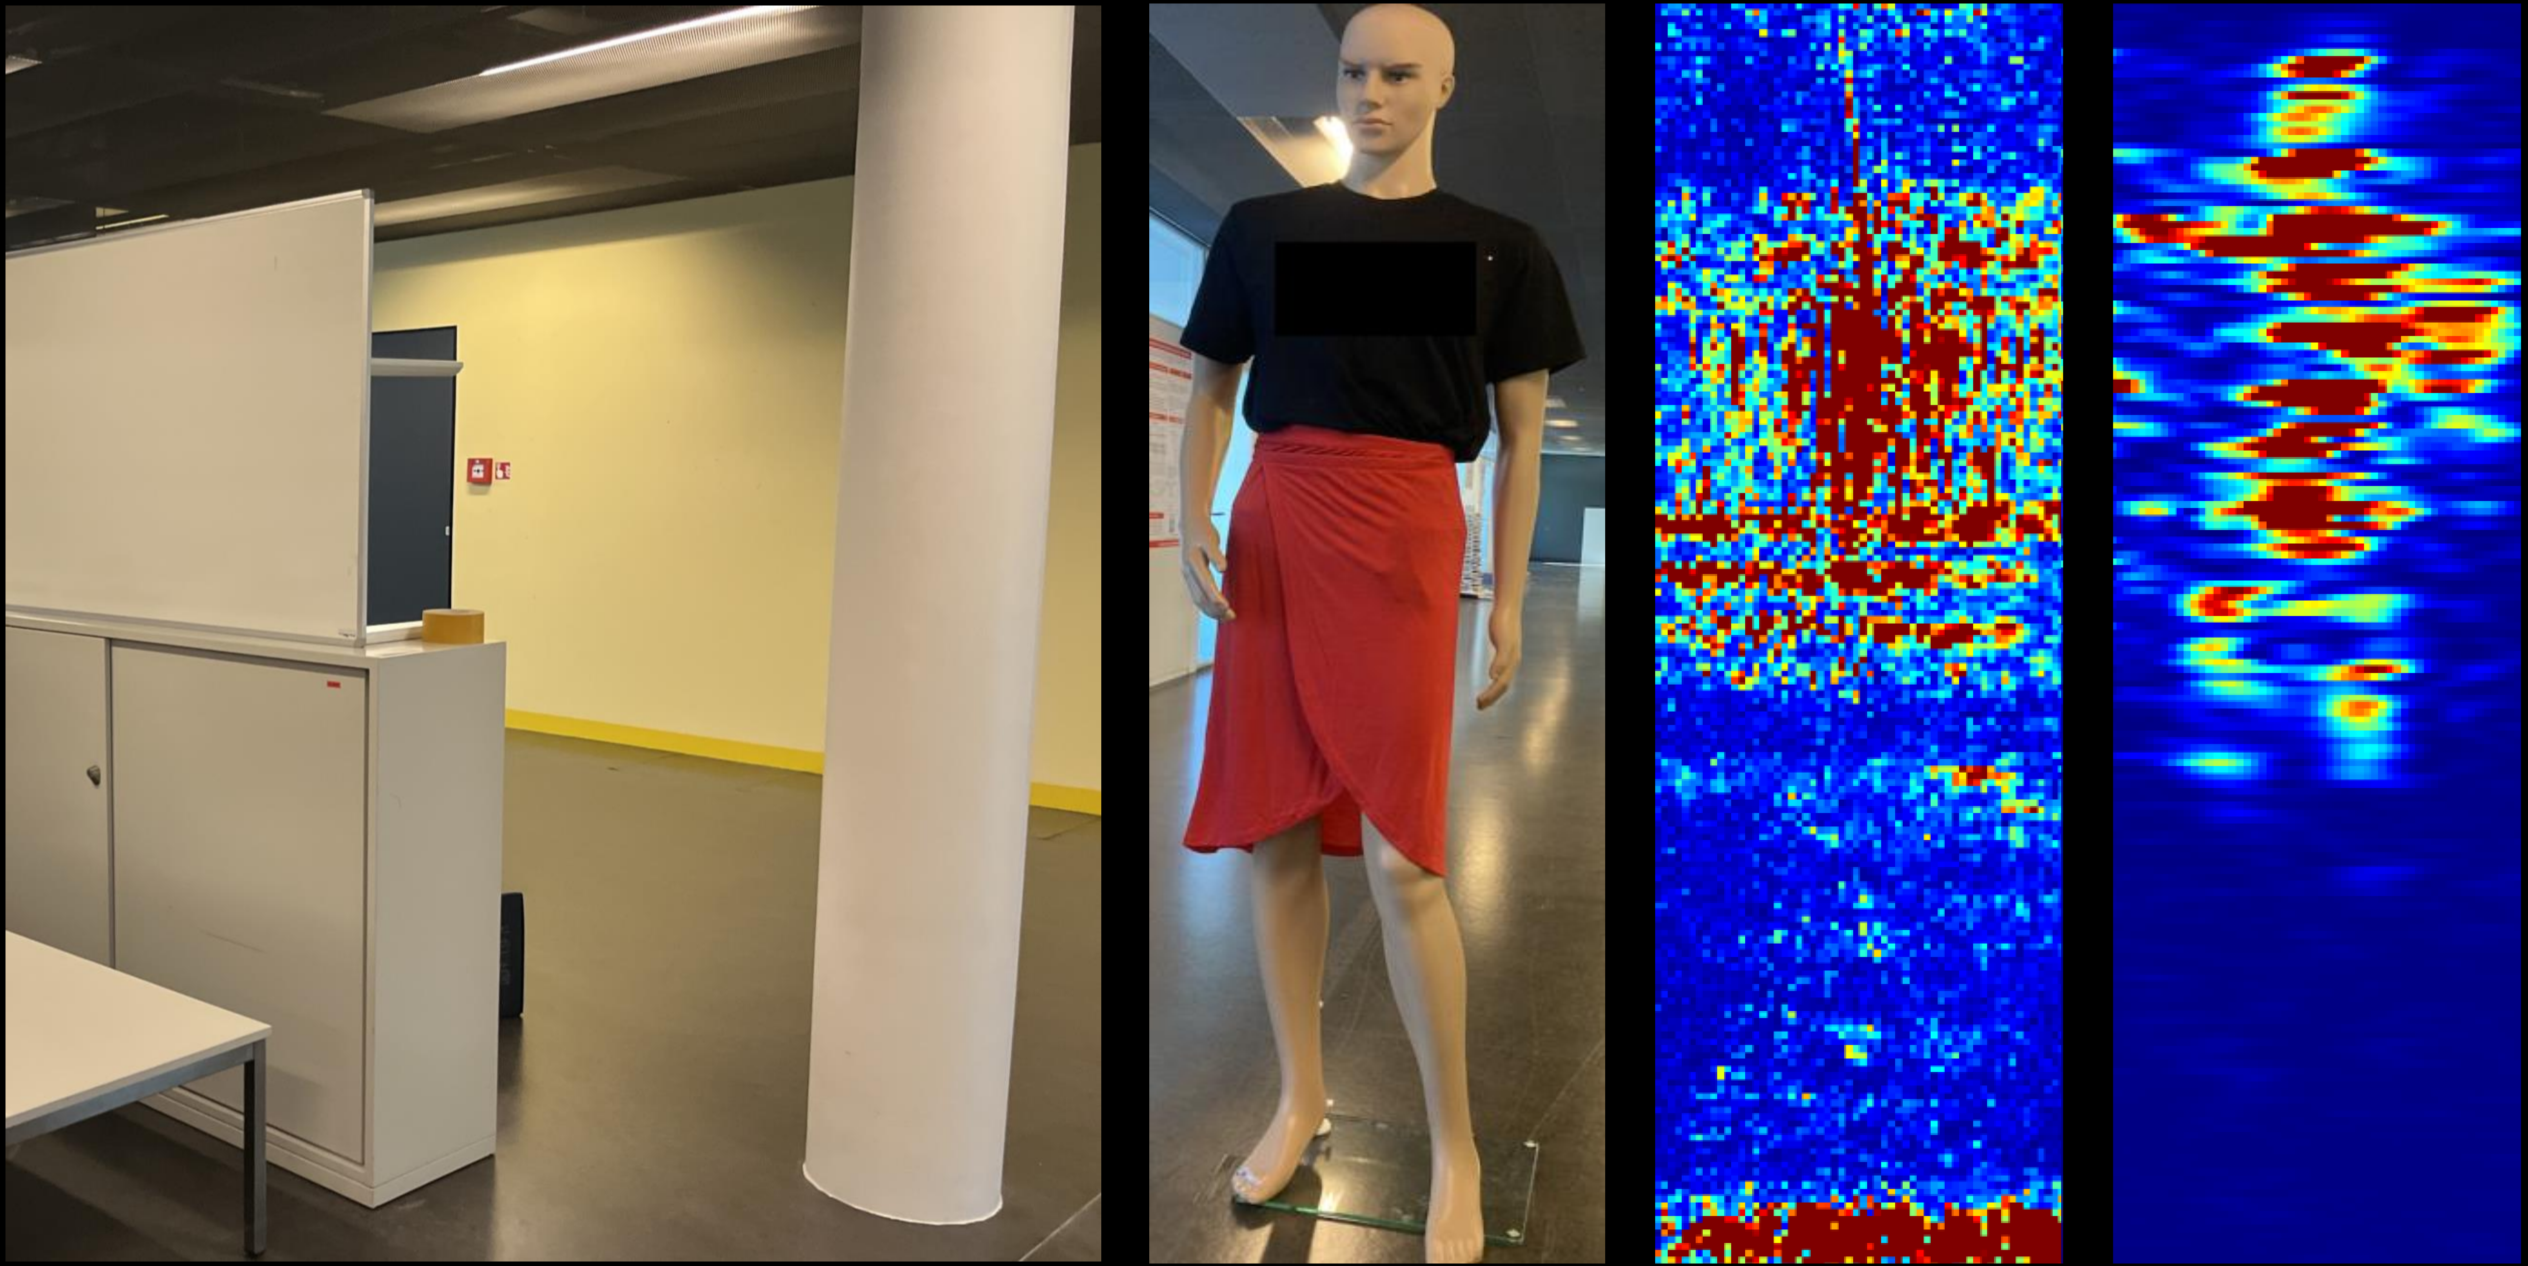

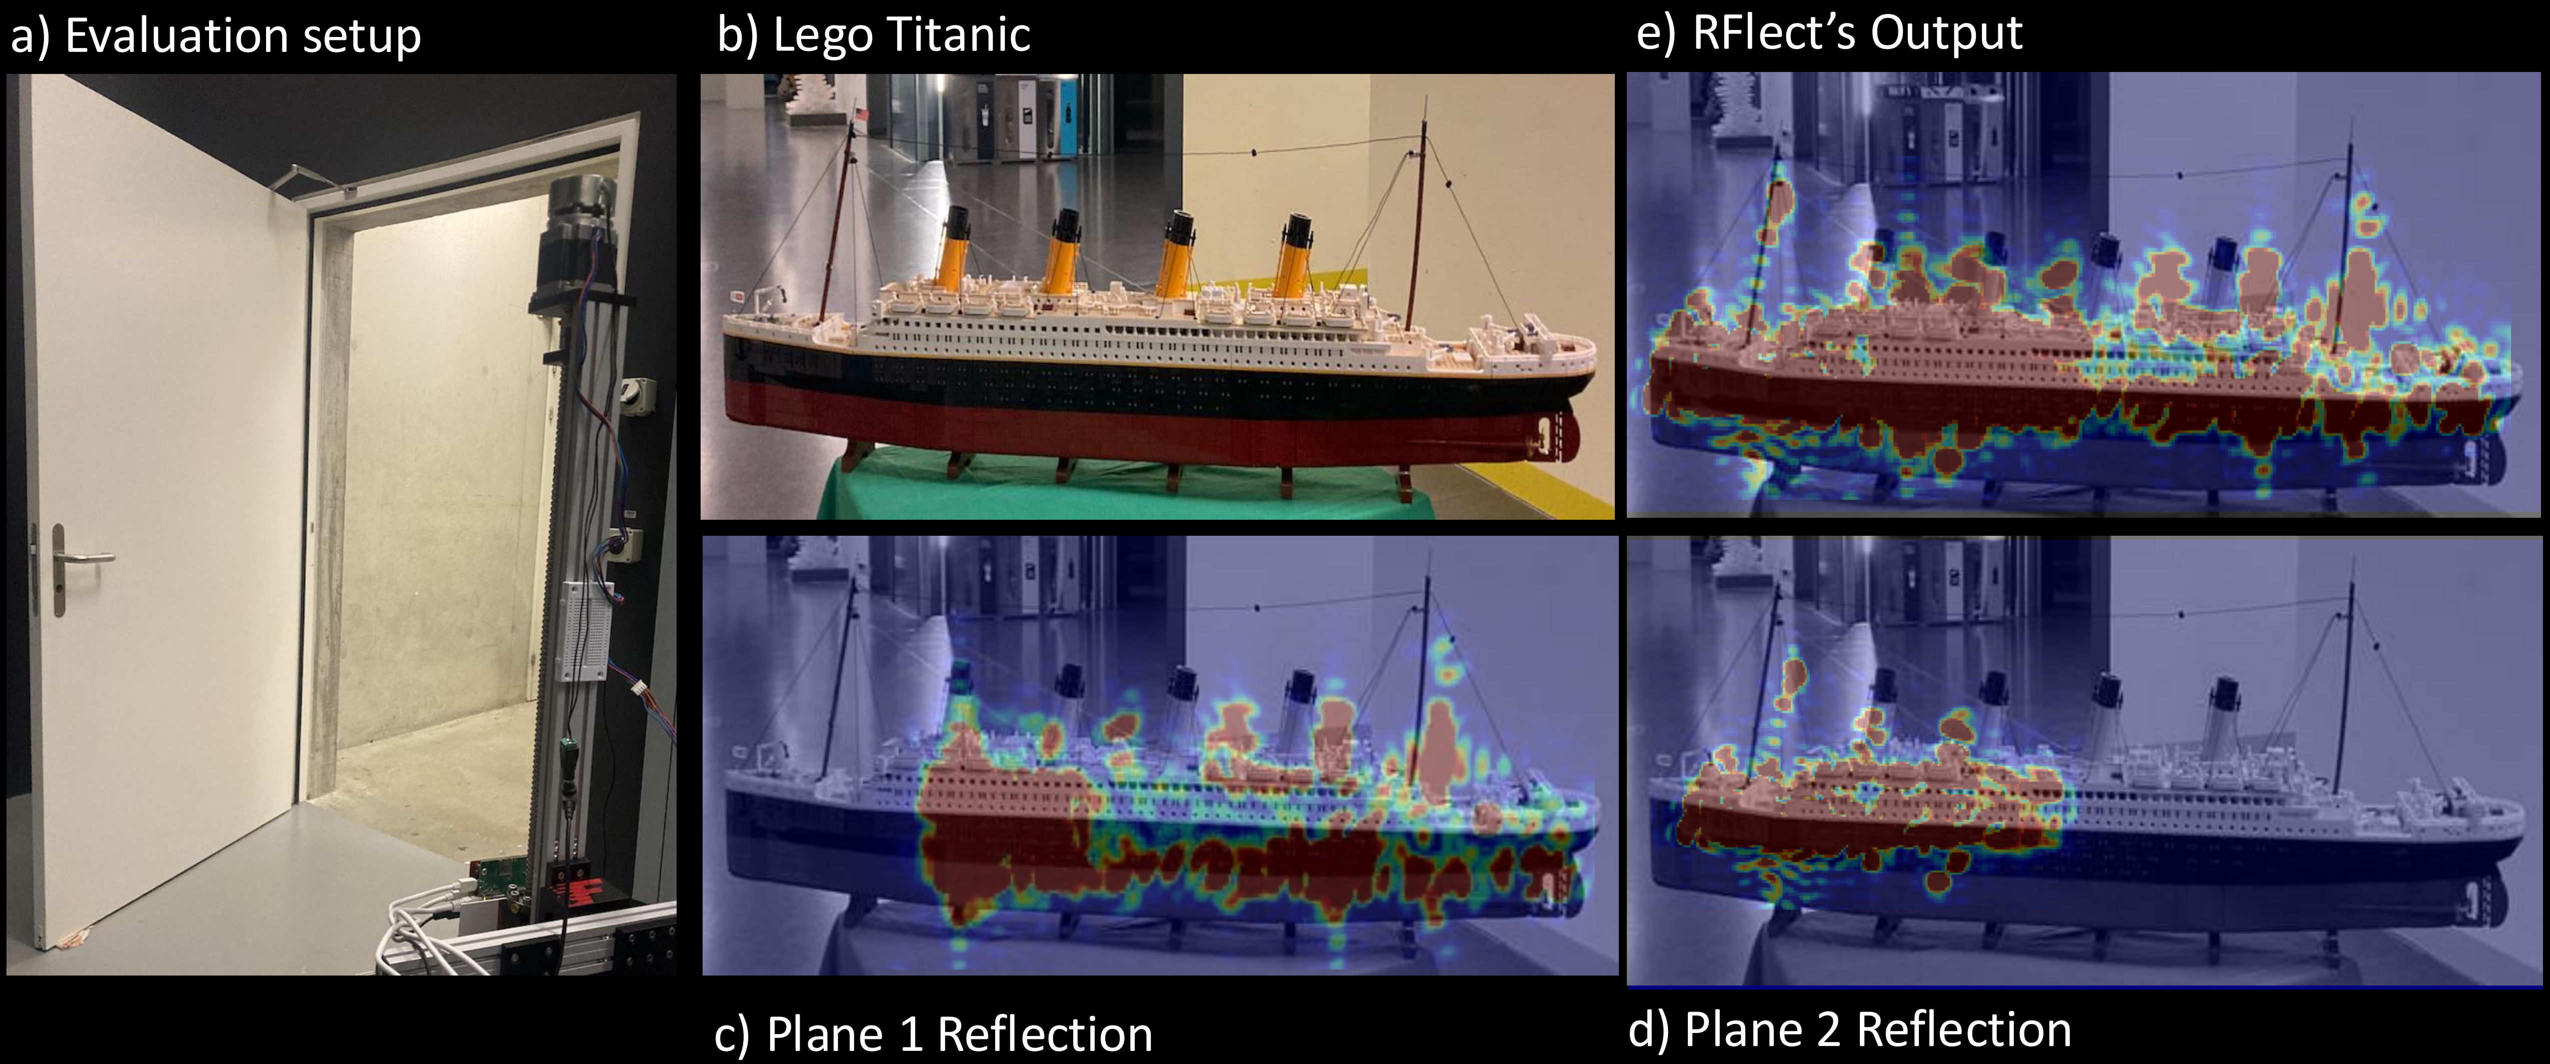
a) Evaluation setup
b) Lego Titanic
e) RFlect’s Output
c) Plane 1 Reflection
d) Plane 2 Reflection

a) Evaluation setup

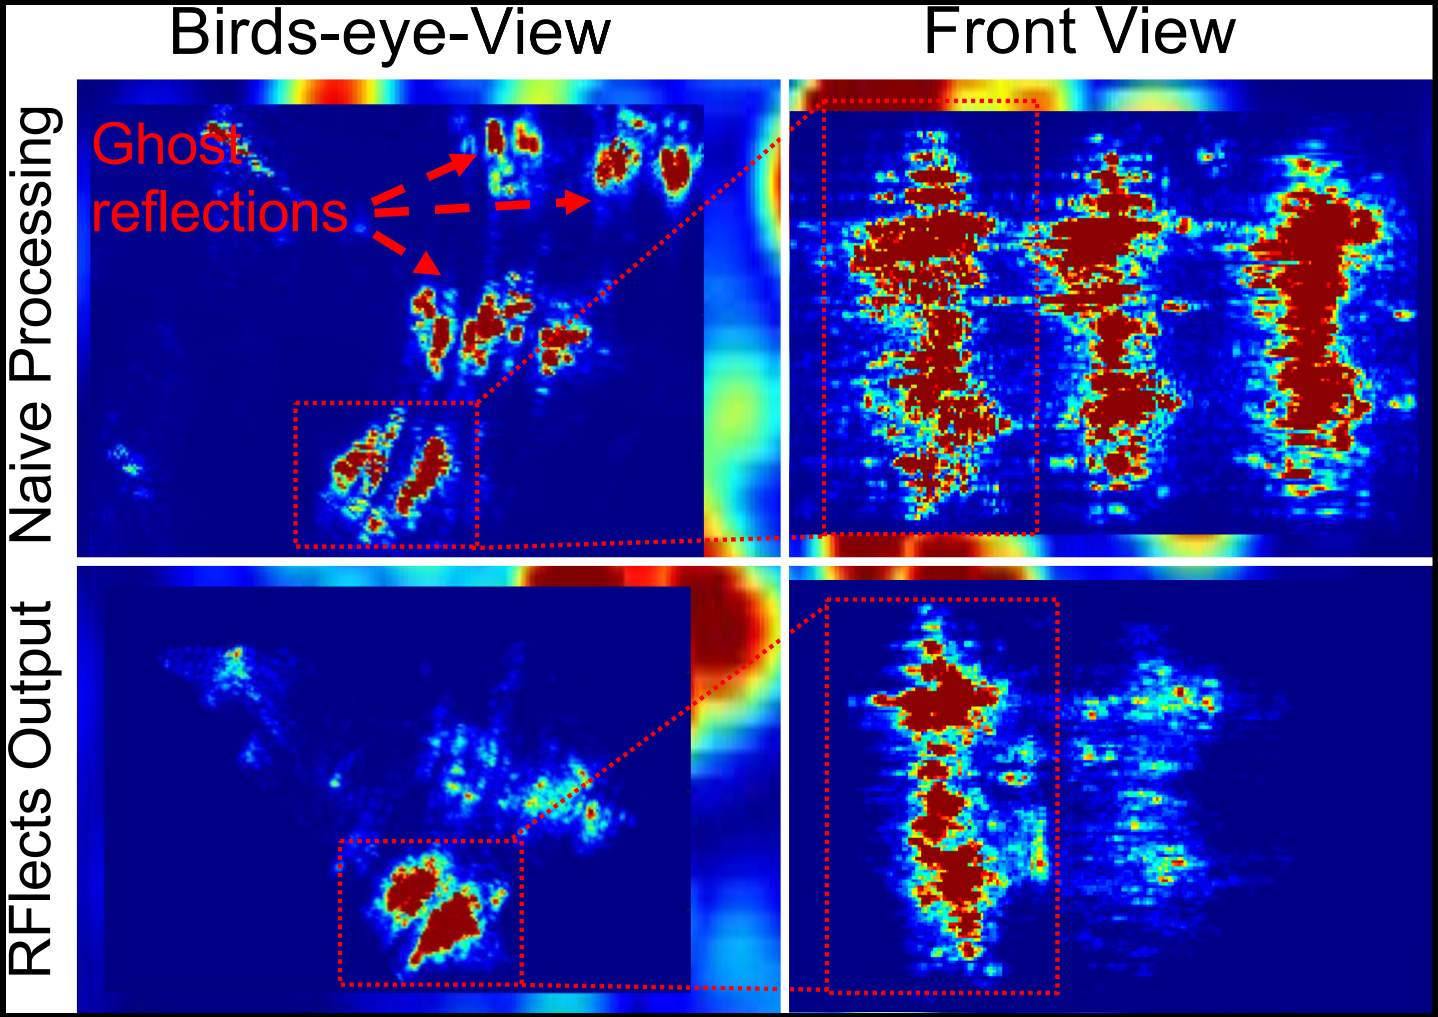

Planar Imaging Results

Target Objects

Baseline: Image objects in line-of-sight of radar

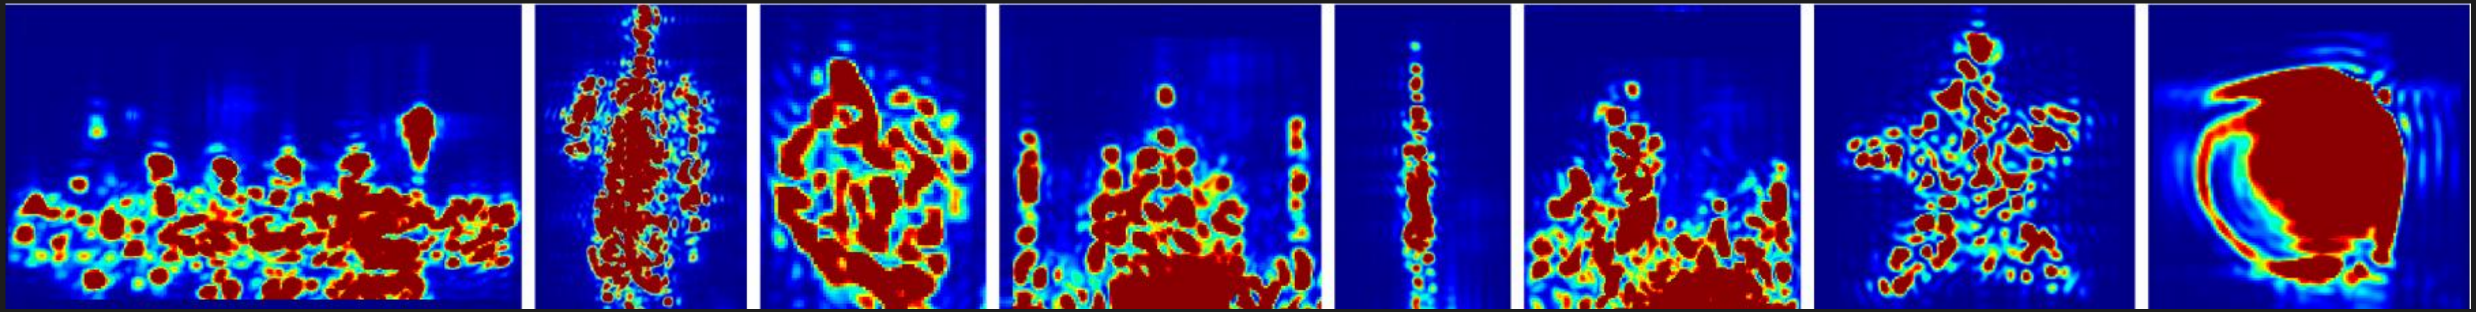

RFlect Around-the-Corner Imaging off Planar Surfaces

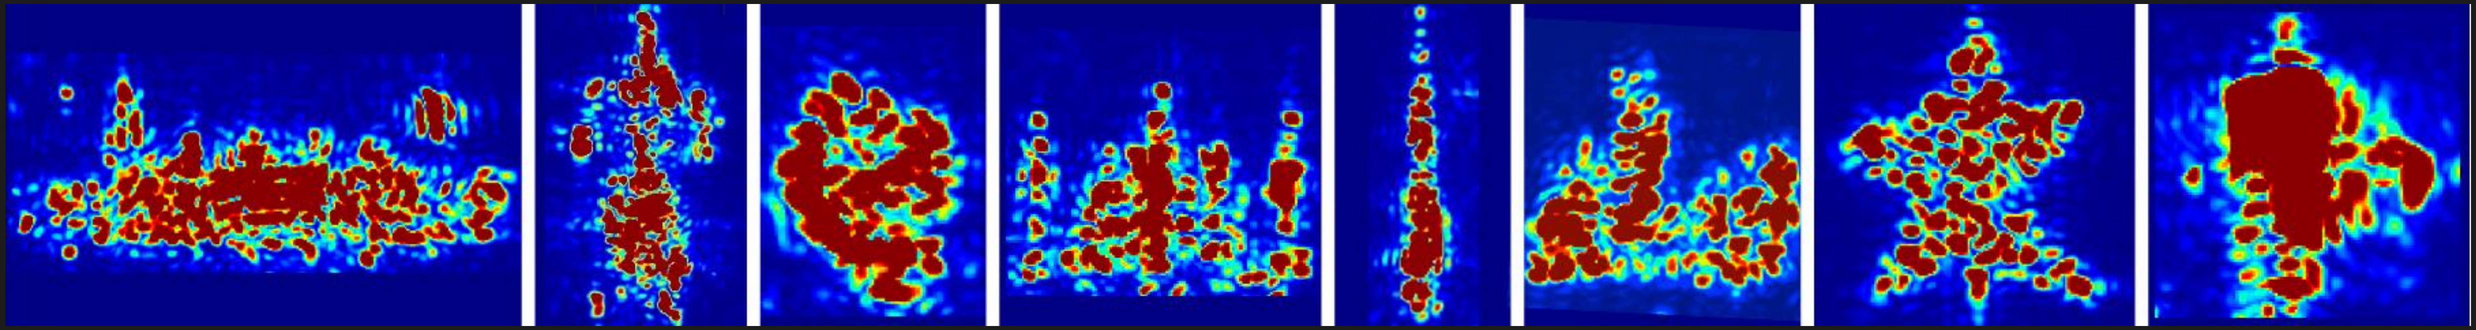

Complex Imaging Results

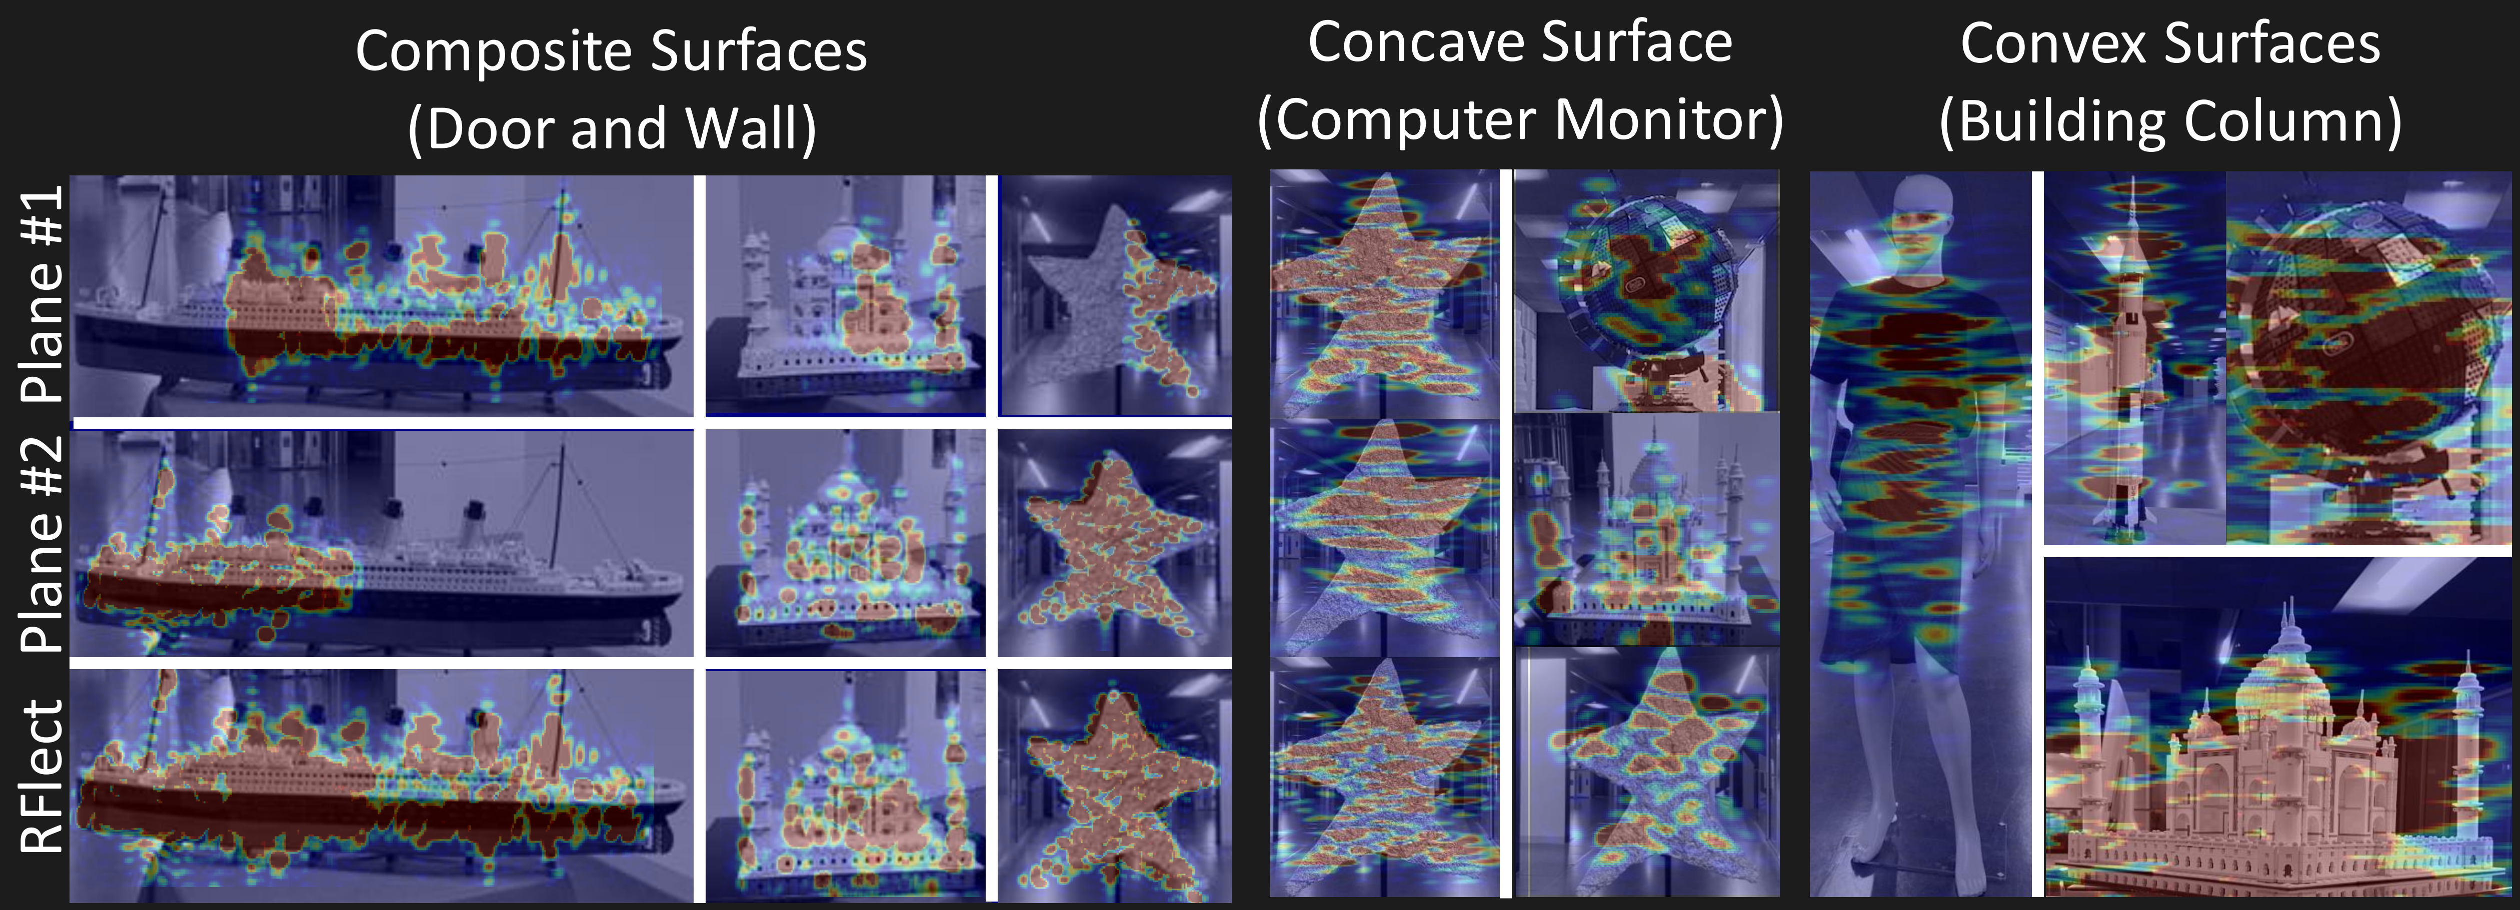

Qualitative Results: Concave Reflectors

Curved Computer Monitor

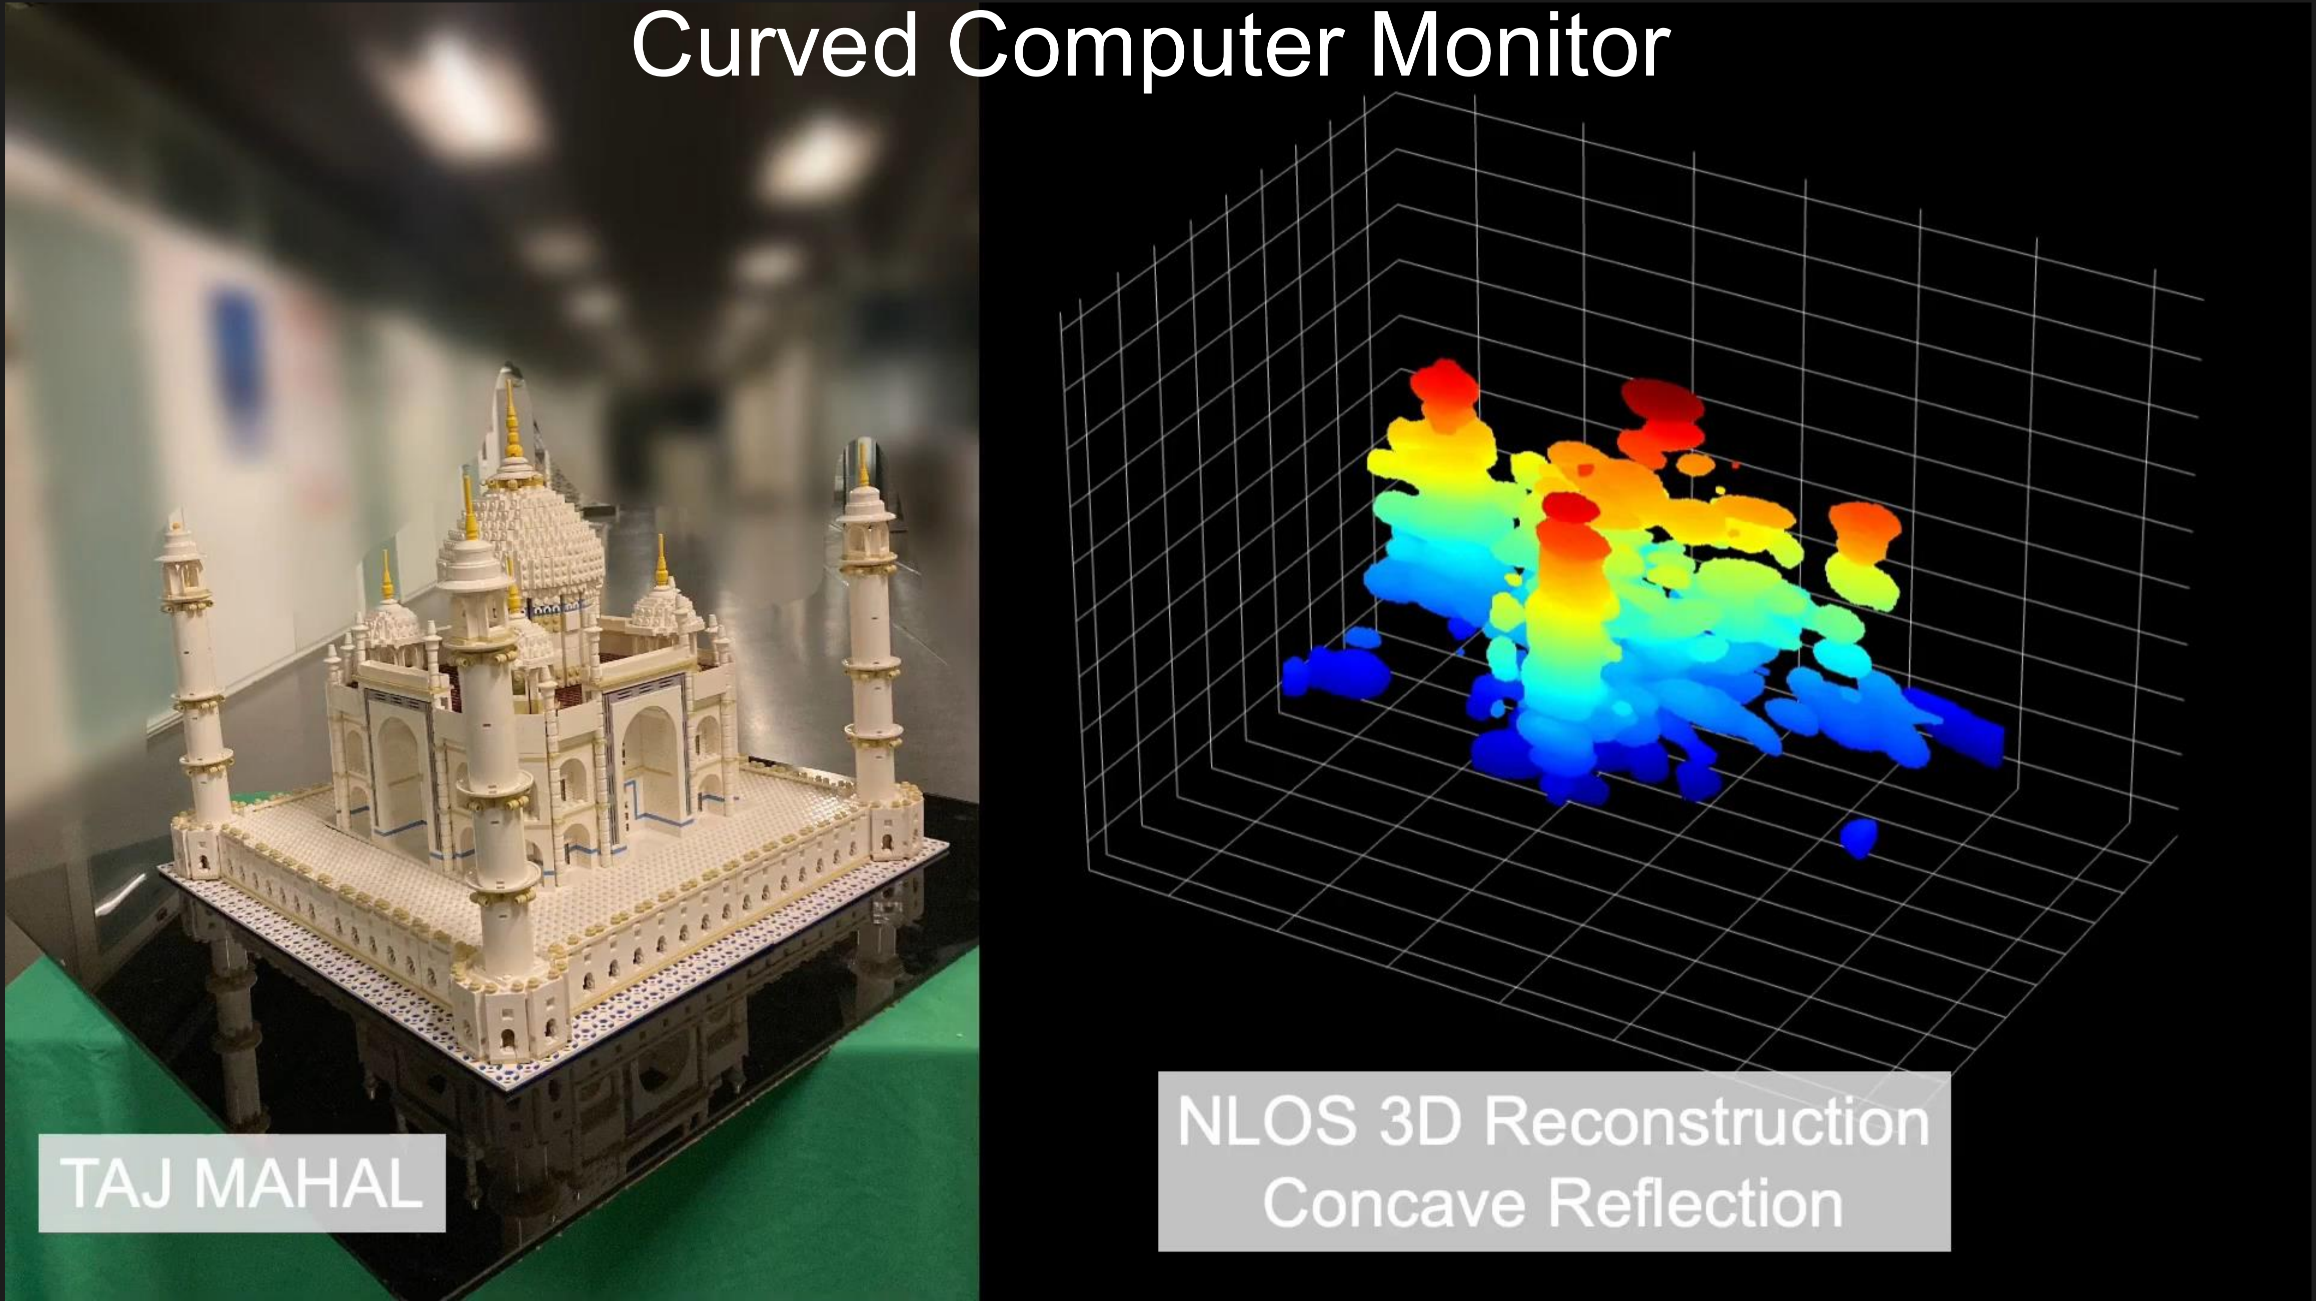

Qualitative Results: Convex Reflectors

Building Column