ciml.media.mit.edu

Data-Driven Computational Imaging

Time	Title	Presenter
08:30 – 08:50	Introduction to Computational Imaging	Guy Satat (MIT)
08:50 – 09:15	Data-Driven Computational Imaging Survey	Tristan Swedish (MIT)
09:15 – 10:00	Data-Driven Non-line-of-sight Imaging and 3D Reconstruction	Guy Satat (MIT)
10:00 – 10:20	Break	
10:20 – 11:00	Rendering and Simulation for Data-Driven Computational Imaging	Tristan Swedish (MIT)
11:00 – 12:00	Visual Sensing Using Machine Learning	Vivek Boominathan (Rice), Ashok Veeraraghavan (Rice)

mit
media
lab

Data Driven Non-line-of-sight Imaging and 3D Reconstruction

Guy Satat

Data-Driven Computational Imaging @ CVPR 2019

Why Non-Line-of-Sight Imaging?

Seeing Around Corners

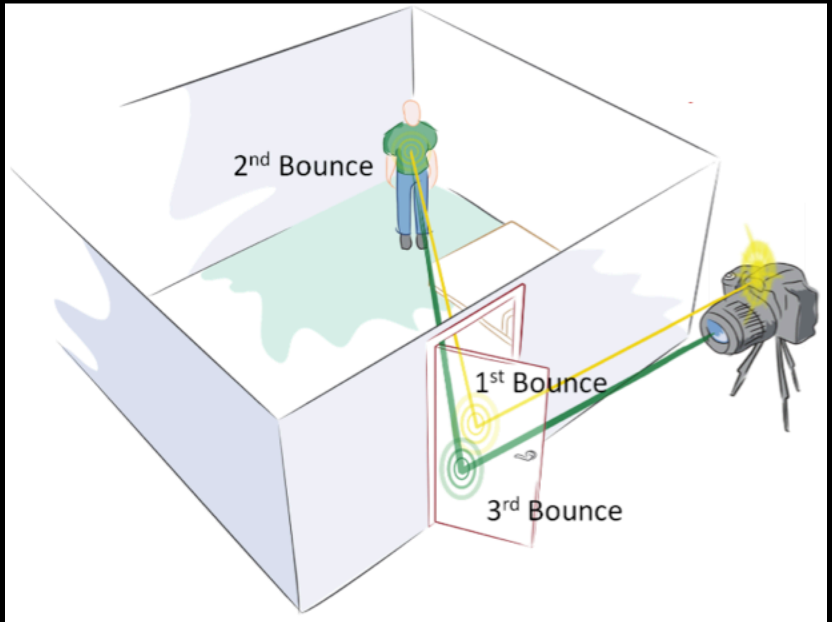

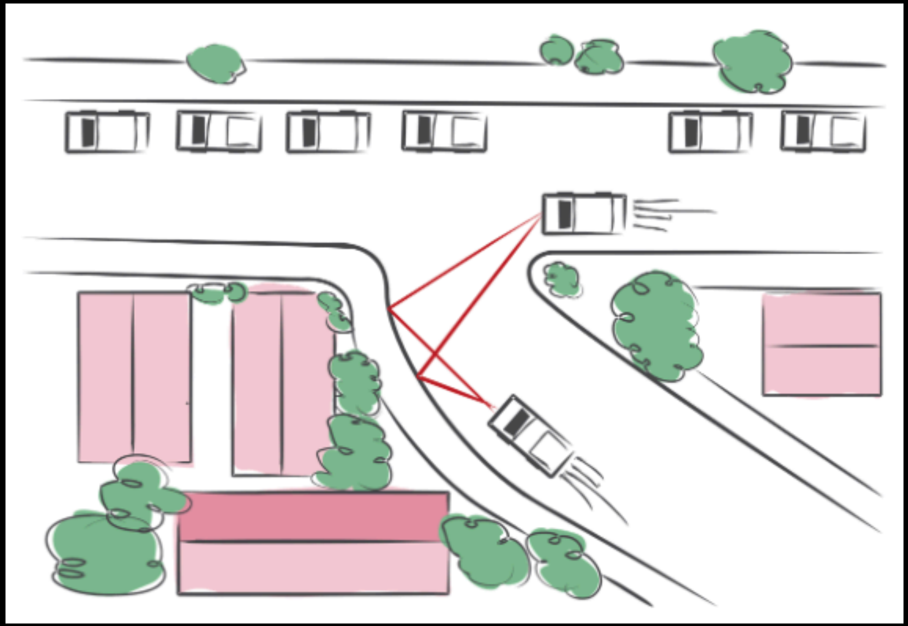

Radar

Camera

Lidar

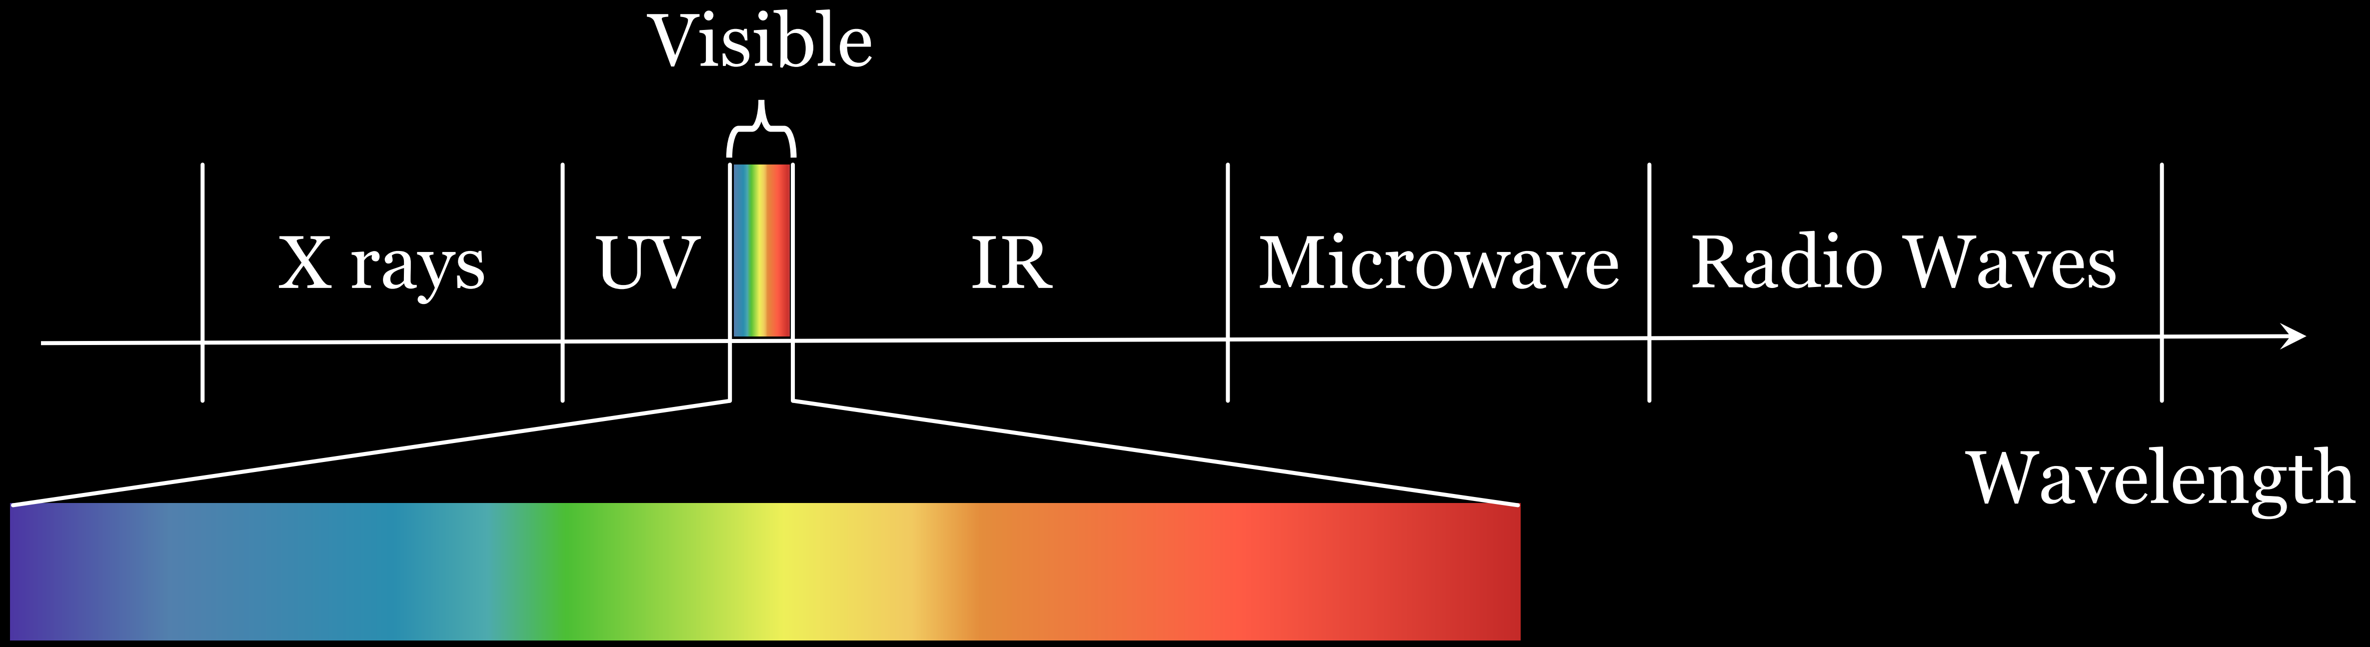

- Resolution
- Optical Contrast

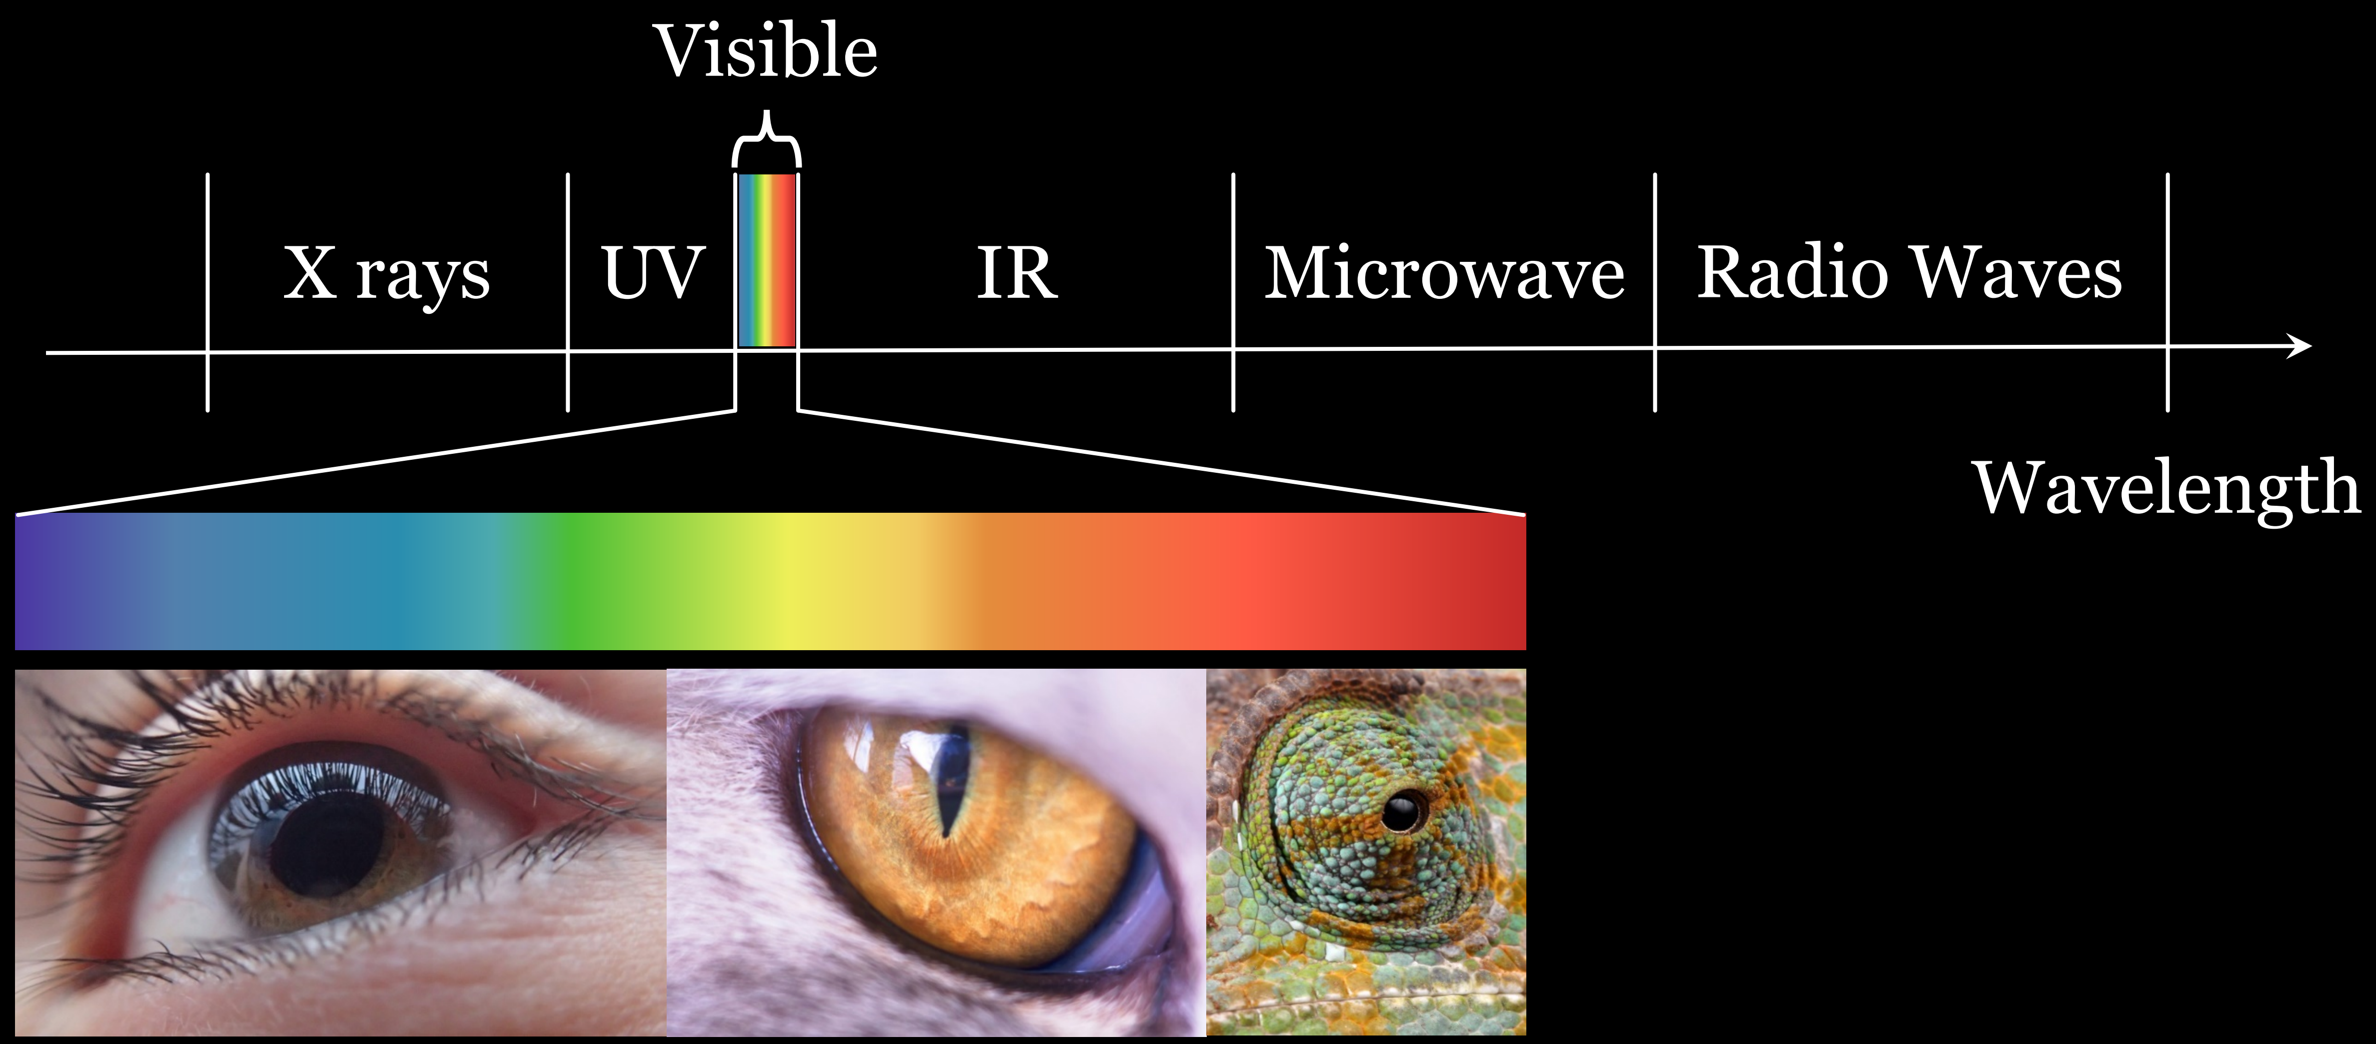

Visible
X rays
UV
IR
Microwave
Radio Waves
Wavelength

Light Scatters

Light and Matter in a Nutshell

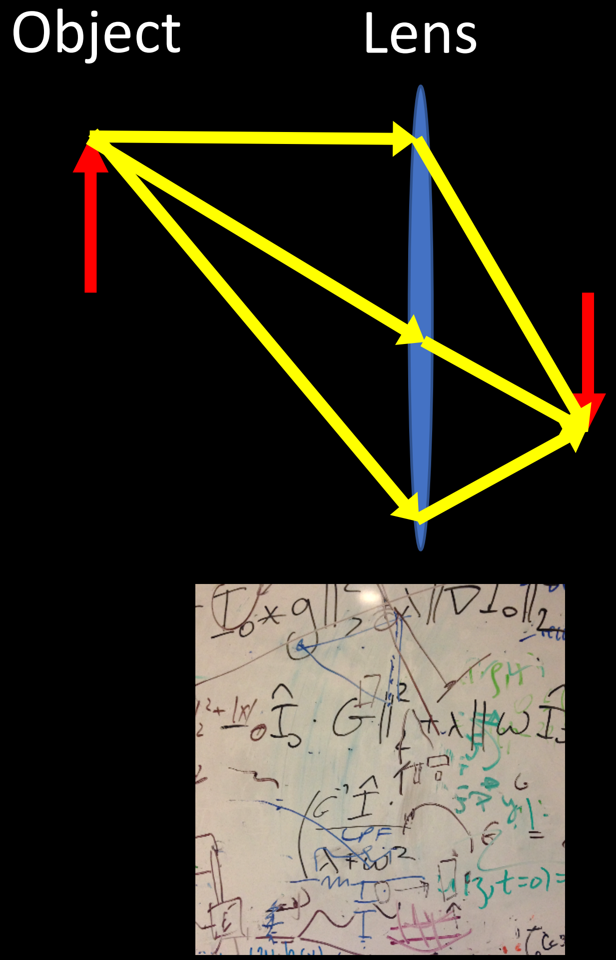

Absorption

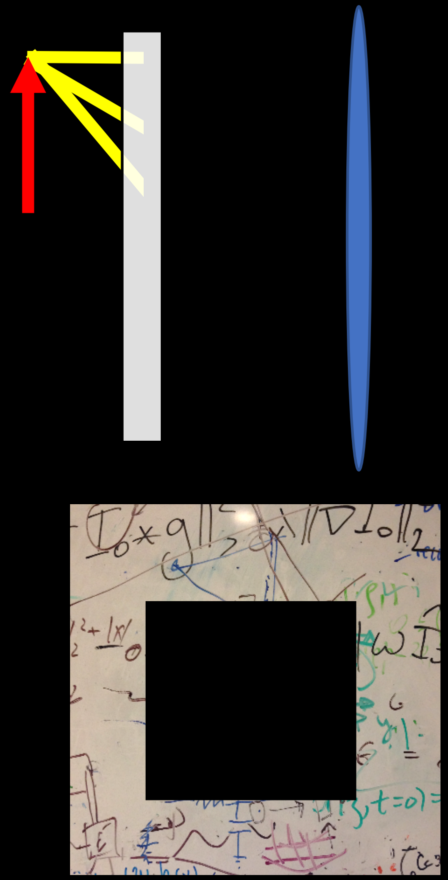

Scattering

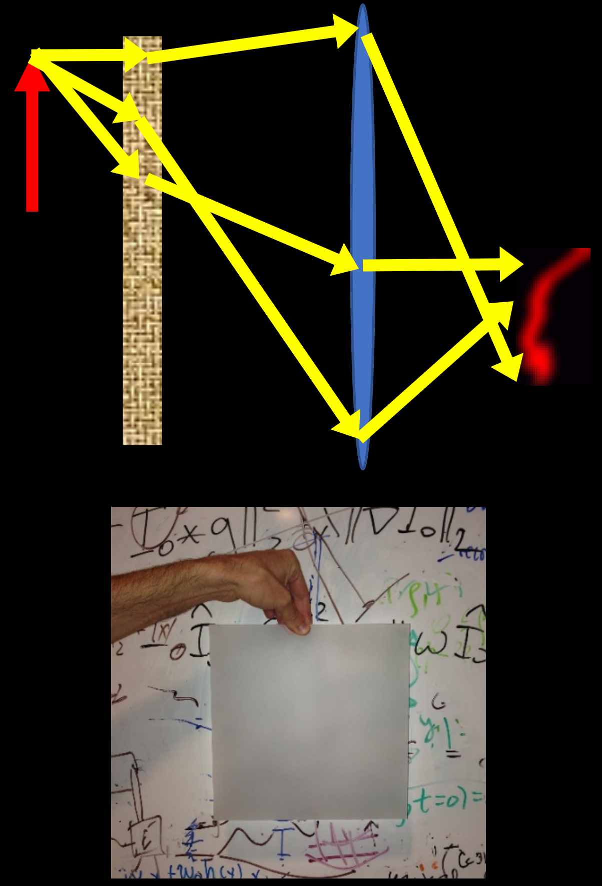

Imaging Geometries

Around the Corner

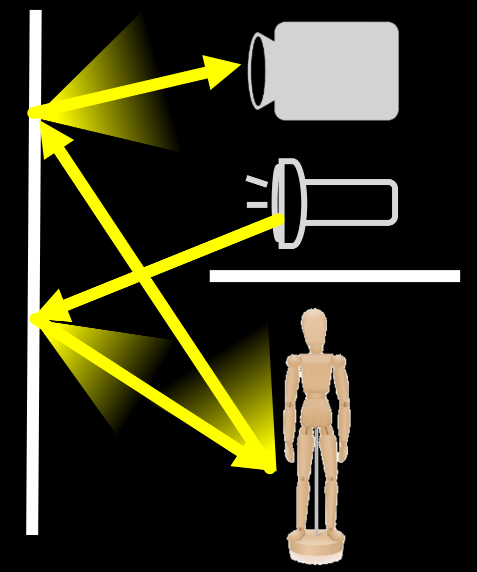

Through Scattering

Transmission Optics

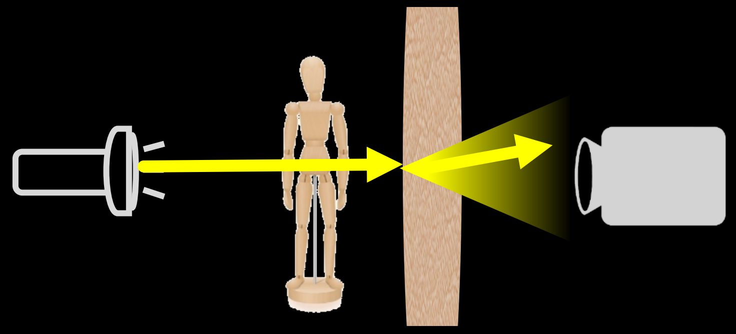

Reflection Optics

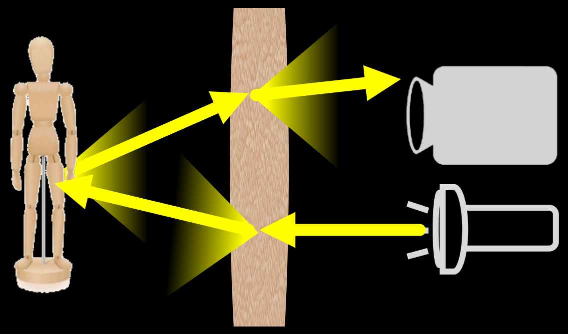

Challenges With Computational Imaging

- Physical modeling (forward model):

$$\underbrace{b}_{\text{Measurement}} = \underbrace{A}_{\text{Forward Model}}(\underbrace{x}_{\text{Unknown Target}})$$

- Inference (Inversion):

$$x = A^{-1}(b)$$

- A can be hard to model
- A is usually non-invertible

Data-Driven Computational Imaging Advantages

- No need of a physical model

- No need to invert a model

  → More robust solution

- Directly learn: $b \rightarrow x$

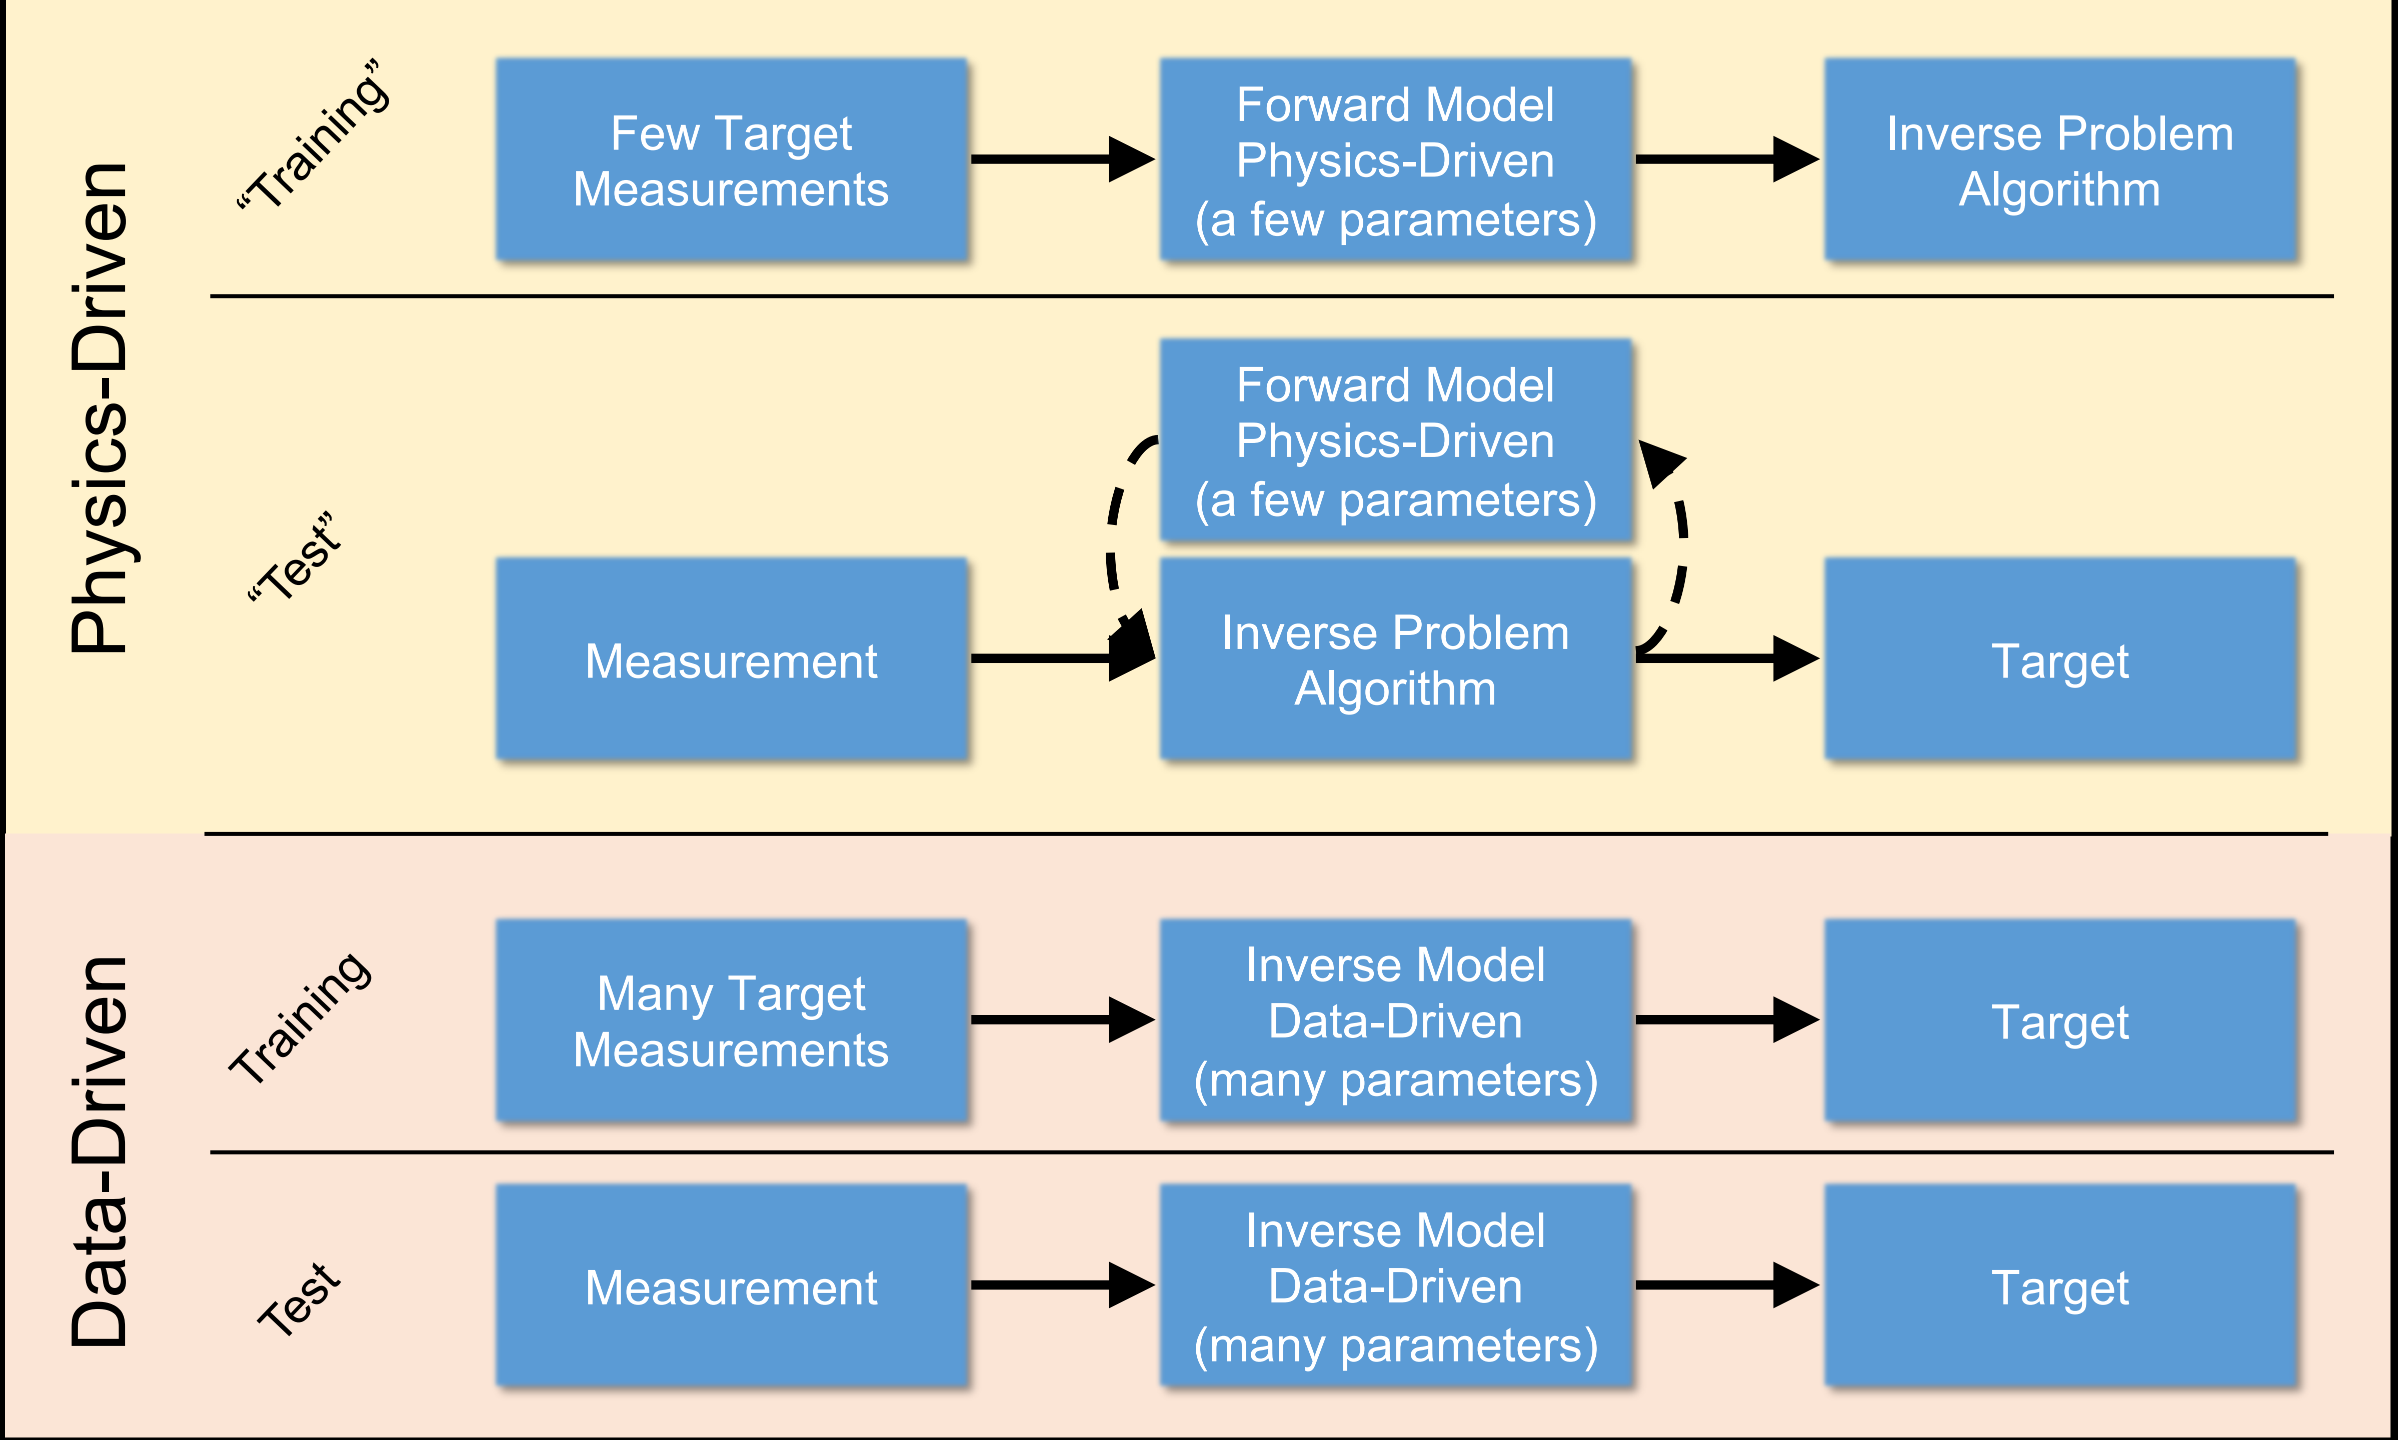
Physics-Driven
"Training"
Few Target Measurements
Forward Model Physics-Driven (a few parameters)
Inverse Problem Algorithm
"Test"
Forward Model Physics-Driven (a few parameters)
Measurement
Inverse Problem Algorithm
Target
Data-Driven
Training
Many Target Measurements
Inverse Model Data-Driven (many parameters)
Target
Test
Measurement
Inverse Model Data-Driven (many parameters)
Target

Challenges with Data-Driven NLOS Imaging

- Lack of available data
- Non-traditional sensors
  - LIDAR, ToF
- Hard to label data
  - Requires additional sensors and experiments
- Generalization

NLOS Computational Imaging

A Few Examples

We Have to Calibrate

Classify Pose **Without** Calibration

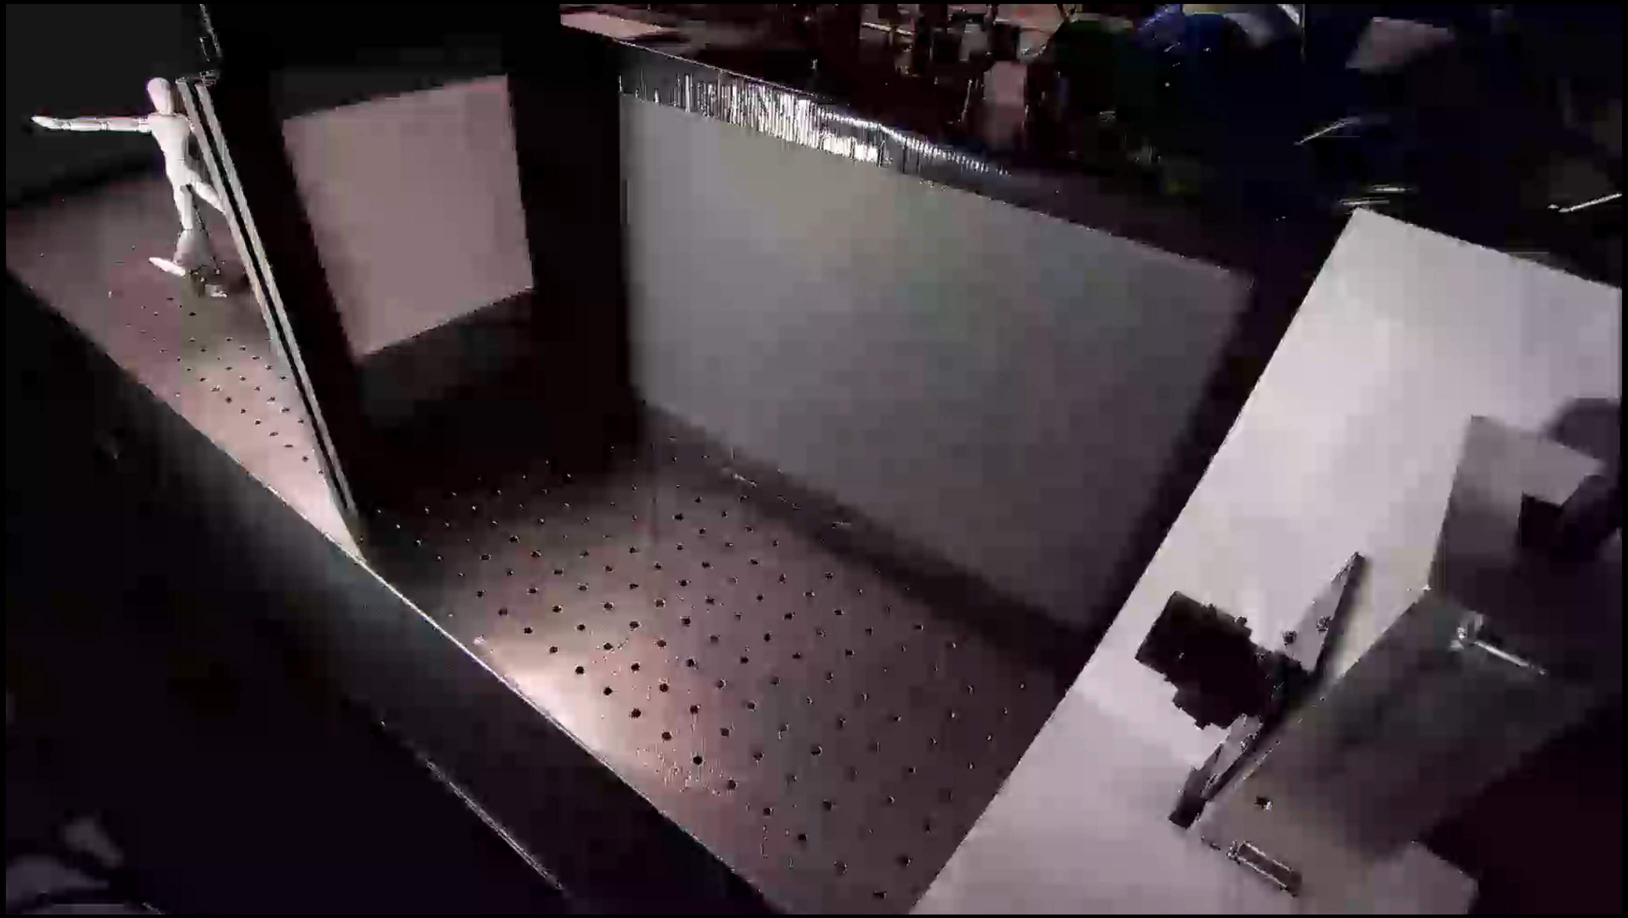

Guy Satat, Matthew Tancik, Otkrist Gupta, Barmak Heshmat, Ramesh Raskar, *Optics Express* 2017

Solution:

Deep Learning

Known to Learn Invariant Models

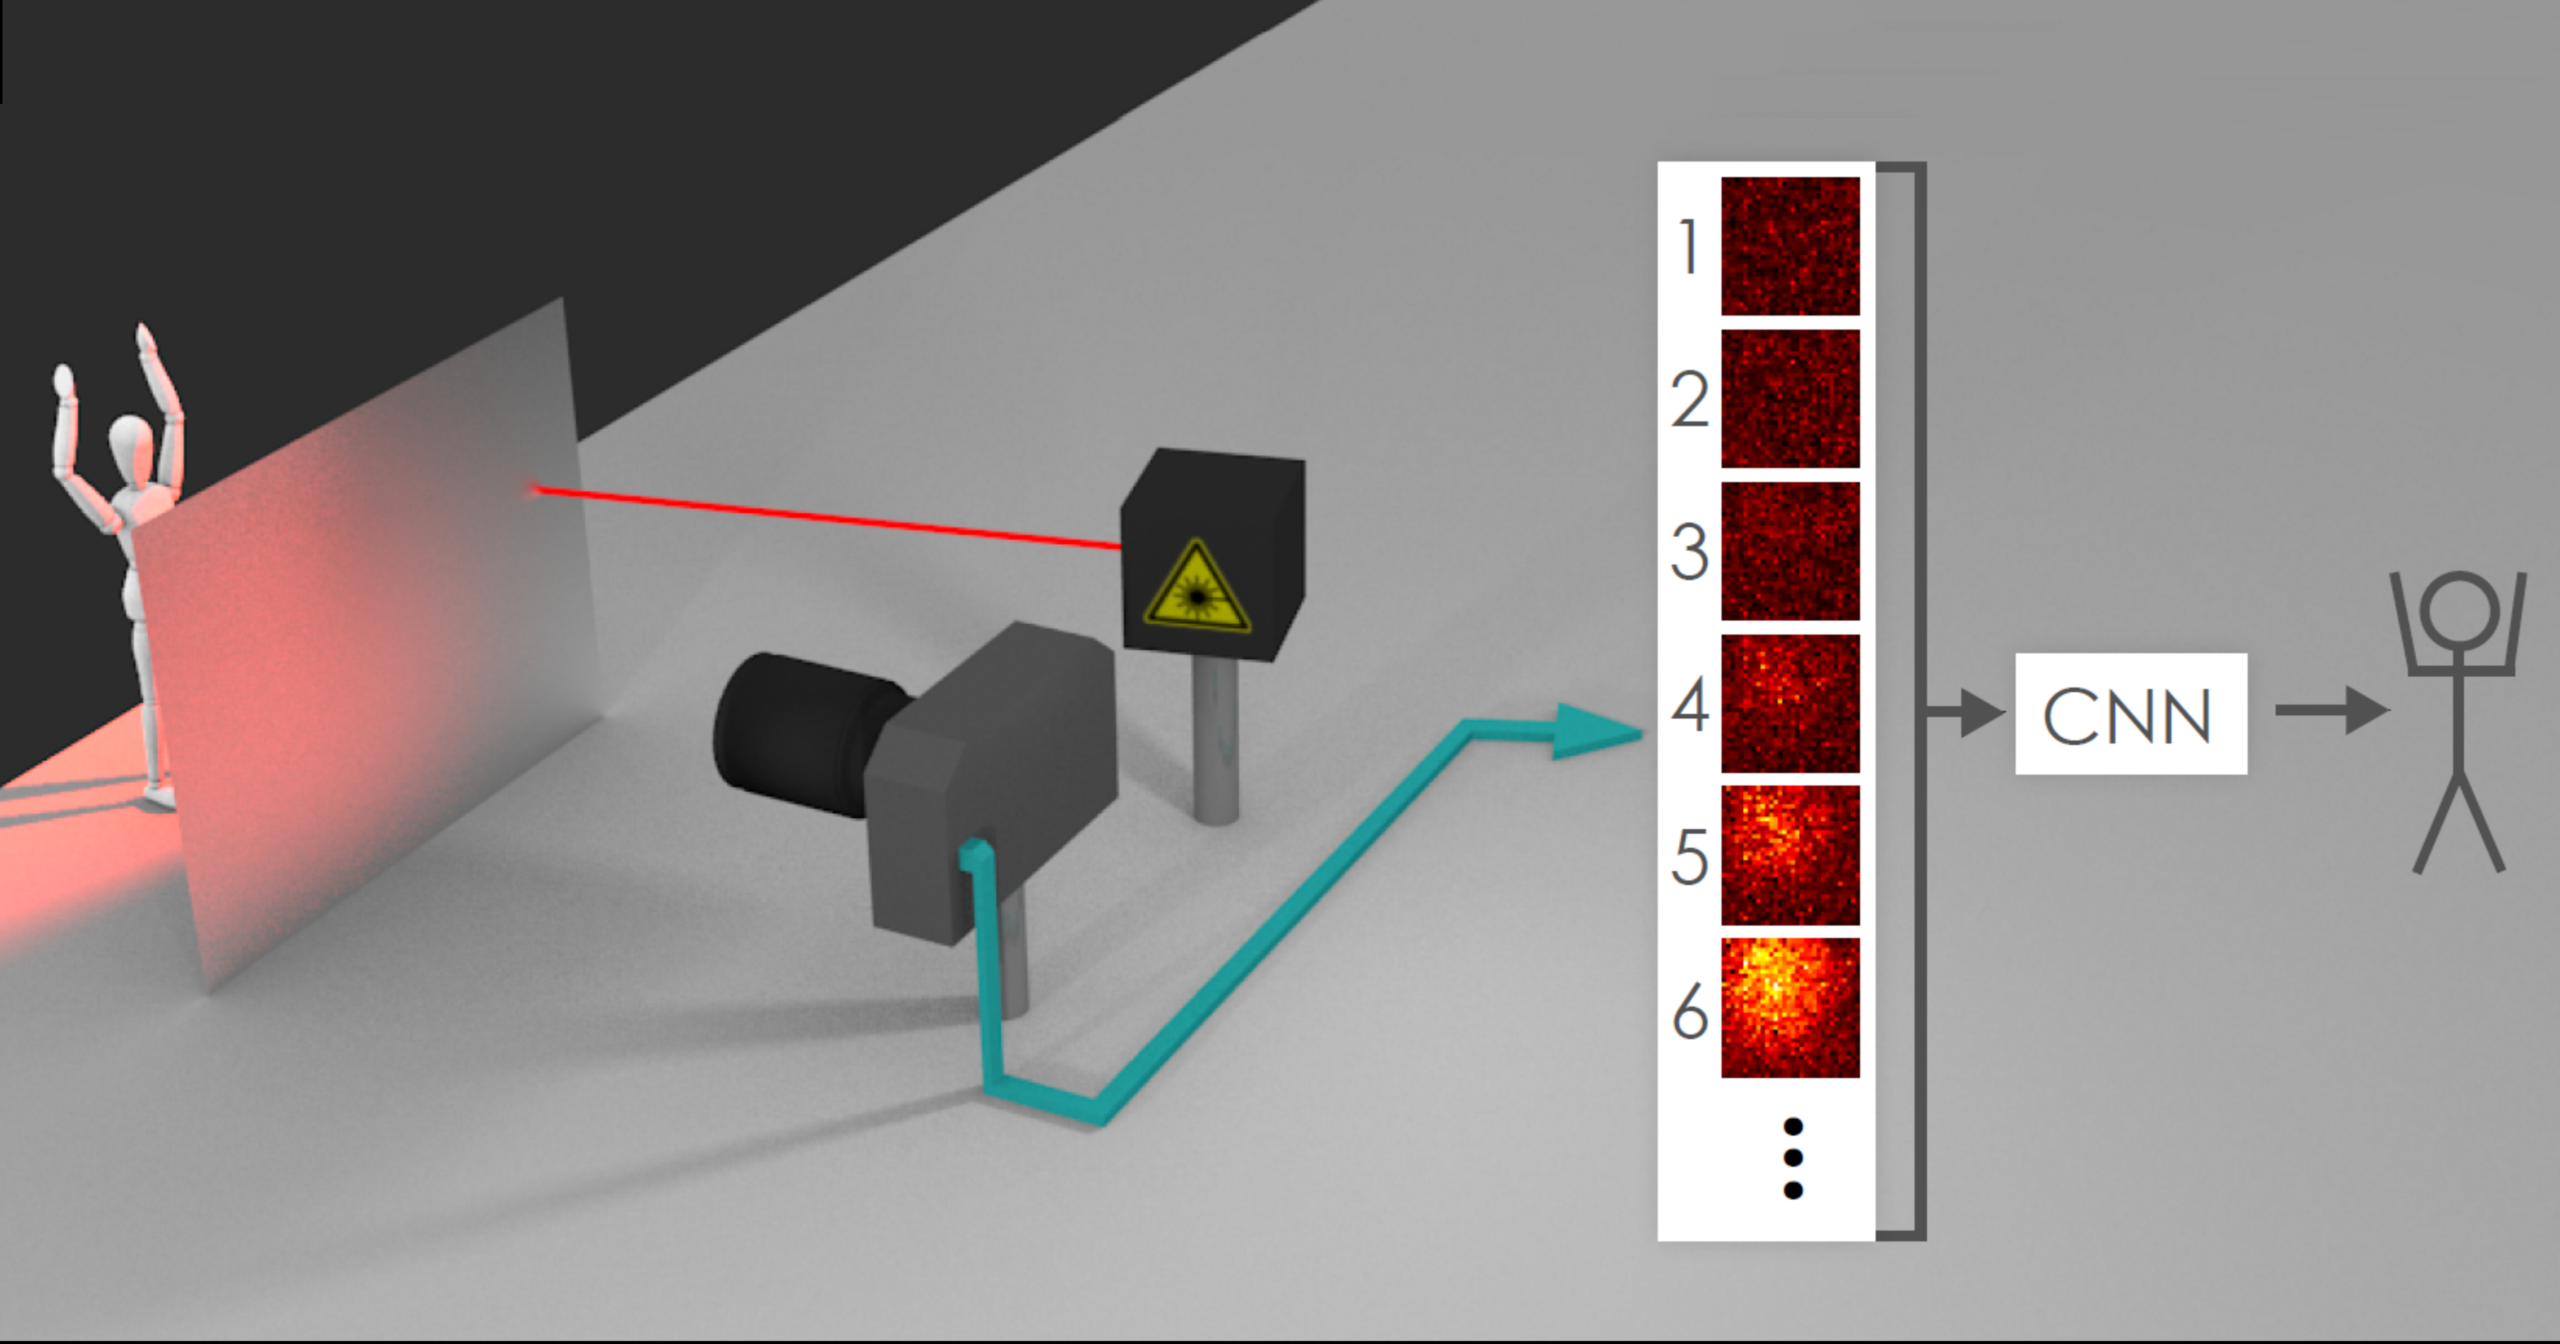
1
2
3
4
5
6
CNN

SPAD Pixel

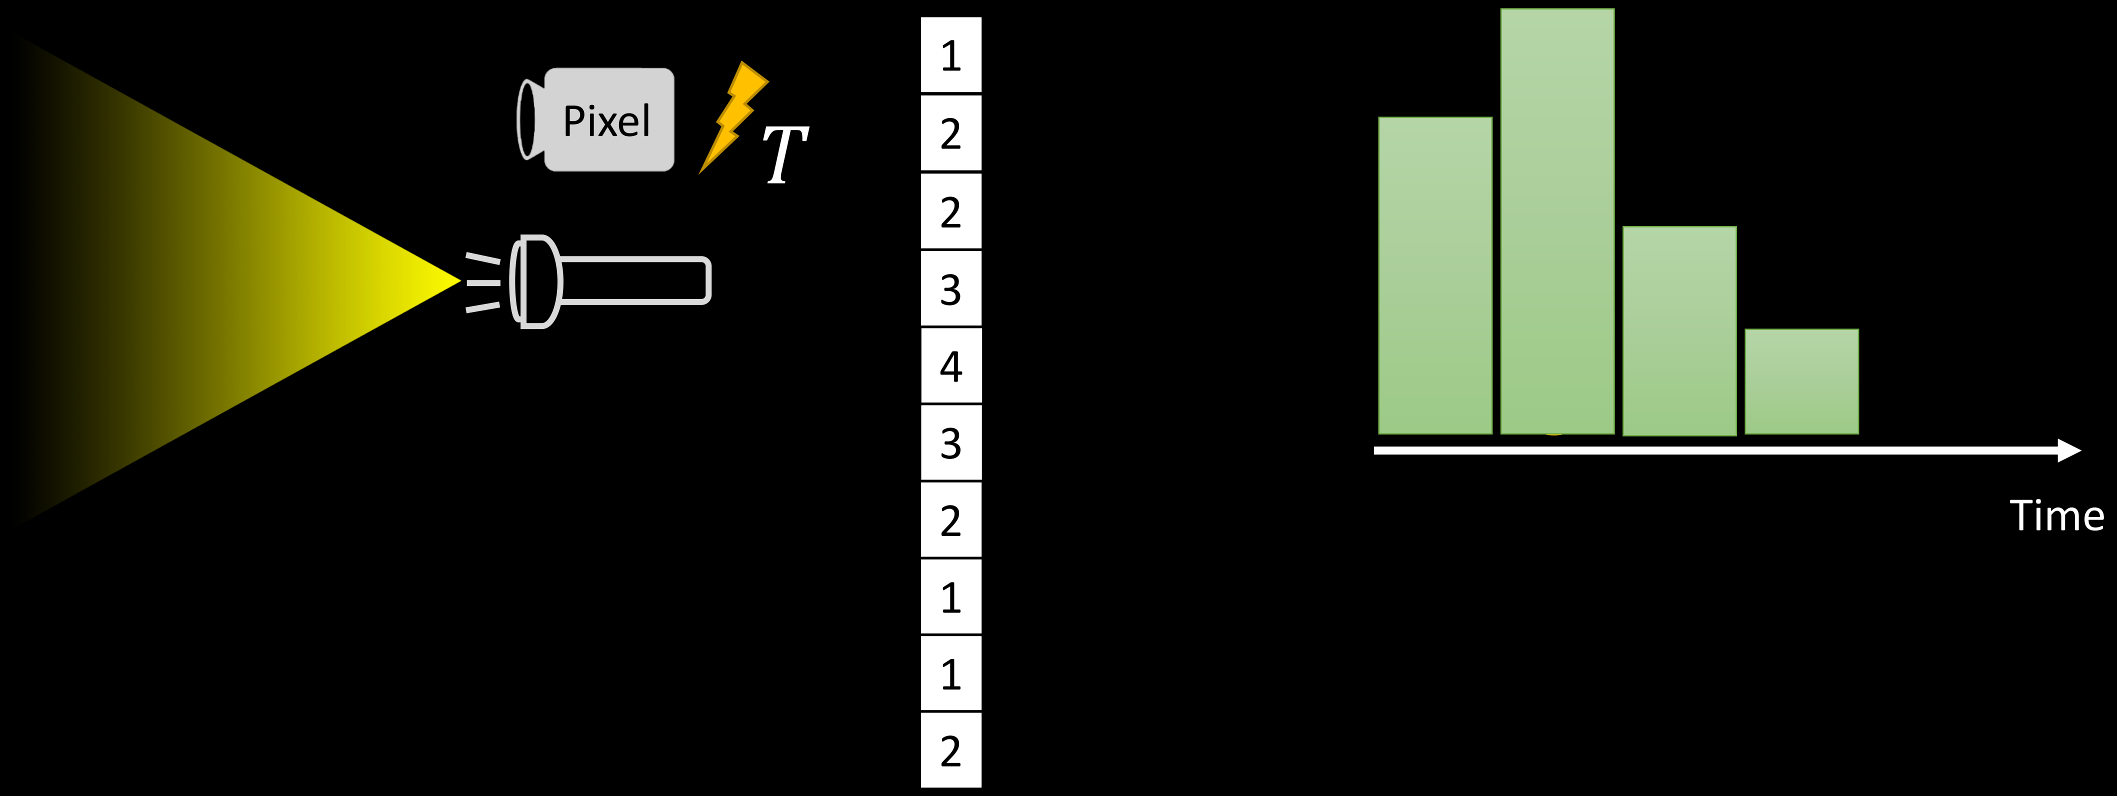

SPAD Camera

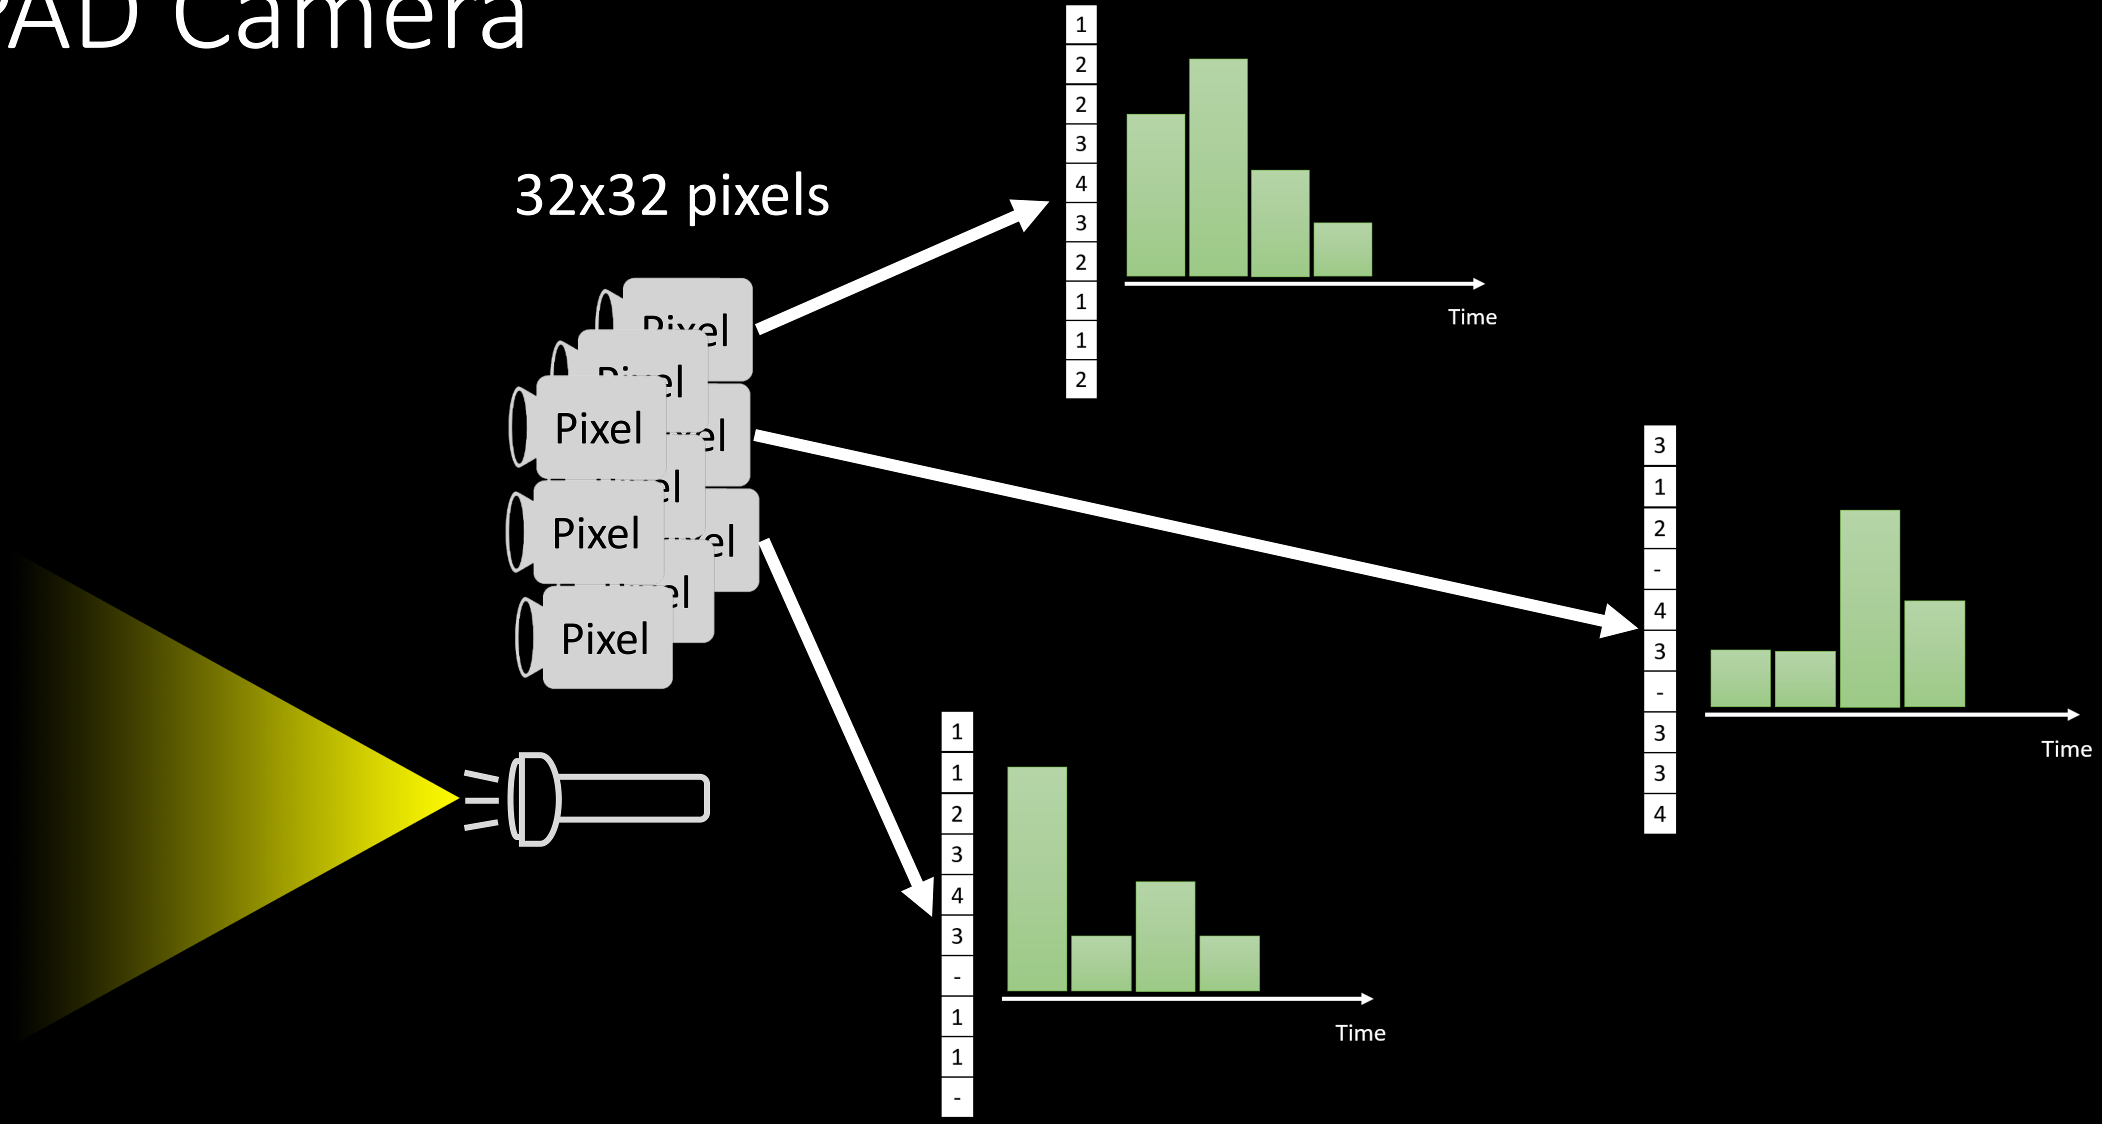
32x32 pixels
Pixel
Pixel
Pixel
1 2 2 3 4 3 2 1 1 2
Time
3 1 2 - 4 3 - 3 3 4
Time
1 1 2 3 4 3 - 1 1 -
Time

How to Train?

Data?

Synthetic Training Dataset

Learning Invariants to Calibration Parameters

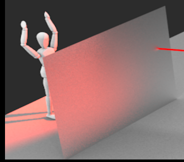

- Create synthetic dataset
- For each example:
  - Randomly sample scene properties
  - Compute forward model
- Train CNN

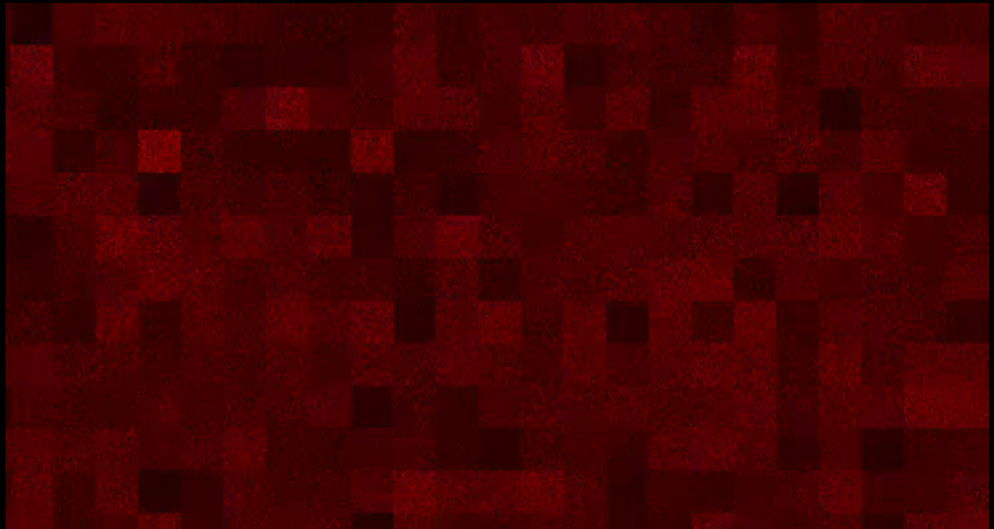

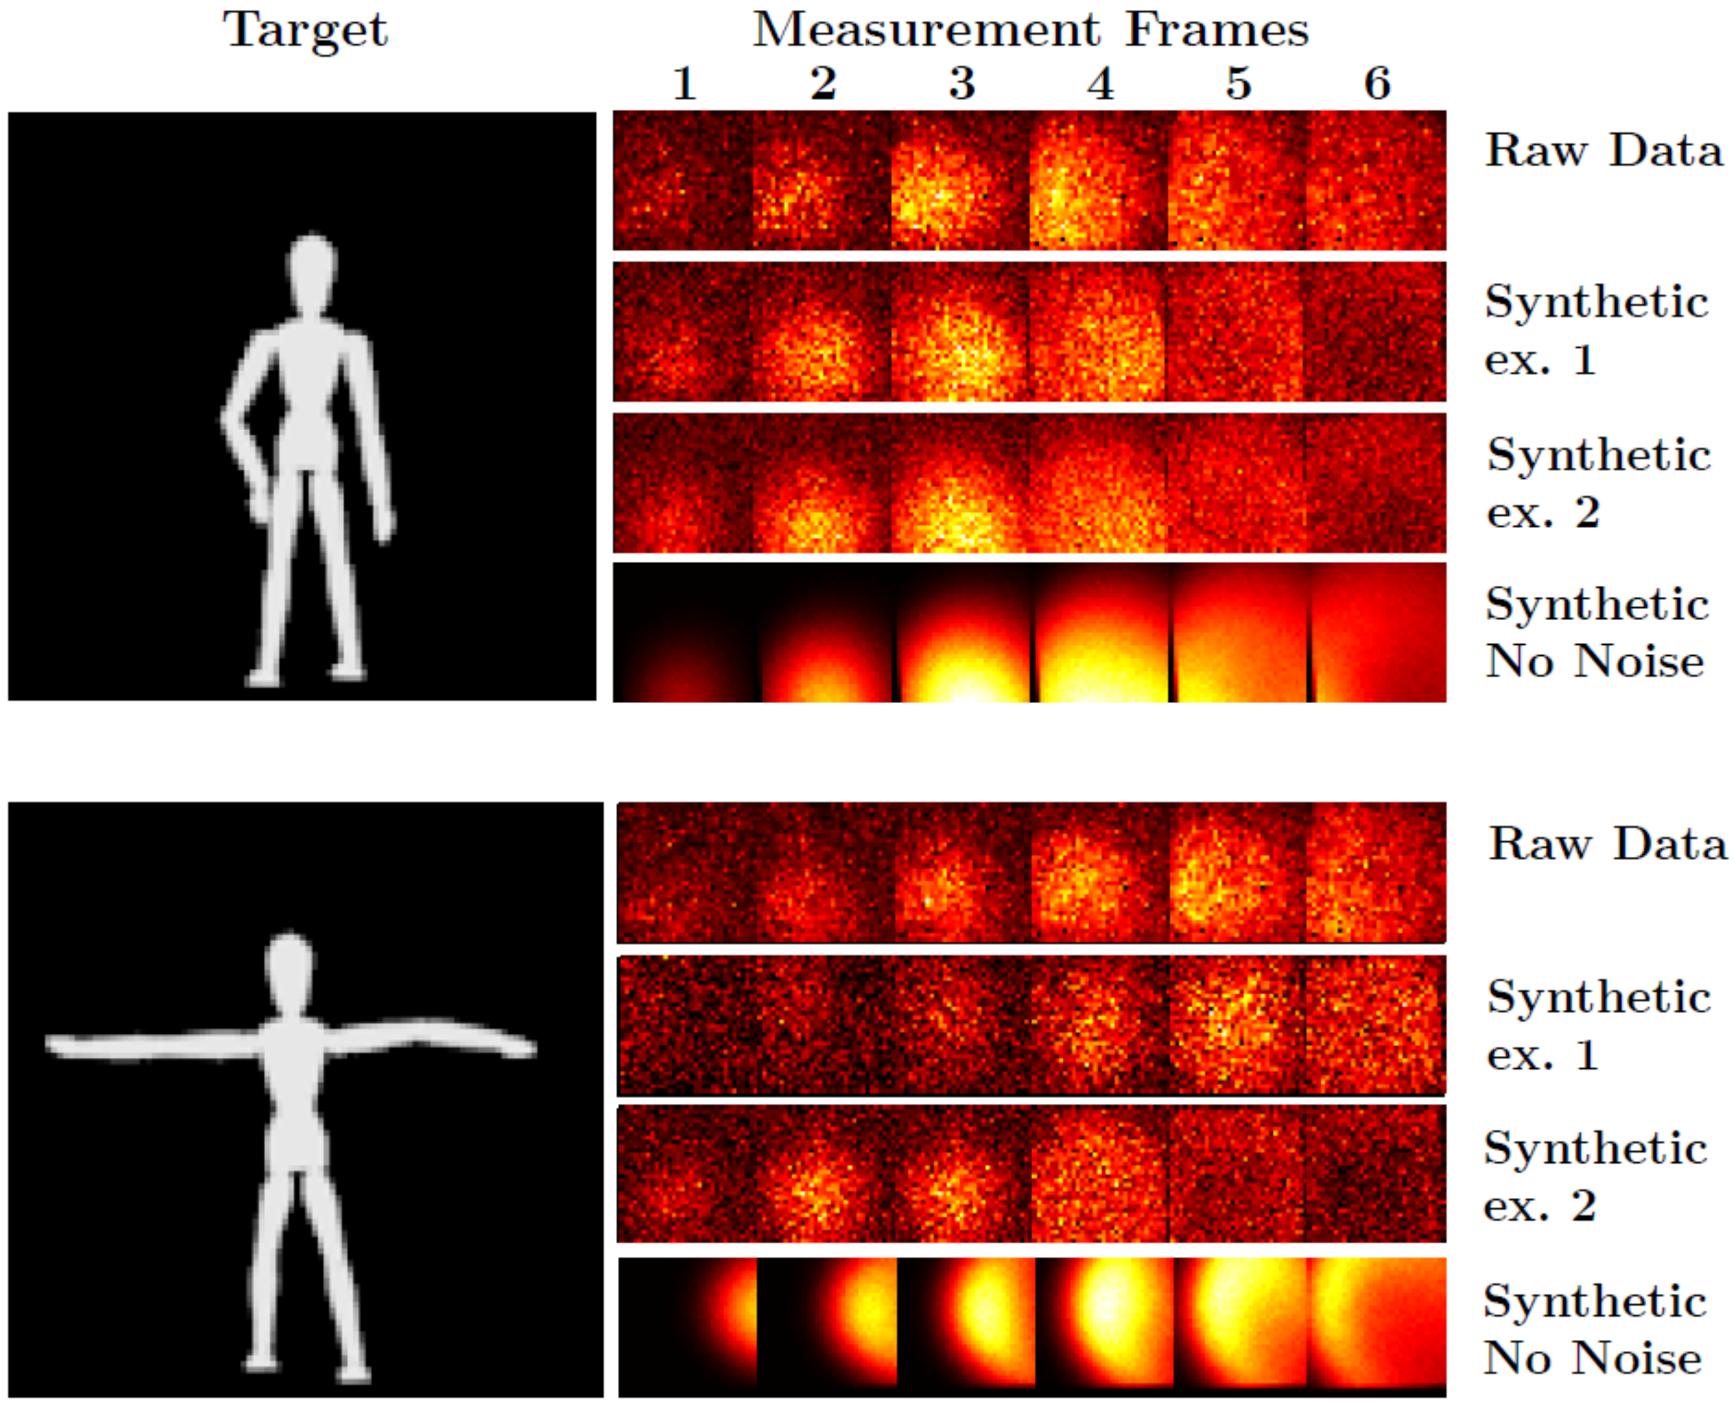

Target
Measurement Frames
1
2
3
4
5
6
Raw Data
Synthetic ex. 1
Synthetic ex. 2
Synthetic No Noise
Raw Data
Synthetic ex. 1
Synthetic ex. 2
Synthetic No Noise

Learning Invariants to Calibration Parameters

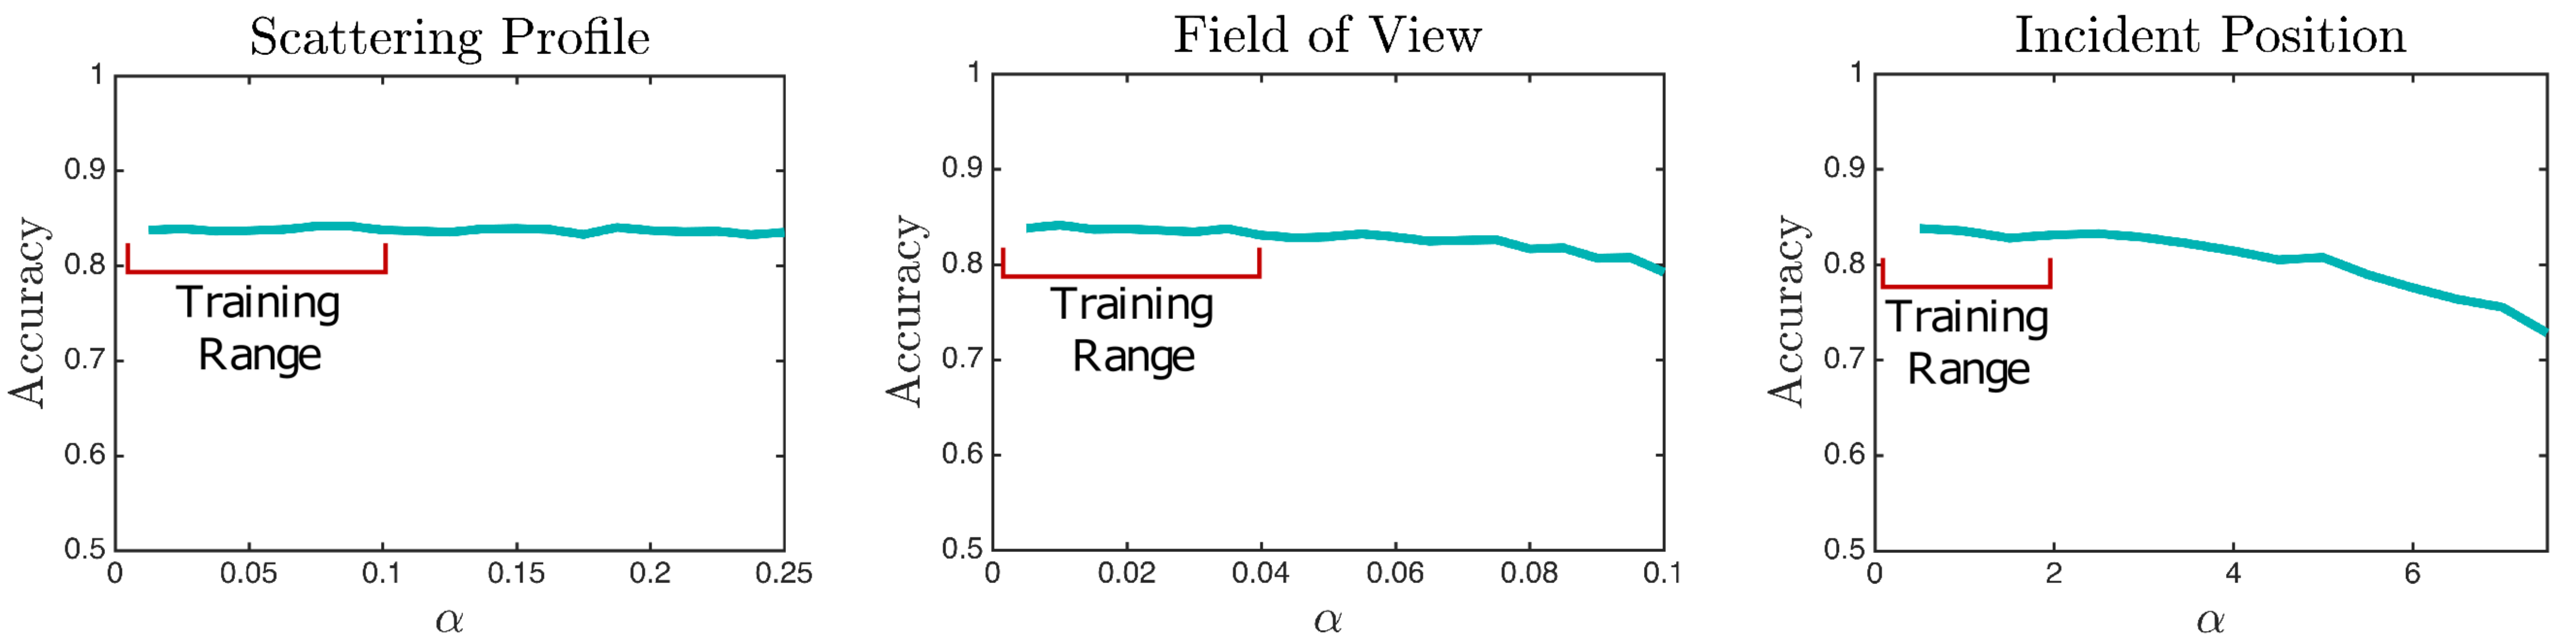

Experimental Results

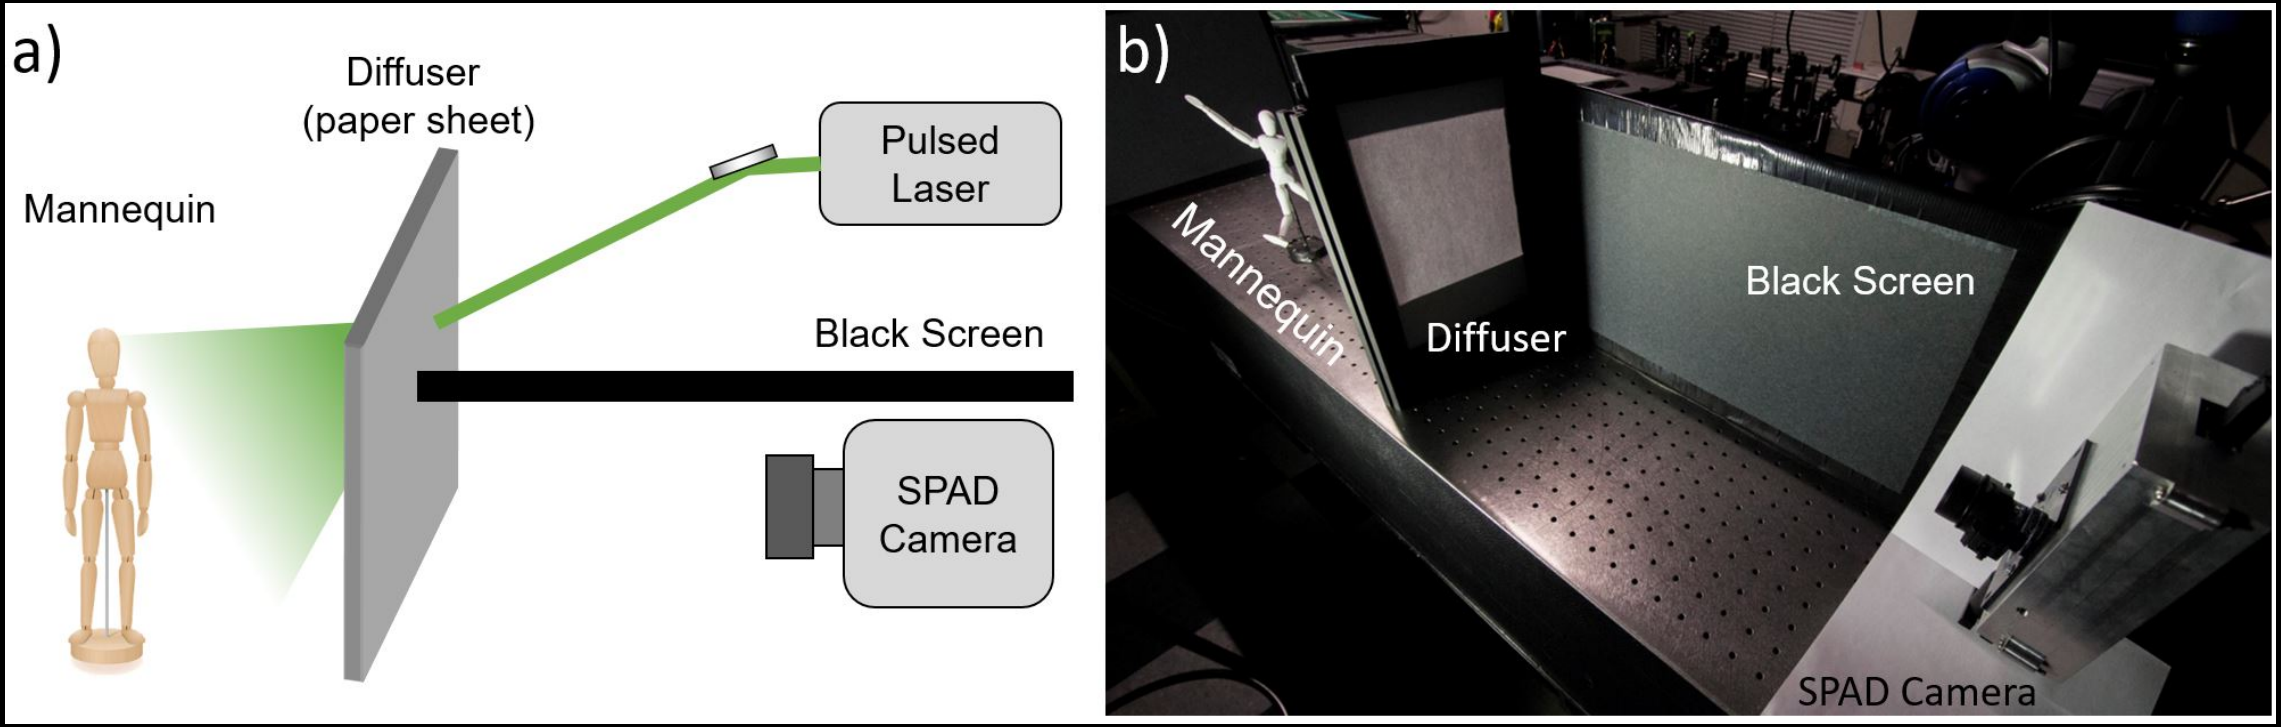

Train on Synthetic Data
Test on Lab Measurement

			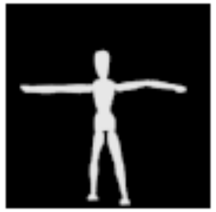
	**0.9**	**0.0**	**0.1**
	**0.3**	**0.7**	**0.0**
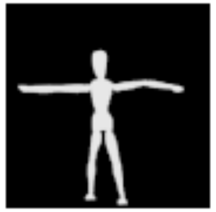	**0.1**	**0.2**	**0.7**

a)

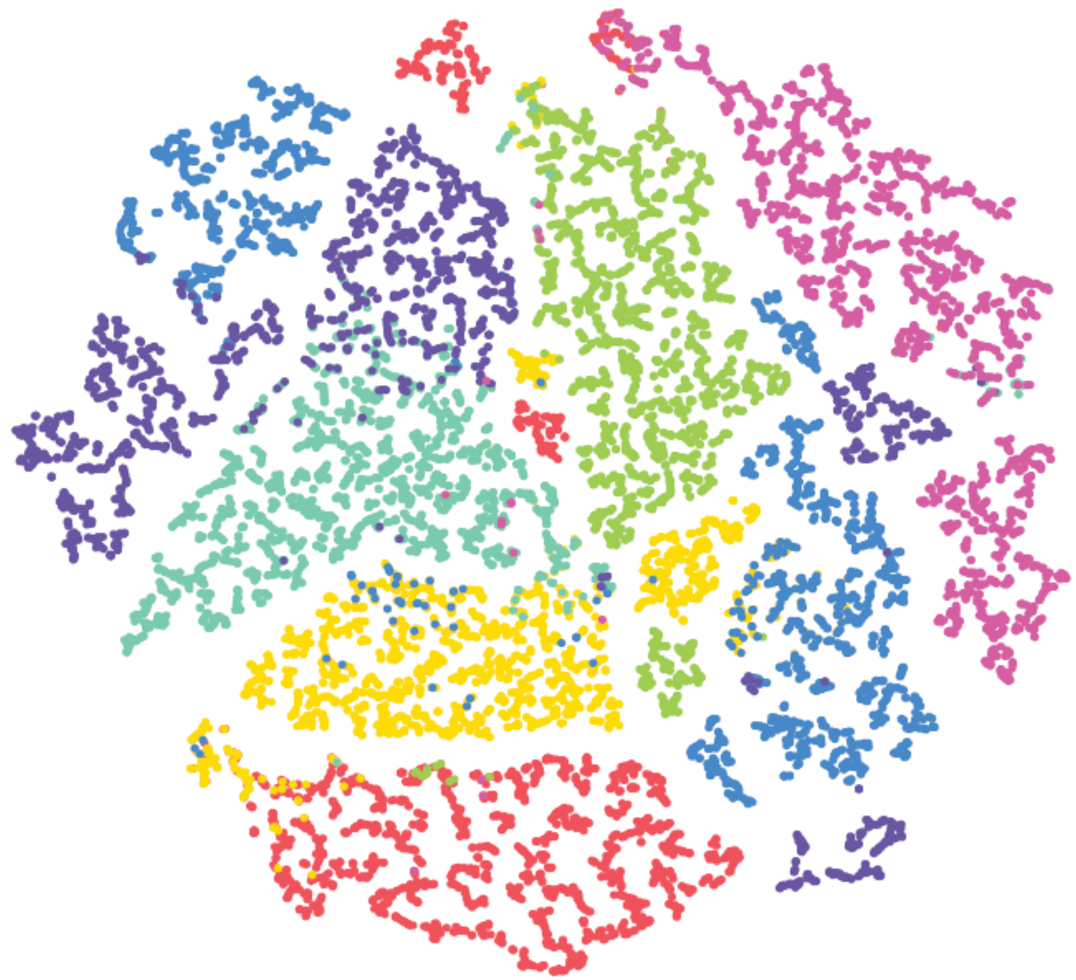

b)

	0.99	0.01	0.00	0.00	0.00	0.00	0.00
	0.02	0.91	0.00	0.01	0.02	0.04	0.00
	0.01	0.00	0.94	0.02	0.03	0.00	0.00
	0.03	0.30	0.02	0.92	0.00	0.00	0.00
	0.02	0.03	0.01	0.00	0.91	0.00	0.03
	0.00	0.08	0.00	0.00	0.00	0.89	0.03
	0.00	0.00	0.00	0.1	0.11	0.02	0.87

Do We Really Need a CNN?

Training set	Clean dataset	Realistic dataset
Mean Example	33.3	33.3
KNN	53.0	30.0
SVM	57.1	20.0
Random forest	68.8	30.0
Single layer network	68.2	23.8
Our CNN	84.0	76.6

What Does The Network Learn?

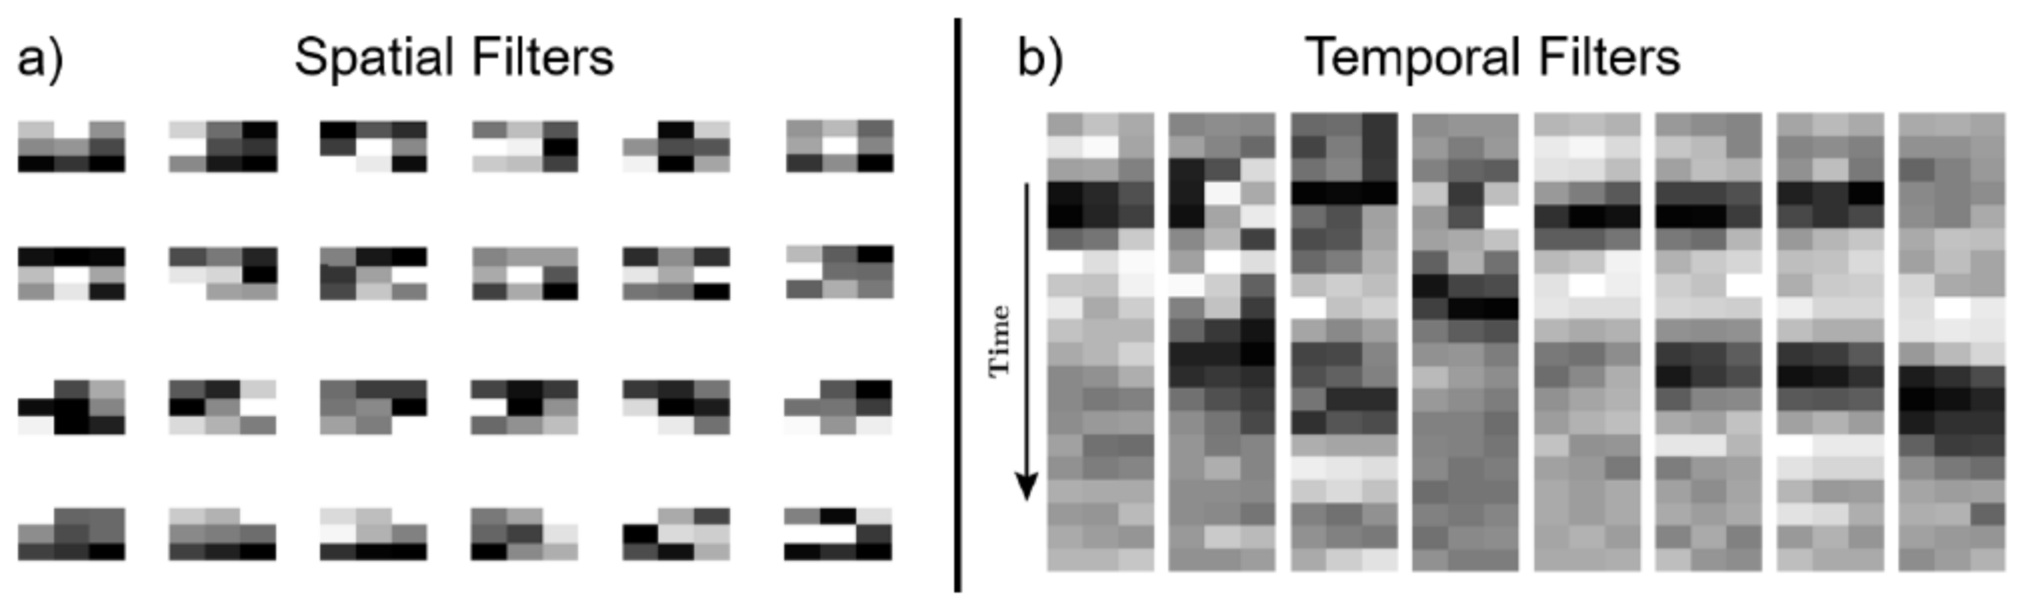

Seeing Around Corners

Flash Photography for NLOS Localization and Identification

Tancik, Satat, Raskar, “Flash Photography for Data-Driven Hidden Scene Recovery,” 2017.
Matthew Tancik, M.Eng Thesis, 2018.

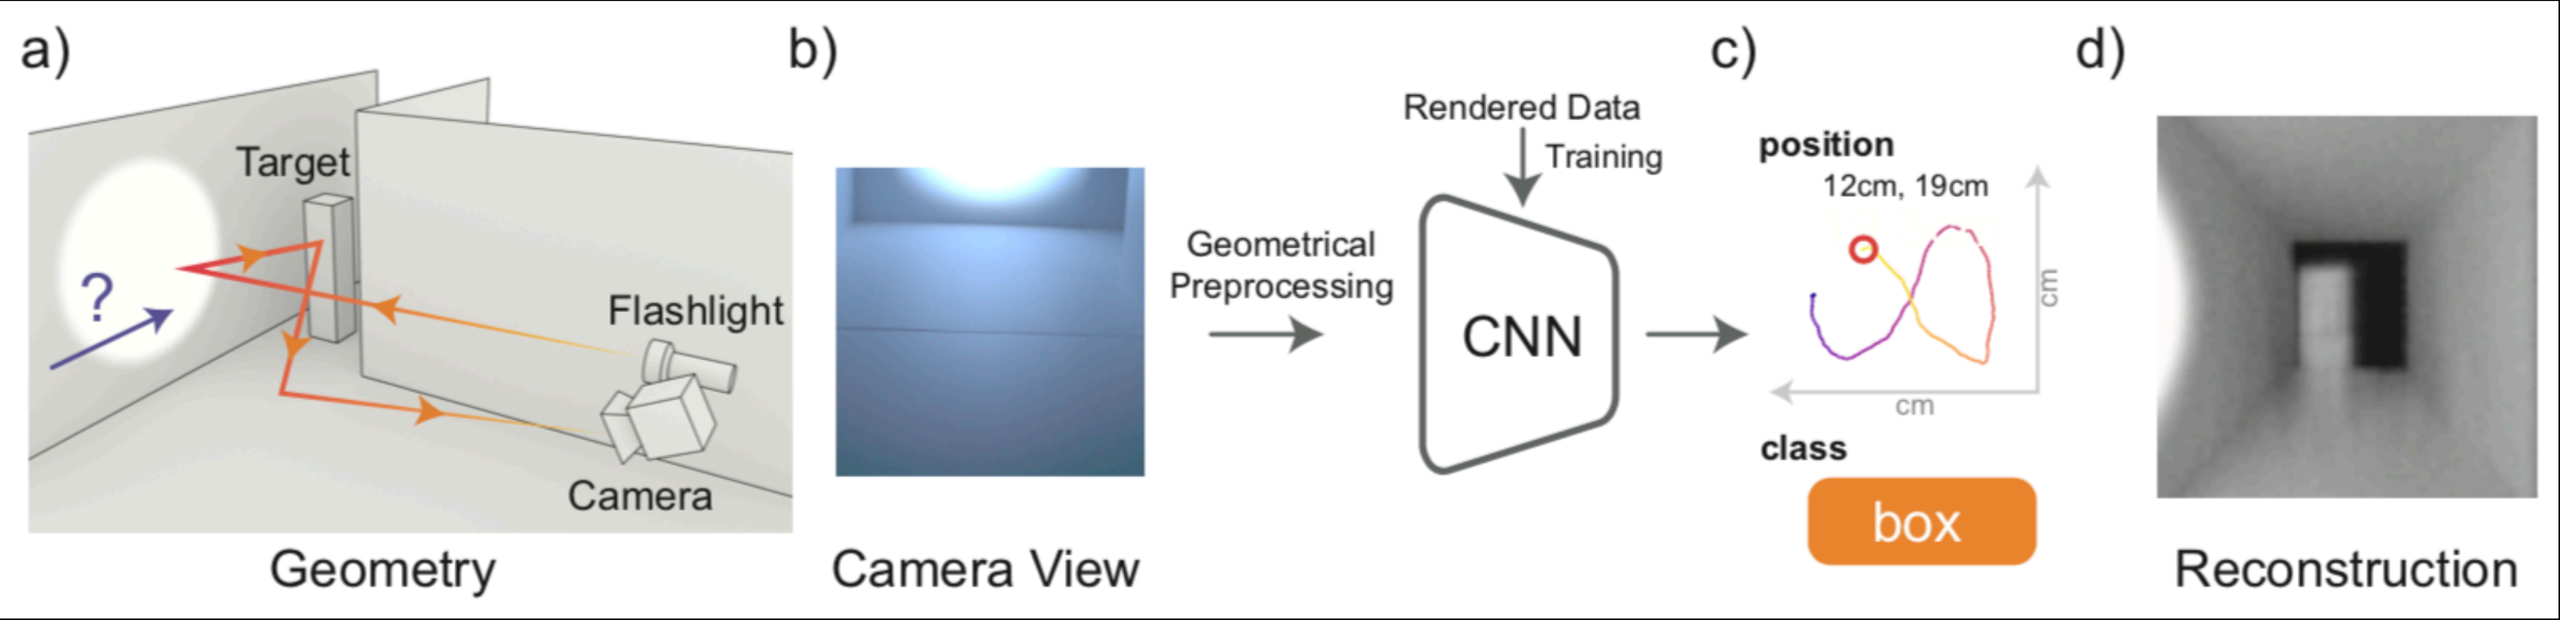
a)
Target
?
Flashlight
Camera
Geometry
b)
Camera View
Geometrical
Preprocessing
Rendered Data
Training
CNN
c)
position
12cm, 19cm
cm
cm
class
box
d)
Reconstruction

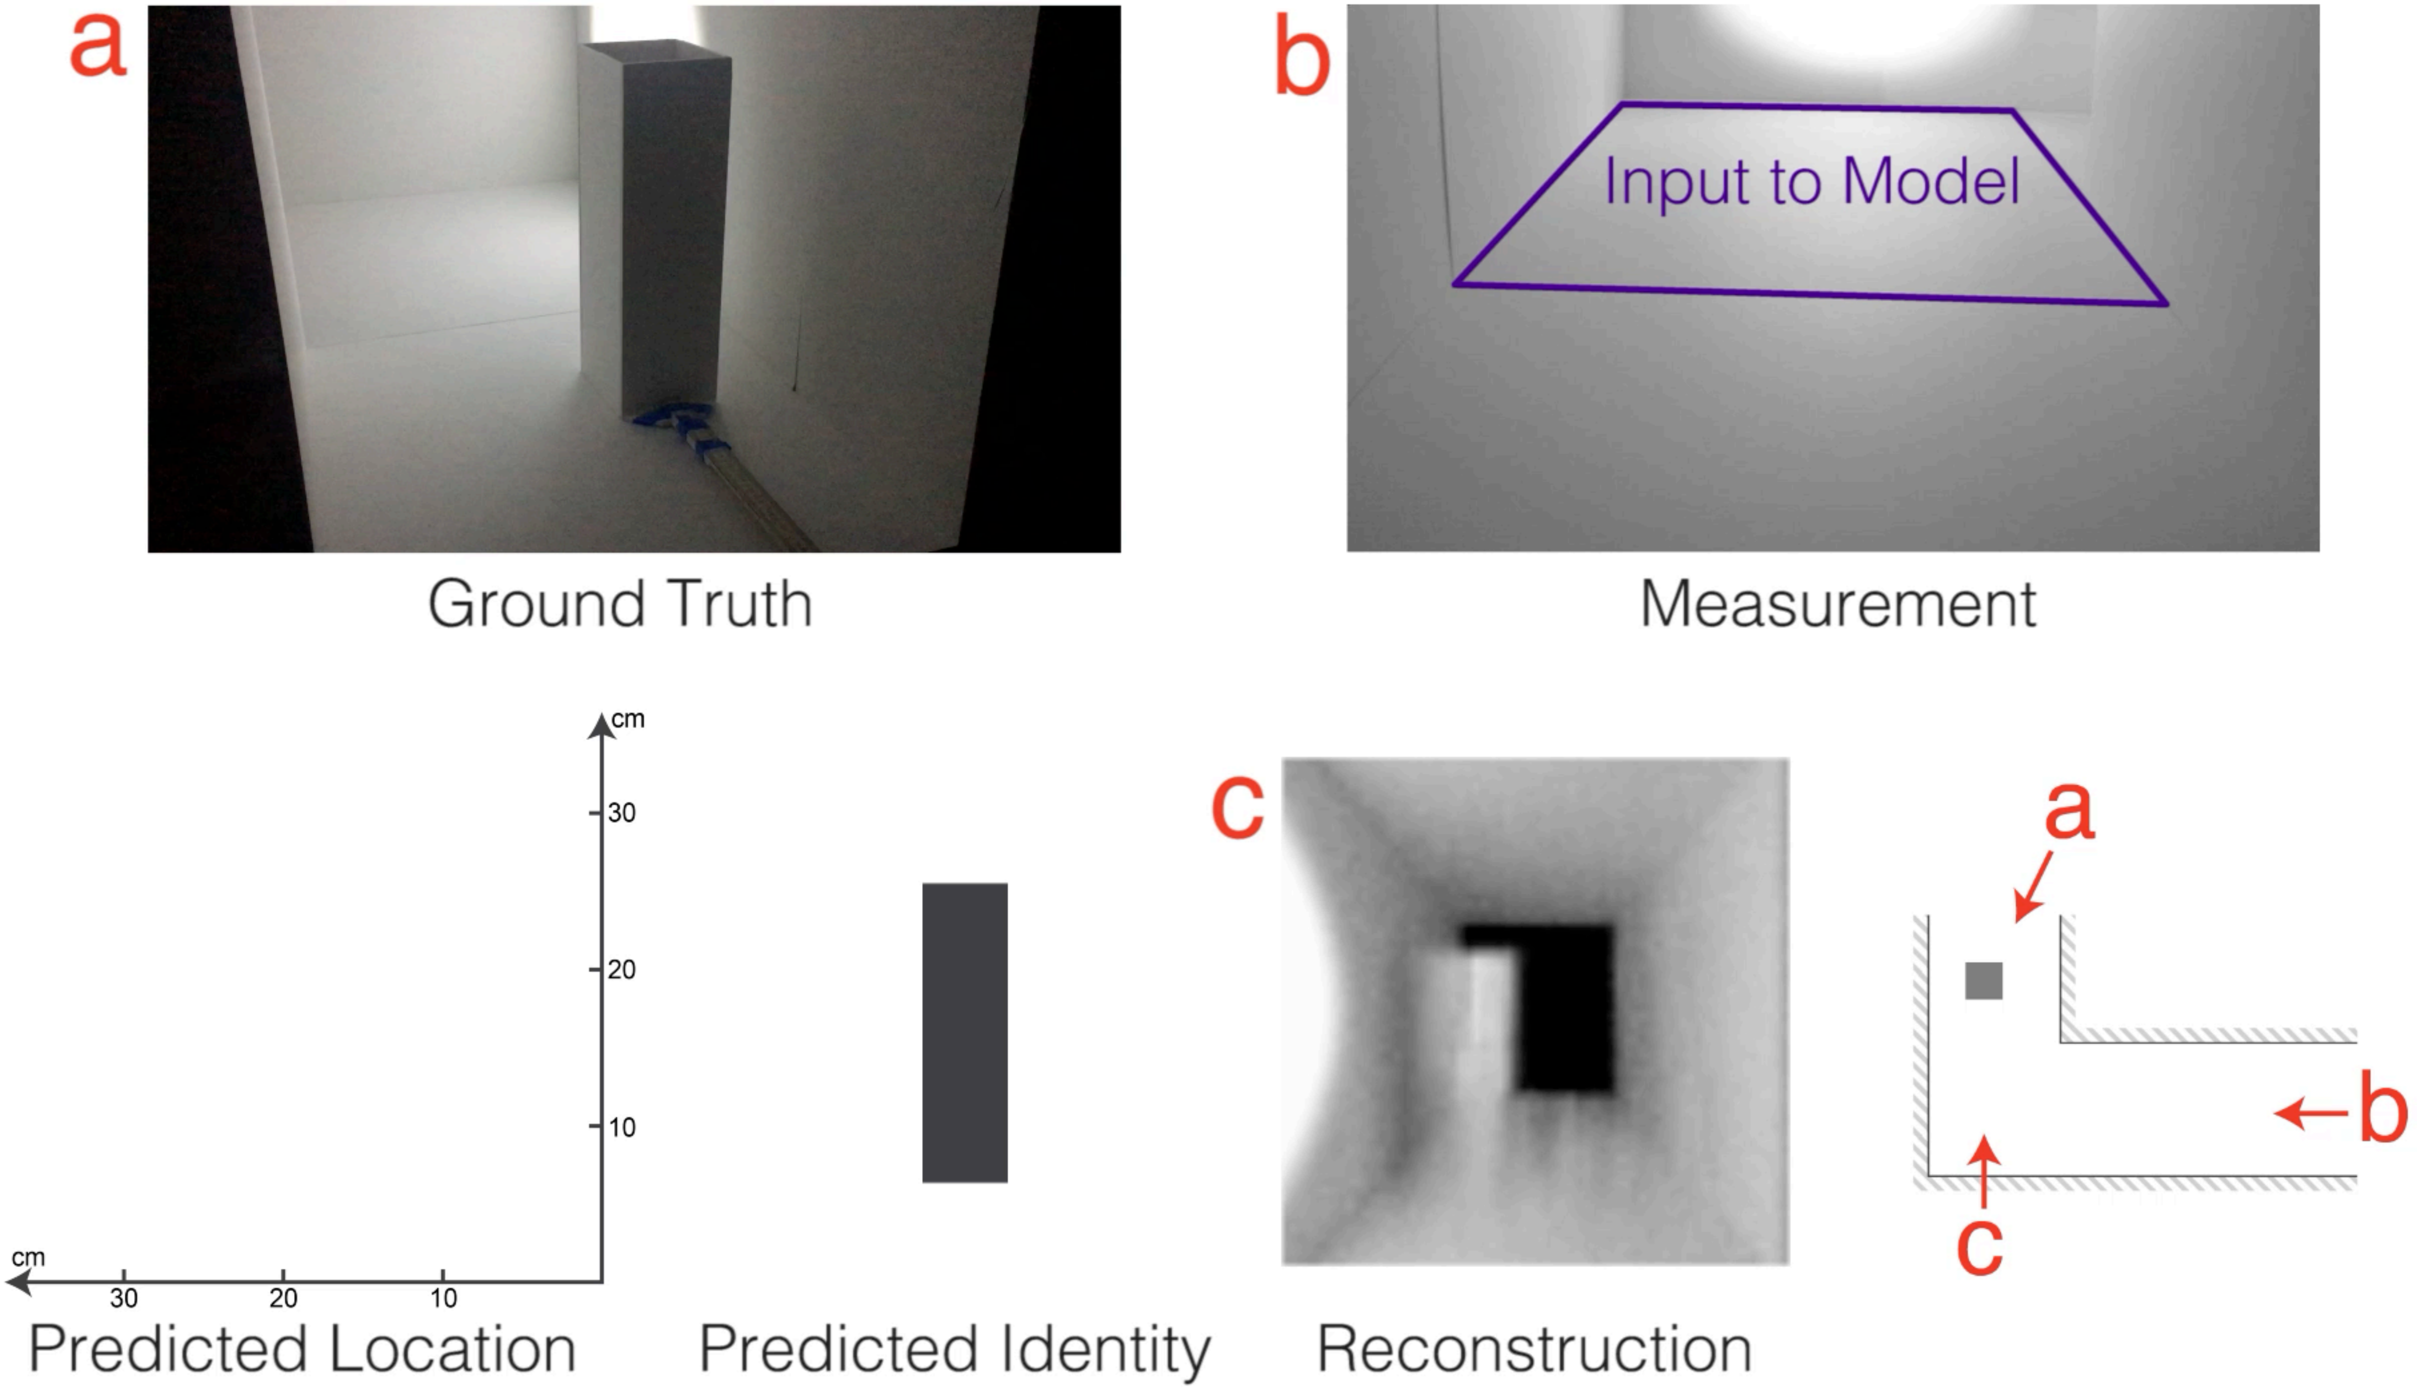
a
Ground Truth
b
Input to Model
Measurement
cm
30
20
10
cm
30
20
10
Predicted Location
Predicted Identity
c
Reconstruction
a
b
c

Preprocessing Steps

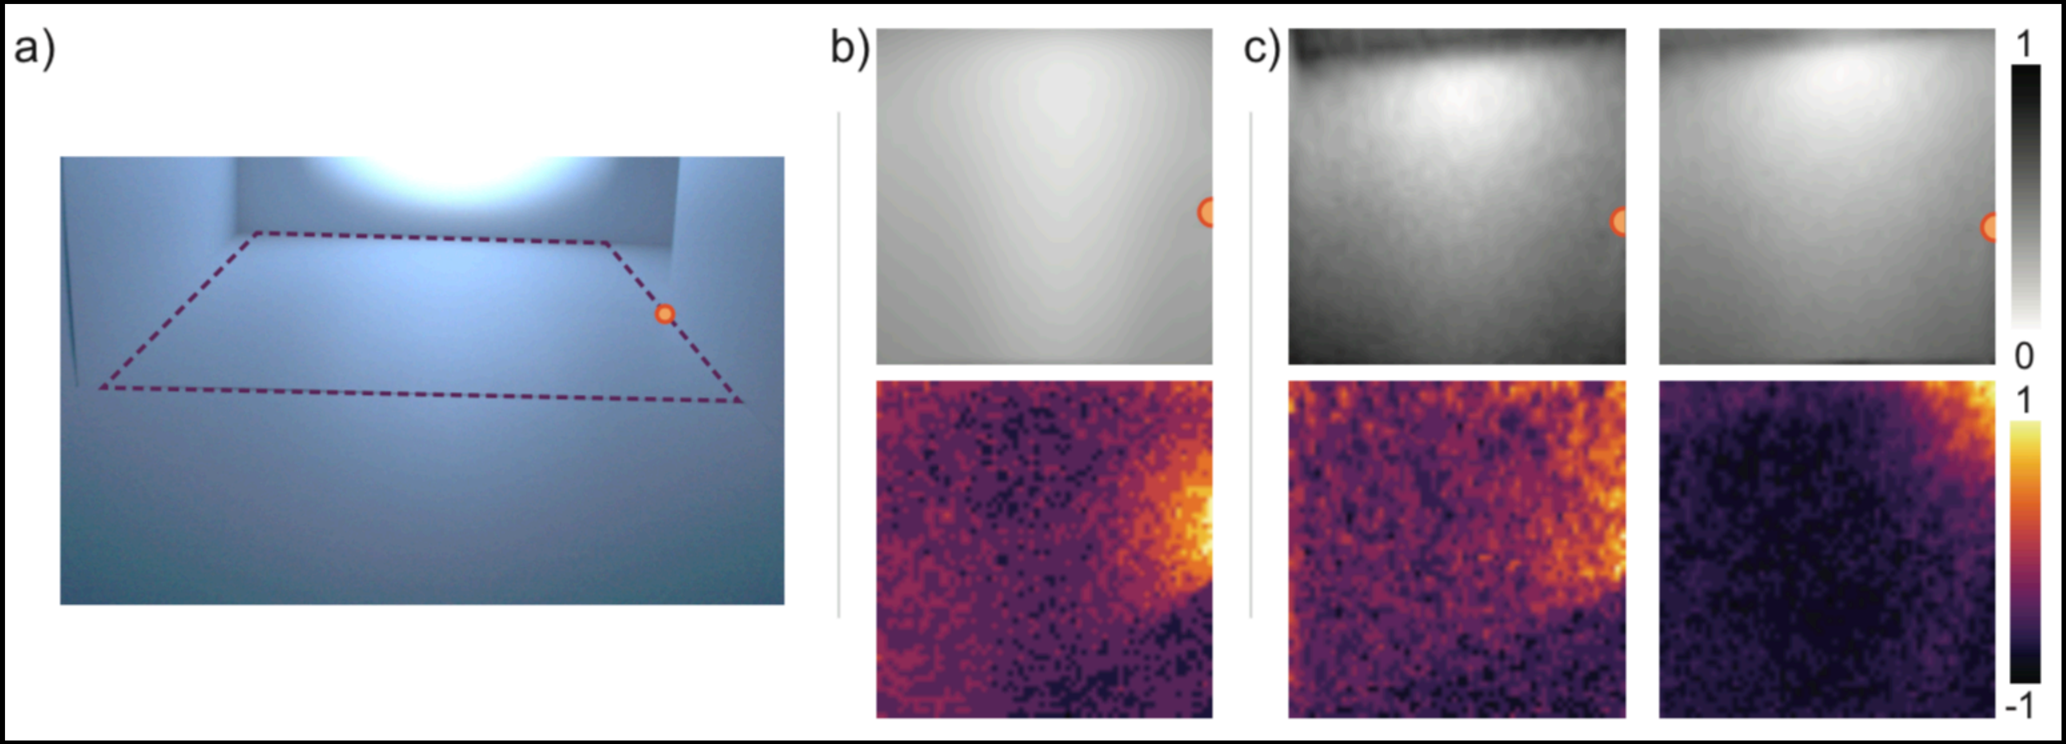

Localization

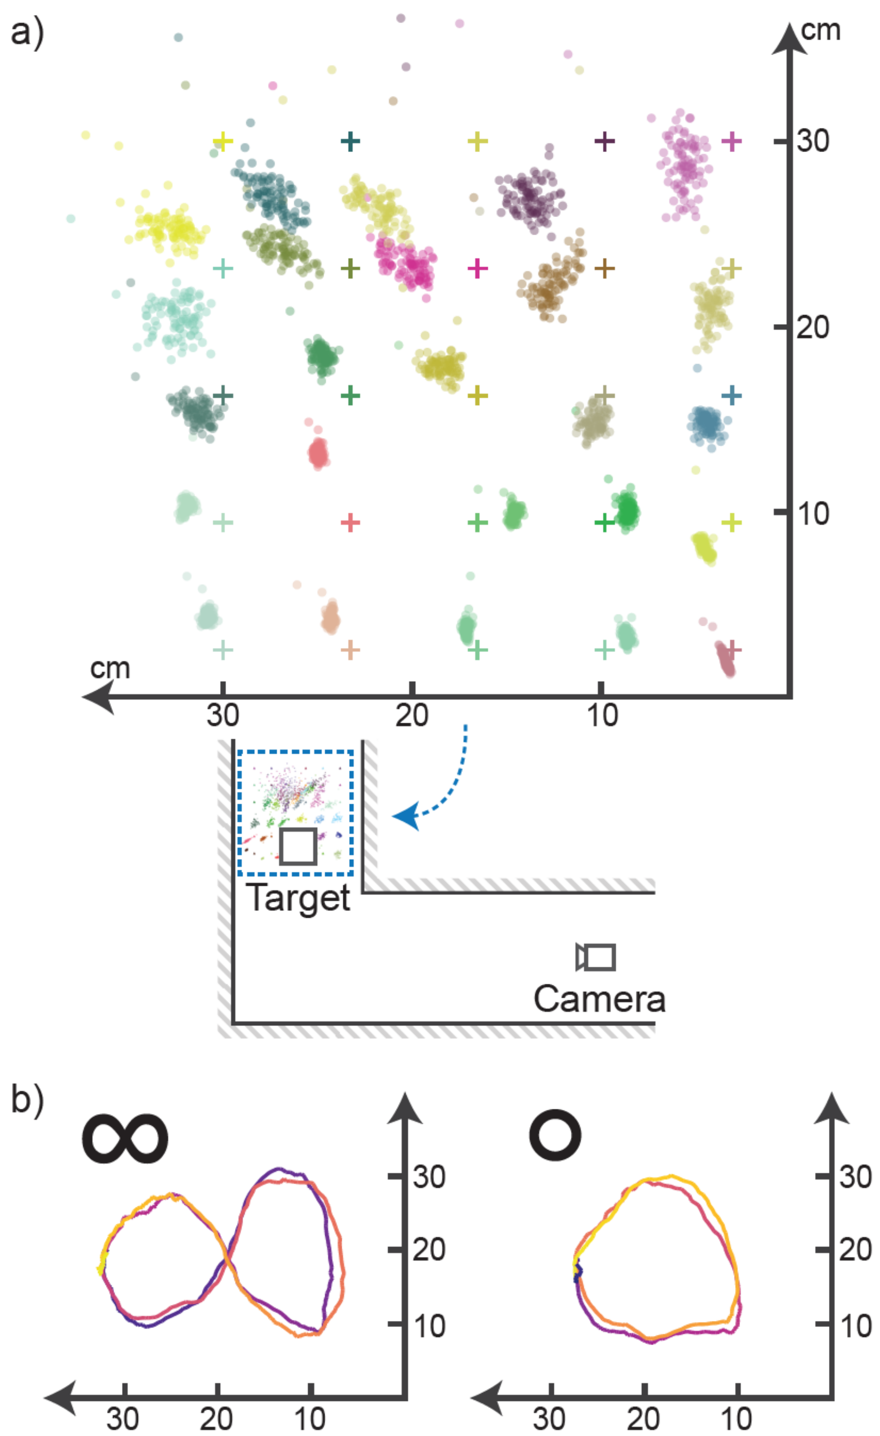

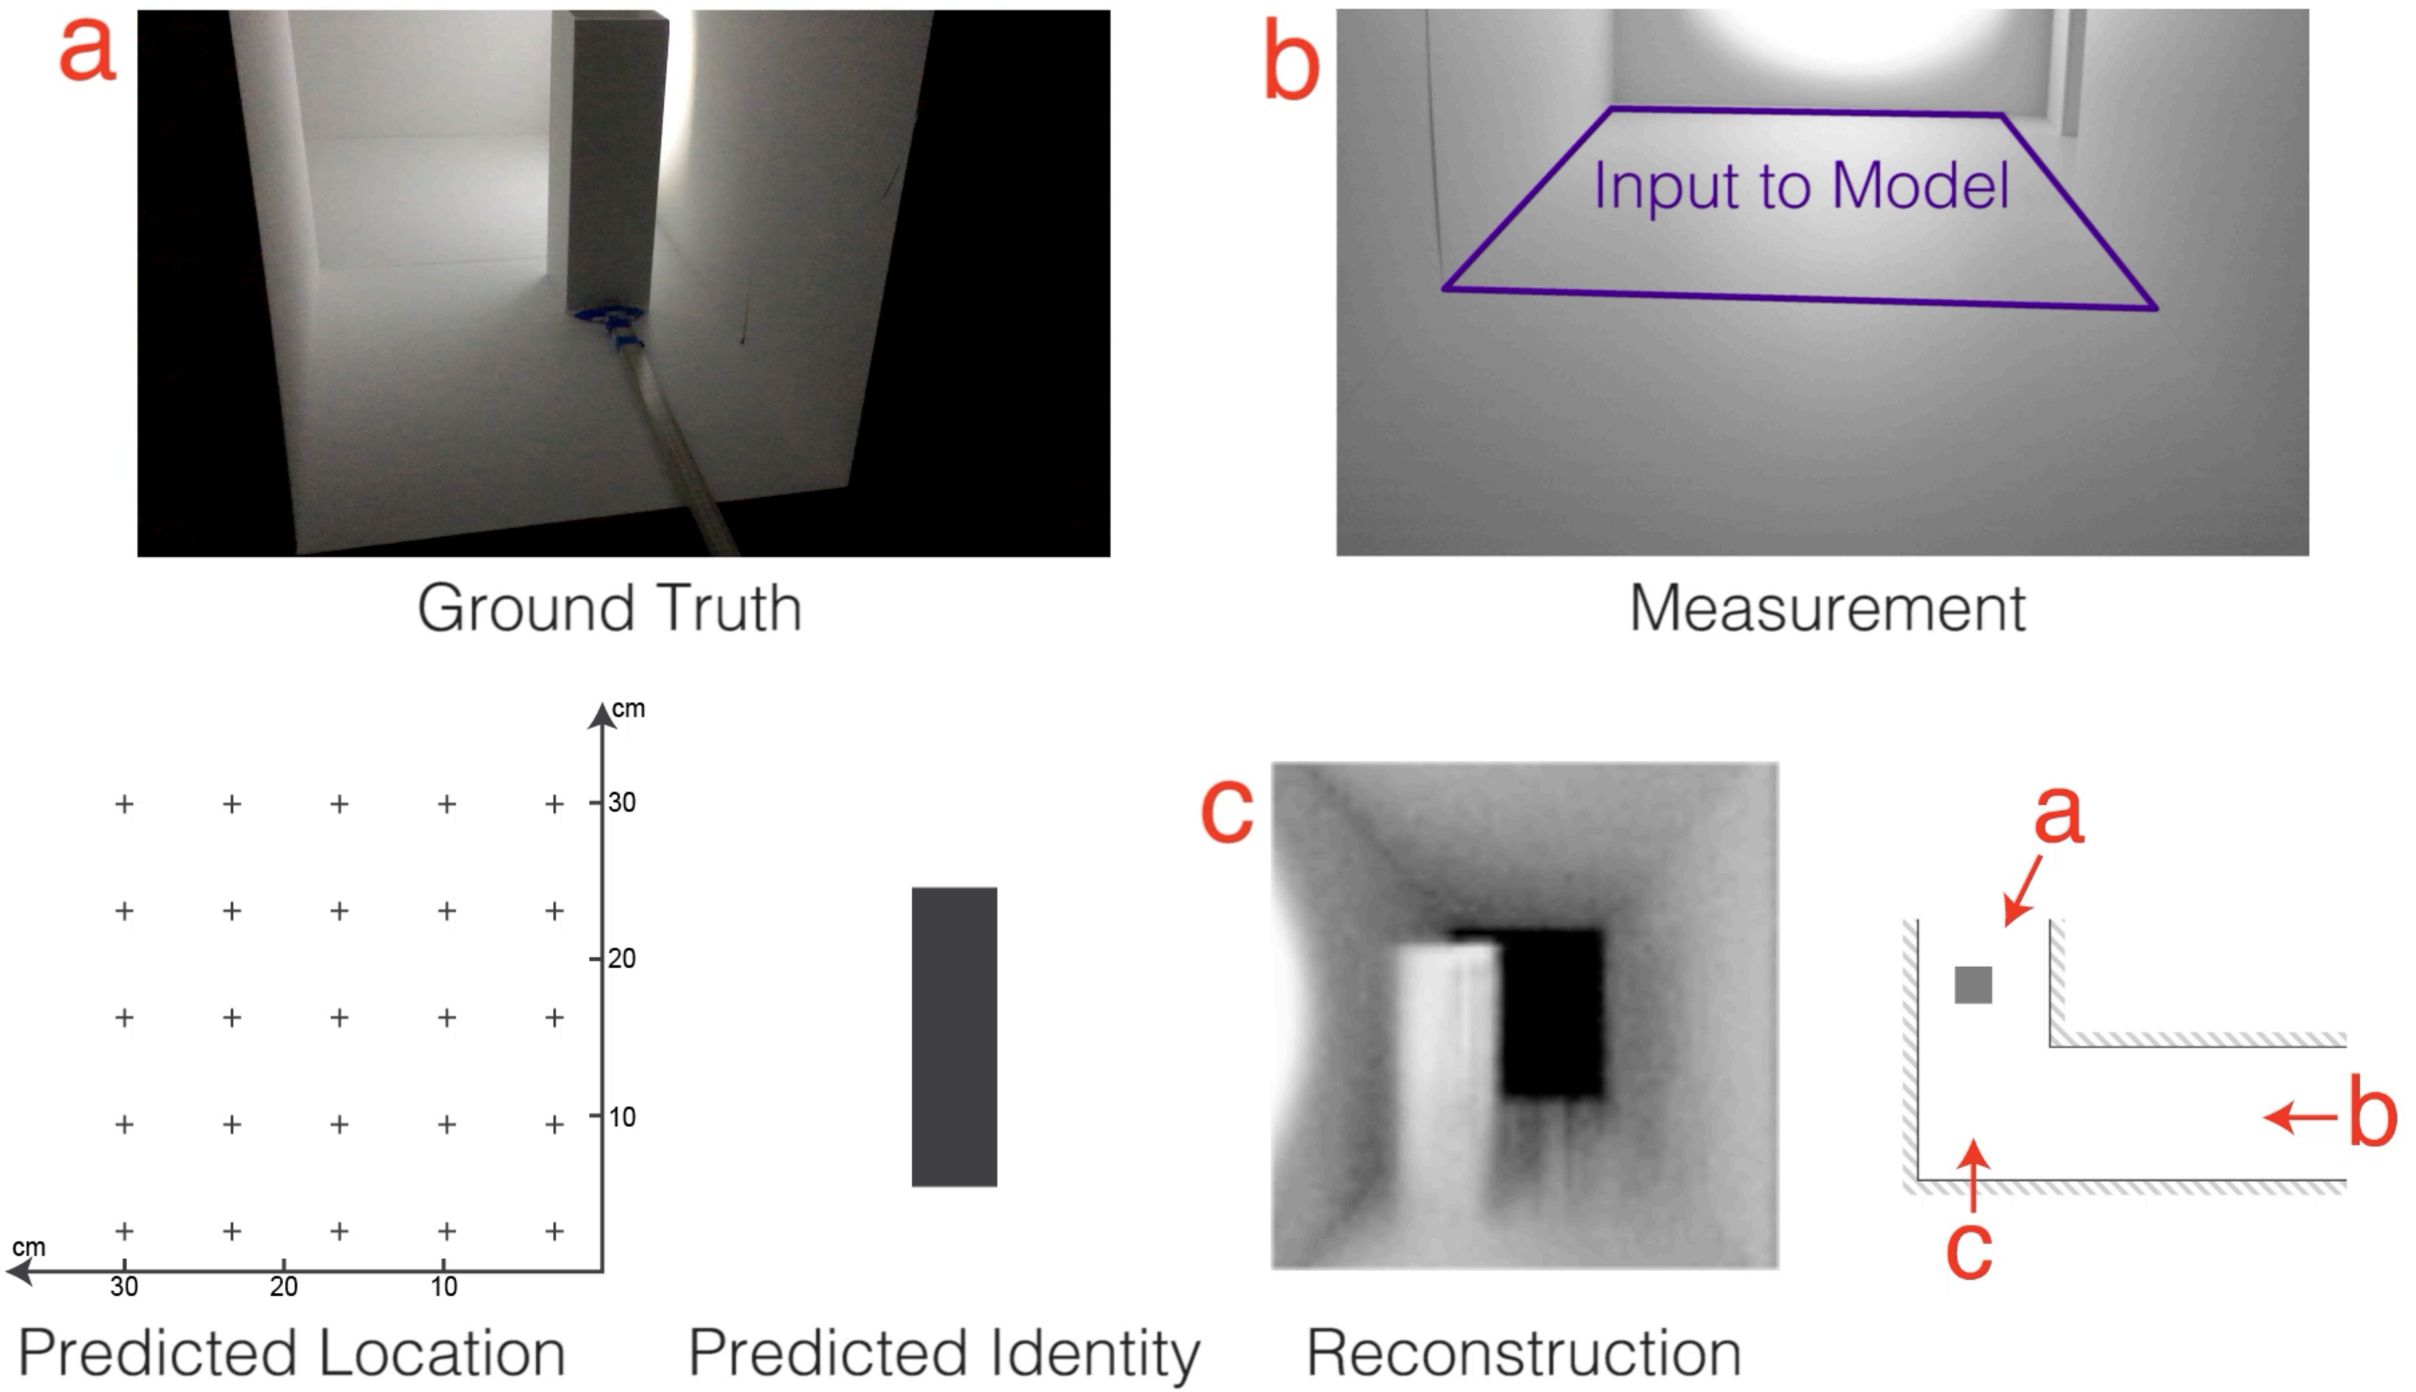
a
Ground Truth
b
Input to Model
Measurement
cm
30
20
10
cm
30
20
10
Predicted Location
Predicted Identity
c
Reconstruction
a
b
c

Identification

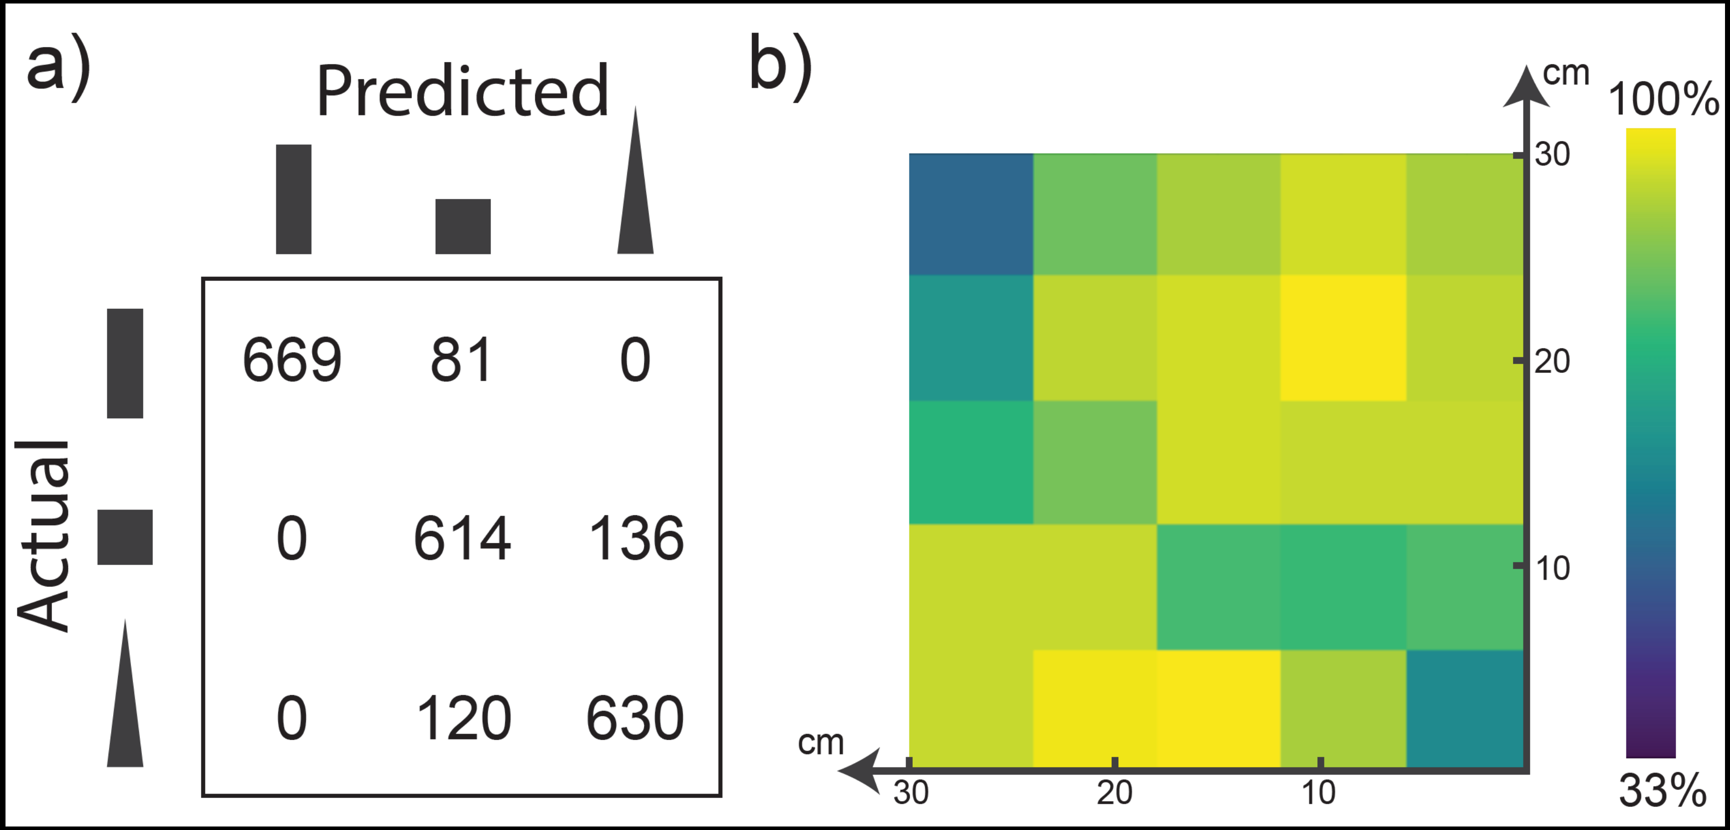

Reconstruction

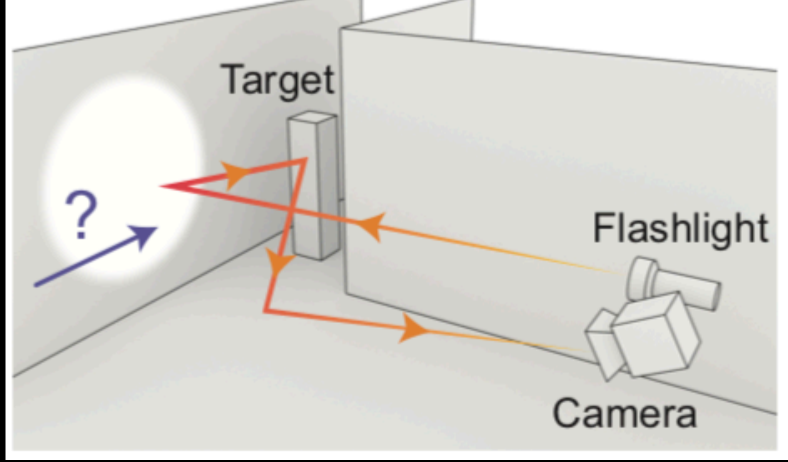

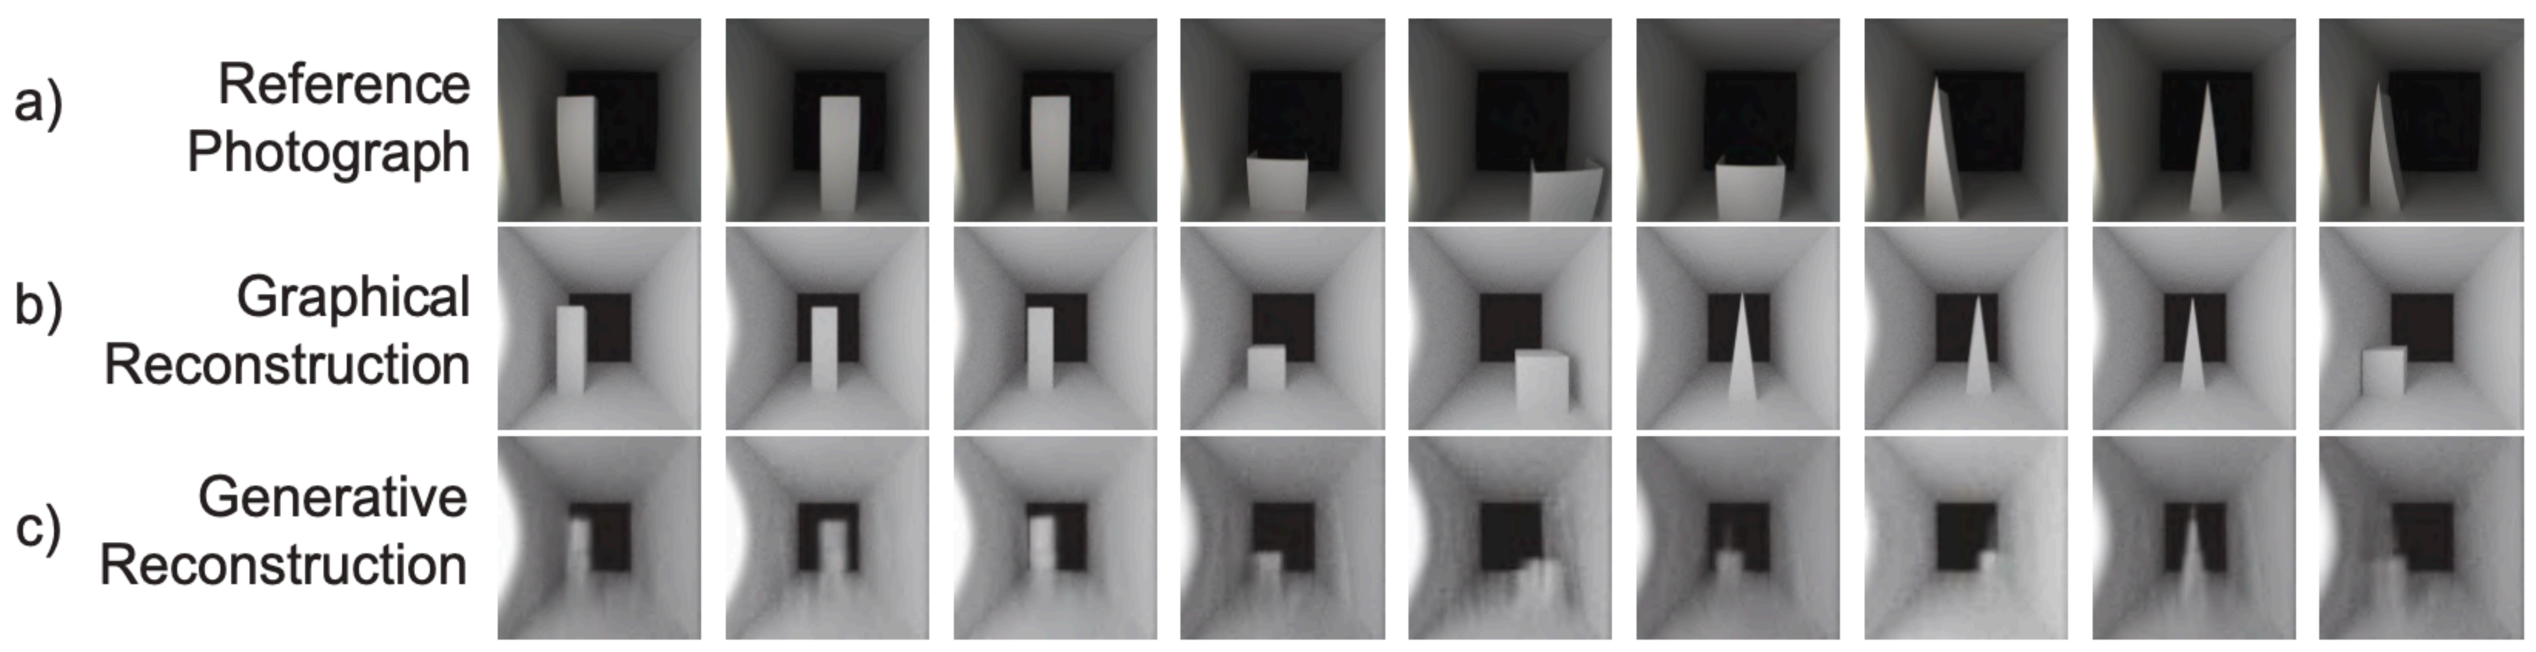

Network Architecture

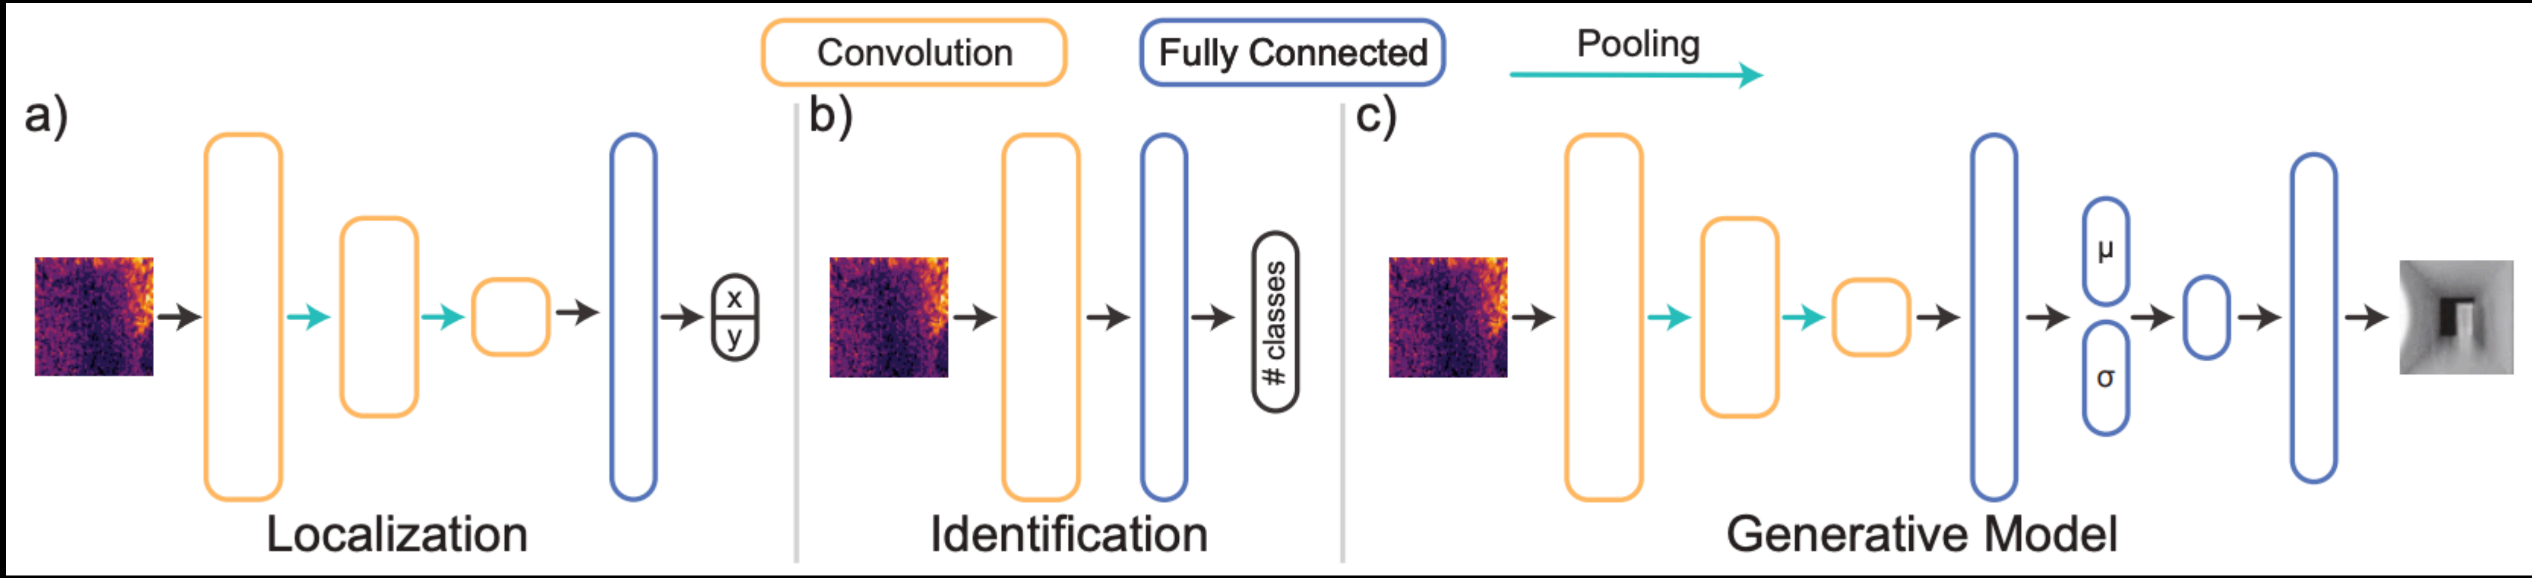

Where is the Information?

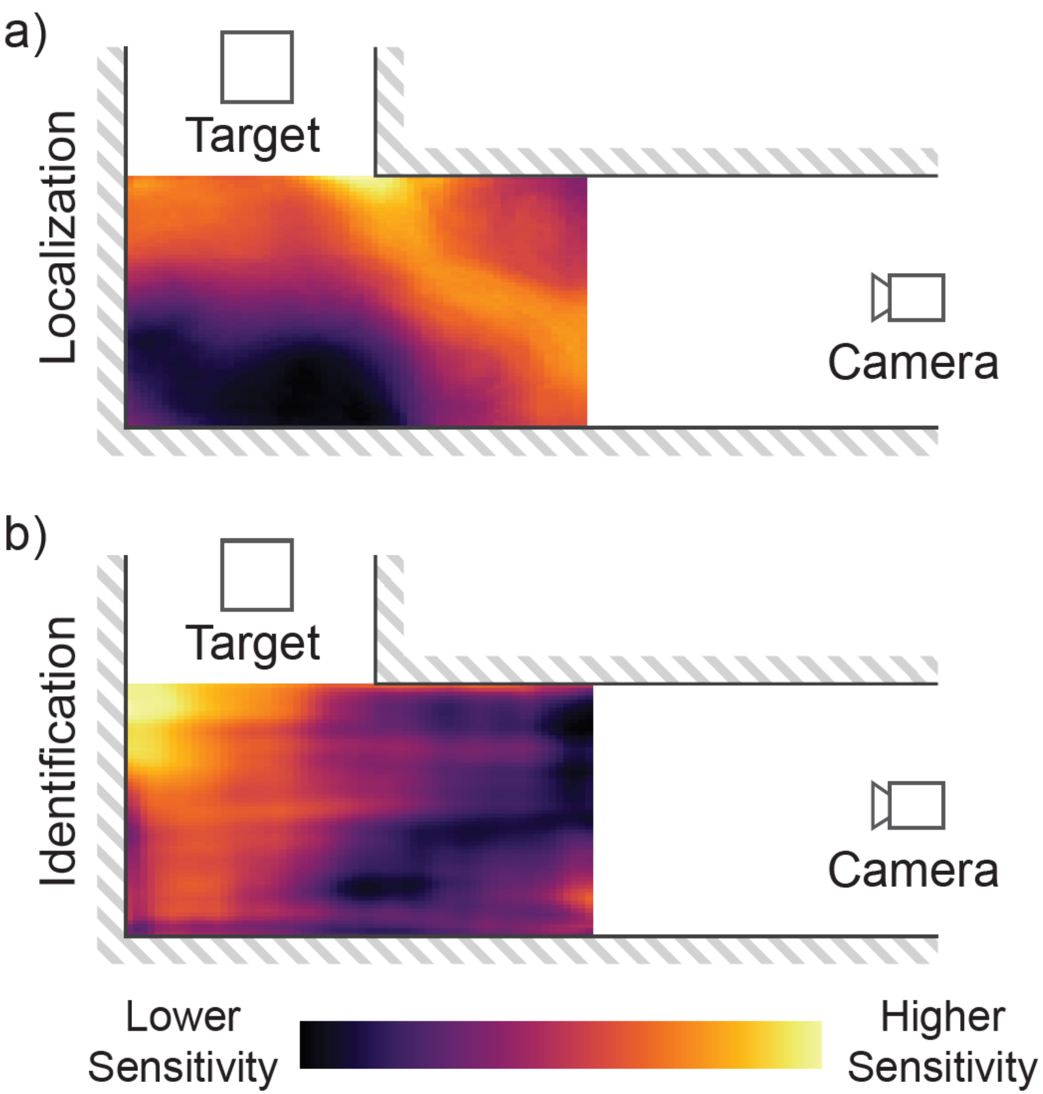

c)

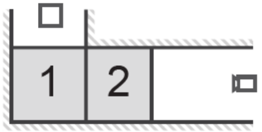

Visible Region	Localization	Identification
1	11.78 cm	51.0%
2	12.15 cm	34.7%
1&2	2.61 cm	87.7%

Training with Real Data

a)

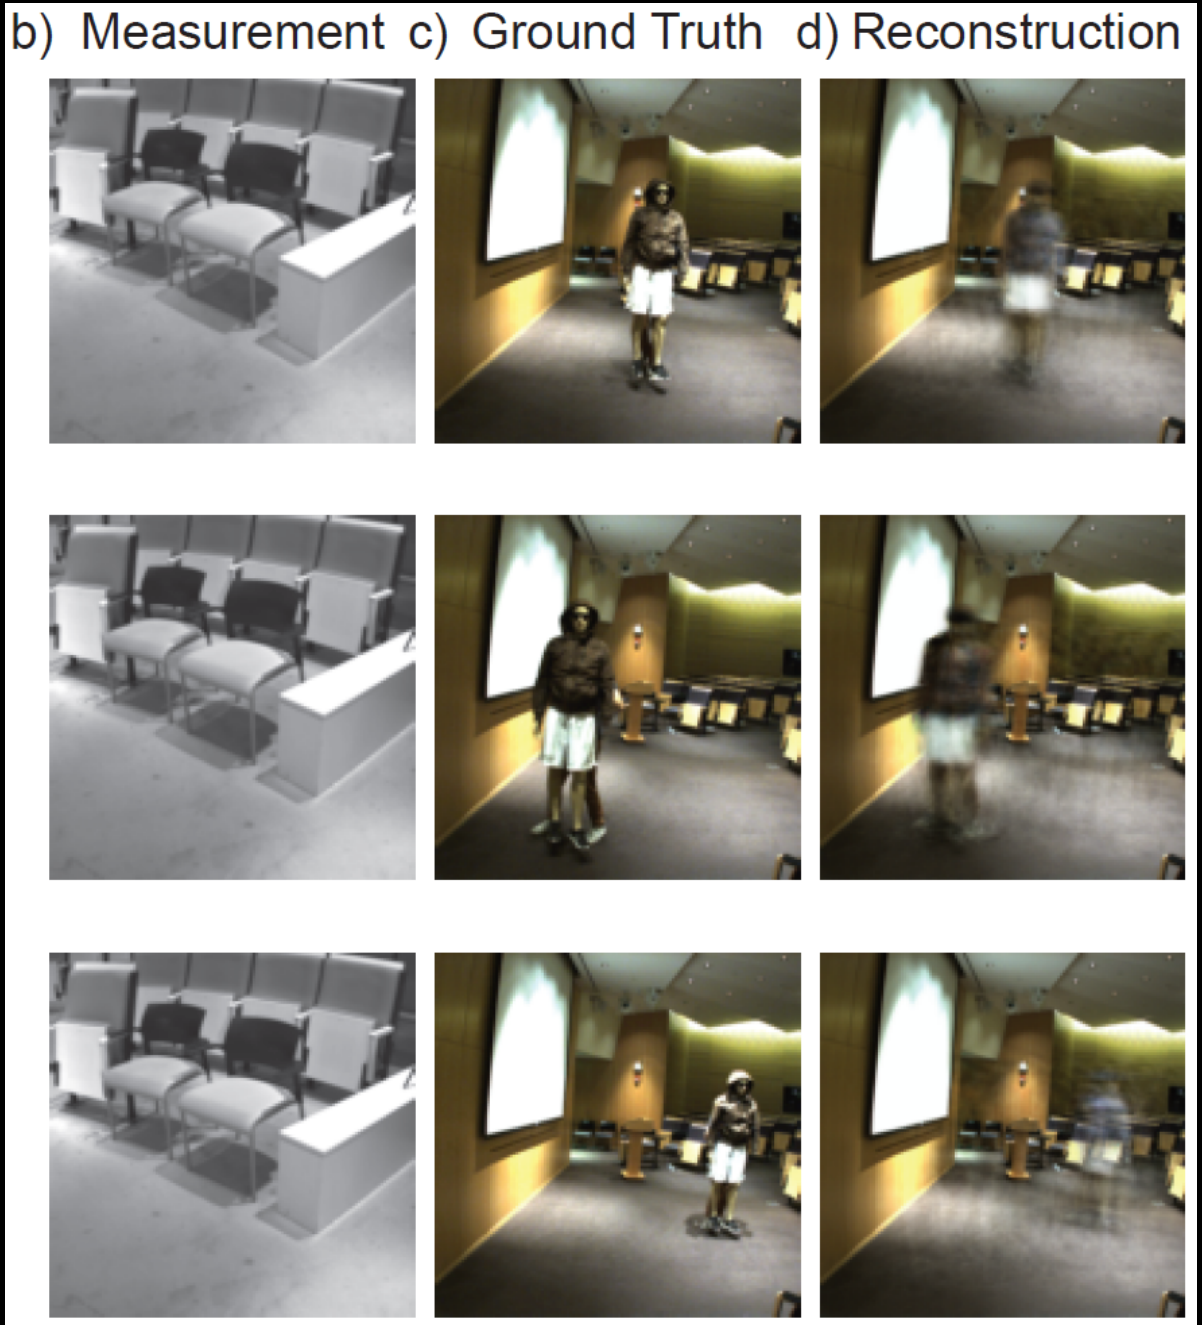

Training Data

Base Camera

Target Camera

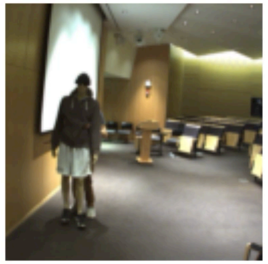

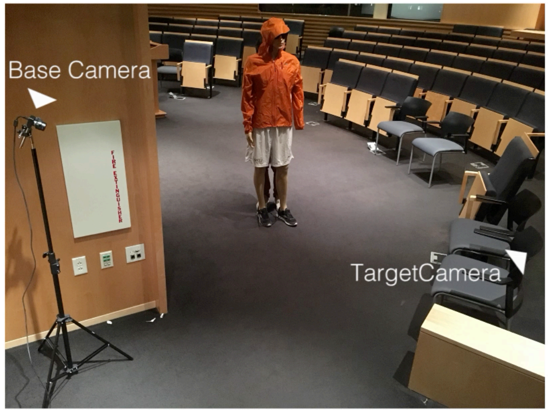

Test Data

Camera Measurement

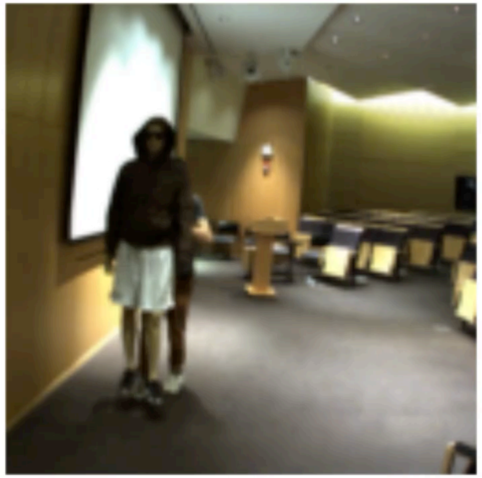

Ground Truth

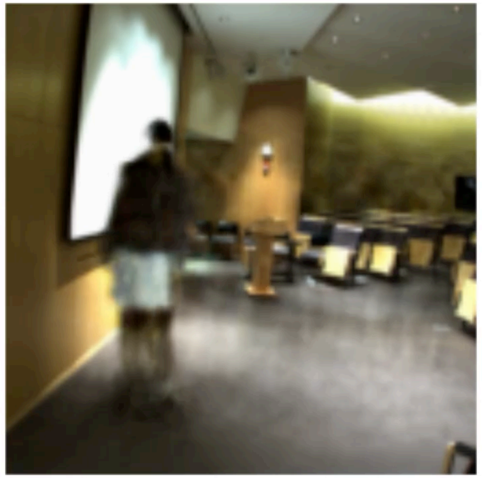

Reconstruction

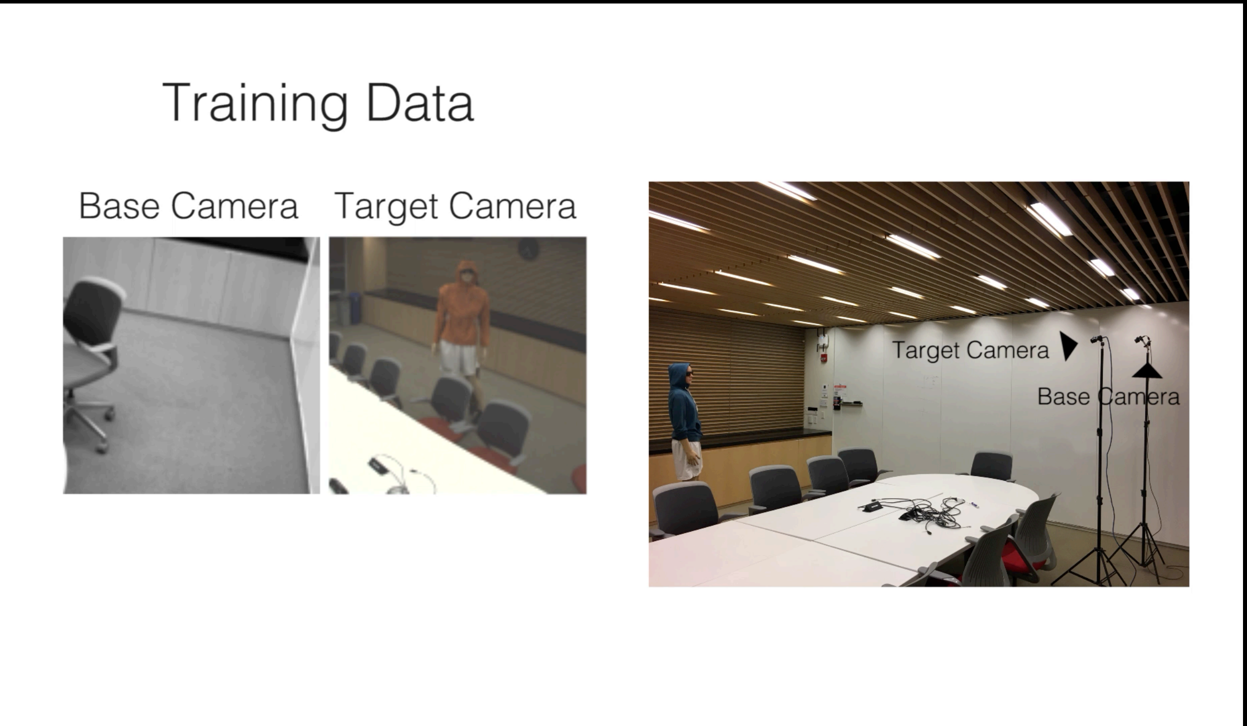
Training Data
Base Camera
Target Camera
Target Camera
Base Camera

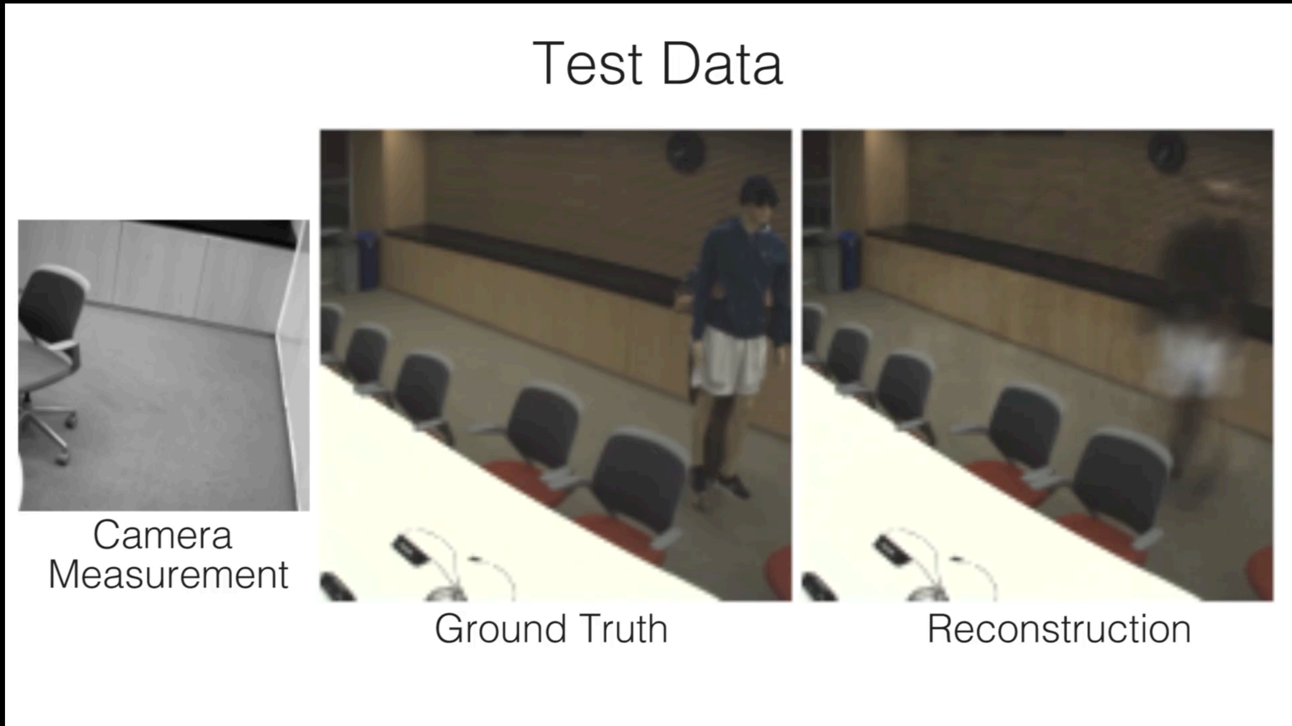
Test Data
Camera Measurement
Ground Truth
Reconstruction

We Don't Really Need A Corner

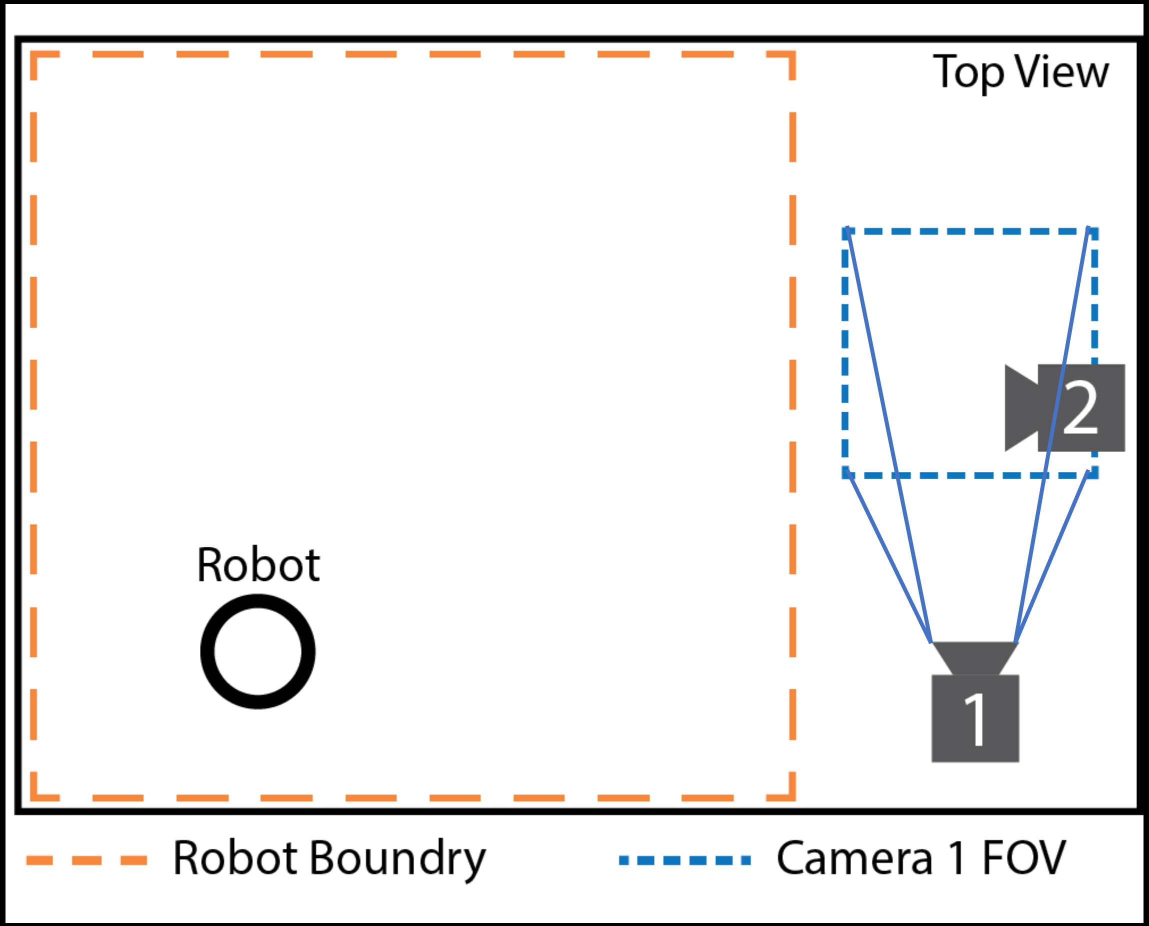

Tancik, Swedish, Satat, Raskar, "Data-Driven Non-Line-of-Sight Imaging With A Traditional Camera," 2018

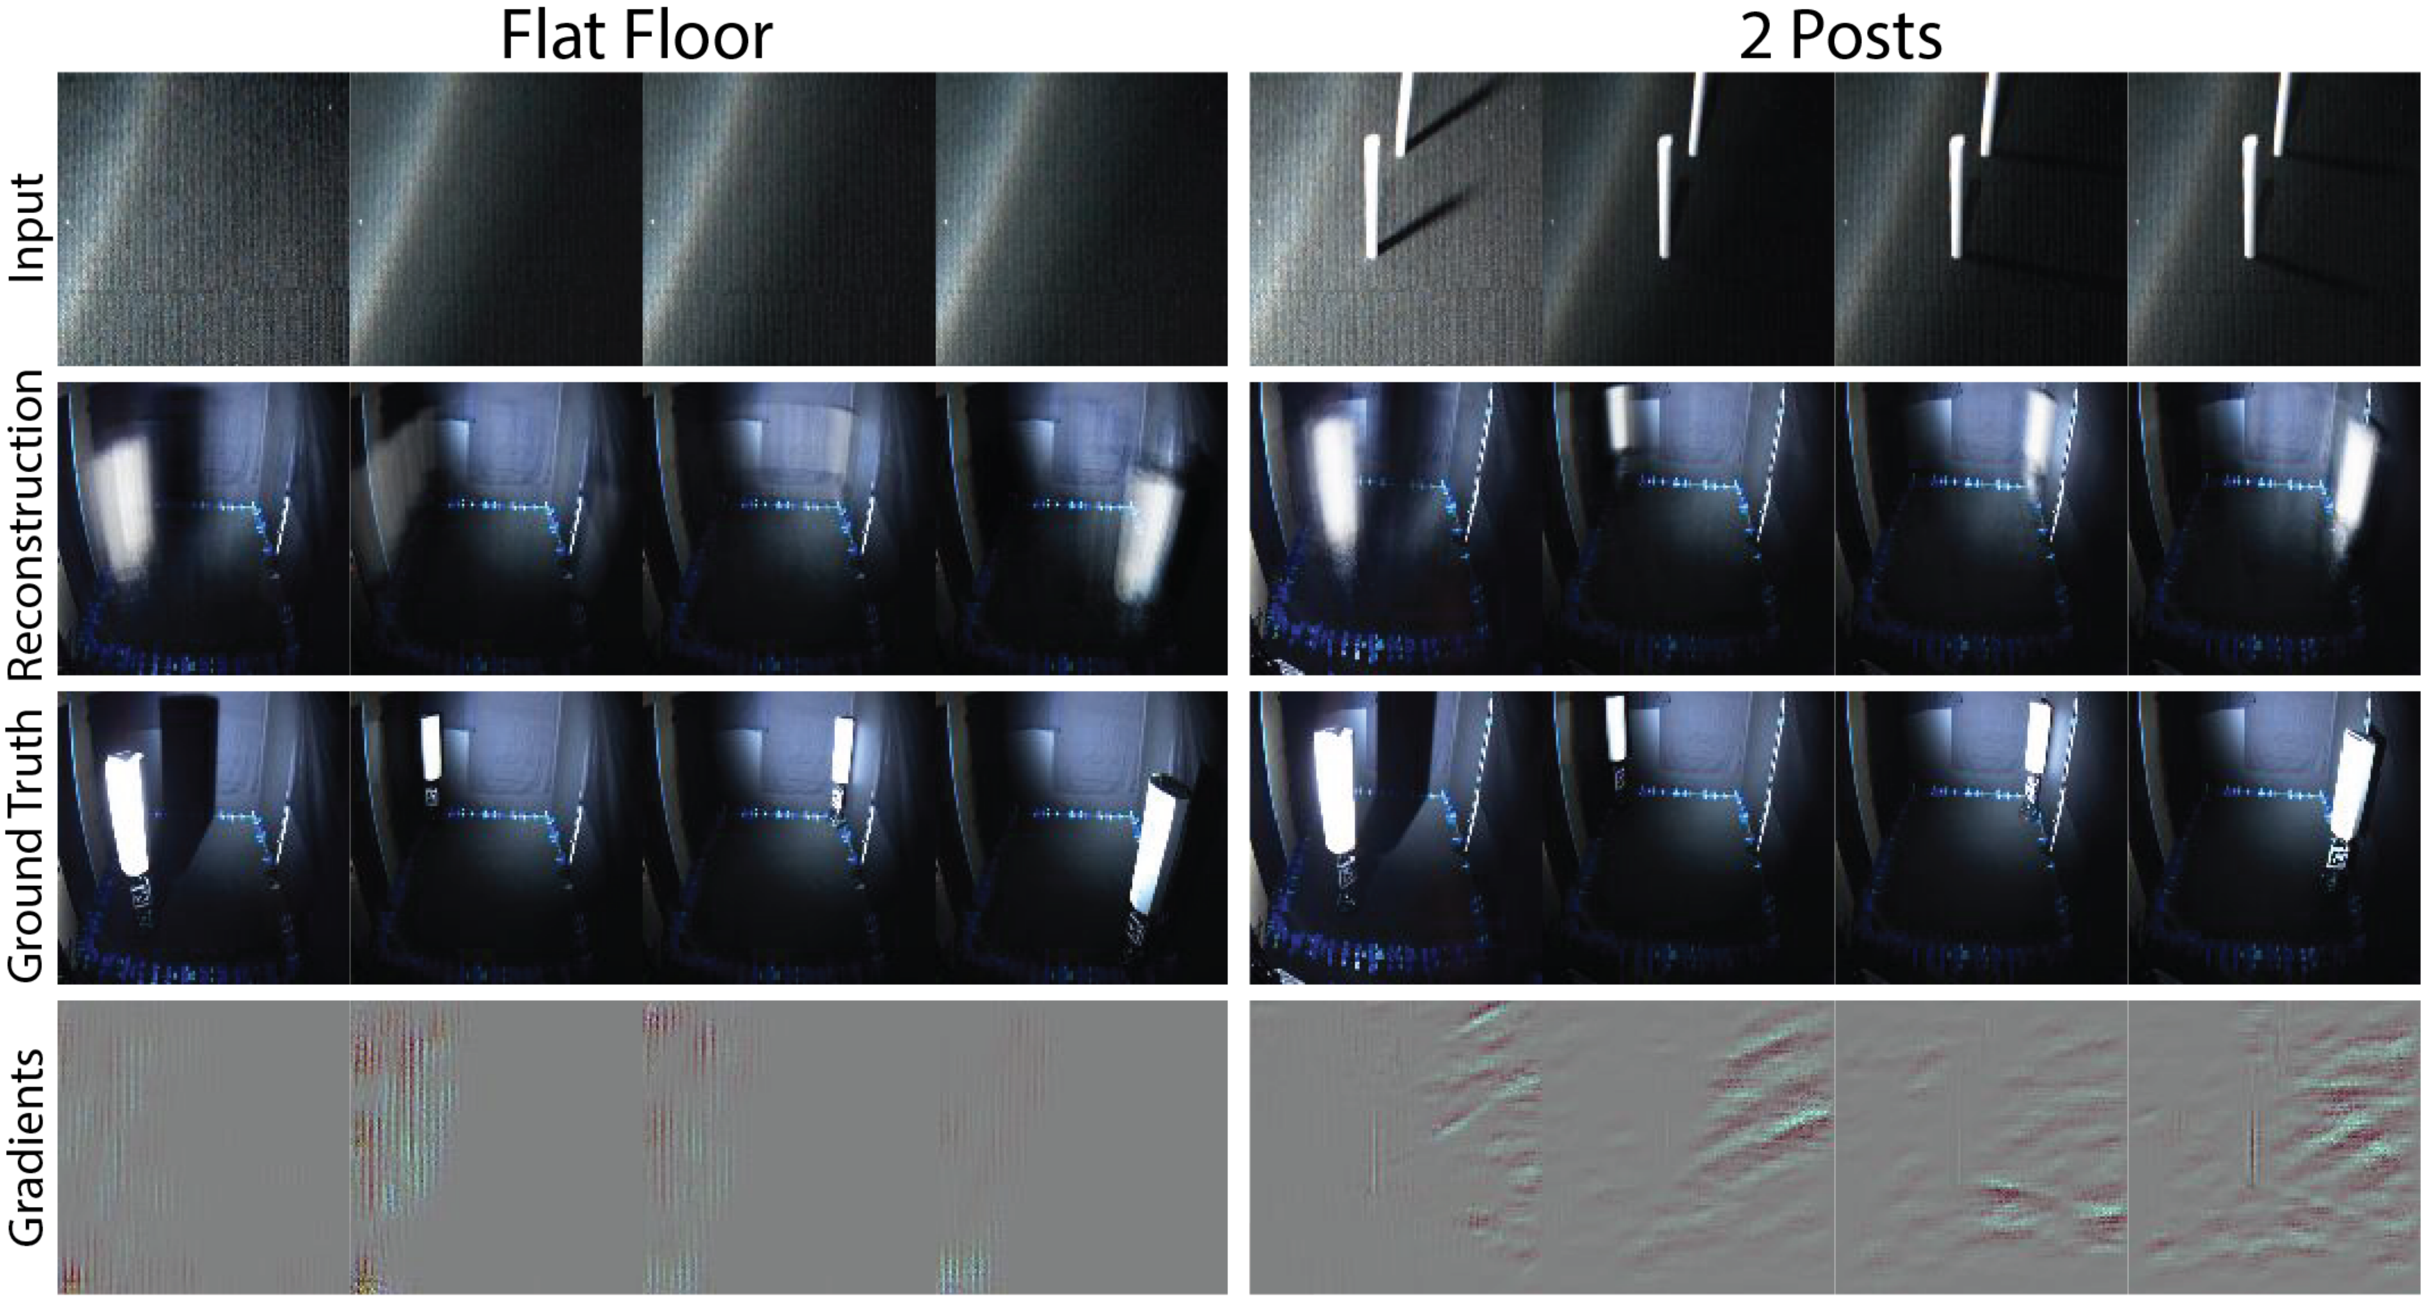

Flat Floor
2 Posts
Input
Reconstruction
Ground Truth
Gradients

Multiple Illumination Points
End-to-end Reconstruction

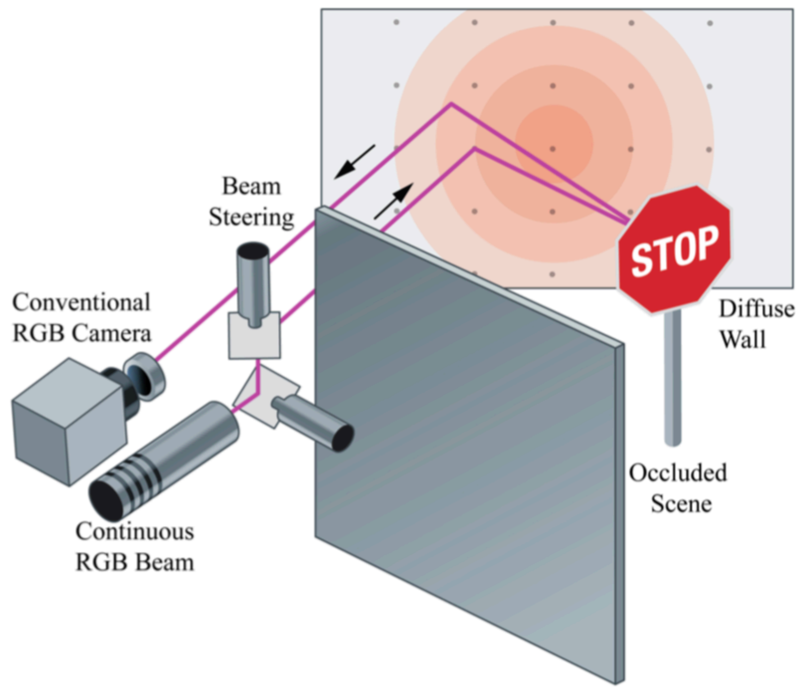

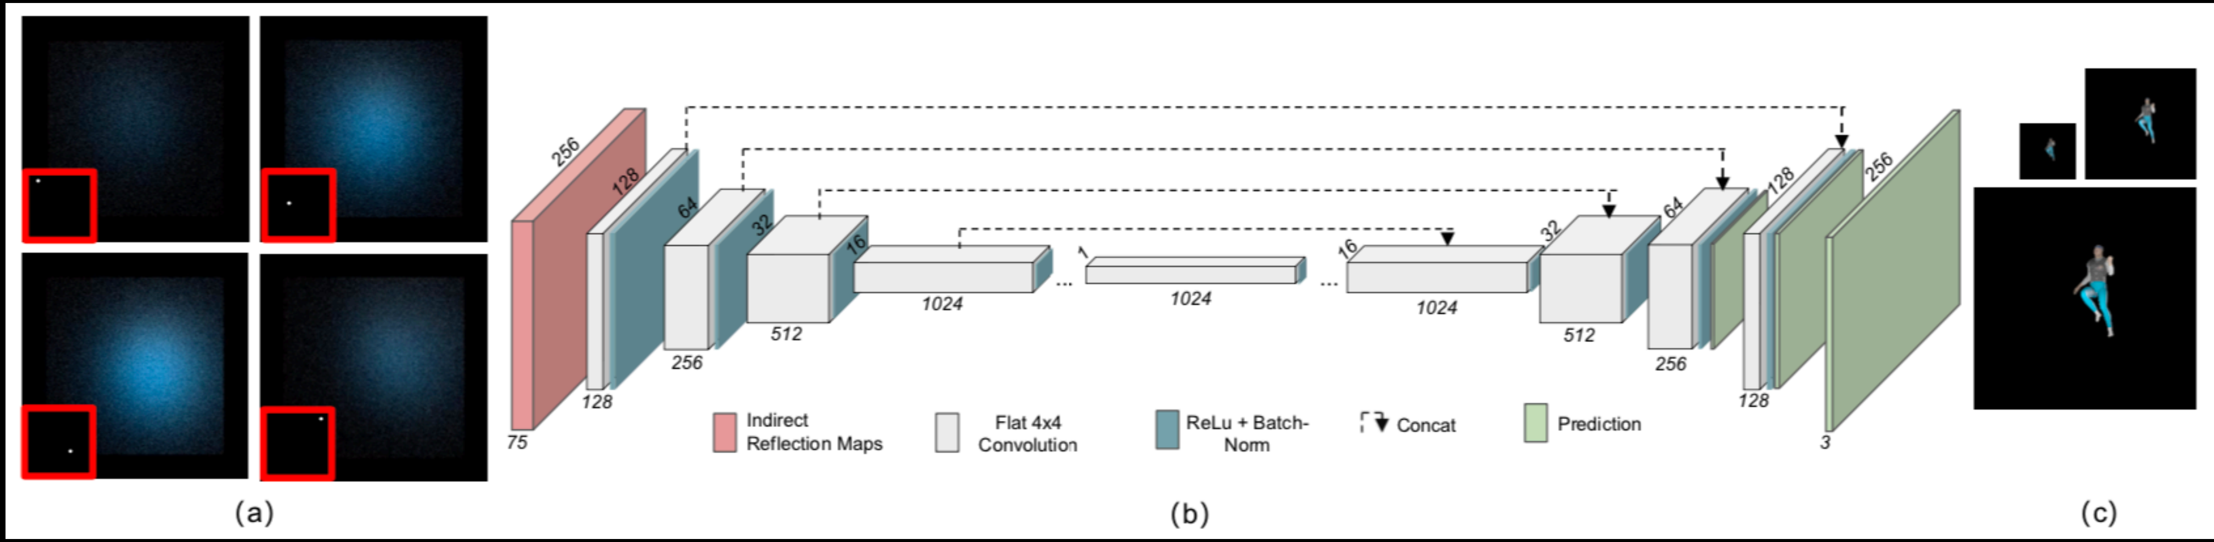

Chen, Daneau, Manna, Heide, “Steady-state Non-Line-of-Sight Imaging,” CVPR, 2019

Multiple Illumination Points
End-to-end Reconstruction

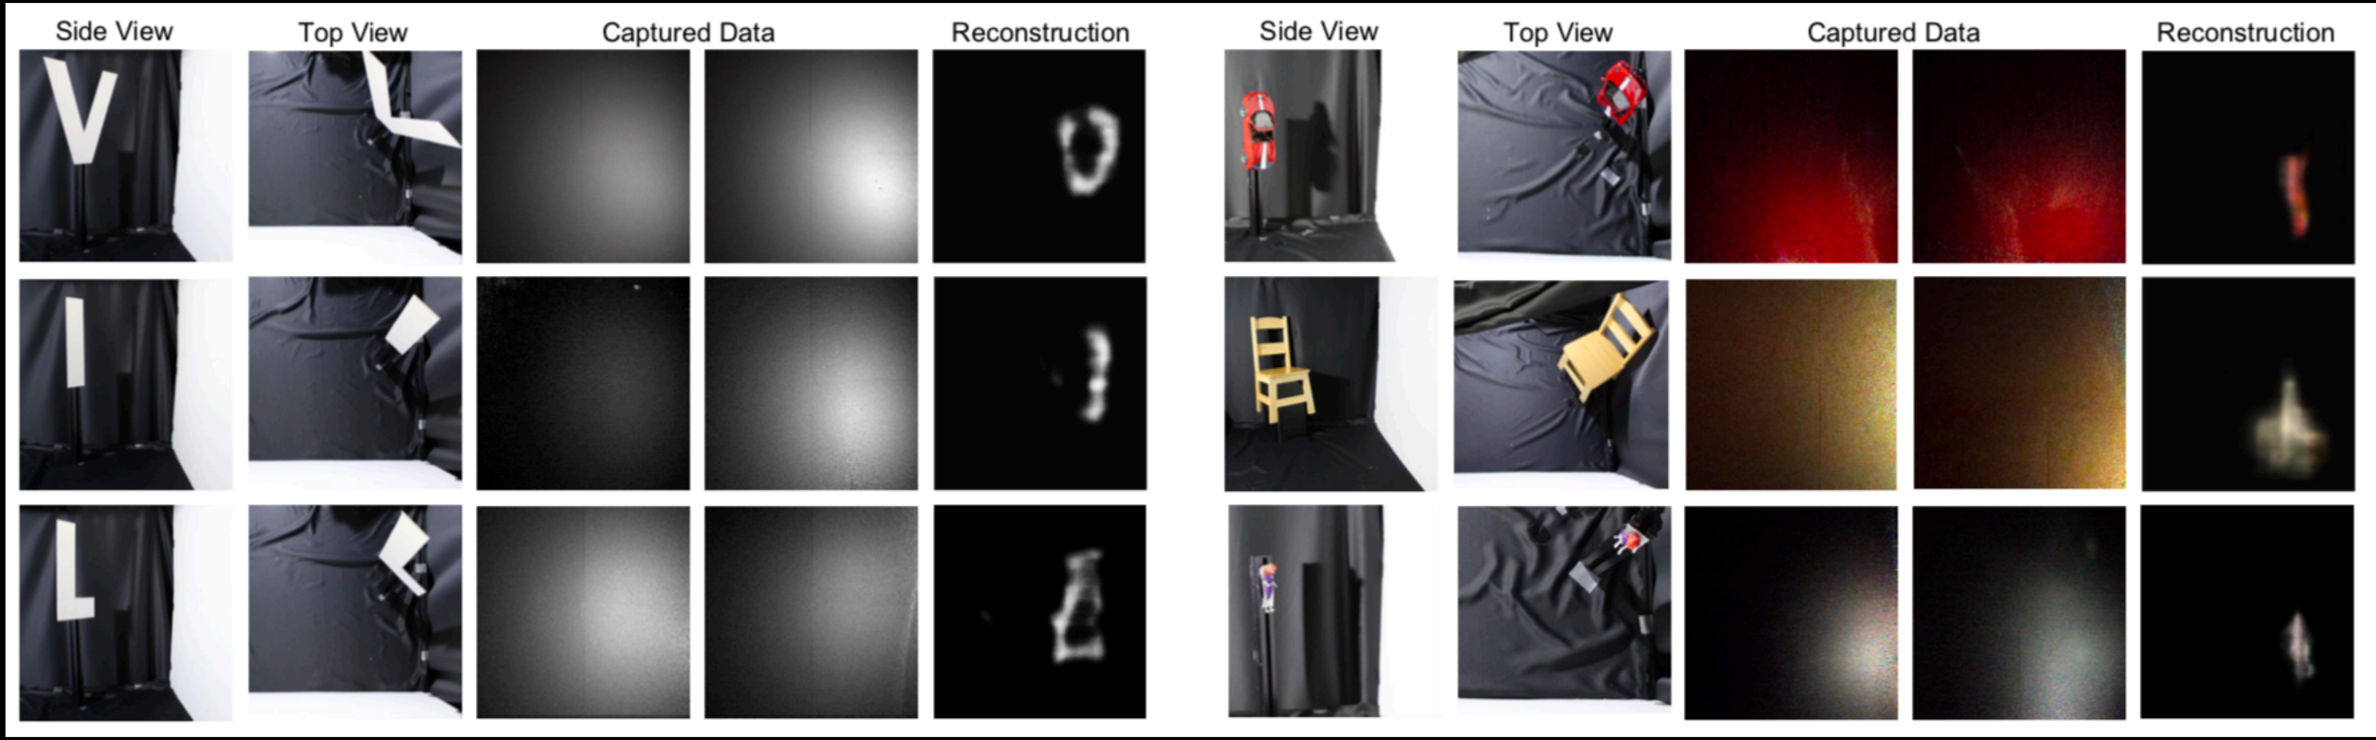

Chen, Daneau, Manna, Heide, “Steady-state Non-Line-of-Sight Imaging,” CVPR, 2019

Adaptive Lighting, 3D localization, Classification

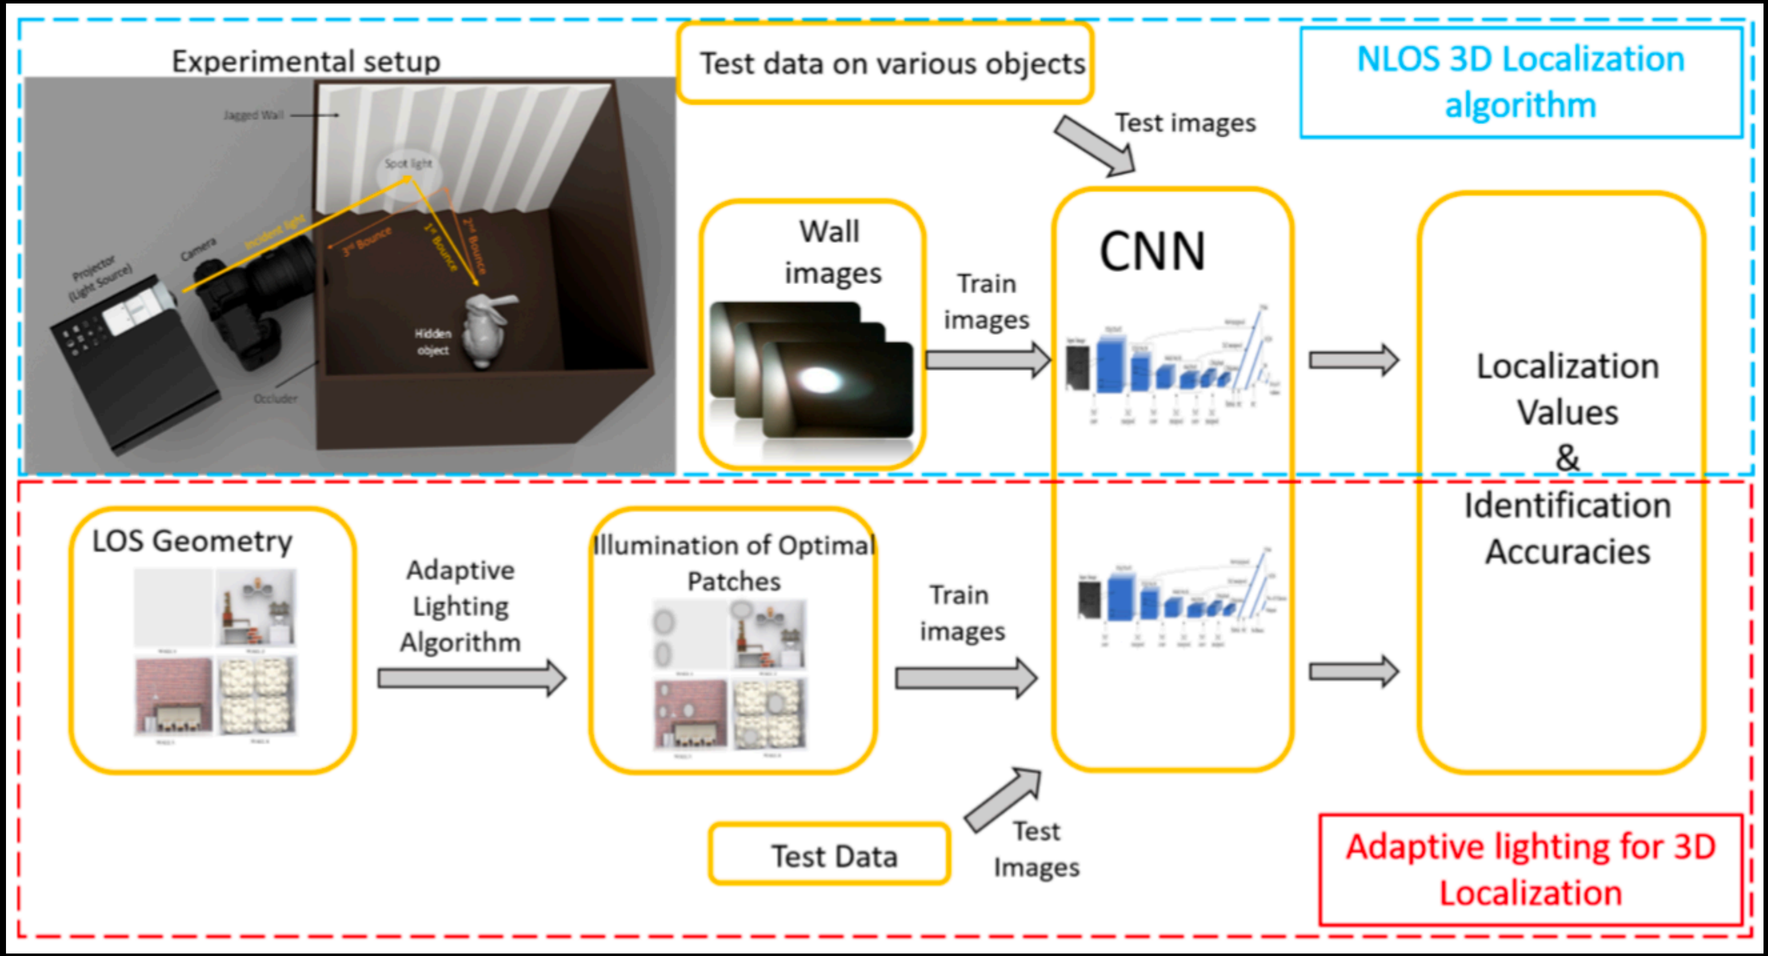

Chandran, Jayasuriya, “Adaptive Lighting for Data-Driven Non-Line-of-Sight 3D Localization and Object Identification,” 2019

Imaging Through Diffusers

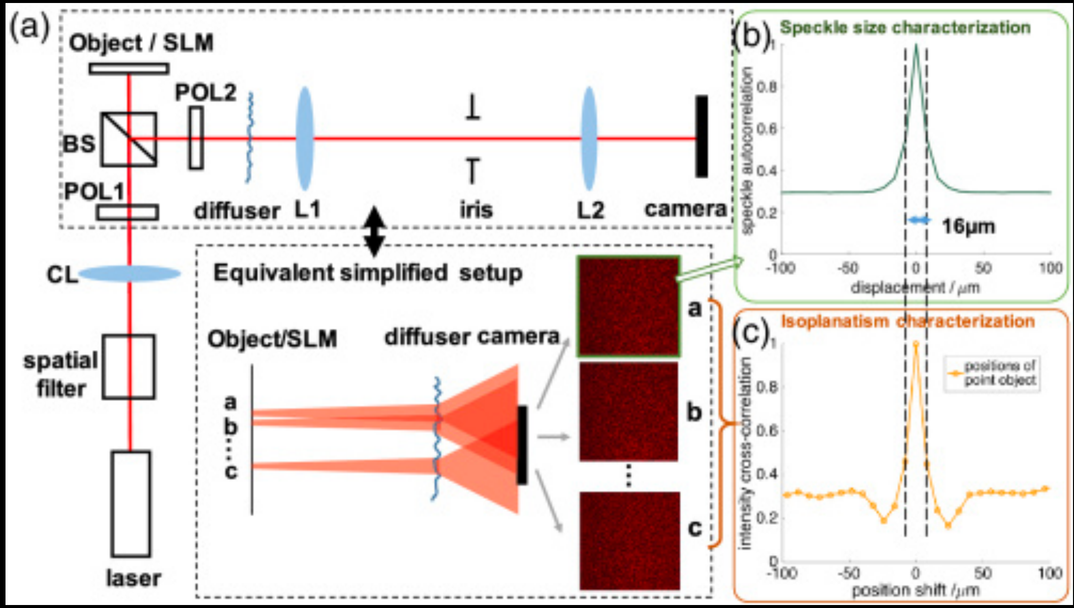

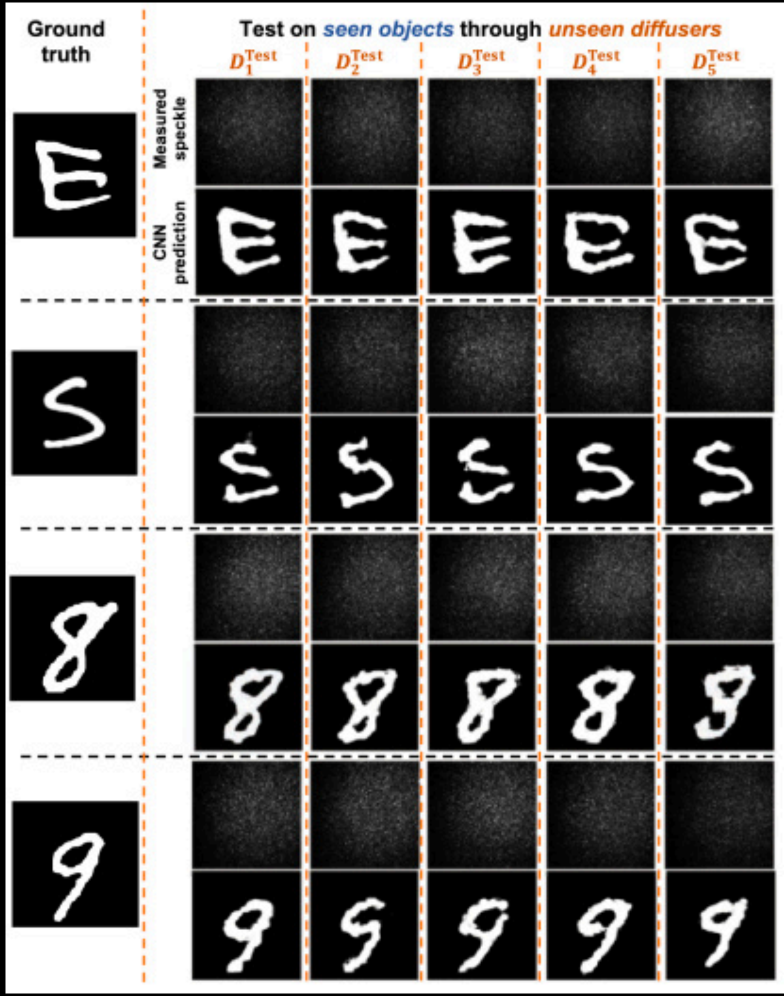

Li, Xue, Tian, “Deep speckle correlation: a deep learning approach toward scalable imaging through scattering media,” 2018

Imaging Through Diffusers

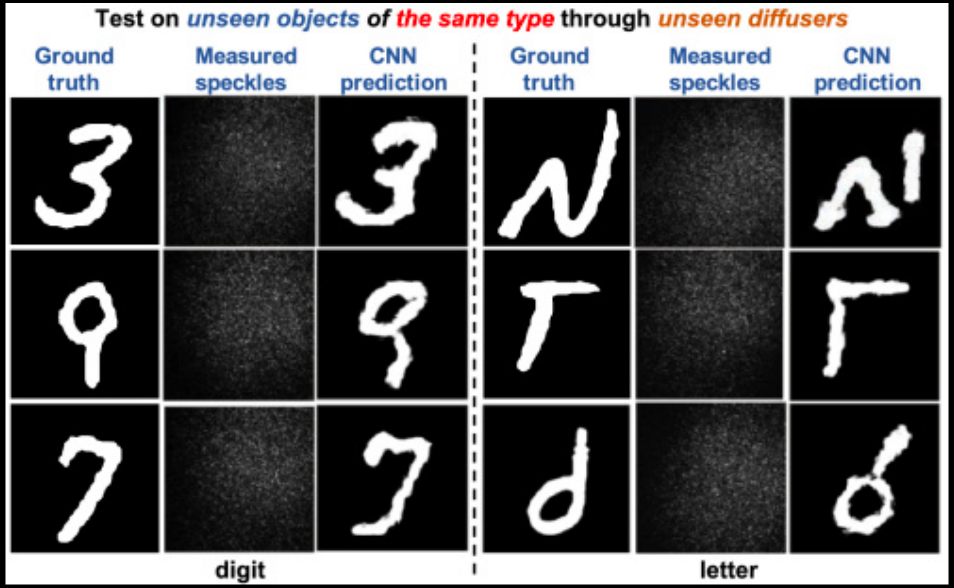

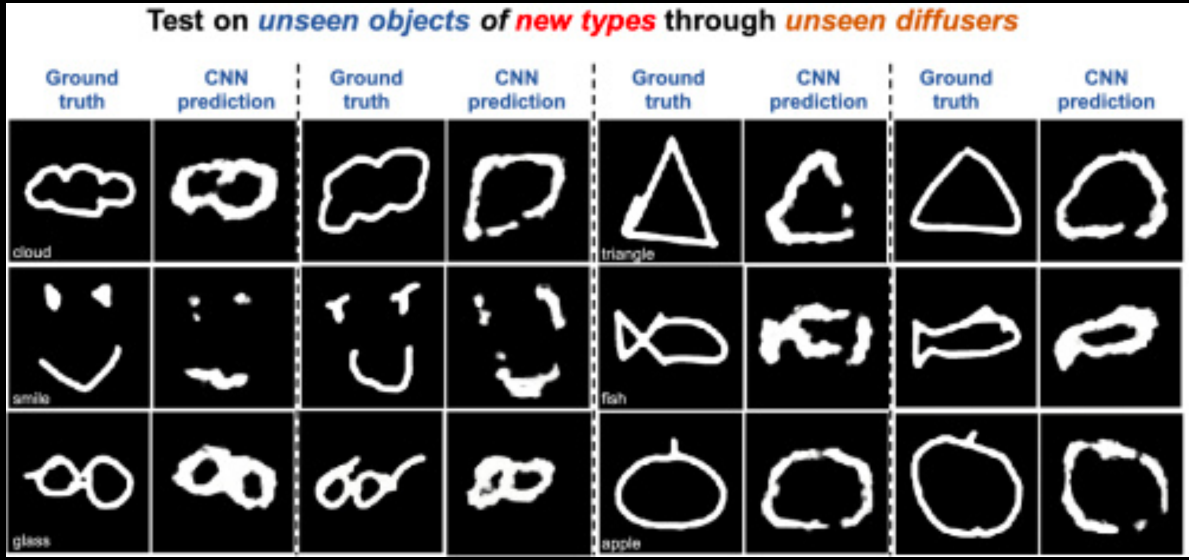

Li, Xue, Tian, "Deep speckle correlation: a deep learning approach toward scalable imaging through scattering media," 2018

Imaging Through Multimode Fibers

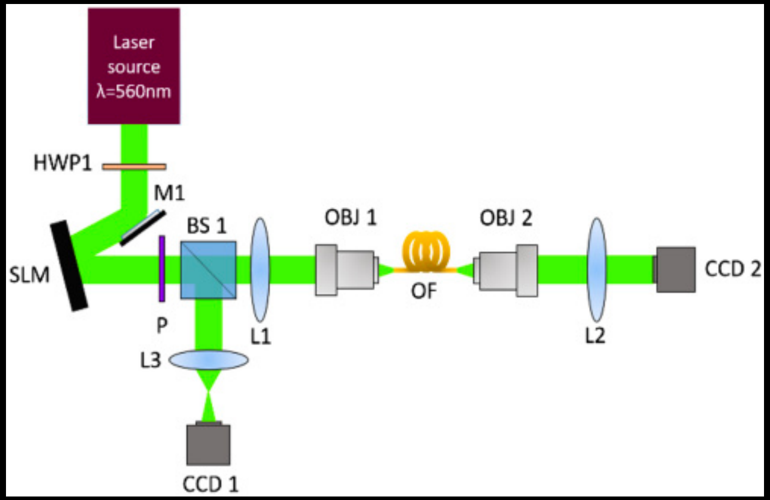

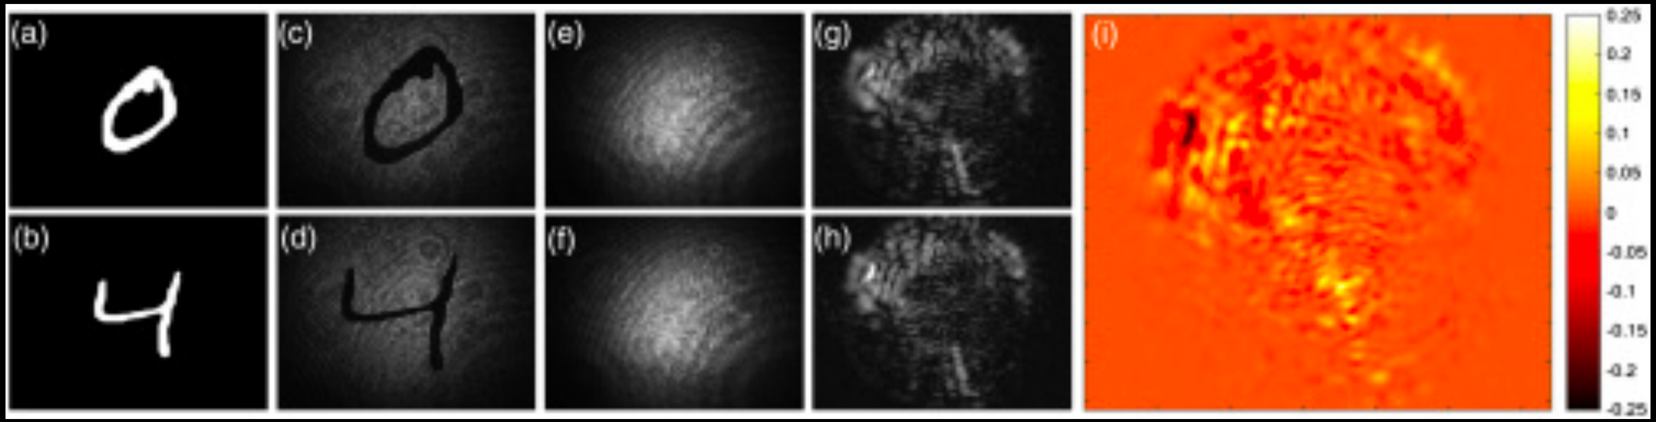

Borhani, Kakkava, Moser, Psaltis, “Learning to see through multimode fibers,” 2018

Limitations and Challenges

- Generalizability vs. Accuracy
- Classification vs. Regression / Reconstruction
- Data Generation Is The Bottleneck
- More complex physics (imaging through fog)

What’s Next?

Model: $$b = Ax$$

Measurement

Forward Model

Unknown Target

Solution: $$x = A^{-1}(b)$$

In Practice: $$x = argmin||Ax - b||^2 + \lambda \mathrm{R}(x)$$

Regularizer

Can we leverage data and our physical knowledge?

Physics + Data

$$x = argmin\|Ax - b\|^2 + \lambda \mathrm{R}(x)$$

1. Brute force: $b \rightarrow x$
2. Learn the regularizer $R(x)$:
    - A is defined by physics.
    - Replace projection step by denoising.
3. Learn A (or its properties)
    - Split the problem to several sub problems.
    - Each sub-problem is learned separately.
    - Physics define the connections.

ADMM-Net

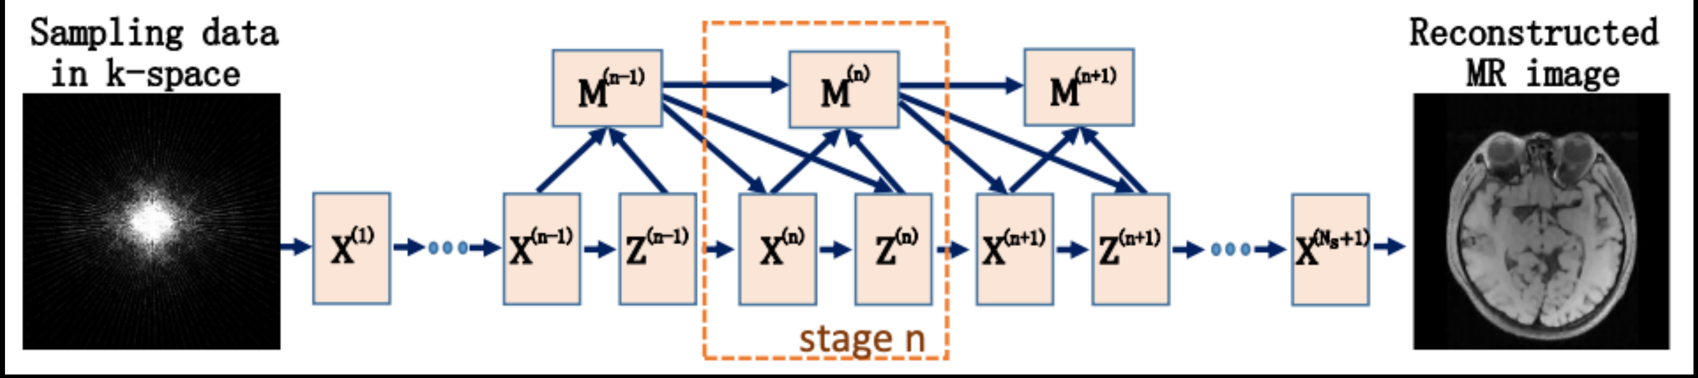

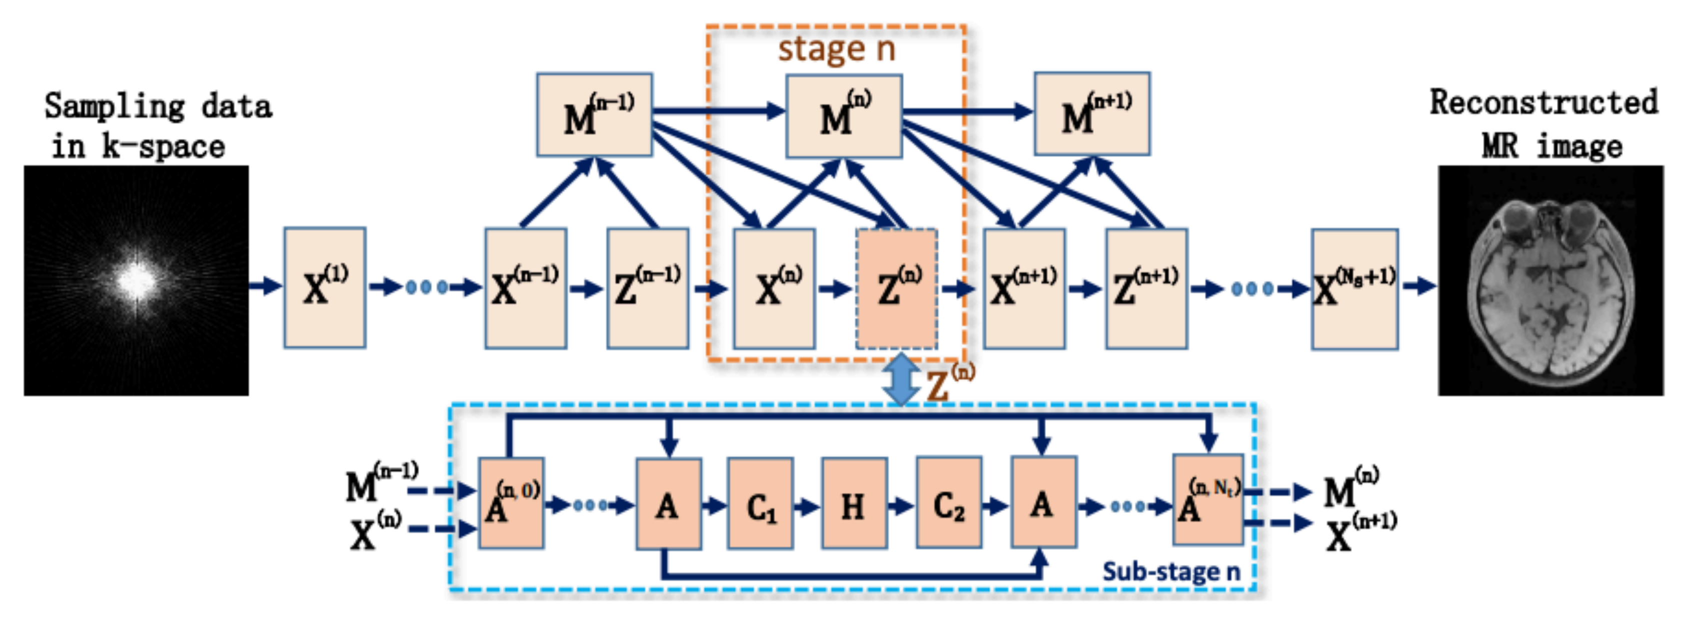

Yang, Sun, Li, Xu, “Deep ADMM-Net for Compressive Sensing MRI,” NIPS, 2016

Unrolled Optimization With Deep Priors

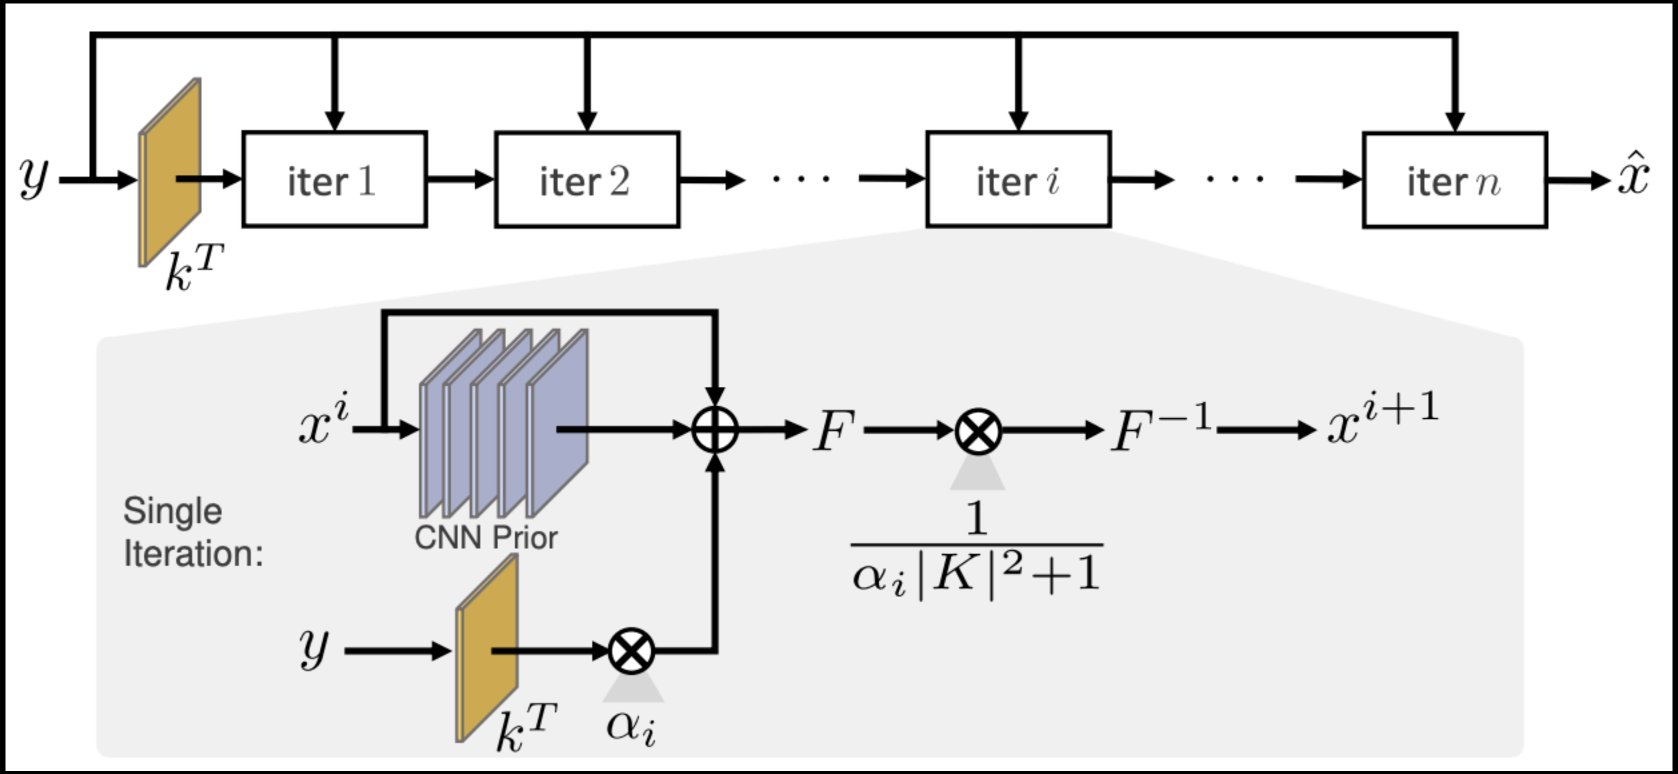

Diamond, Sitzmann, Heide, Wetzstein, “Unrolled Optimization with Deep Priors,” 2018

Summary

media.mit.edu/~guysatat
guysatat@mit.edu

- NLOS Computational Imaging
- Data Driven NLOS Computational Imaging
  - Through scattering
  - Around corners
  - SPAD sensor / traditional camera
- Incorporating Physics