deep imaging

Lecture 19: Physics-based CNN examples

Machine Learning and Imaging

BME 548L

Roarke Horstmeyer

What physical parameters effect image formation?

- Illumination
  - Spatial pattern
  - Angle of incidence
  - Color, polarization
- Lens and optics
  - Position/orientation
  - Shape
  - Focus
  - Transparency
- Detector
  - Pixel size
  - Pixel shape & fill factor
  - Color filters
  - Other filters
- Digitization
  - E to P curves
  - Digitization schemes/thresholds
  - Data transmission, multiplexing
- Physical object

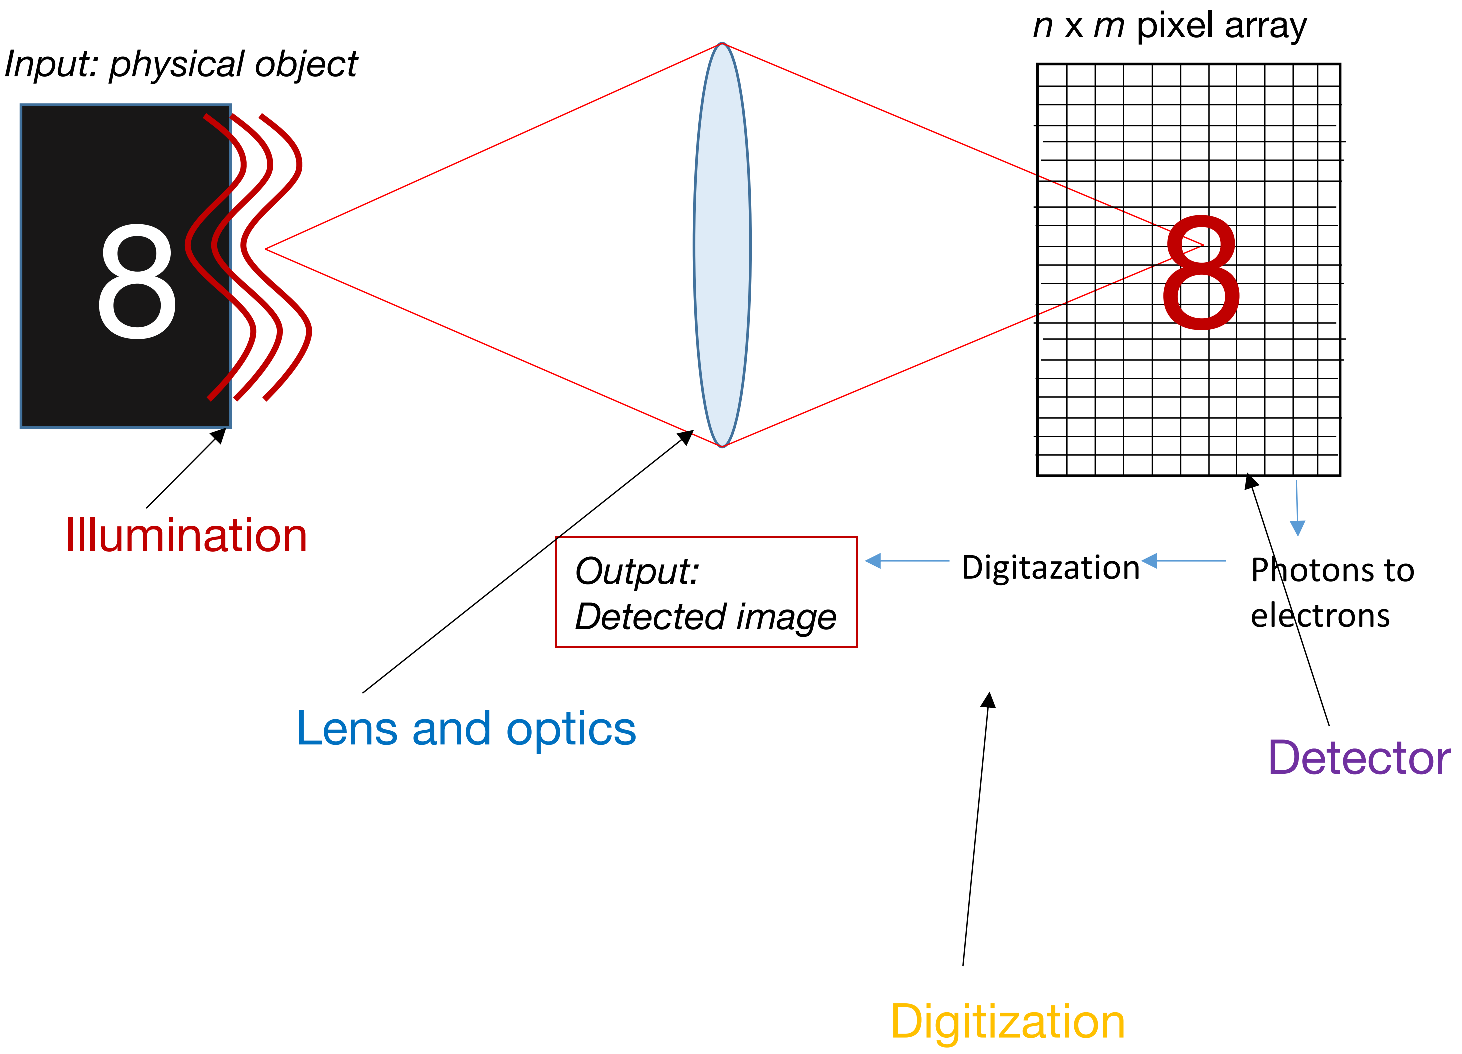

What physical parameters effect image formation?

deep imaging

- Illumination
  - Spatial pattern
  - Angle of incidence
  - Color, polarization
- Lens and optics
  - Position/orientation
  - Shape
  - Focus
  - Transparency
- Detector
  - Pixel size
  - Pixel shape & fill factor
  - Color filters
  - Other filters
- Digitization
  - E to P curves
  - Digitization schemes/thresholds
  - Data transmission, multiplexing
- Physical object

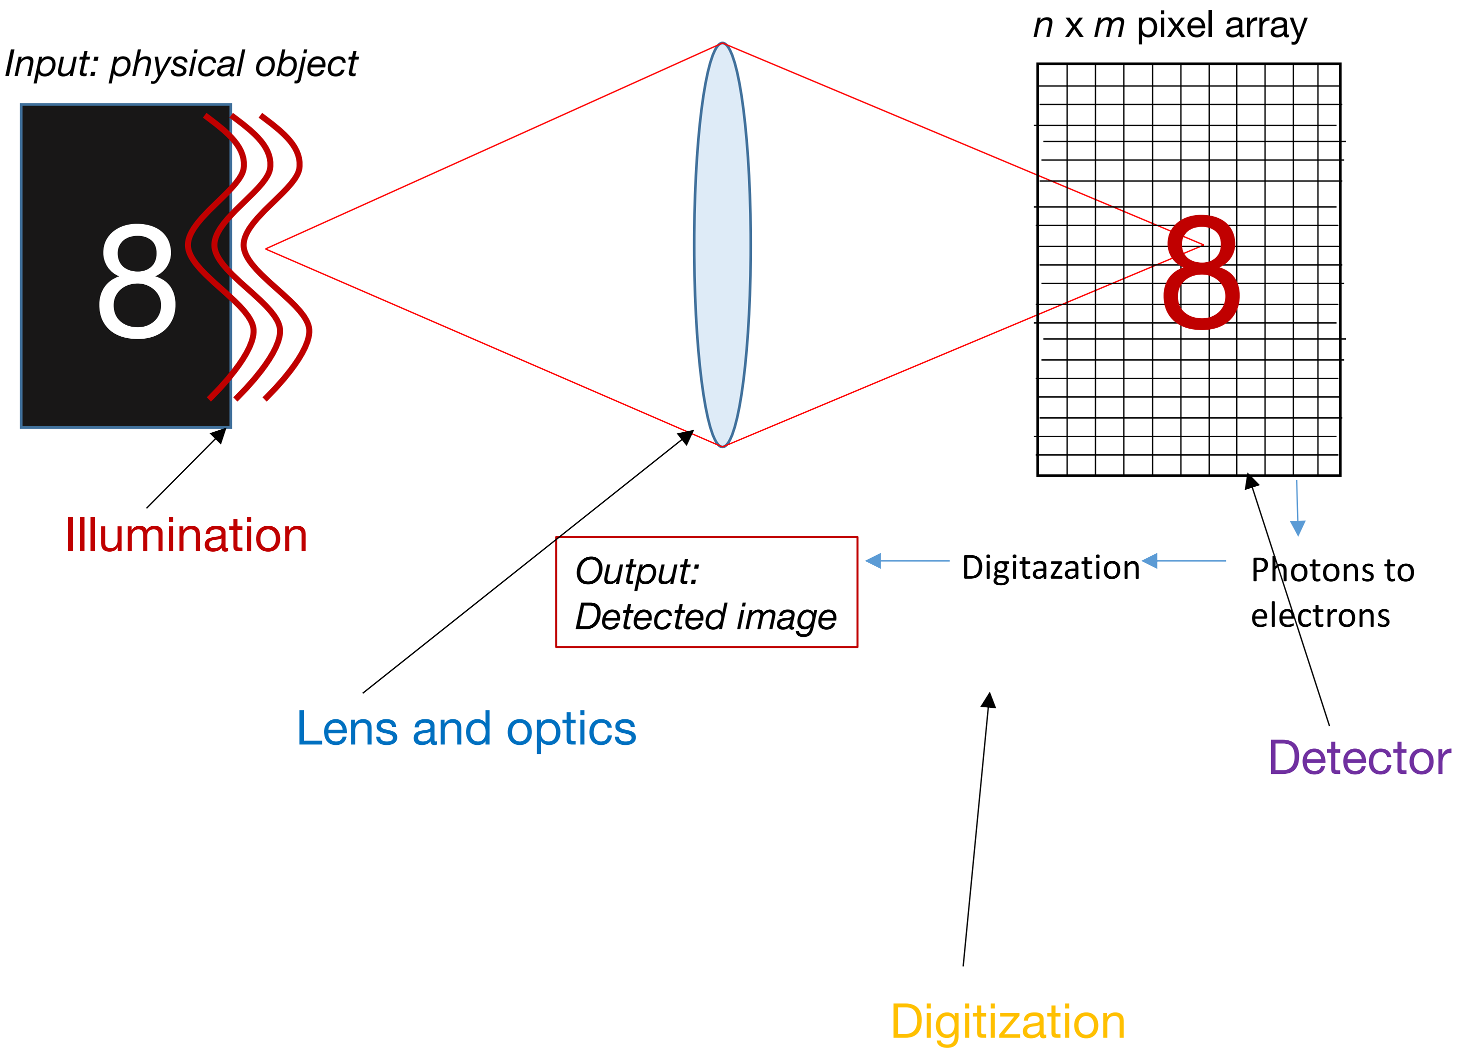

deep imaging

Examples: Lenses and optics

Hybrid optical-electronic convolutional neural networks with optimized diffractive optics for image classification

Julie Chang[1], Vincent Sitzmann[2], Xiong Dun[3], Wolfgang Heidrich[3] & Gordon Wetzstein[2]

deep imaging

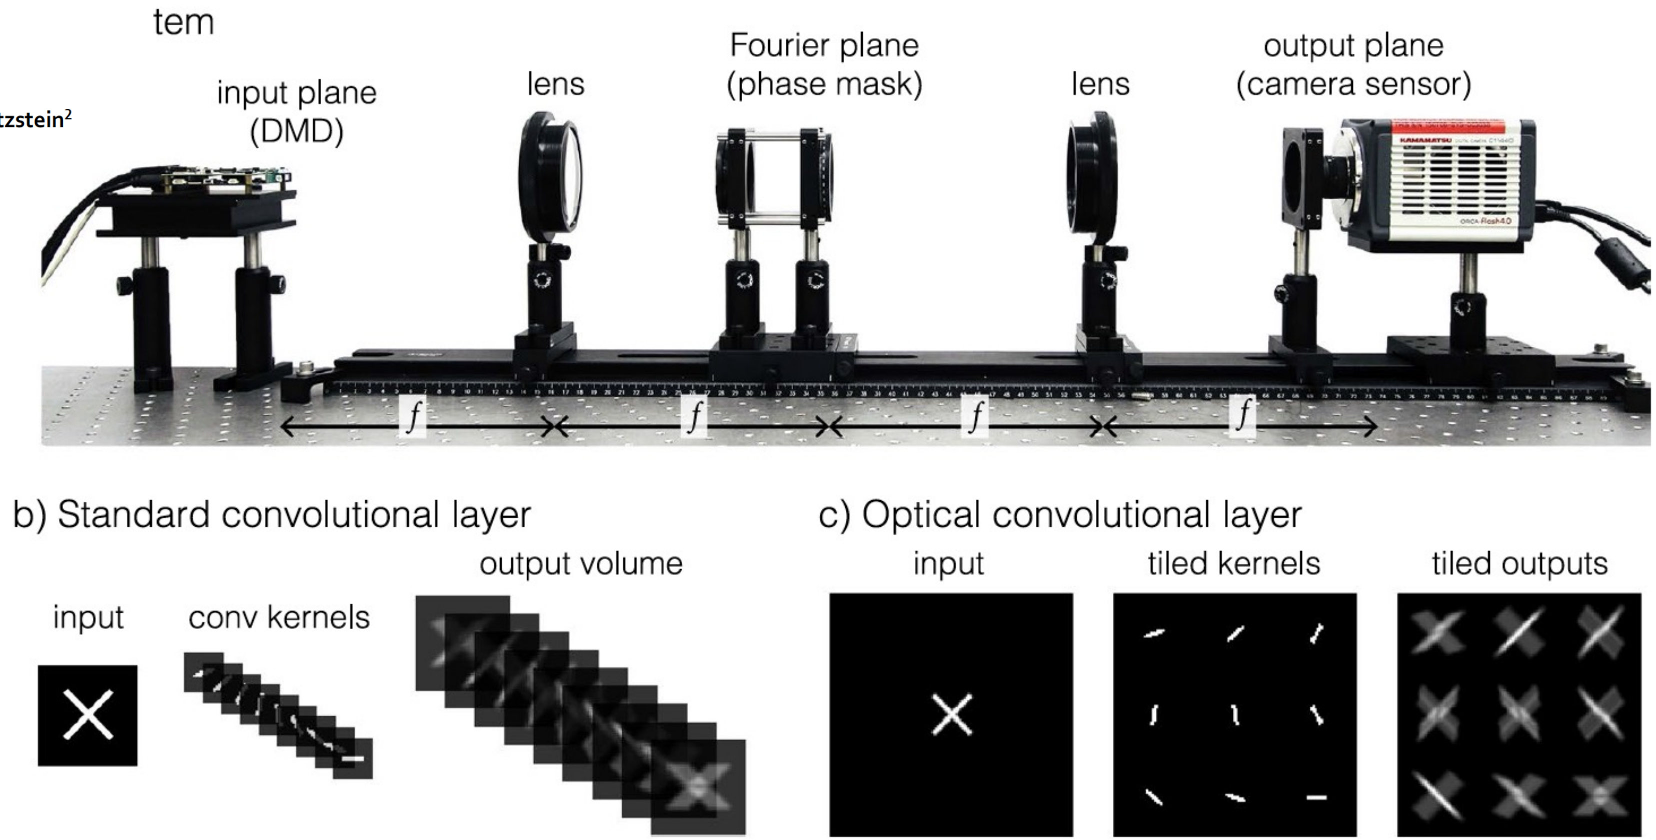

Figure 1. Optical convolutional layer design. (**a**) Diagram of a 4*f* system that could be adapted to implement optical convolutional (opt-conv) layers by placing a phase mask in the Fourier plane. (**b**) The standard components of a digital convolutional layer, including an input image, a stack of convolutional kernels, and a corresponding output volume. (**c**) The equivalent components in an opt-conv layer, where the kernels and outputs are tiled in a 2D array instead of stacked in the depth dimension.

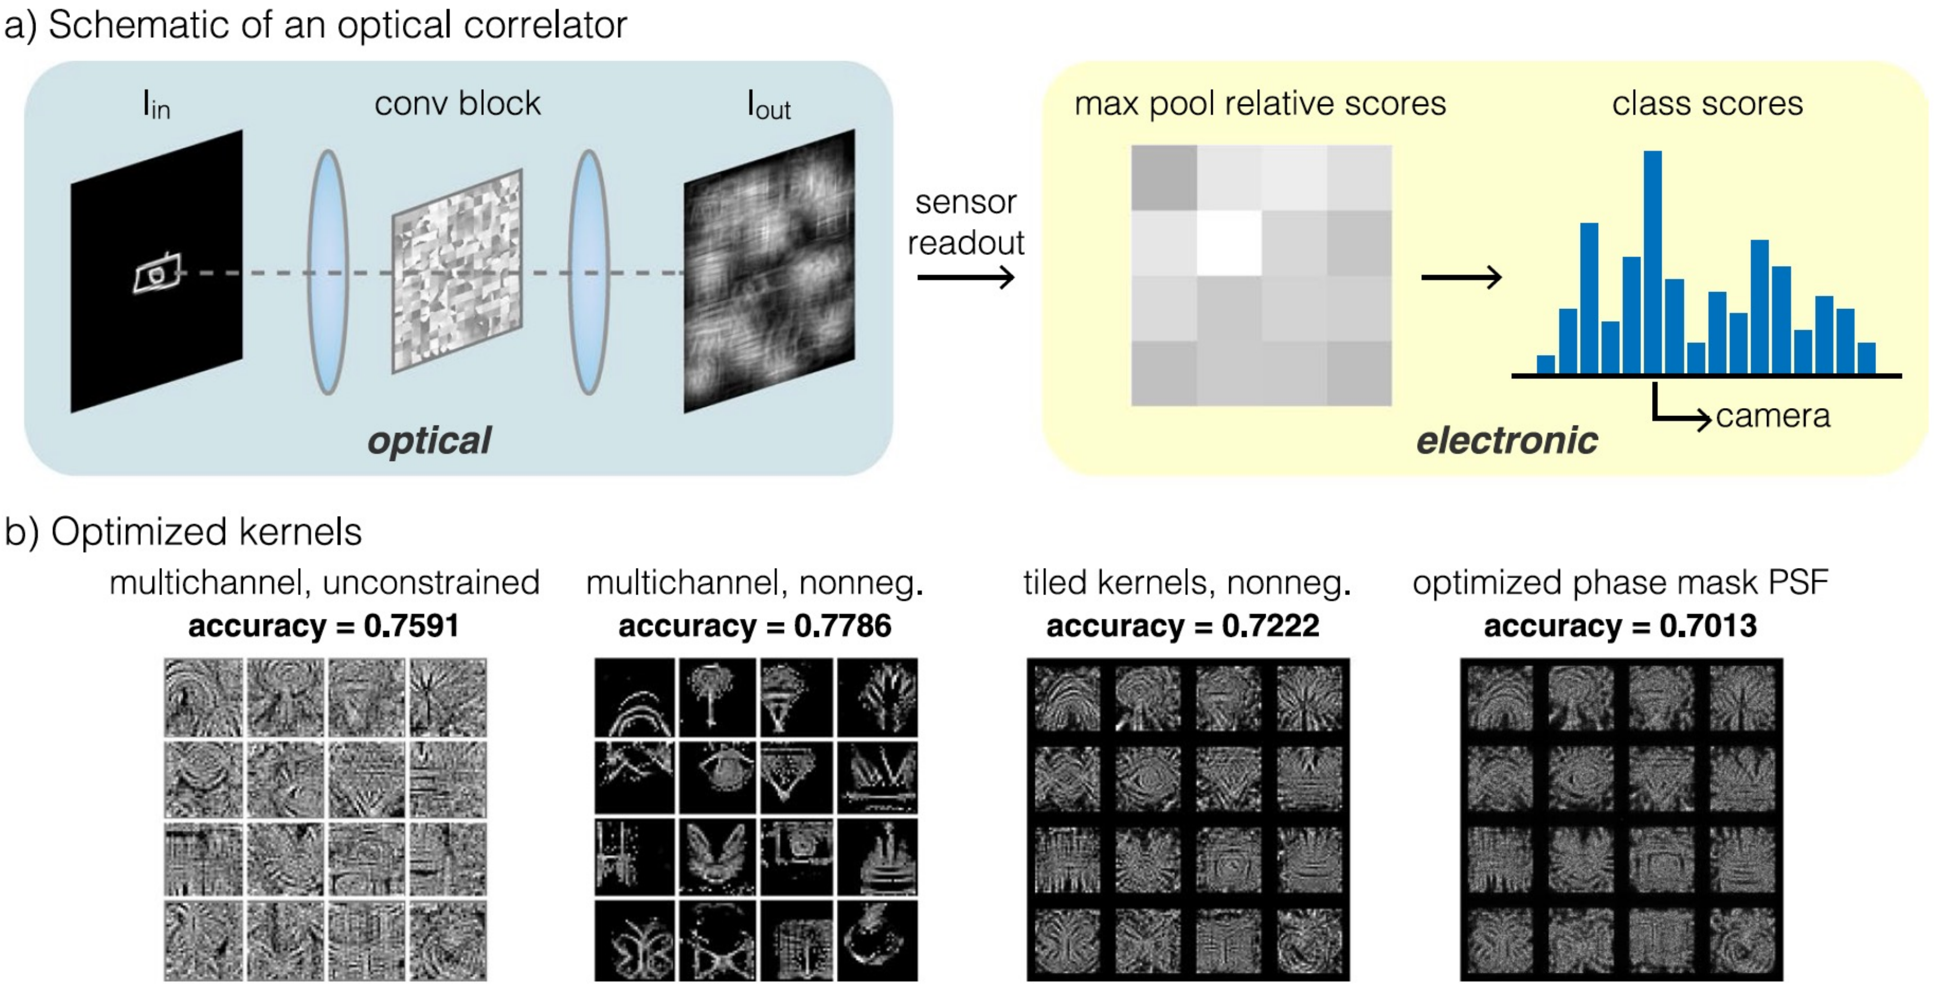

Figure 2. Learned optical correlator. (**a**) Schematic of an optical correlator, where the conv block consists of the 4f system shown in Fig. 1. (**b**) Characteristic optimized kernels of a multichannel unconstrained digital convolutional layer, a multichannel nonnegative digital convolutional layer, a single channel opt-conv layer with tiled kernels, and the PSF produced by phase mask optimization with the previous optimized tiled kernels as the target.

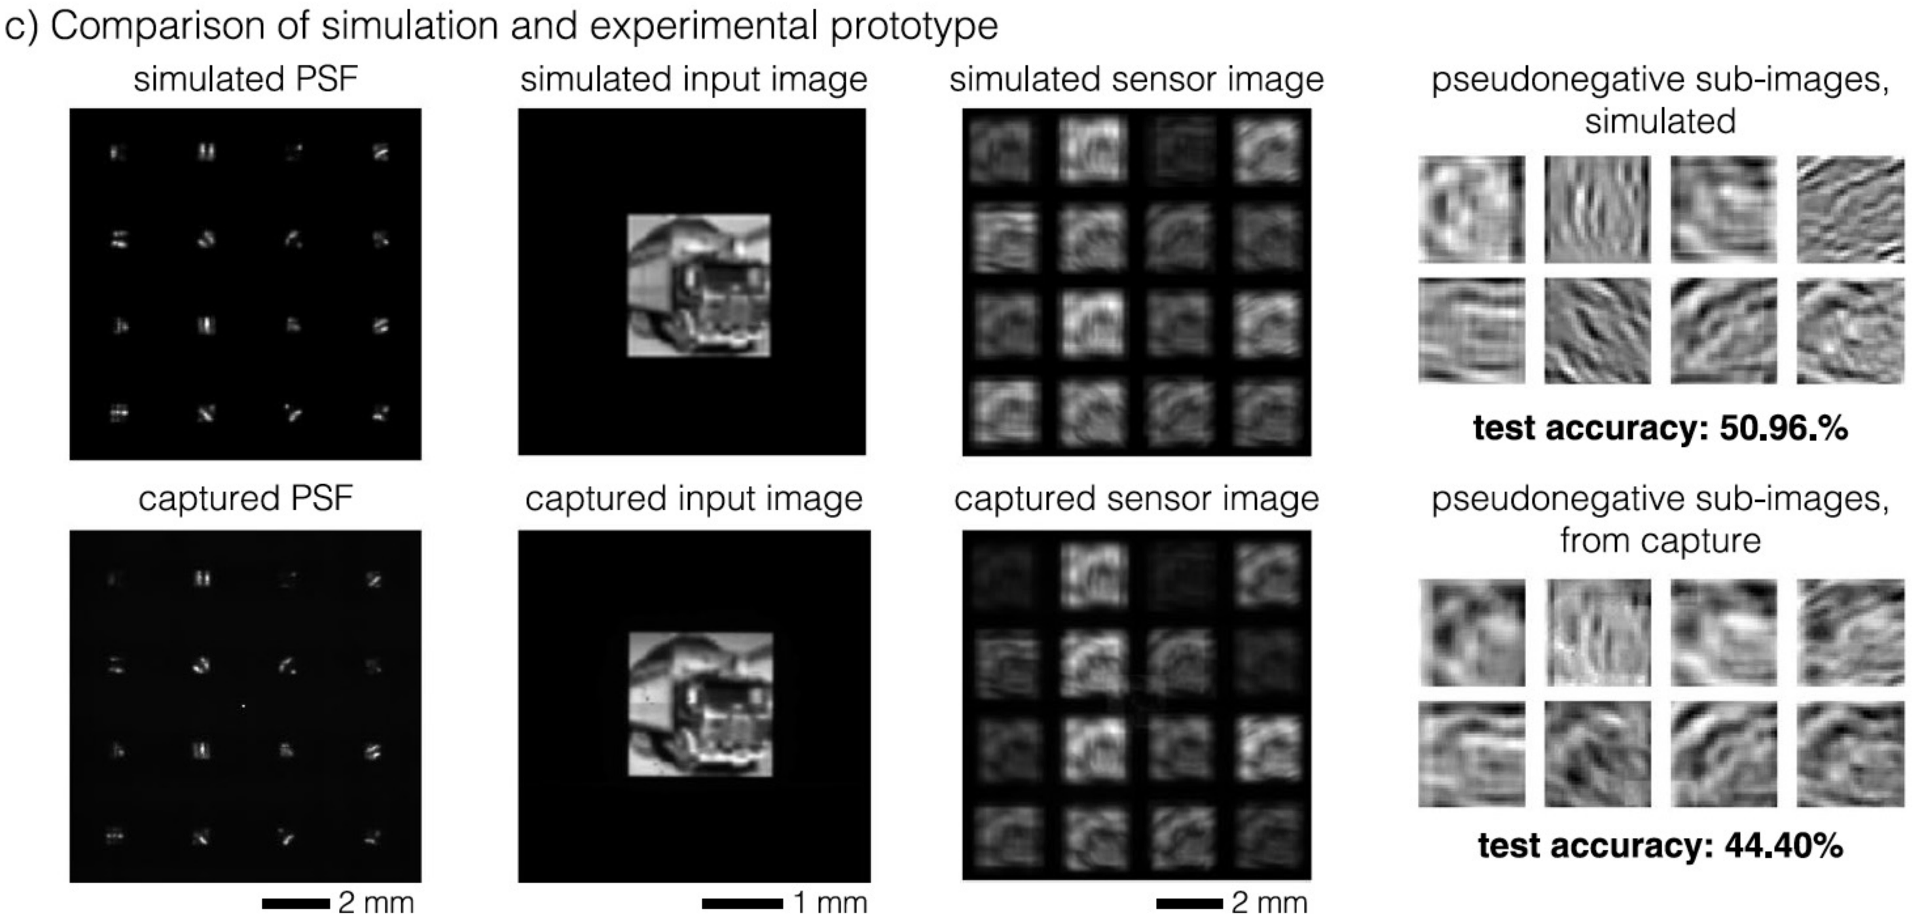

Figure 3. Hybrid optoelectronic CNN. (**a**) Schematic of a model with a single opt-conv layer, after which the sensor image is processed and fed into subsequent digital CNN layers. (**b**) The optimized phase mask template and microscope images of the fabricated phase mask, at different zoom levels. (**c**) Comparison of simulated and captured versions of the PSF produced by the phase mask, a sample input image, the respective sensor image, and pseudonegative sub-images after subtraction of corresponding positive (top two rows) and negative (bottom two rows) sub-images.

deep imaging

End-to-end Optimization of Optics and Image Processing for Achromatic Extended Depth of Field and Super-resolution Imaging

deep imaging

VINCENT SITZMANN*, Stanford University, USA
STEVEN DIAMOND*, Stanford University, USA
YIFAN PENG*, The University of British Columbia, Canada and Stanford University, USA
XIONG DUN, King Abdullah University of Science and Technology, Saudi Arabia
STEPHEN BOYD, Stanford University, USA
WOLFGANG HEIDRICH, King Abdullah University of Science and Technology, Saudi Arabia
FELIX HEIDE, Stanford University, USA
GORDON WETZSTEIN, Stanford University, USA

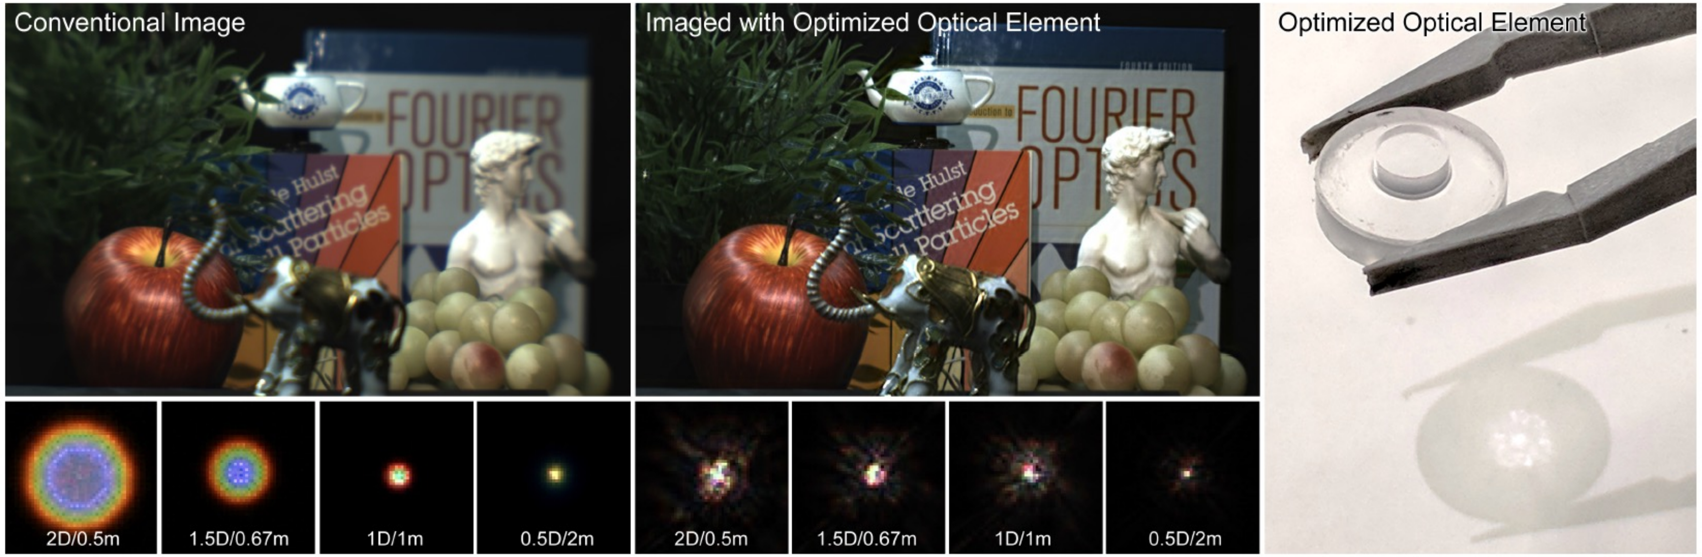

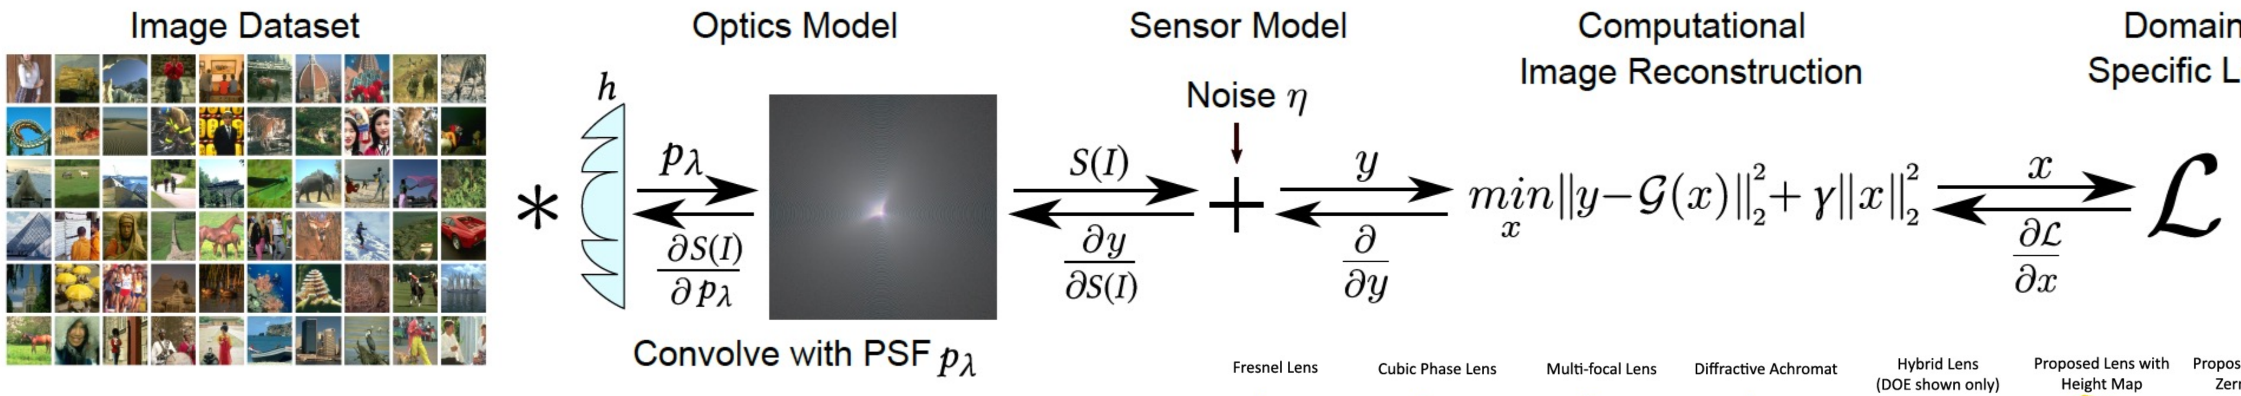
Image Dataset
Optics Model
Sensor Model
Computational Image Reconstruction
Domain-Specific Loss
h
p_λ
$\frac{\partial S(I)}{\partial p_\lambda}$
Convolve with PSF p_λ
Noise η
$S(I)$
$\frac{\partial y}{\partial S(I)}$
y
$\frac{\partial}{\partial y}$
$\min_x \|y-\mathcal{G}(x)\|_2^2 + \gamma\|x\|_2^2$
x
$\frac{\partial \mathcal{L}}{\partial x}$
$\mathcal{L}$

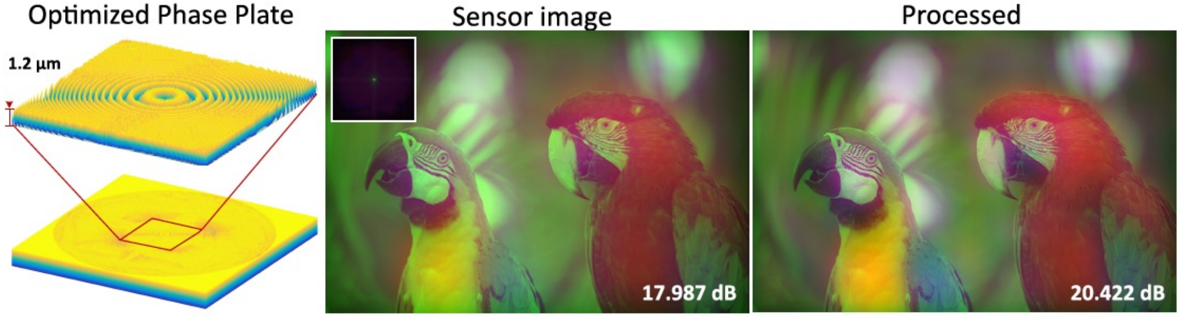
Optimized Phase Plate
1.2 µm
Sensor image
Processed
17.987 dB
20.422 dB

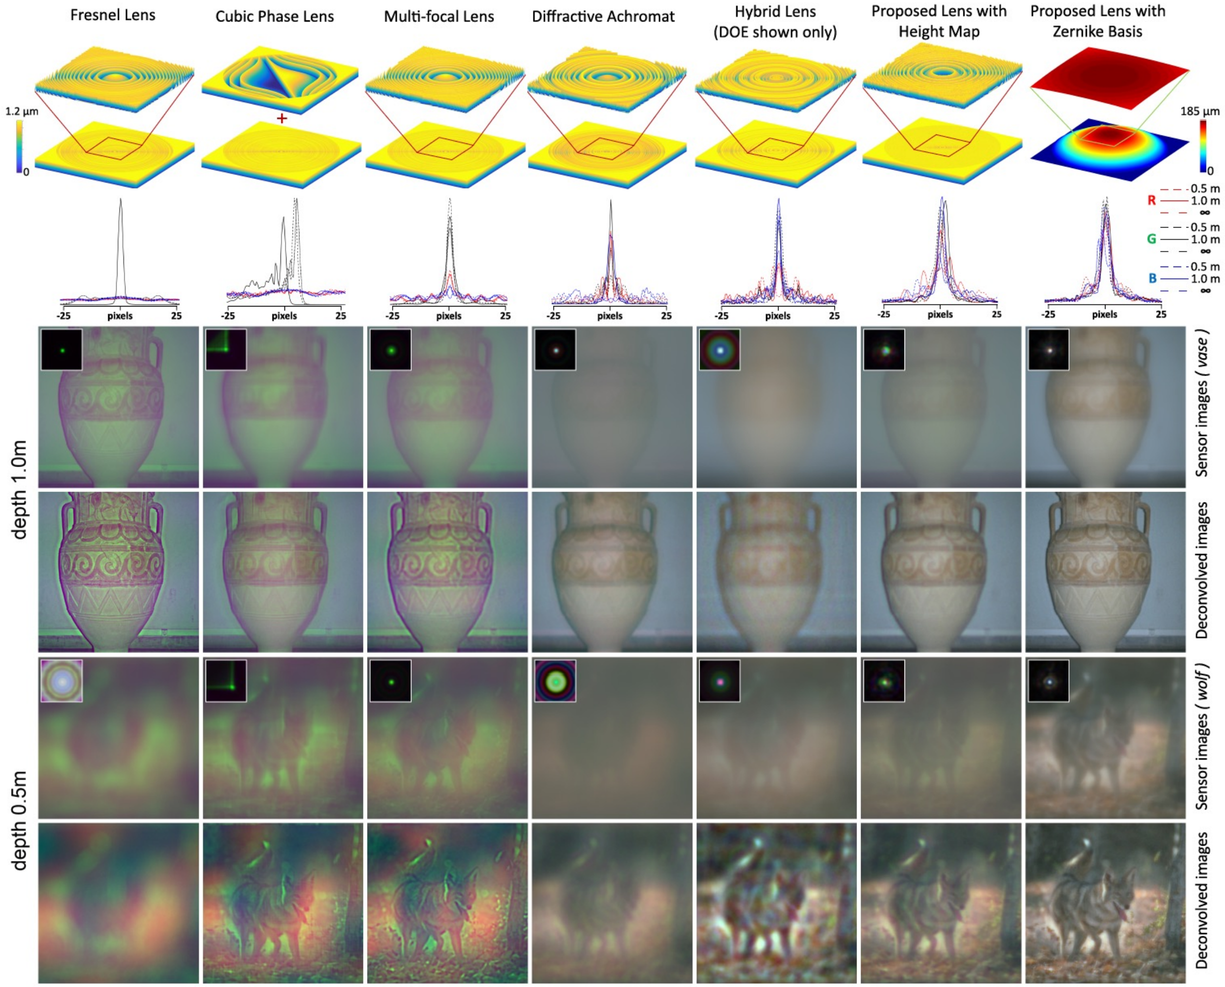
Fresnel Lens
Cubic Phase Lens
Multi-focal Lens
Diffractive Achromat
Hybrid Lens (DOE shown only)
Proposed Lens with Height Map
Proposed Lens with Zernike Basis
1.2 µm
185 µm
depth 1.0m
depth 0.5m
Sensor images (vase)
Deconvolved images
Sensor images (wolf)
Deconvolved images

Learned phase coded aperture for the benefit of depth of field extension

SHAY ELMALEM,[*] RAJA GIRYES, AND EMANUEL MAROM

School of Electrical Engineering, The Iby and Aladar Fleischman Faculty of Engineering, Tel Aviv University, Tel Aviv, Israel
[*]*shay.elmalem@gmail.com*

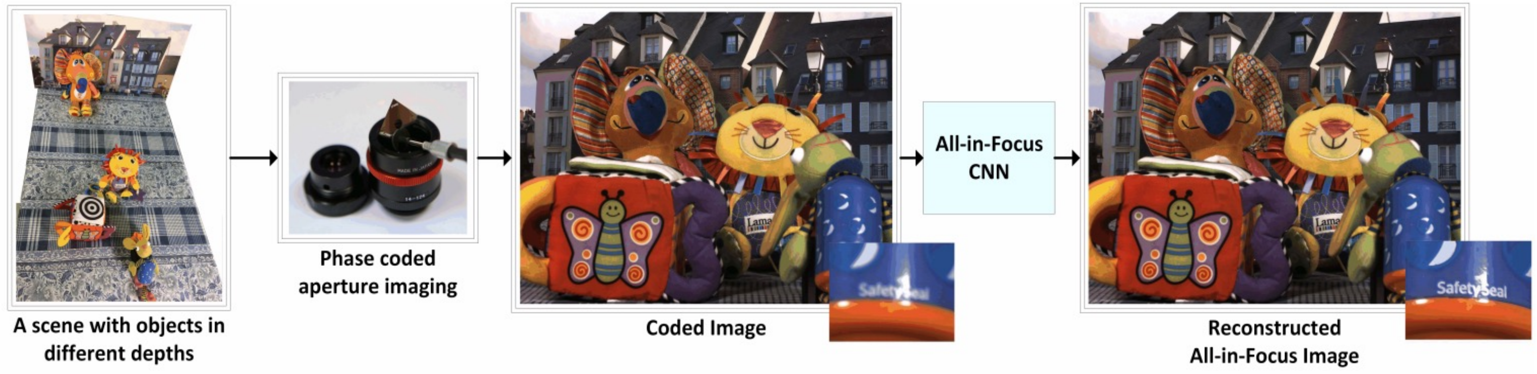

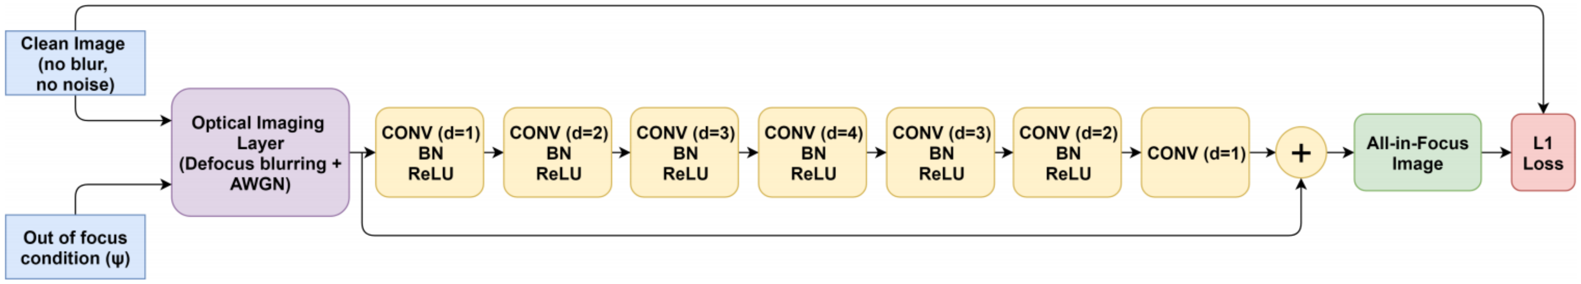

(a) (b) (c)

Multicolor localization microscopy and point-spread-function engineering by deep learning

ERAN HERSHKO,[1,2,3] LUCIEN E. WEISS,[2,3] TOMER MICHAELI,[1] AND YOAV SHECHTMAN,[2,*]

[1]*Electrical Engineering Department, Technion, 32000 Haifa, Israel*
[2]*Biomedical Engineering Department, Technion, 32000 Haifa, Israel*
[3]*Equal contribution*
**yoavsh@bm.technion.ac.il*

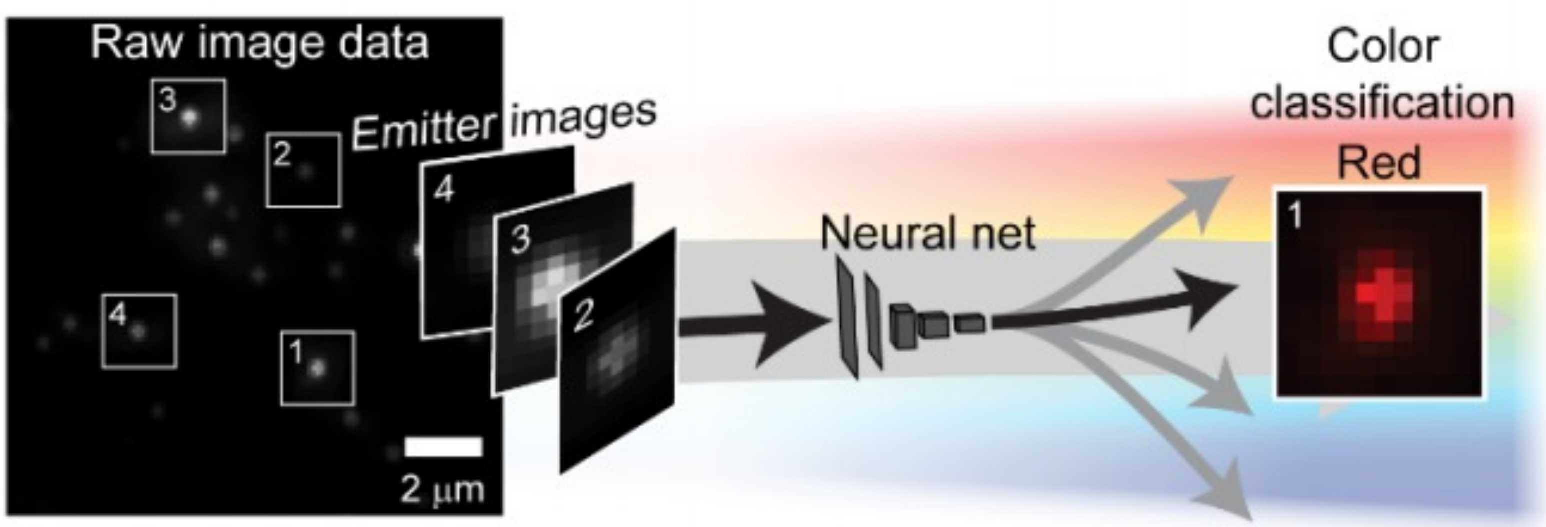

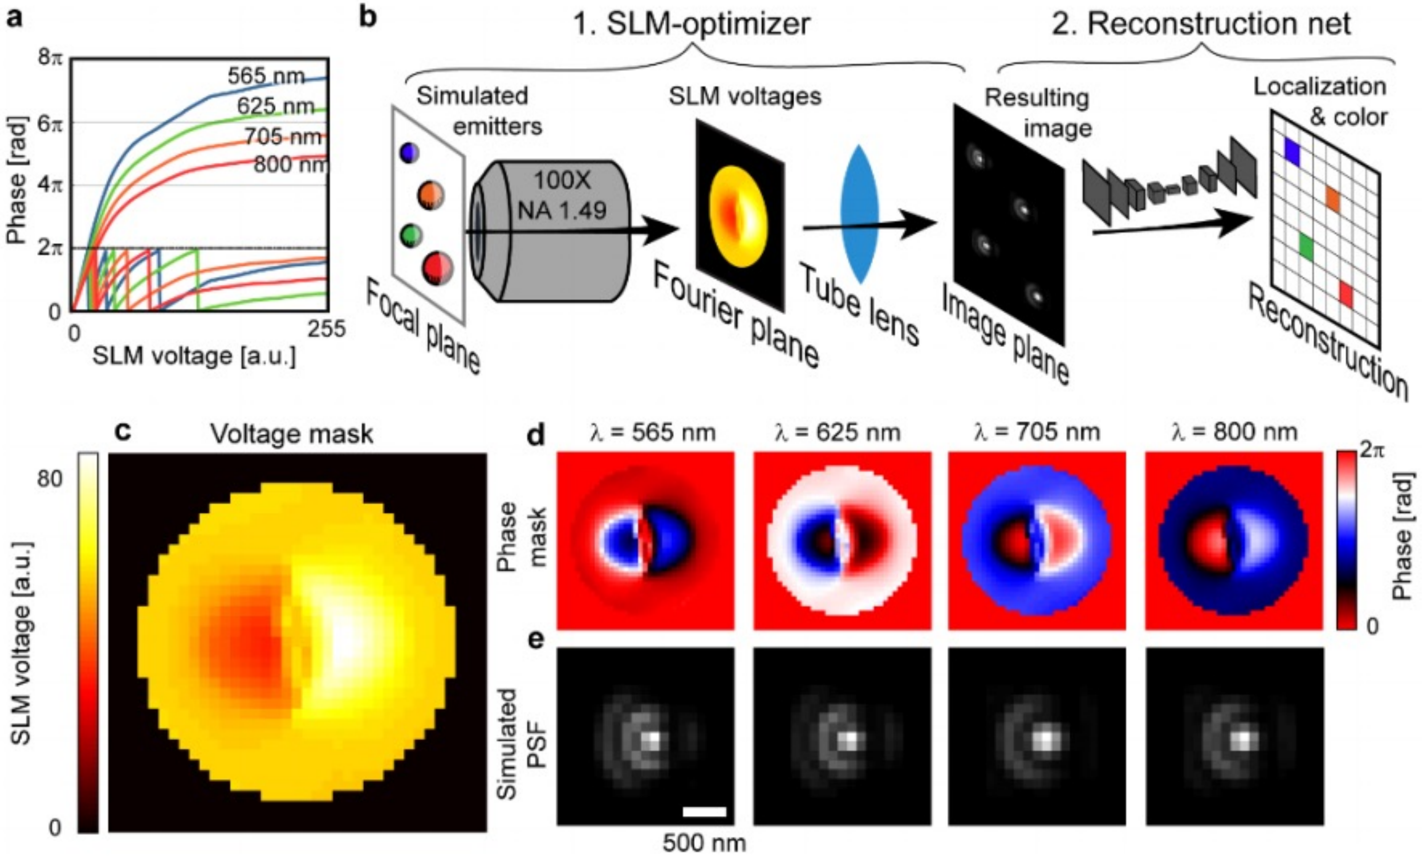

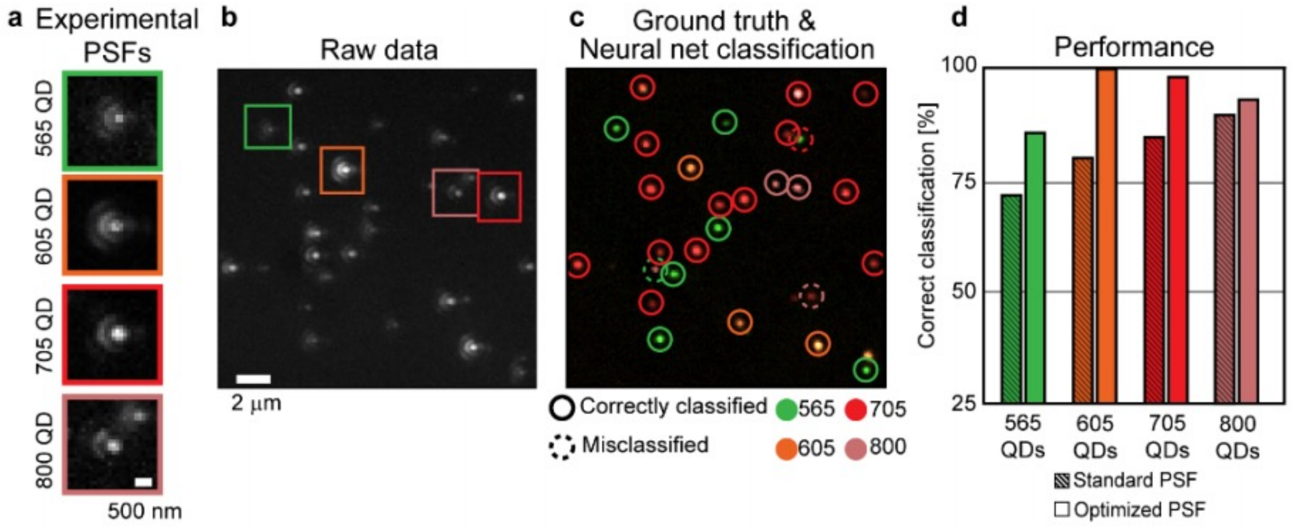

deep imaging

DeepSTORM3D: dense three dimensional localization microscopy and point spread function design by deep learning

Elias Nehme[1,2], Daniel Freedman[3], Racheli Gordon[2], Boris Ferdman[2,4], Lucien E. Weiss[2], Onit Alalouf[2], Reut Orange[2,4], Tomer Michaeli[1], and Yoav Shechtman[2,4,*]

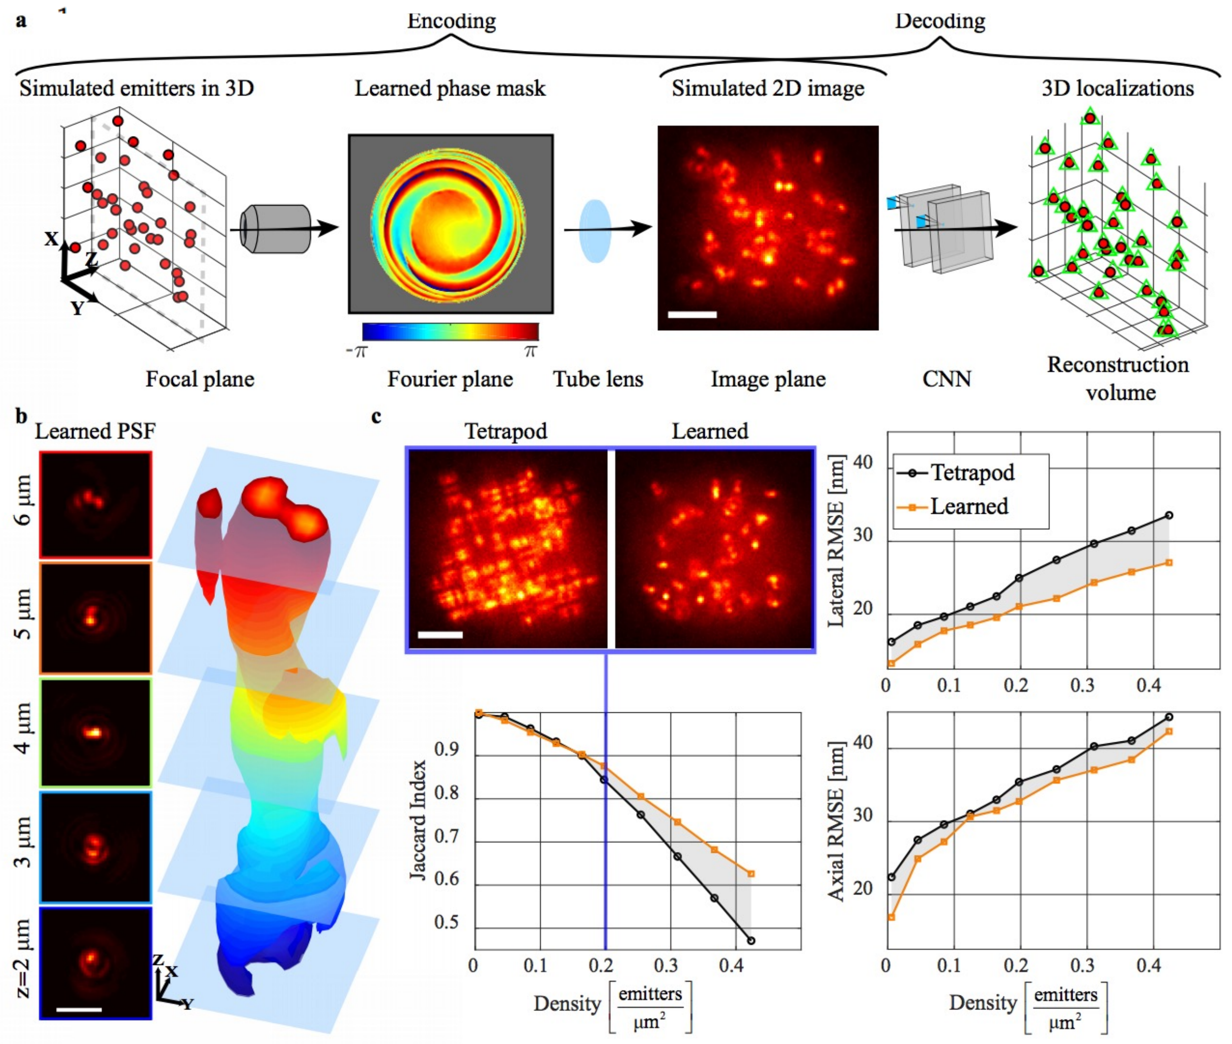

deep imaging

What physical parameters effect image formation?

- Illumination
    - Spatial pattern
    - Angle of incidence
    - Color, polarization
- Lens and optics
    - Position/orientation
    - Shape
    - Focus
    - Transparency
- Detector
    - Pixel size
    - Pixel shape & fill factor
    - Color filters
    - Other filters
- Digitization
    - E to P curves
    - Digitization schemes/thresholds
    - Data transmission, multiplexing
- Physical object

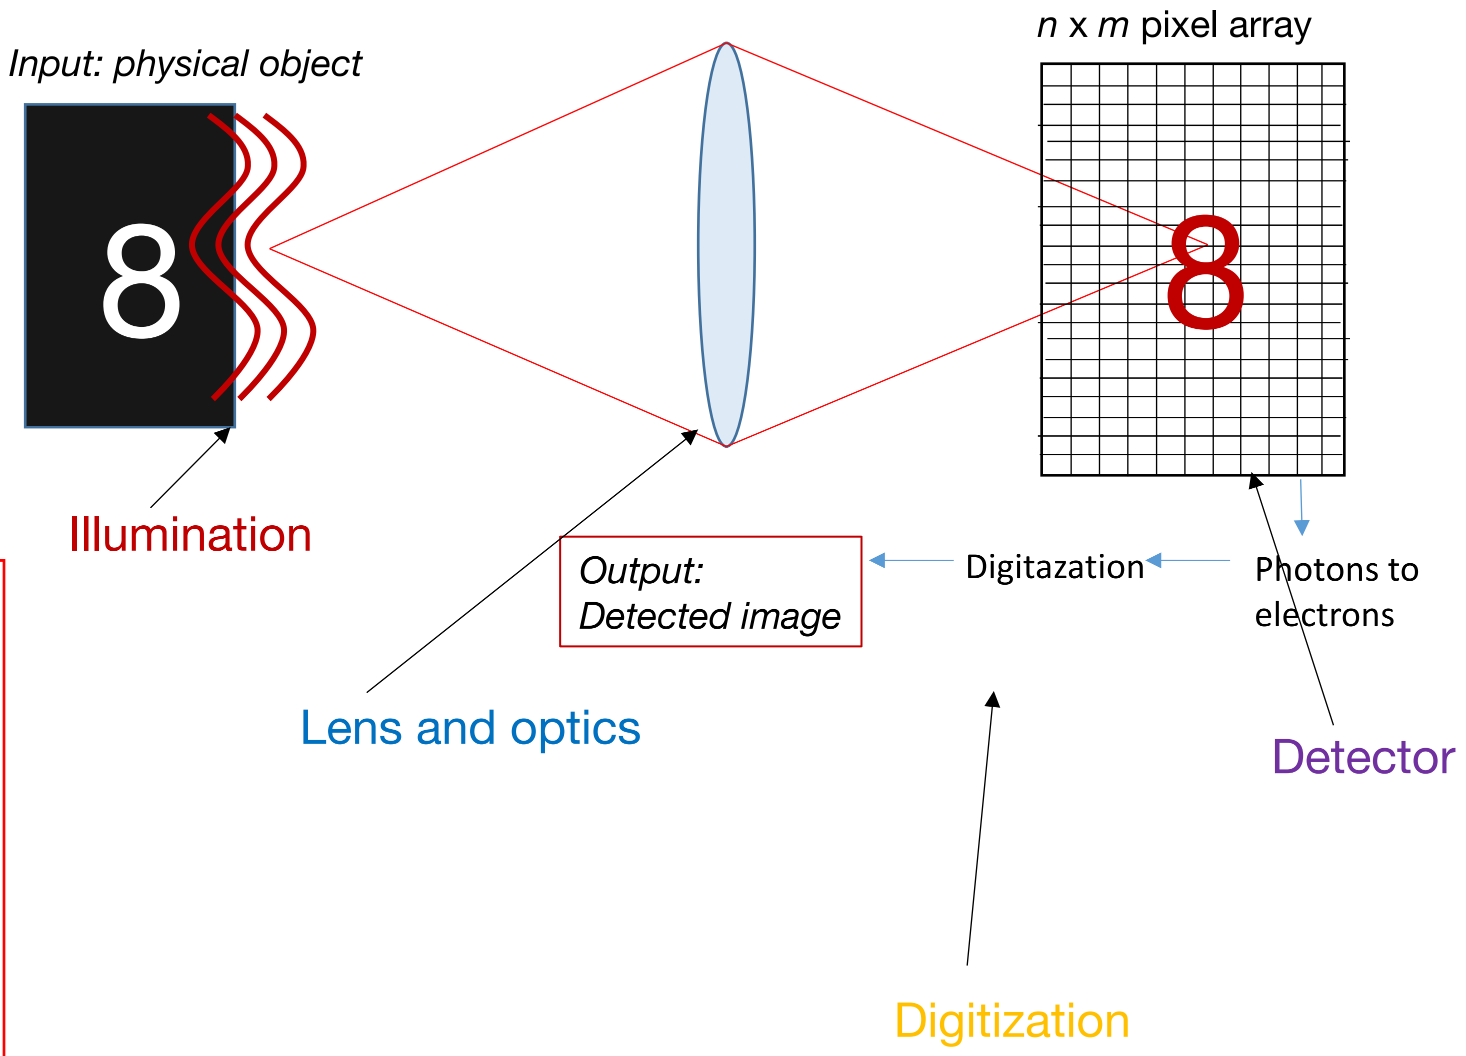

Examples: Detection and sampling

ReconNet: Non-Iterative Reconstruction of Images from Compressively Sensed Random Measurements

Kuldeep Kulkarni[1,2], Suhas Lohit[1], Pavan Turaga[1,2], Ronan Kerviche[3], and Amit Ashok[3]

[1]School of Electrical, Computer, and Energy Engineering, Arizona State University, Tempe, AZ
[2]School of Arts, Media and Engineering, Arizona State University, Tempe, AZ
[3]College of Optical Sciences, University of Arizona, Tucson, AZ

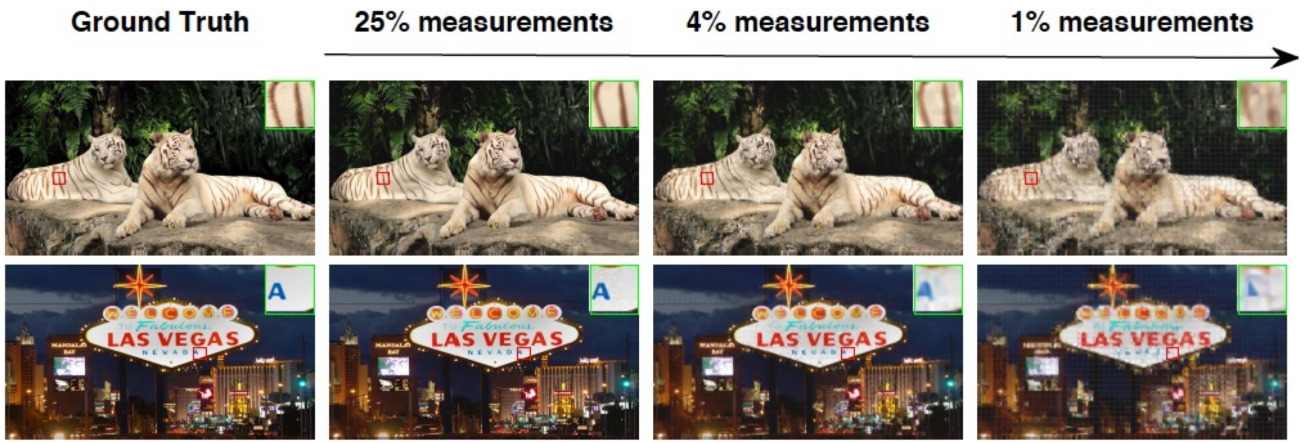

DEEP LEARNING SPARSE TERNARY PROJECTIONS FOR COMPRESSED SENSING OF IMAGES

Duc Minh Nguyen, Evaggelia Tsiligianni, Nikos Deligiannis

Vrije Universiteit Brussel, Pleinlaan 2, B-1050 Brussels, Belgium
imec, Kapeldreef 75, B-3001 Leuven, Belgium
Email: {mdnguyen, etsiligi, ndeligia}@etrovub.be

DeepBinaryMask: Learning a Binary Mask for Video Compressive Sensing

Michael Iliadis, *Member, IEEE,* Leonidas Spinoulas, *Member, IEEE,* and Aggelos K. Katsaggelos, *Fellow, IEEE*

imaging

ReconNet: Non-Iterative Reconstruction of Images from Compressively Sensed Random Measurements

Kuldeep Kulkarni[1,2], Suhas Lohit[1], Pavan Turaga[1,2], Ronan Kerviche[3], and Amit Ashok[3]

[1]School of Electrical, Computer, and Energy Engineering, Arizona State University, Tempe, AZ
[2]School of Arts, Media and Engineering, Arizona State University, Tempe, AZ
[3]College of Optical Sciences, University of Arizona, Tucson, AZ

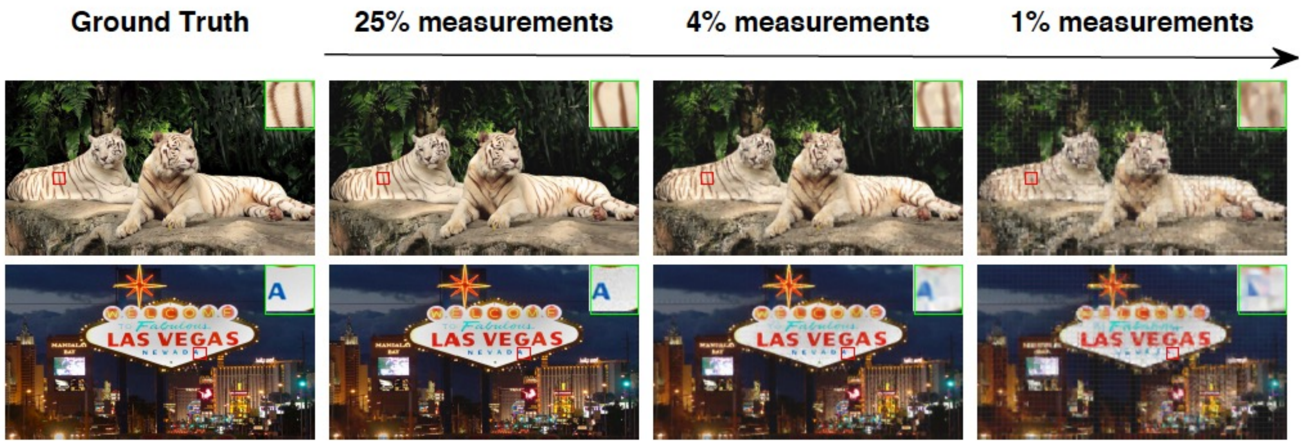

DEEP LEARNING SPARSE TERNARY PROJECTIONS FOR COMPRESSED SENSING OF IMAGES

Duc Minh Nguyen, Evaggelia Tsiligianni, Nikos Deligiannis

Vrije Universiteit Brussel, Pleinlaan 2, B-1050 Brussels, Belgium
imec, Kapeldreef 75, B-3001 Leuven, Belgium
Email: {mdnguyen, etsiligi, ndeligia}@etrovub.be

DeepBinaryMask: Learning a Binary Mask for Video Compressive Sensing

Michael Iliadis, *Member, IEEE,* Leonidas Spinoulas, *Member, IEEE,* and Aggelos K. Katsaggelos, *Fellow, IEEE*

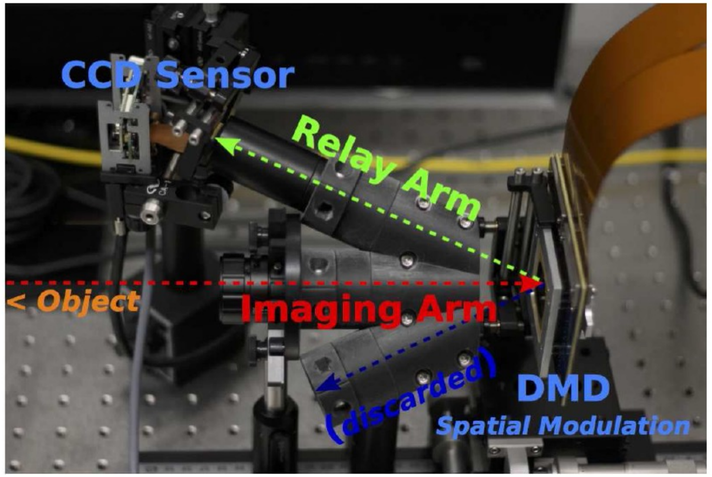

Figure 6: Compressive imager testbed layout with the object imaging arm in the center, the two DMD imaging arms are on the sides.

Standard compressive sensing problem:

$$\min_{\mathbf{x}} \quad ||\mathbf{\Psi x}||_1 \qquad s.t \qquad ||\mathbf{y} - \mathbf{\Phi x}||_2 \leq \epsilon.$$

Use iterative solvers to determine **x**

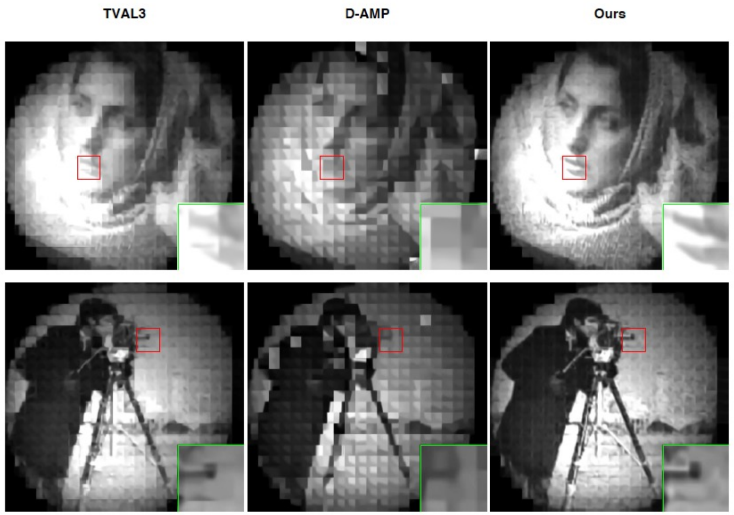

Figure 8: The figure shows reconstruction results on 3 images collected using our block SPC operating at measurement rate of 0.04. The reconstructions of our algorithm are qualitatively better than those of TVAL3 and D-AMP.

Proposed reconstruction method via CNN:

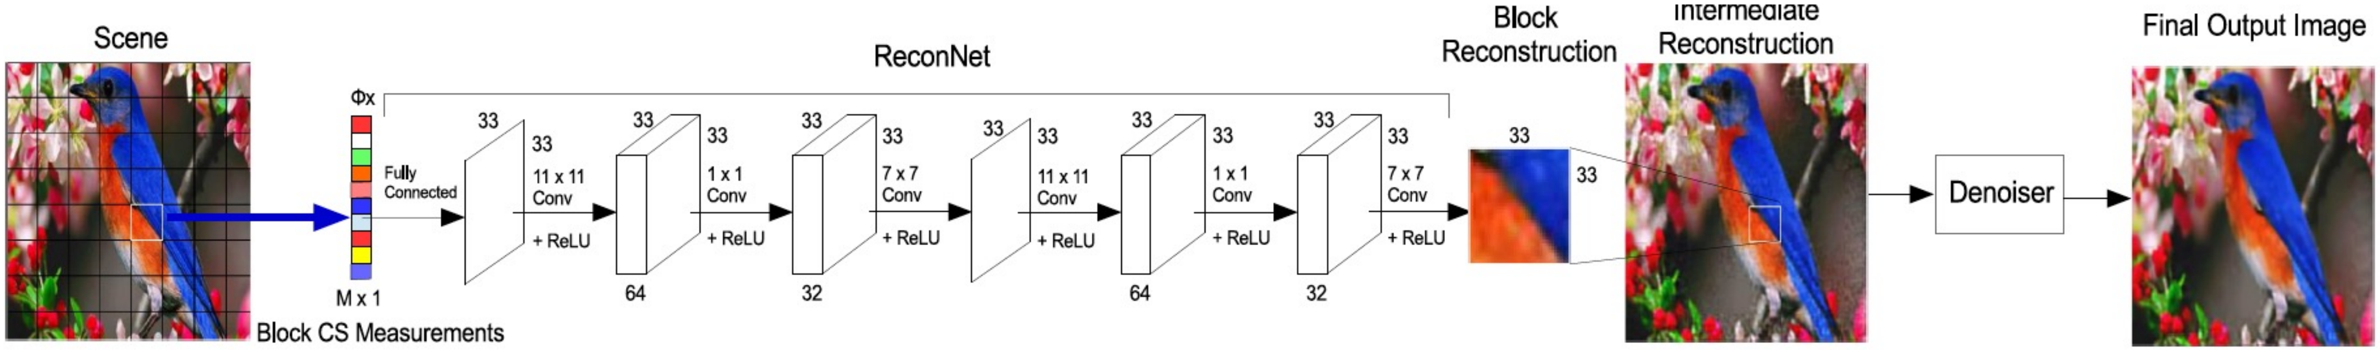

Learning Sensor Multiplexing Design through Back-propagation

Ayan Chakrabarti
Toyota Technological Institute at Chicago
6045 S. Kenwood Ave., Chicago, IL
ayanc@ttic.edu

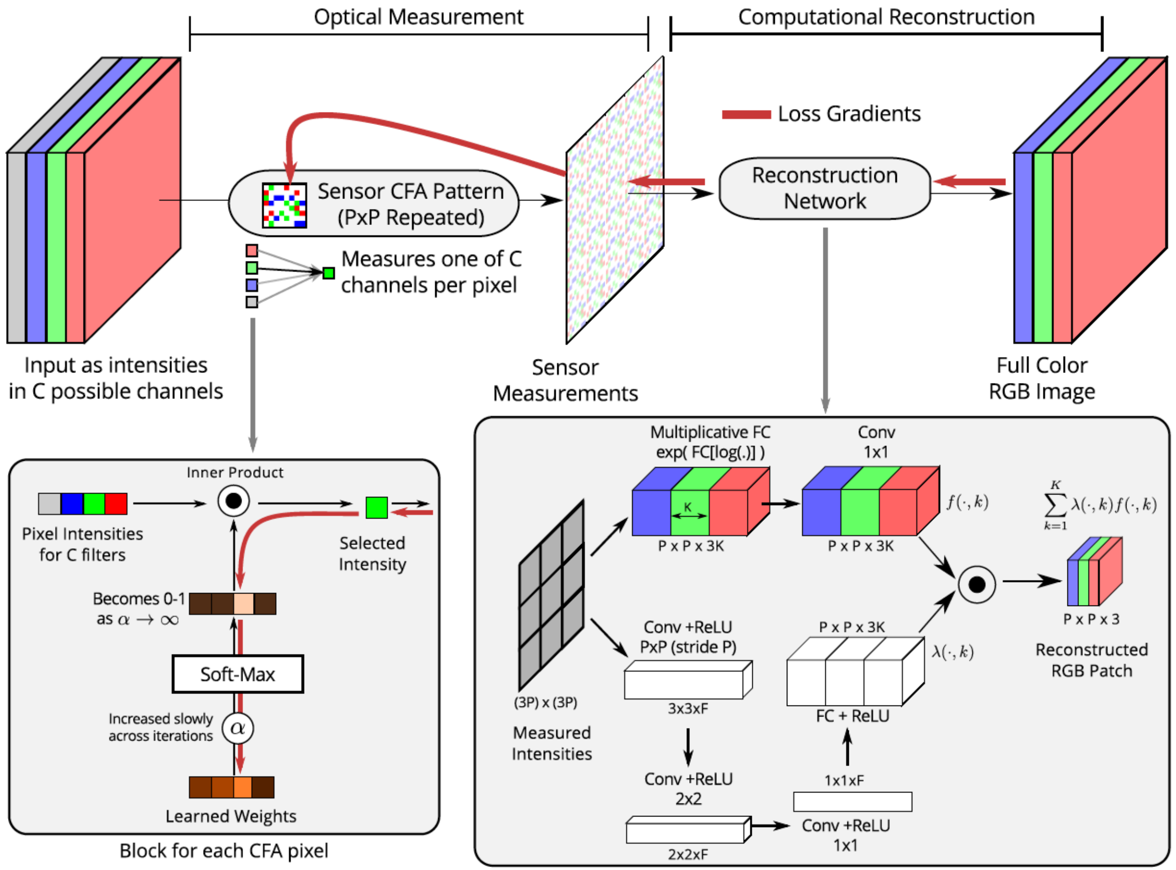

Learning Sensor Multiplexing Design through Back-propagation

Ayan Chakrabarti
Toyota Technological Institute at Chicago
6045 S. Kenwood Ave., Chicago, IL
ayanc@ttic.edu

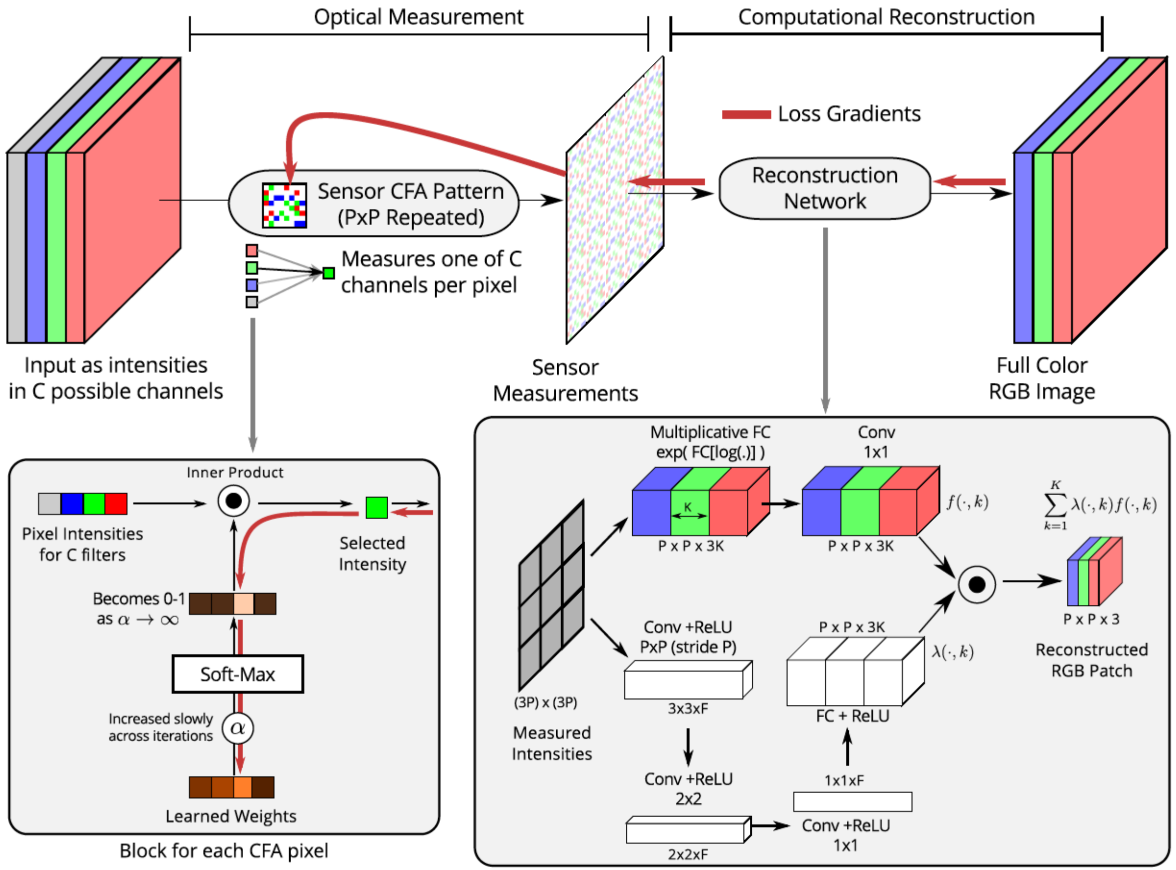

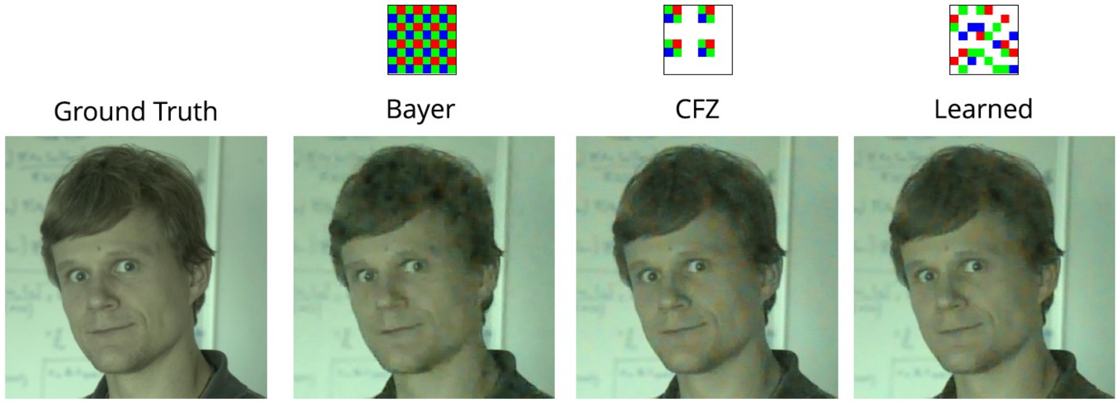

Noise STD	Percentile	Bayer [2]	CFZ [4]	Learned
0	25%	47.62	**48.04**	47.97
	50%	51.72	**52.17**	52.12
	75%	54.97	**55.32**	55.30
0.0025	25%	44.61	46.05	**46.08**
	50%	47.55	49.08	**49.17**
	75%	50.52	51.57	**51.76**
0.0050	25%	42.55	44.33	**44.37**
	50%	45.63	47.01	**47.19**
	75%	48.73	49.68	**49.94**

Adaptive Image Sampling using Deep Learning and its Application on X-Ray Fluorescence Image Reconstruction

Qiqin Dai, Henry Chopp, Emeline Pouyet, Oliver Cossairt, Marc Walton, and Aggelos K. Katsaggelos, *Fellow, IEEE*

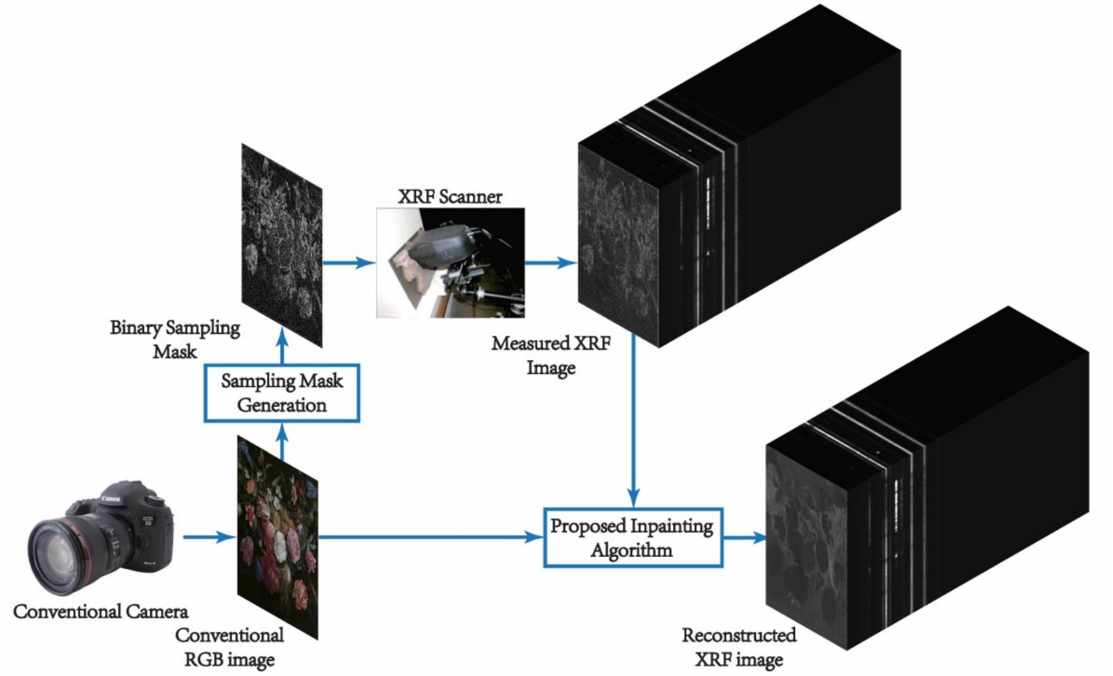

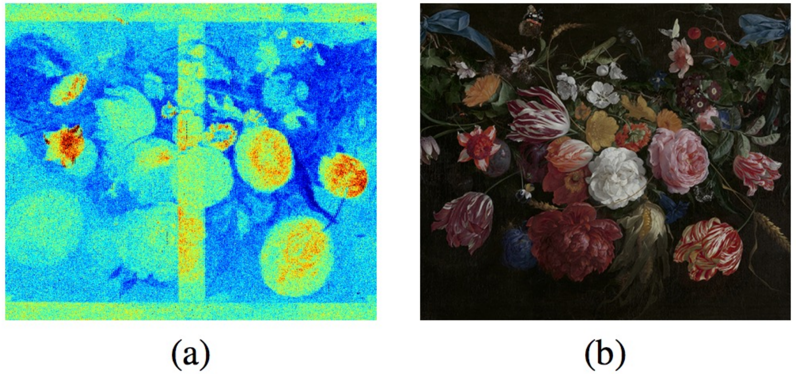

Fig. 1. (a) XRF map showing the distribution of $Pb\ L\eta$ XRF emission line (sum of channel #582 - 602) of the "Bloemen en insecten" (ca 1645), by Jan Davidsz. de Heem, in the collection of Koninklijk Museum voor Schone Kunsten (KMKSA) Antwerp and (b) the HR RGB image.

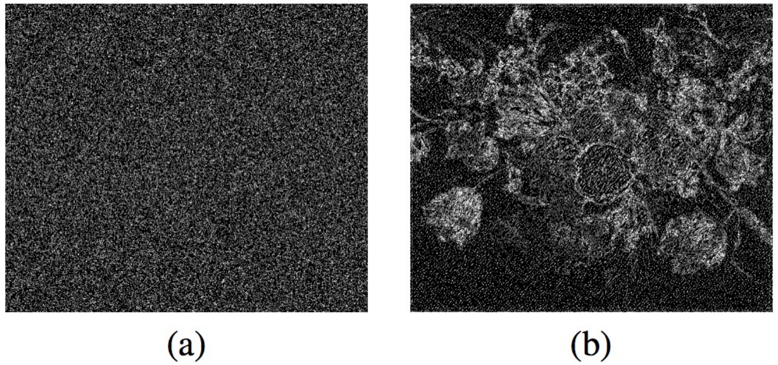

Fig. 2. (a) Random binary sampling mask that skips 80% of pixels and (b) Adaptive binary sampling mask that skips 80% of pixels based on the input RGB images in Fig 1 (b).

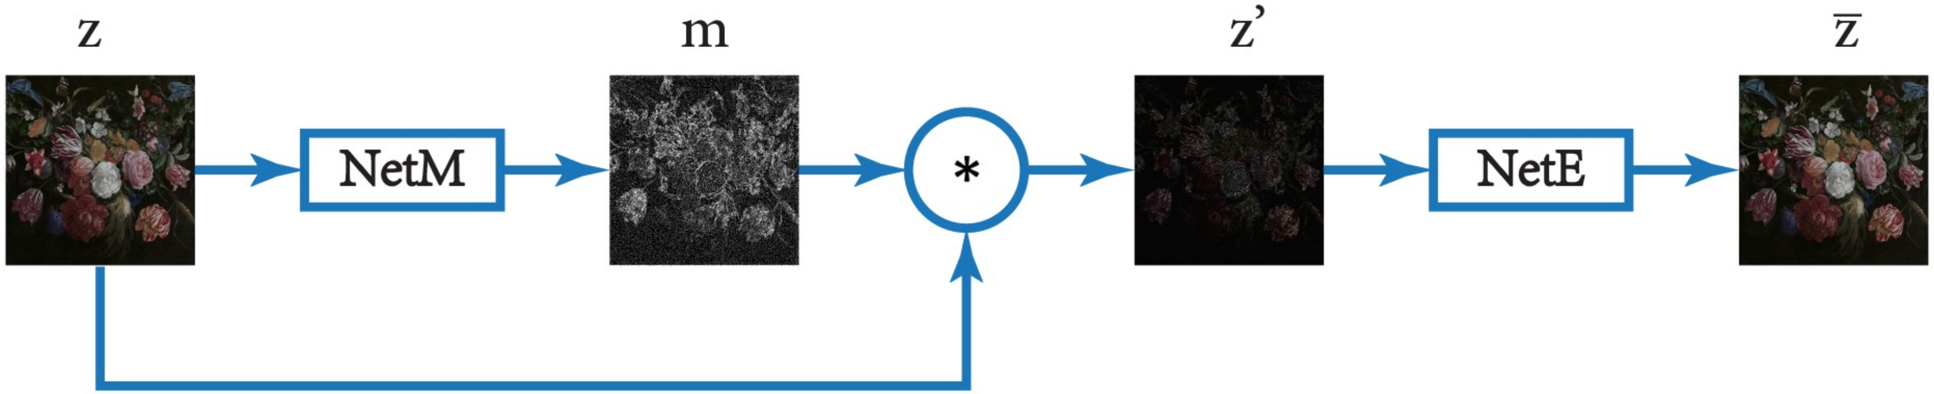
z
m
z'
$\bar{z}$
NetM
*
NetE

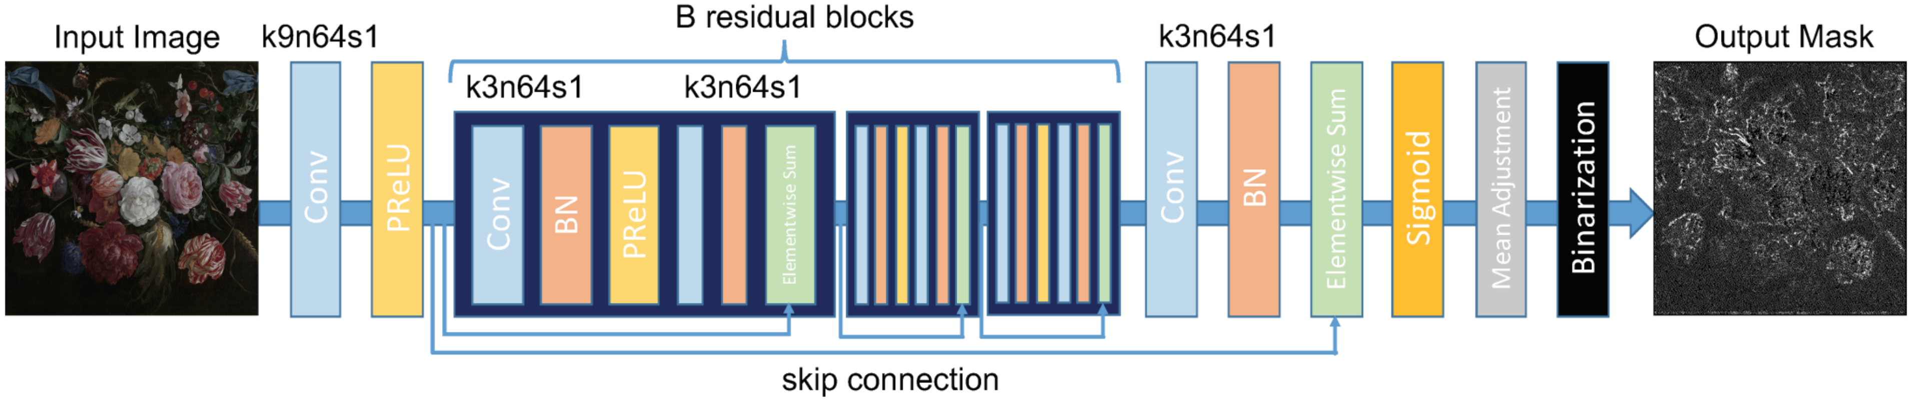
Input Image
k9n64s1
B residual blocks
k3n64s1
k3n64s1
k3n64s1
Output Mask
Conv
PReLU
Conv
BN
PReLU
Elementwise Sum
Conv
BN
Elementwise Sum
Sigmoid
Mean Adjustment
Binarization
skip connection

Learning a Variational Network for Reconstruction of Accelerated MRI Data

Kerstin Hammernik[1*], Teresa Klatzer[1], Erich Kobler[1], Michael P Recht[2,3], Daniel K Sodickson[2,3], Thomas Pock[1,4] and Florian Knoll[2,3]

[1] Institute of Computer Graphics and Vision, Graz University of Technology, Graz, Austria

[2] Center for Biomedical Imaging, Department of Radiology, NYU School of Medicine, New York, NY, United States

[3] Center for Advanced Imaging Innovation and Research (CAI²R), NYU School of Medicine, New York, NY, United States

[4] Center for Vision, Automation & Control, AIT Austrian Institute of Technology GmbH, Vienna, Austria

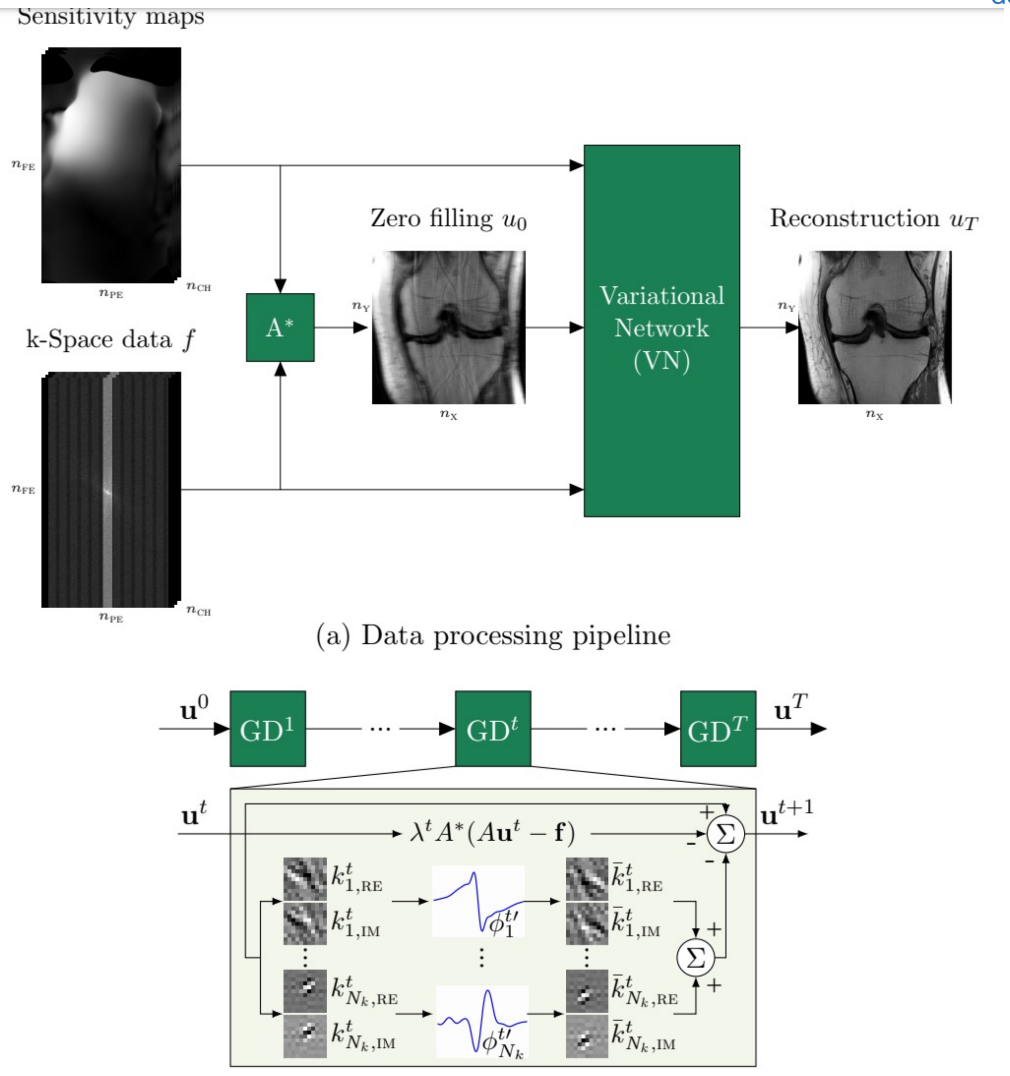

(a) Data processing pipeline

(b) Structure of the variational network (VN)

Facebook Fast MRI Challenge:
https://ai.facebook.com/blog/using-reinforcement-learning-to-personalize-ai-accelerated-mri-scans/

Reducing Uncertainty in Undersampled MRI Reconstruction with Active Acquisition

Zizhao Zhang[1,2,*] Adriana Romero[2] Matthew J. Muckley[3] Pascal Vincent[2] Lin Yang[1] Michal Drozdzal[2]
[1] University of Florida [2] Facebook AI Research [3] NYU School of Medicine

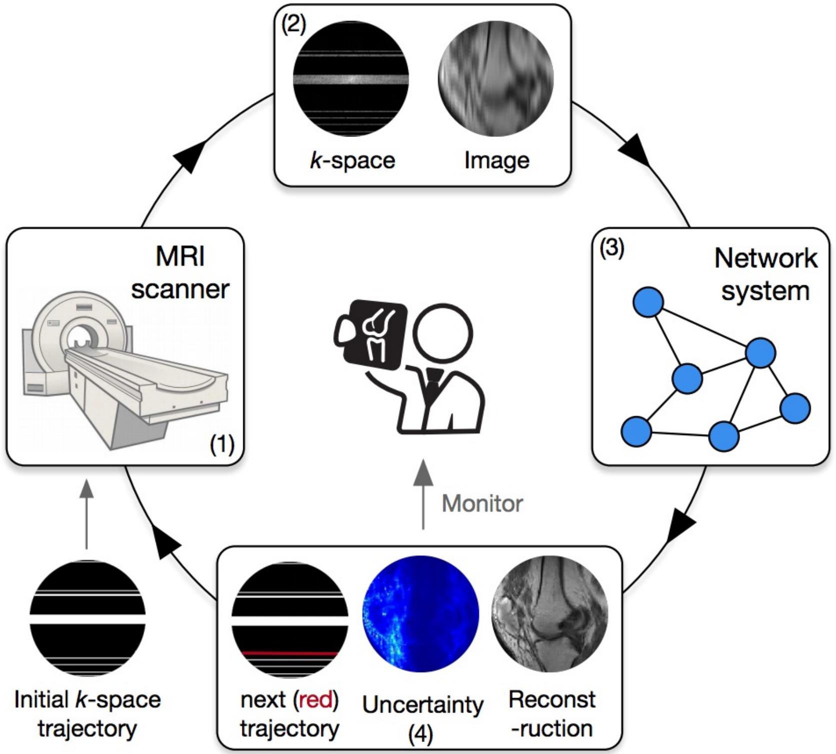

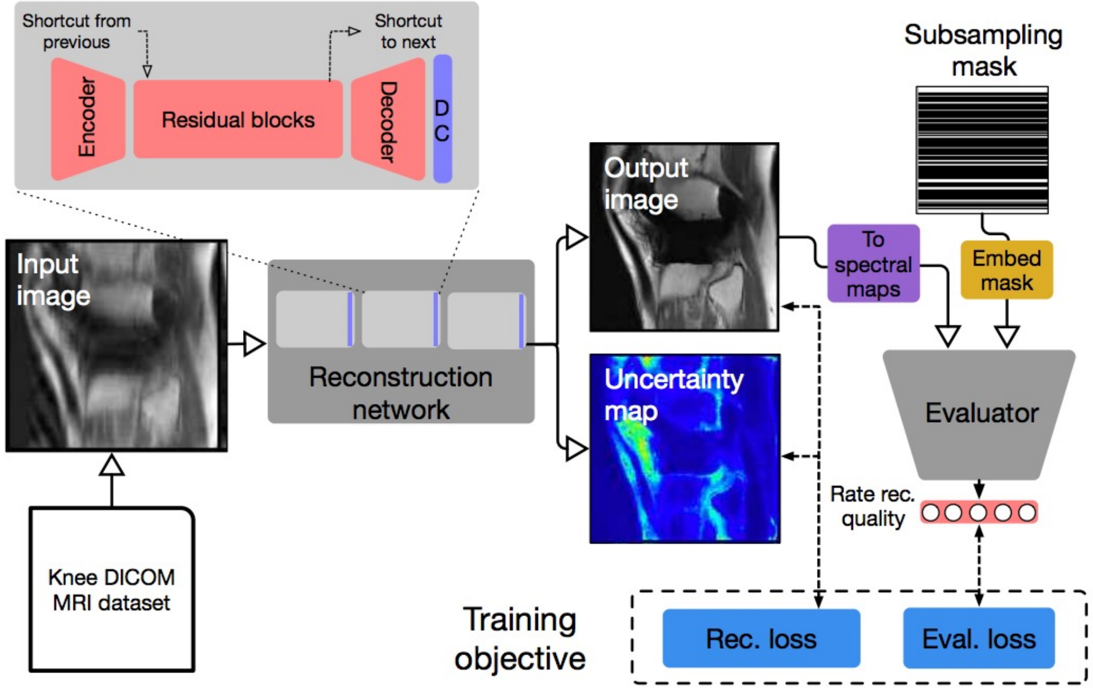

Figure 3: The training pipeline of the proposed method.

Learning a Probabilistic Strategy for Computational Imaging Sensor Selection

deep imaging

He Sun, *Member, IEEE,* and Adrian V. Dalca, *Member, IEEE,* and Katherine L. Bouman, *Member, IEEE*

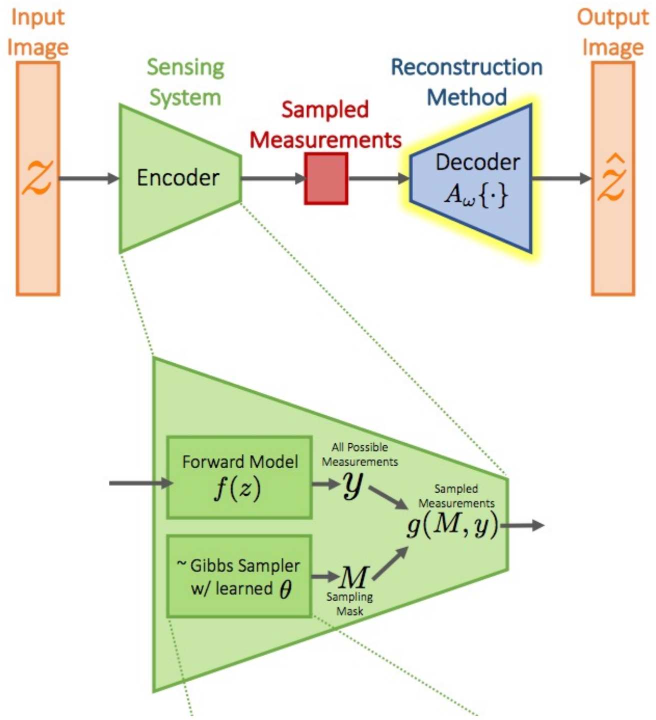

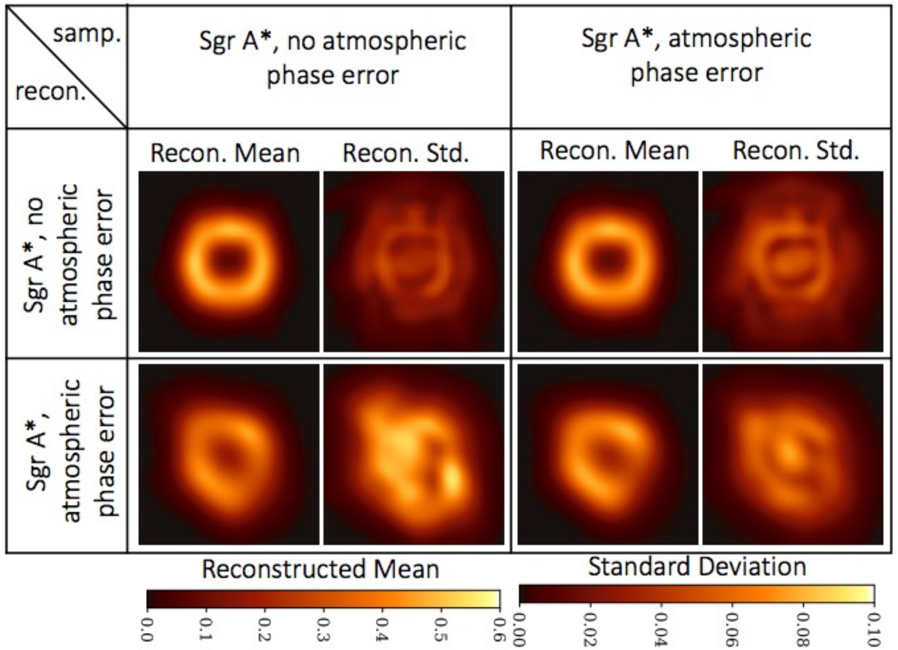

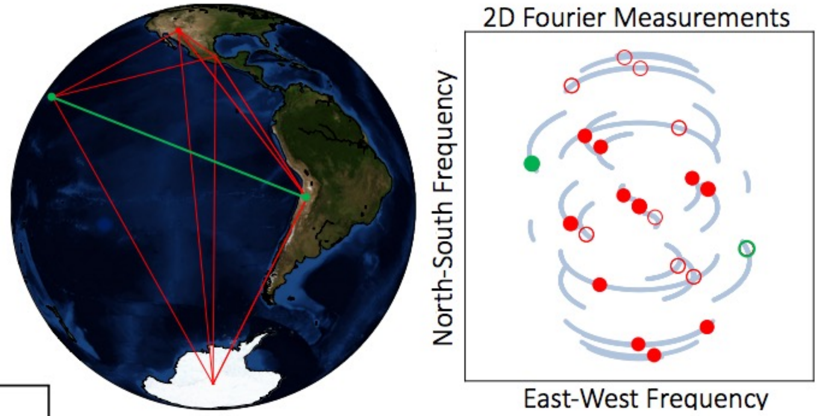

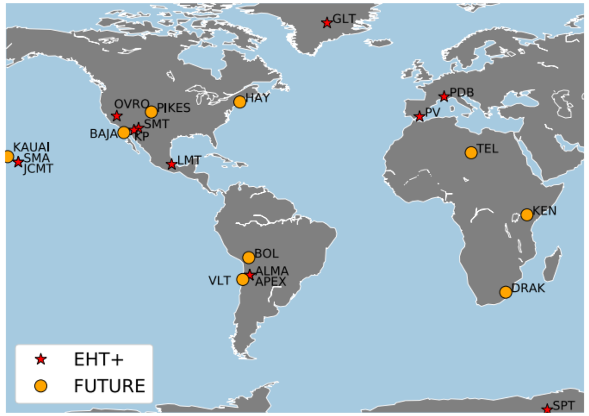

Fig. 3. Site map of potential future EHT telescope locations. Twelve sites ("EHT+") marked with blue stars are existing telescopes currently participating in or planning to join the EHT. The other nine sites ("FUTURE"), marked with orange dots, are potential locations where new telescopes could be added. "FUTURE" sites are selected as locations that can observe at the necessary 230 GHz (1.3 mm wavelength) observed by the EHT.

What physical parameters effect image formation?

- Illumination
  - Spatial pattern
  - Angle of incidence
  - Color, polarization
- Lens and optics
  - Position/orientation
  - Shape
  - Focus
  - Transparency
- Detector
  - Pixel size
  - Pixel shape & fill factor
  - Color filters
  - Other filters
- Digitization
  - E to P curves
  - Digitization schemes/thresholds
  - Data transmission, multiplexing
- Physical object

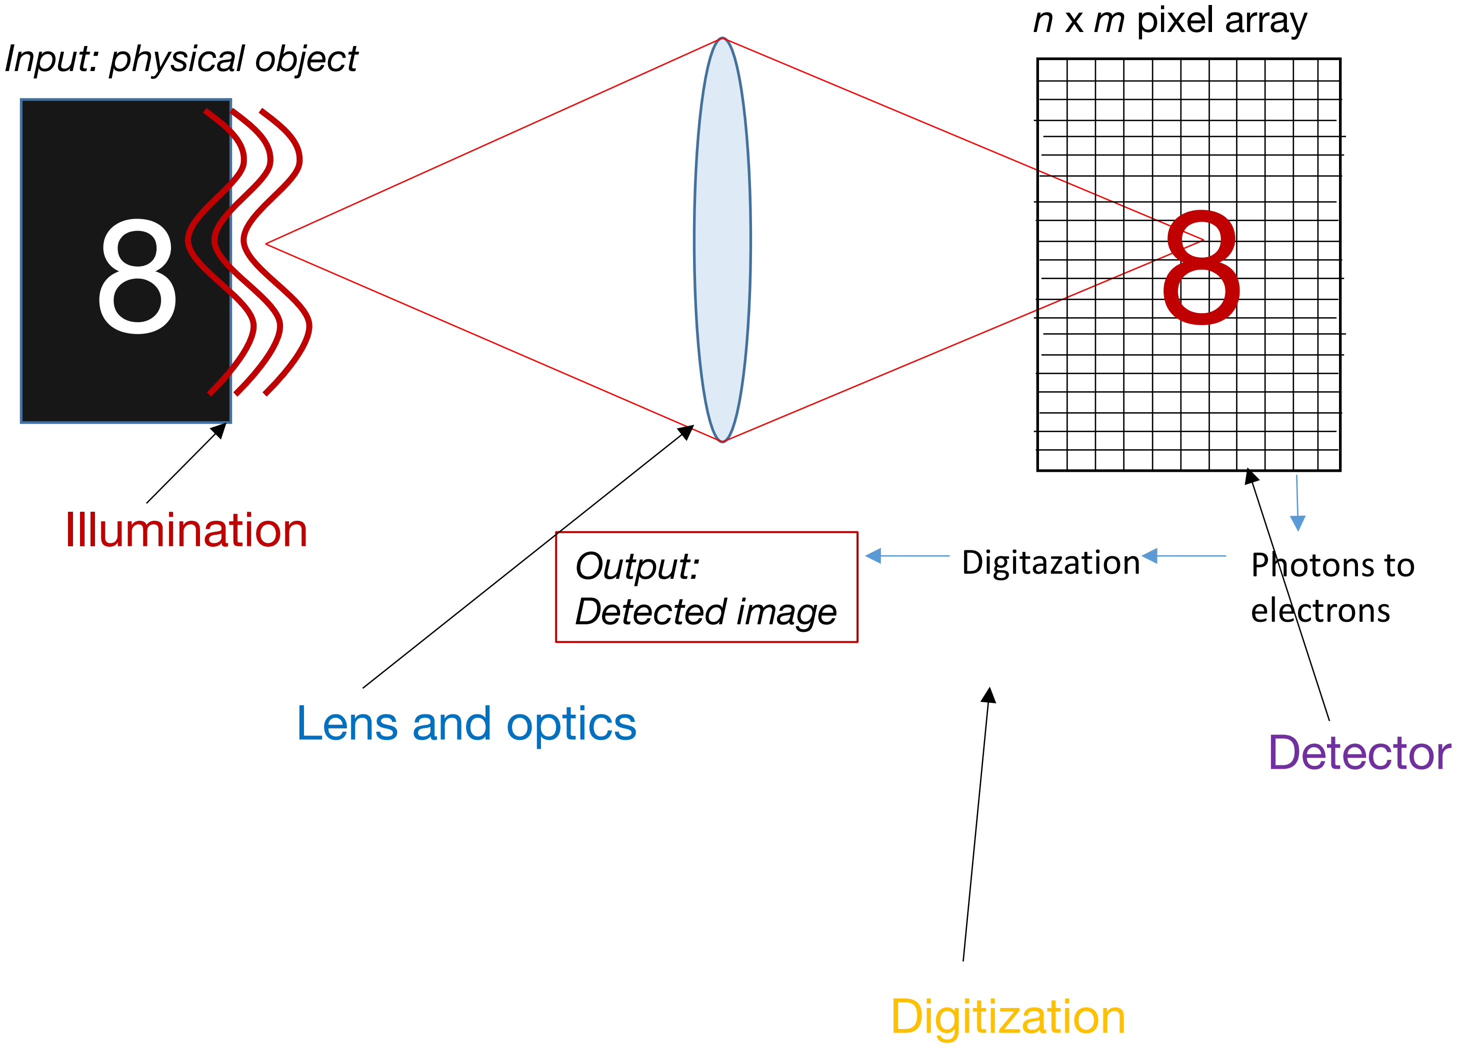

Examples: Illumination

Accurate and efficient classification with LED illumination

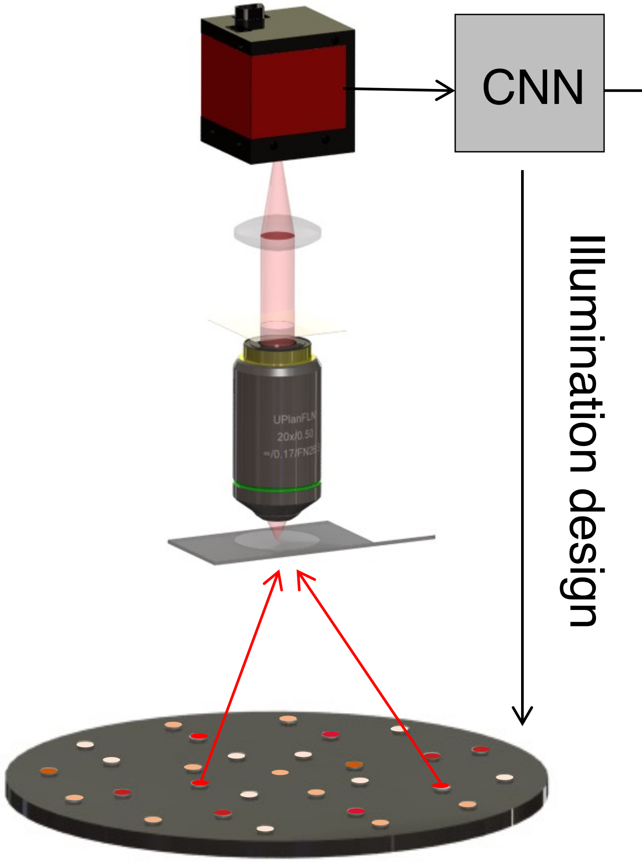

Classifier: Bright-field (1 image) | FPM (29 images) | Learned (1-2 images)

(0) Not infected

(1) Infected with malaria

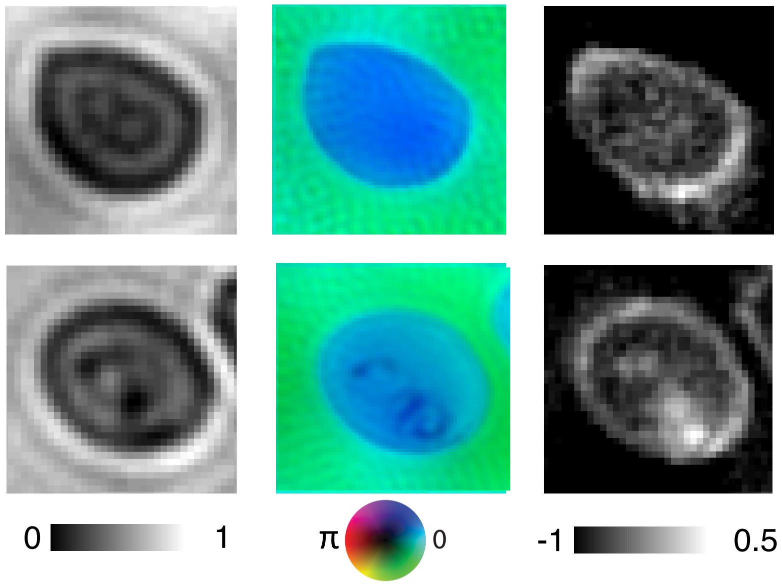

LED Pattern:

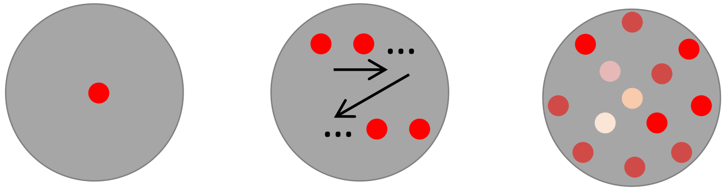

Center LED | Scanned | Learned

Classification Accuracy: **75%** | **96.5%** | **95%**

R. Horstmeyer et al., "Convolutional neural networks that teach microscopes how to image," arXiv (2017)

Physics-based Learned Design: Optimized Coded-Illumination for Quantitative Phase Imaging

Michael R. Kellman*§, Emrah Bostan*, Nicole Repina,‡
Michael Lustig,* Laura Waller*

August 13, 2018

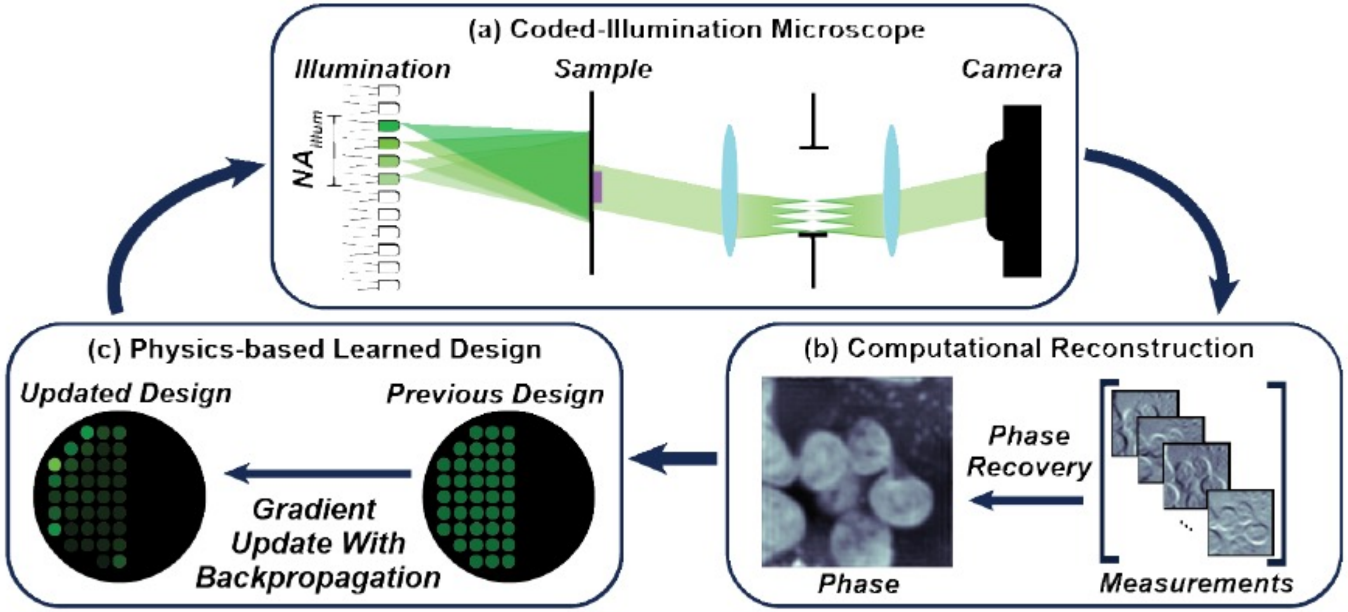

Figure 1: Learning Coded-Illumination Design for Quantitative Phase Imaging: (a) Schematic of the LED-illumination microscope where multiple intensity measurements are captured under unique coded-illumination patterns, (b) Computational phase reconstruction of the sample's optical phase with coded-illumination measurements. (c) Optimization for learning of coded-illumination design based on the non-linear iterative reconstruction.

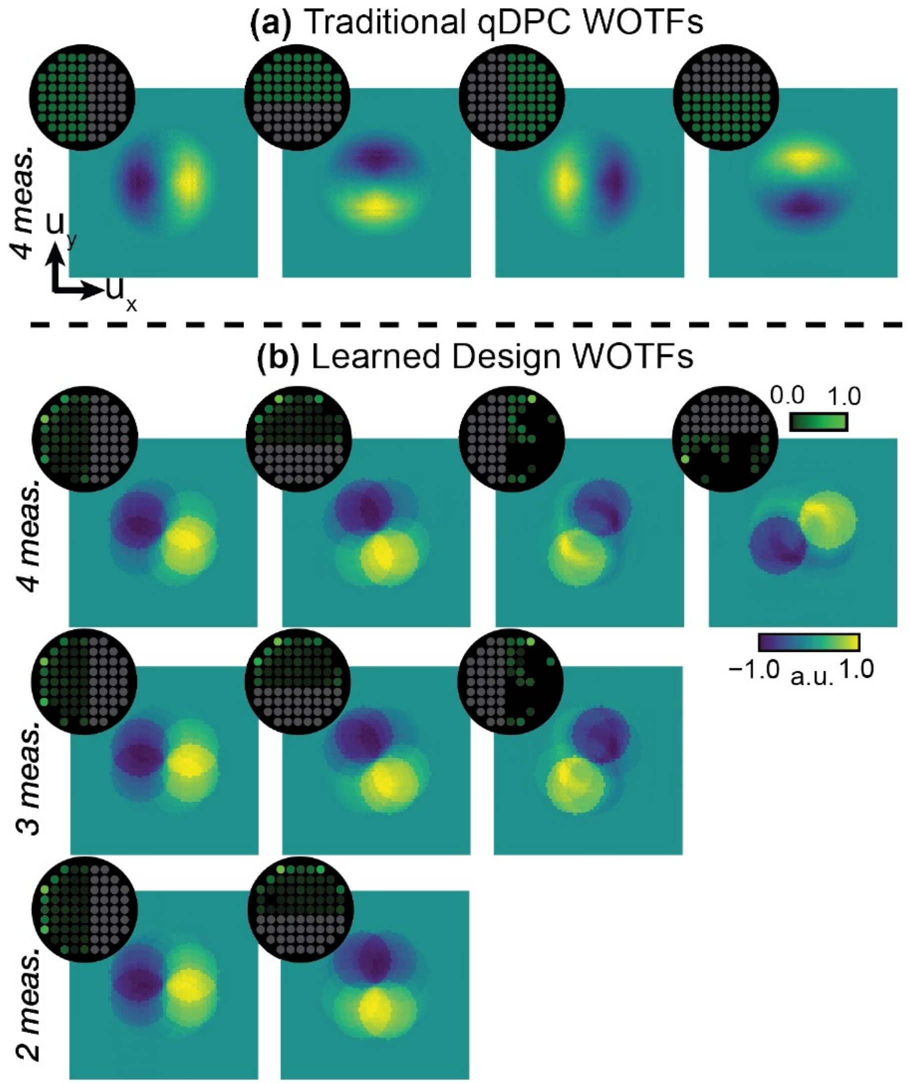
(a) Traditional qDPC WOTFs
4 meas.
u_y
u_x
(b) Learned Design WOTFs
0.0 1.0
−1.0 a.u. 1.0
4 meas.
3 meas.
2 meas.

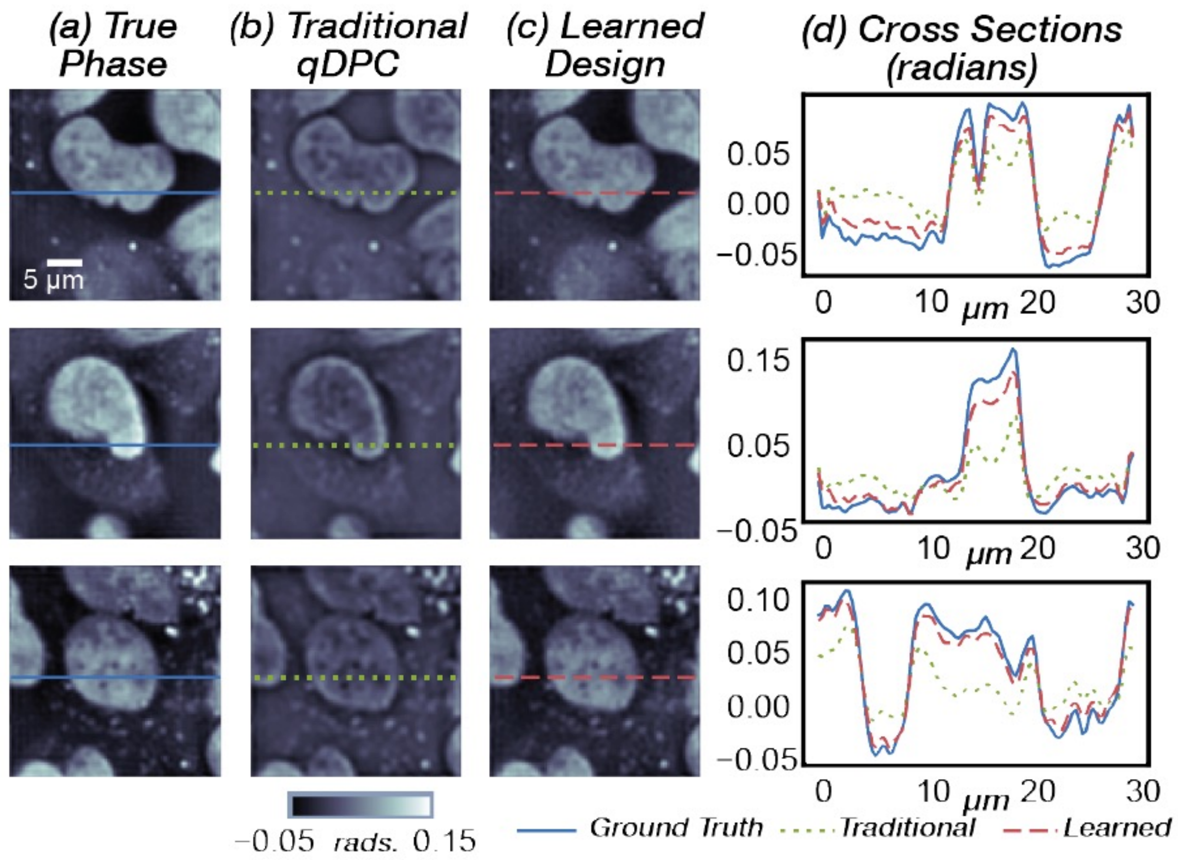
(a) True Phase
(b) Traditional qDPC
(c) Learned Design
(d) Cross Sections (radians)
5 µm
−0.05 rads. 0.15
μm
Ground Truth
Traditional
Learned

What physical parameters effect image formation?

- Illumination
  - Spatial pattern
  - Angle of incidence
  - Color, polarization
- Lens and optics
  - Position/orientation
  - Shape
  - Focus
  - Transparency
- Detector
  - Pixel size
  - Pixel shape & fill factor
  - Color filters
  - Other filters
- Digitization
  - E to P curves
  - Digitization schemes/thresholds
  - Data transmission, multiplexing
- Physical objects

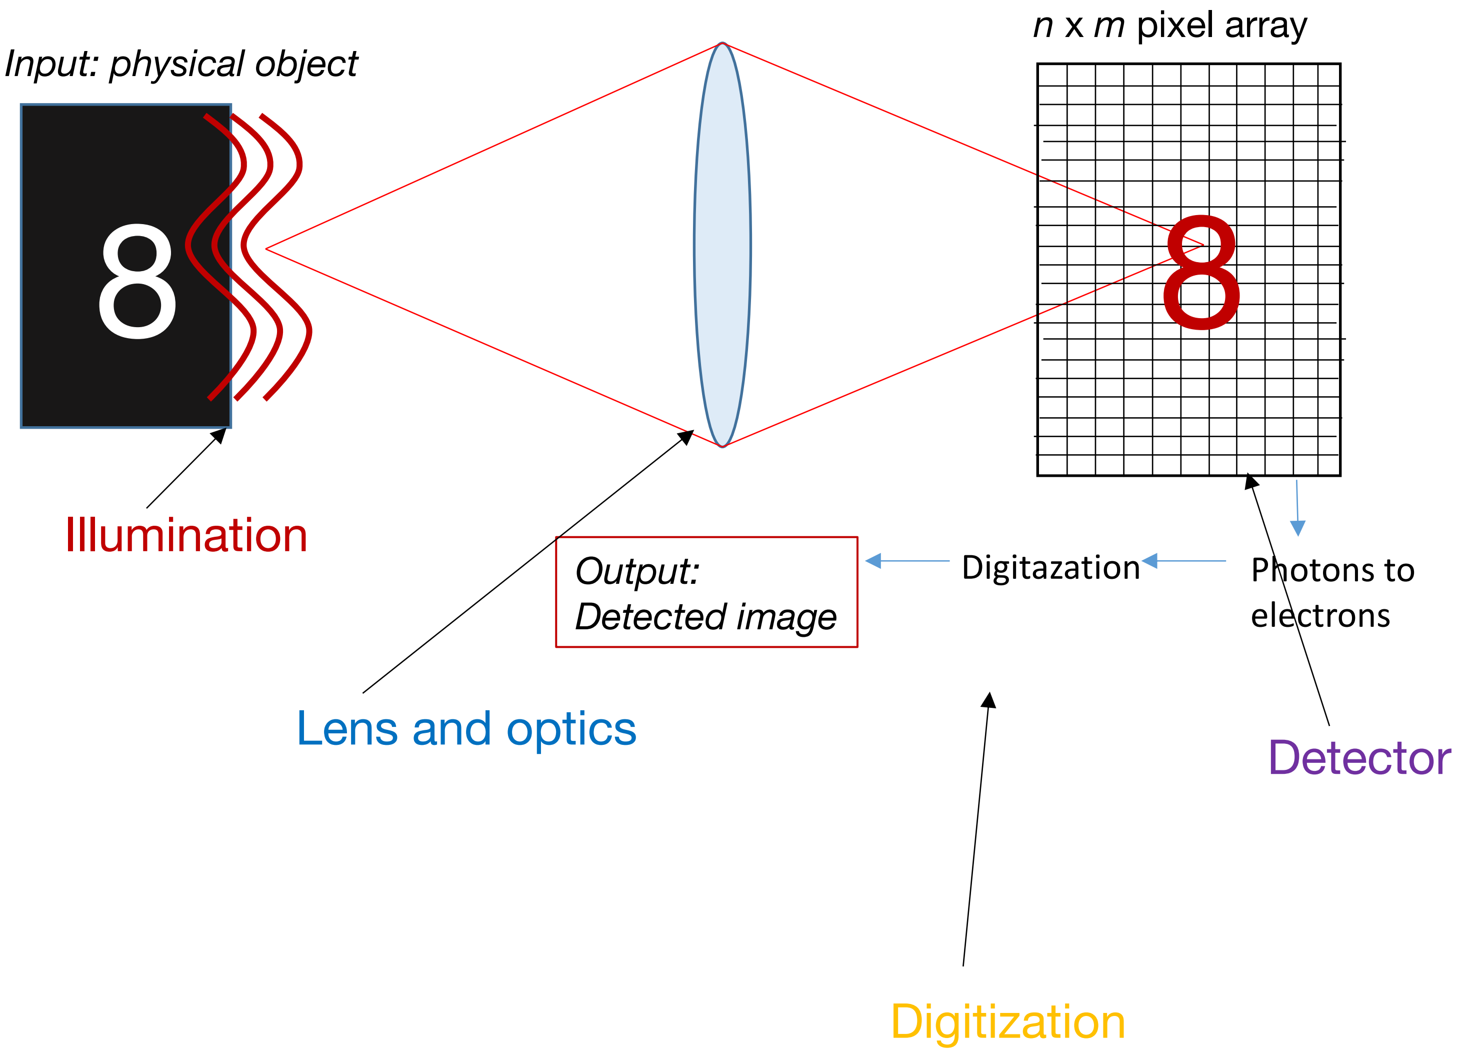

Deep learning with coherent nanophotonic circuits

Yichen Shen ✉, Nicholas C. Harris ✉, Scott Skirlo, Mihika Prabhu, Tom Baehr-Jones, Hochberg, Xin Sun, Shijie Zhao, Hugo Larochelle, Dirk Englund & Marin Soljačić

Nature Photonics **11**, 441–446 (2017) | Download Citation

Figure 1: General architecture of the ONN.

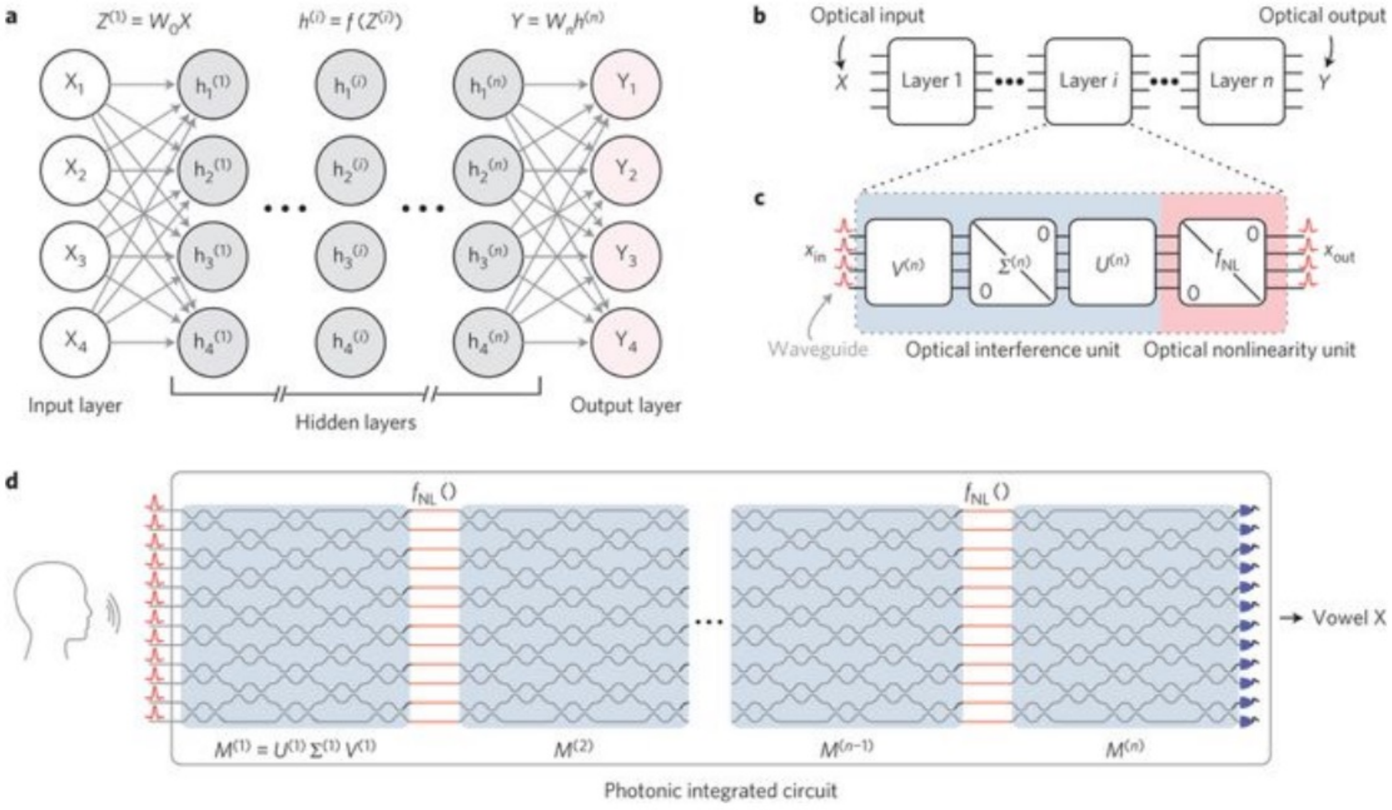

a, General artificial neural network architecture composed of an input layer, a number of hidden layers and an output layer. **b**, Decomposition of the general neural network into individual layers. **c**, Optical interference and nonlinearity units that compose each layer of the artificial neural network. **d**, Proposal for an all-optical, fully integrated neural network.

deep imaging

Figure 2: Illustration of OIU.

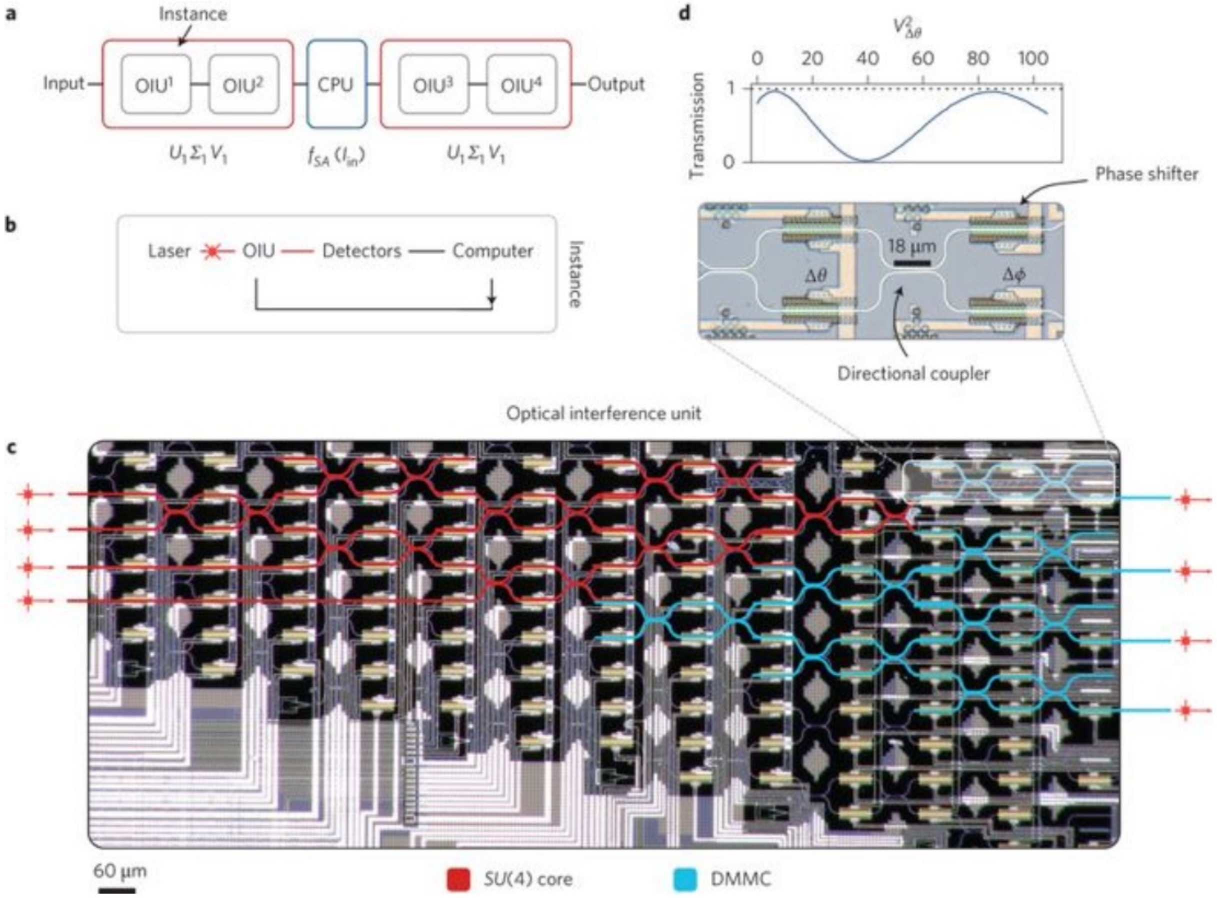

Figure 3: Vowel recognition.

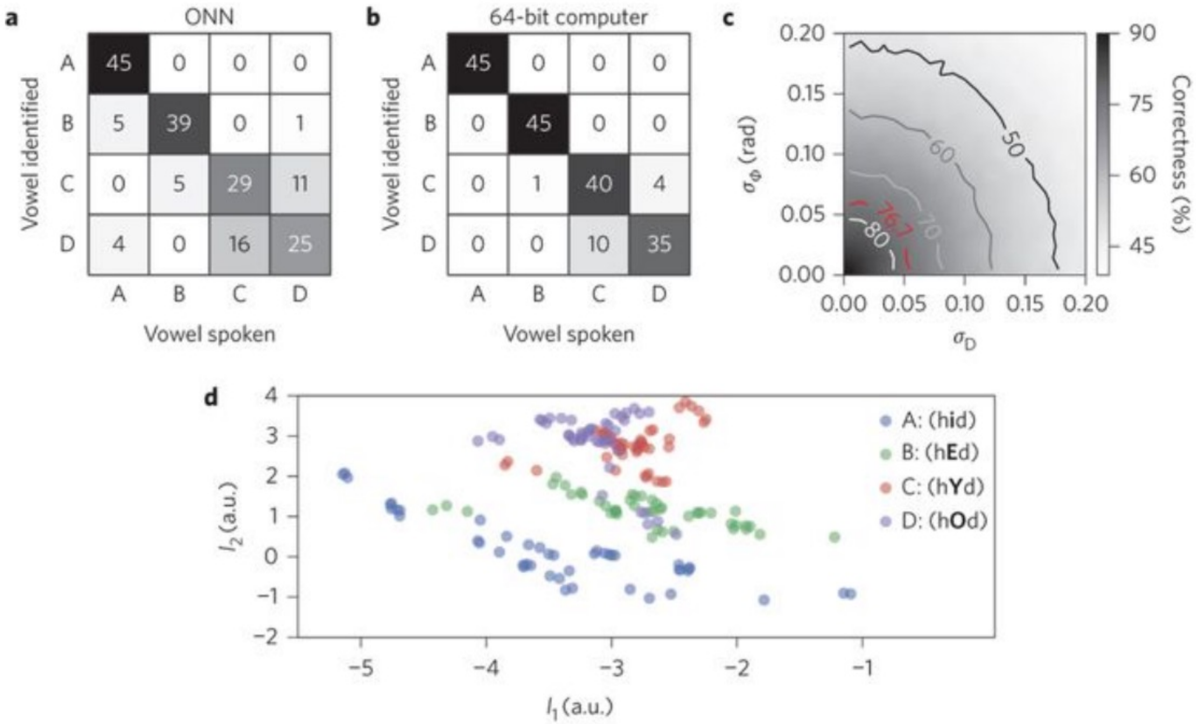

nature photonics

REVIEW ARTICLE

https://doi.org/10.1038/s41566-018-0246-9

Inverse design in nanophotonics

Sean Molesky[1], Zin Lin[2], Alexander Y. Piggott[3], Weiliang Jin[1], Jelena Vučković[3] and Alejandro W. Rodriguez[1]*

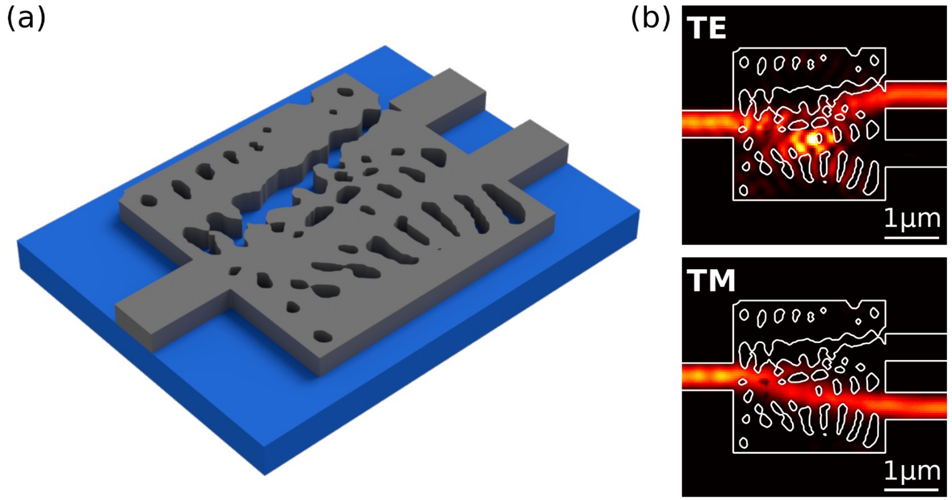

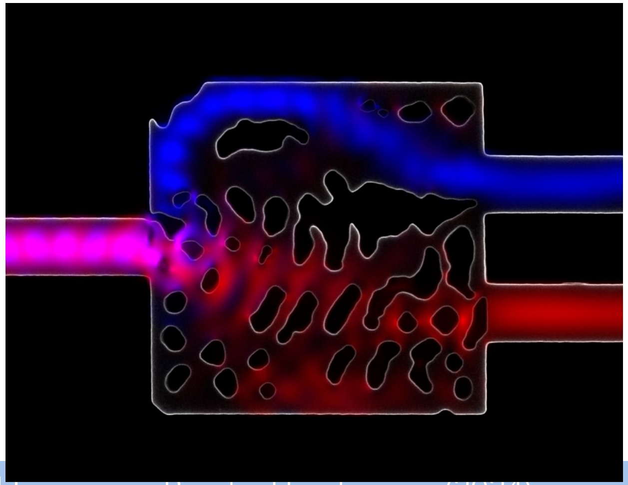

deep imaging

Machine learning based compact photonic structure design for strong light confinement

MIRBEK TURDUEV[1,5,*], CAGRI LATIFOGLU[2,5], IBRAHIM HALIL GIDEN[3,6], and Y. SINAN HANAY[4,5]

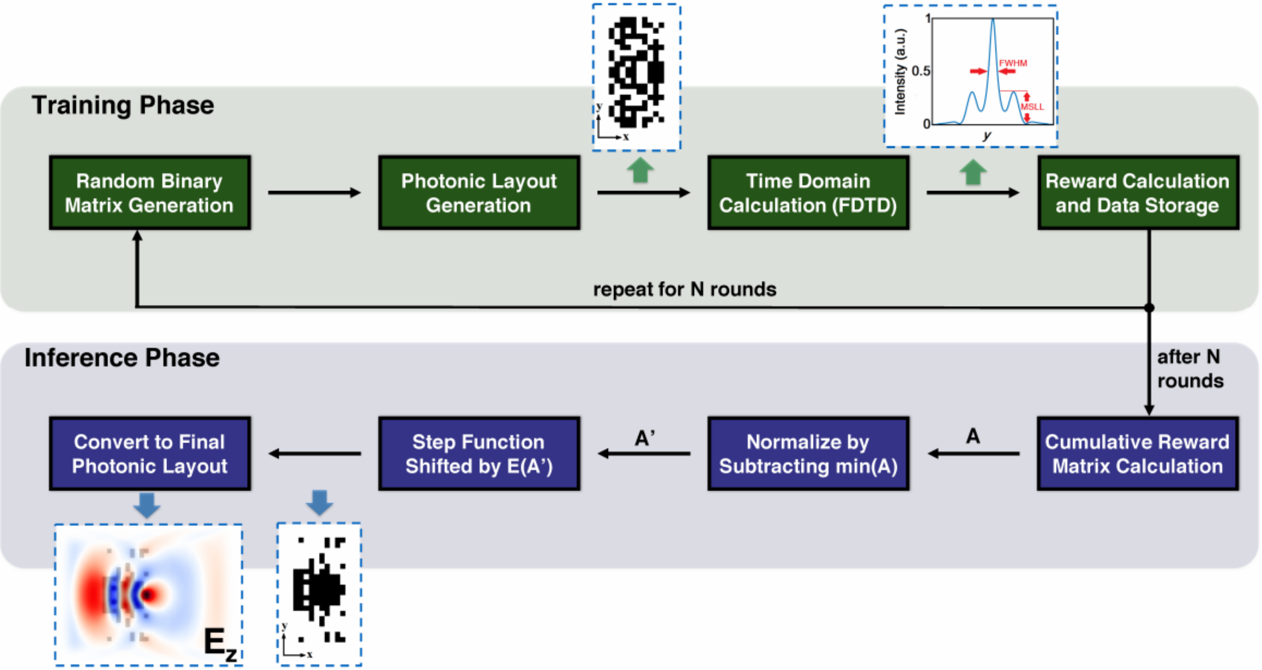

deep imaging

Ethical questions surrounding deep convolutional networks

1. What are your expectations for an image reconstruction algorithm used in a clinical setting?

Ethical questions surrounding deep convolutional networks

1. What are your expectations for an image reconstruction algorithm used in a clinical setting?

2. What types of “guarantees” should we be able to make, if any, to a patient?

deep imaging

Ethical questions surrounding deep convolutional networks

1. What are your expectations for an image reconstruction algorithm used in a clinical setting?

2. What types of "guarantees" should we be able to make, if any, to a patient?

3. How should we guide future development of ML software to meet any guarantees?

Ethical questions surrounding deep convolutional networks

1. What are your expectations for an image reconstruction algorithm used in a clinical setting?

2. What types of "guarantees" should we be able to make, if any, to a patient?

3. How should we guide future development of ML software to meet any guarantees?

4. How should we guide future development of ML-designed hardware to meet any guarantees?

5. Thoughts towards a system of checks and balances?

The Machine Learning in Imaging Ethics Questionnaire

Situation 1: In 5 years, you walk into a clinic because you have a spot on your skin that you are concerned about. The clinician is too busy, so you step over to a terminal with a standard microscope and it images your arm. It says you are fine.

Are you comfortable with leaving the office?

Yes:

No:

Why or why not? What might change how you feel?

deep imaging

The Machine Learning in Imaging Ethics Questionnaire

Situation 2: The same thing happens. But this time, the machine reports that it is 99% confident in its diagnosis, given previous examples of skin marks that have been verified by doctors as benign. It also gives you the opportunity to take a look at some of these previous example images it is basing its decision on. You notice that they don’t look 100% like the mark on your arm, as is expected, but they look pretty similar.

Are you comfortable with leaving the office?

Yes:

No:

Why or why not? What might change how you feel?

The Machine Learning in Imaging Ethics Questionnaire

Situation 3: In 5 years, the same thing happens. But this time, a doctor comes up after the machine makes its suggested diagnosis. He takes a very cursory look (10 seconds) and then confirms the machine's opinion.

Are you now comfortable with leaving the office?

Yes:

No:

The Machine Learning in Imaging Ethics Questionnaire

Situation 4: In 10 years, you go up to a modified microscope, “the Tissue Scanner 3000”, that has a number of fancy lenses and lights. As a machine learning expert by now, you’re aware that this microscope is optimized for looking at skin lesions. It performs a scan with a particular lighting configuration and reports a score of 98% confident that the lesion is benign, allowing you to look through other examples. It asks If you’d like another scan for additional confidence or a different outcome, at which point the illumination changes and it does some more scanning and reports a 99% confidence level. You can continue with another scan, but…

Are you now comfortable with leaving the office?

Yes:

No: