

Lecture 16: Introduction to Physical Layers in Machine Learning

Machine Learning and Imaging

BME 548L

Roarke Horstmeyer

Notes

- Homework 3 due Monday 3/25 at 11pm
- Homework 4 posted most likely next Monday, due 2 weeks later

- I'll review final project proposals this week

Project proposal review meetings

I would like to have a short (15 min.) 1-on-1 meeting with all final project teams this week to provide feedback on your project proposals. It is also an opportunity for you all to ask me questions, and to refine your strategy for your final project as needed early on, to maximize chances for success

Please sign up for a 15 minute time to chat here:


<https://calendly.com/rwh4/15min>

We can use the course zoom link for the meetings:

<https://duke.zoom.us/j/93342721843>

Other Computer Vision Tasks

Semantic Segmentation




GRASS, CAT,
TREE, SKY

No objects, just pixels

Balanced equation

Classification + Localization




CAT

Single Object

Over-determined

Object Detection




DOG, DOG, CAT

Multiple Object

Over-determined

Instance Segmentation



DOG, DOG, CAT

Over-determined

Super-resolution


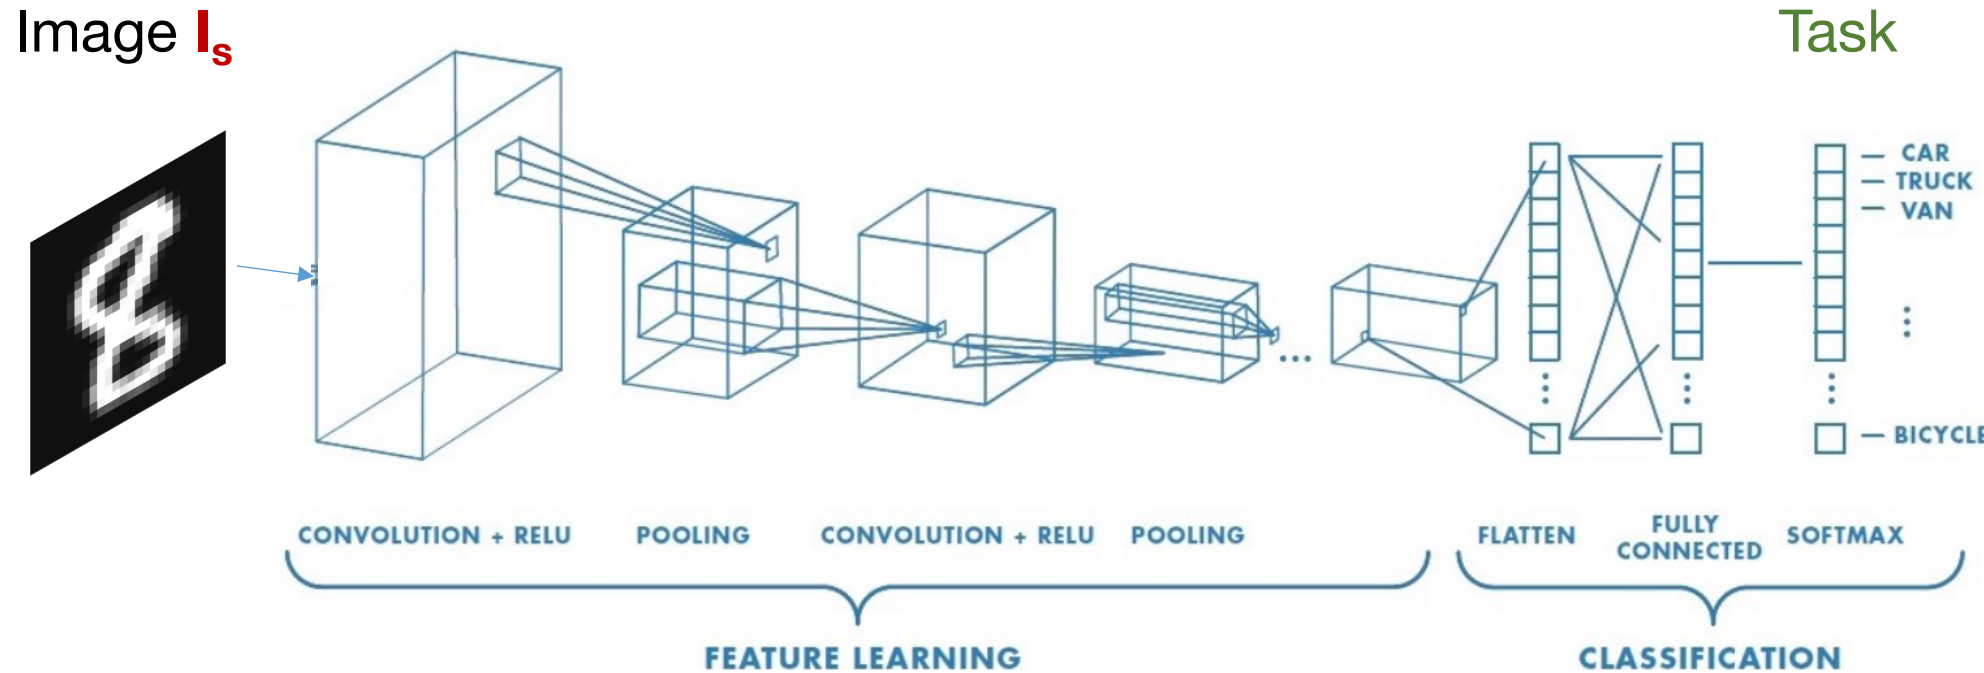


Figure 1: Super-resolved image (left) is almost indistinguishable from original (right). [4x upscaling]

Under-determined


Bringing together physical and digital image representations




$$\text{Task} = \mathbf{W}_n \dots \text{ReLU}[\mathbf{W}_1 \text{ReLU}[\mathbf{W}_0 \mathbf{I}_s] \dots]$$

Simple model of image formation

Continuous $I_0(x_0, y_0)$



$n \times m$ pixel array



Discrete signal


$$I_s(x, y), (x, y) \in \mathbb{Z}^{n \times m}$$

Digitization

Photons to electrons


What does the Sampling Theorem mean for us?

Continuous functions




*conditions

Discretize vectors
(and matrices)




Simple model of image formation

Discrete $I_0(x_0, y_0)$



$n \times m$ pixel array




Discrete signal

$$I_s(x, y), (x, y) \in \mathbb{Z}^{n \times m}$$

Digitization

Photons to electrons



Bringing together physical and digital image representations

Physical Layers

Digital Layers


Task

Physical world

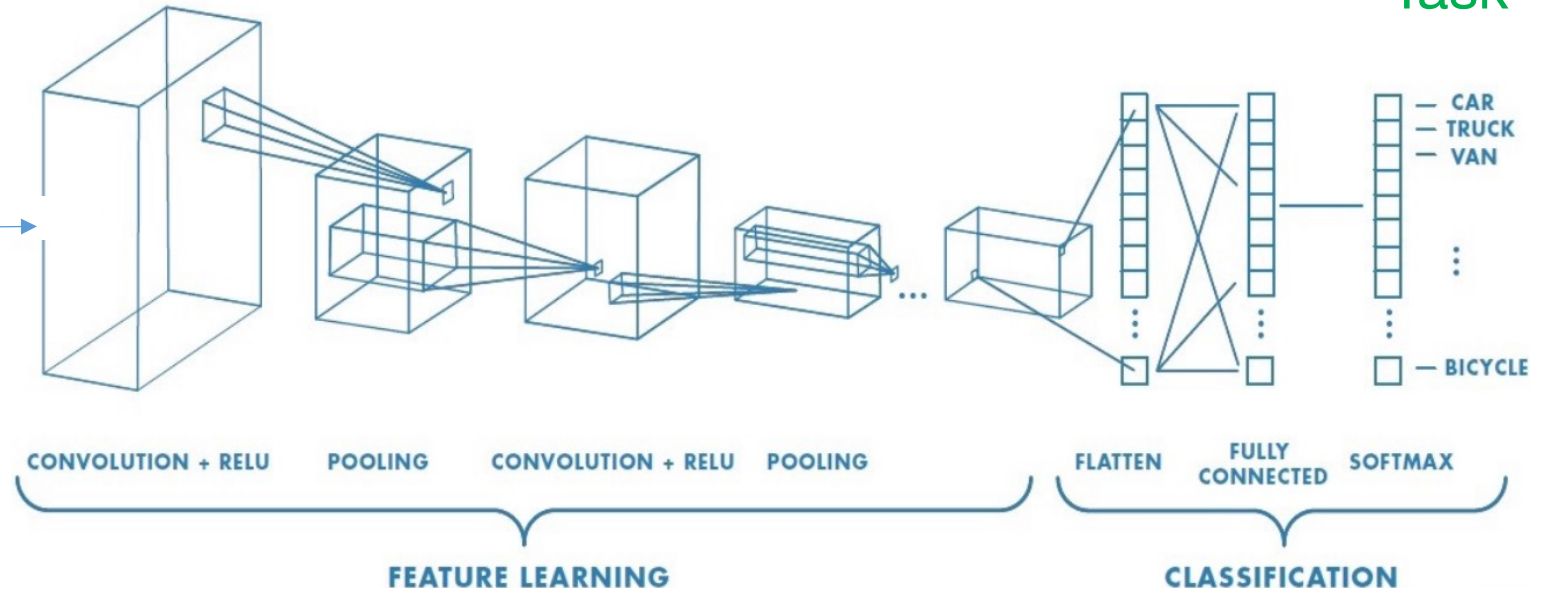
Digital Image

$I_0(x_0, y_0)$

$I_s(x, y)$




$$I_s = f[I_0]$$



Digital layers

Physical layers

$$\text{Task} = \mathbf{W}_n \dots \text{ReLU}[\mathbf{W}_1 \text{ReLU}[\mathbf{W}_0 f[I_0]] \dots]$$




Bringing together physical and digital image representations

Physical Layers

Digital Layers

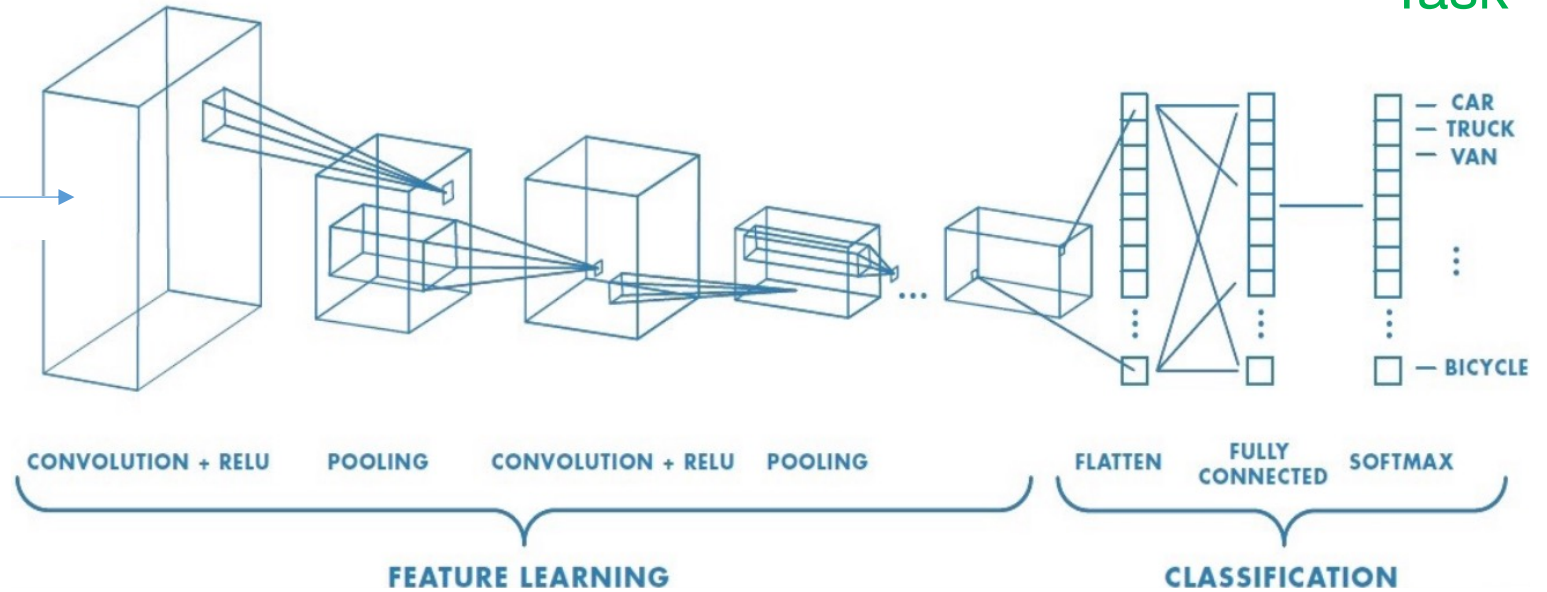
Task

Physical Function I_0



Digitized I_s

$$I_s = f[I_0]$$



Digital layers

Physical layers


$$\text{Task} = \mathbf{W}_n \dots \text{ReLU}[\mathbf{W}_1 \text{ReLU}[\mathbf{W}_0 f[I_0]] \dots]$$

Required properties of physical mapping $f[\]$ for DNN optimization?

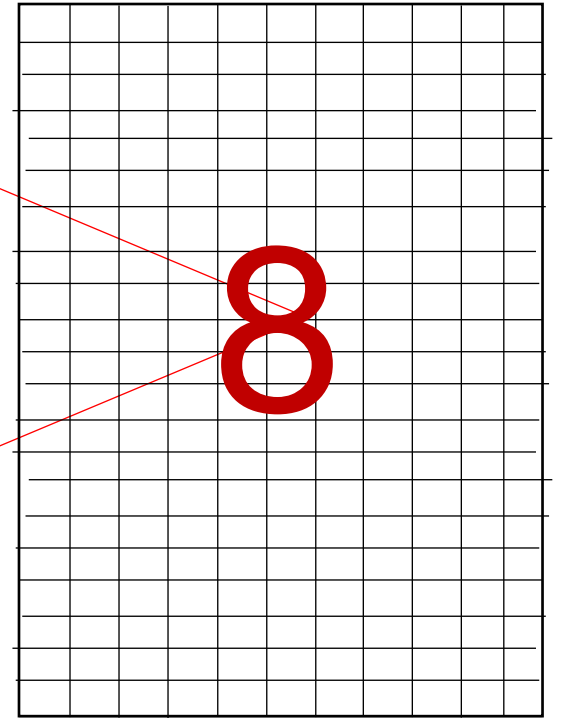
- Finite
- Non-zero gradients
- Differentiable*
- Known structure (for now...)
- Anything else?

What physical parameters effect image formation?

Input: physical object



$n \times m$ pixel array




Output: Detected image

Digitazation


Photons to electrons

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





Simple mathematical model of image formation



First - what is light and how can we model it?

- Interpretation #1: Radiation (*Incoherent*)
- Model: Rays





- Real, non-negative
- Models absorption and brightness

$$I_{\text{tot}} = I_1 + I_2$$

First - what is light and how can we model it?



- Interpretation #1: Radiation (*Incoherent*)
- Model: Rays



- Real, non-negative
- Models absorption and brightness

$$I_{\text{tot}} = I_1 + I_2$$

- Interpretation #2: Electromagnetic wave (*Coherent*)
- Model: Waves




- Complex field
- Models Interference

$$E_{\text{tot}} = E_1 + E_2$$

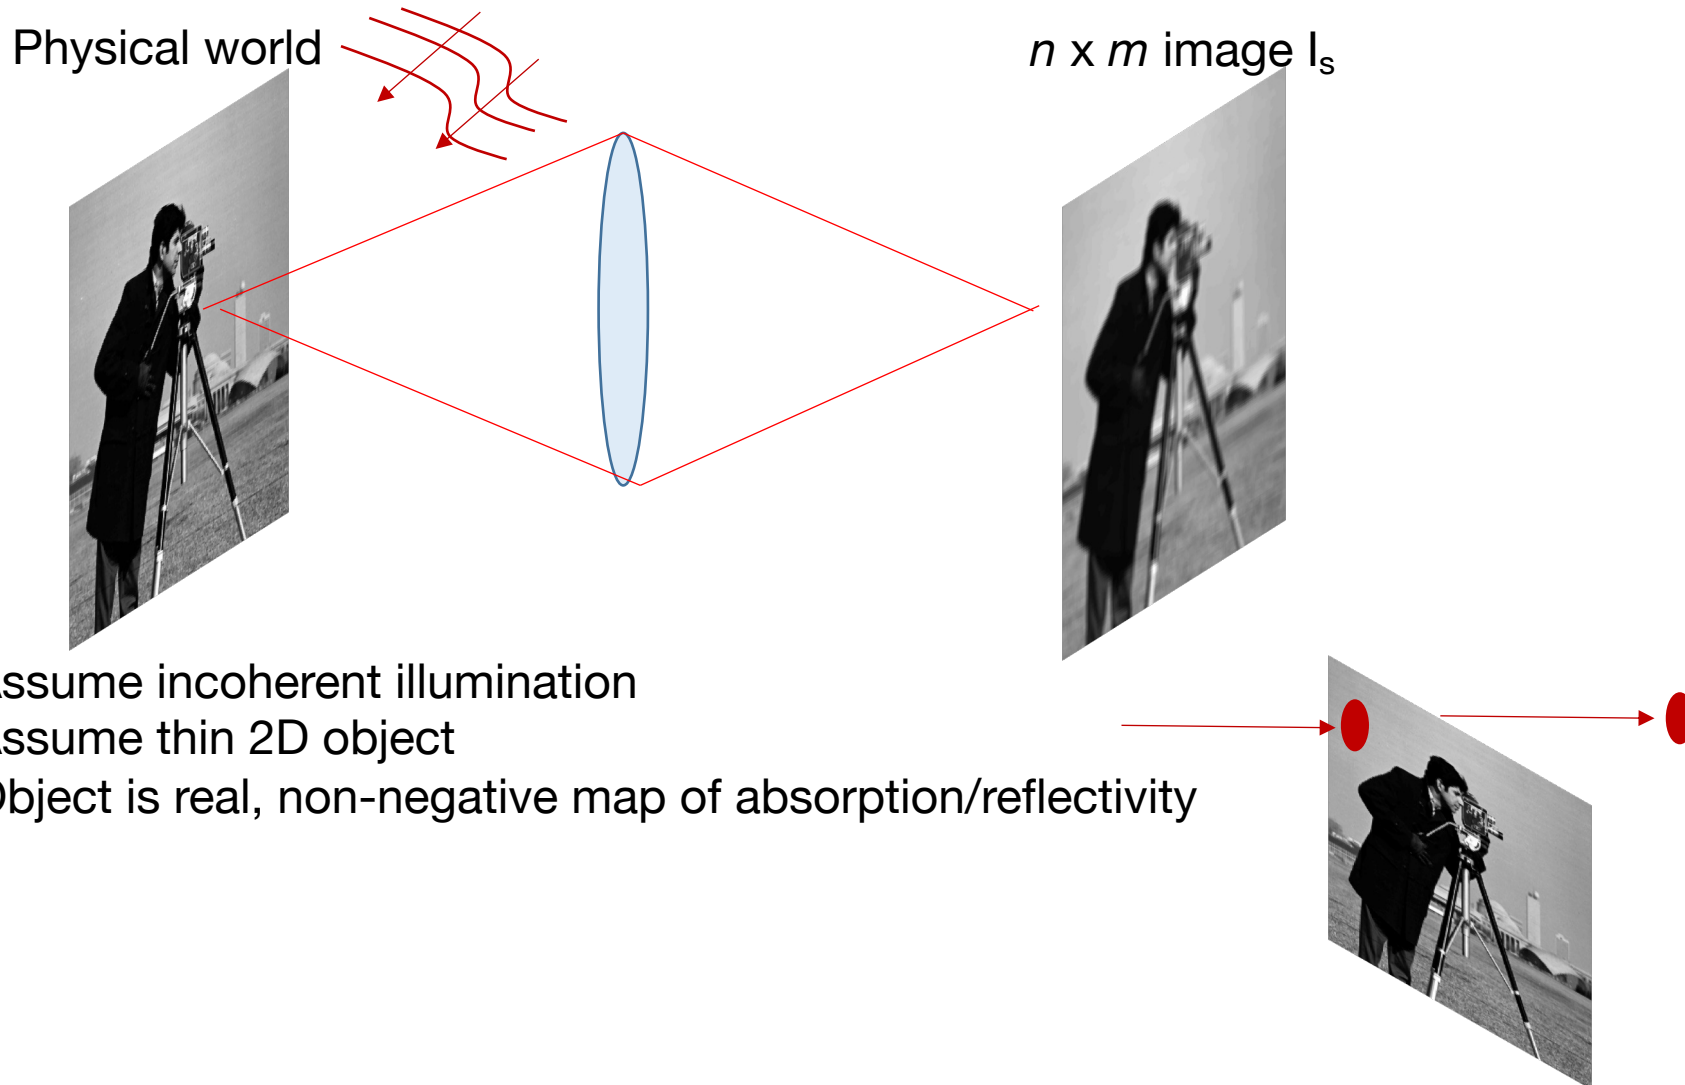
- Interpretation #3: Particle
- Model: Photons

Simple mathematical model of incoherent image formation



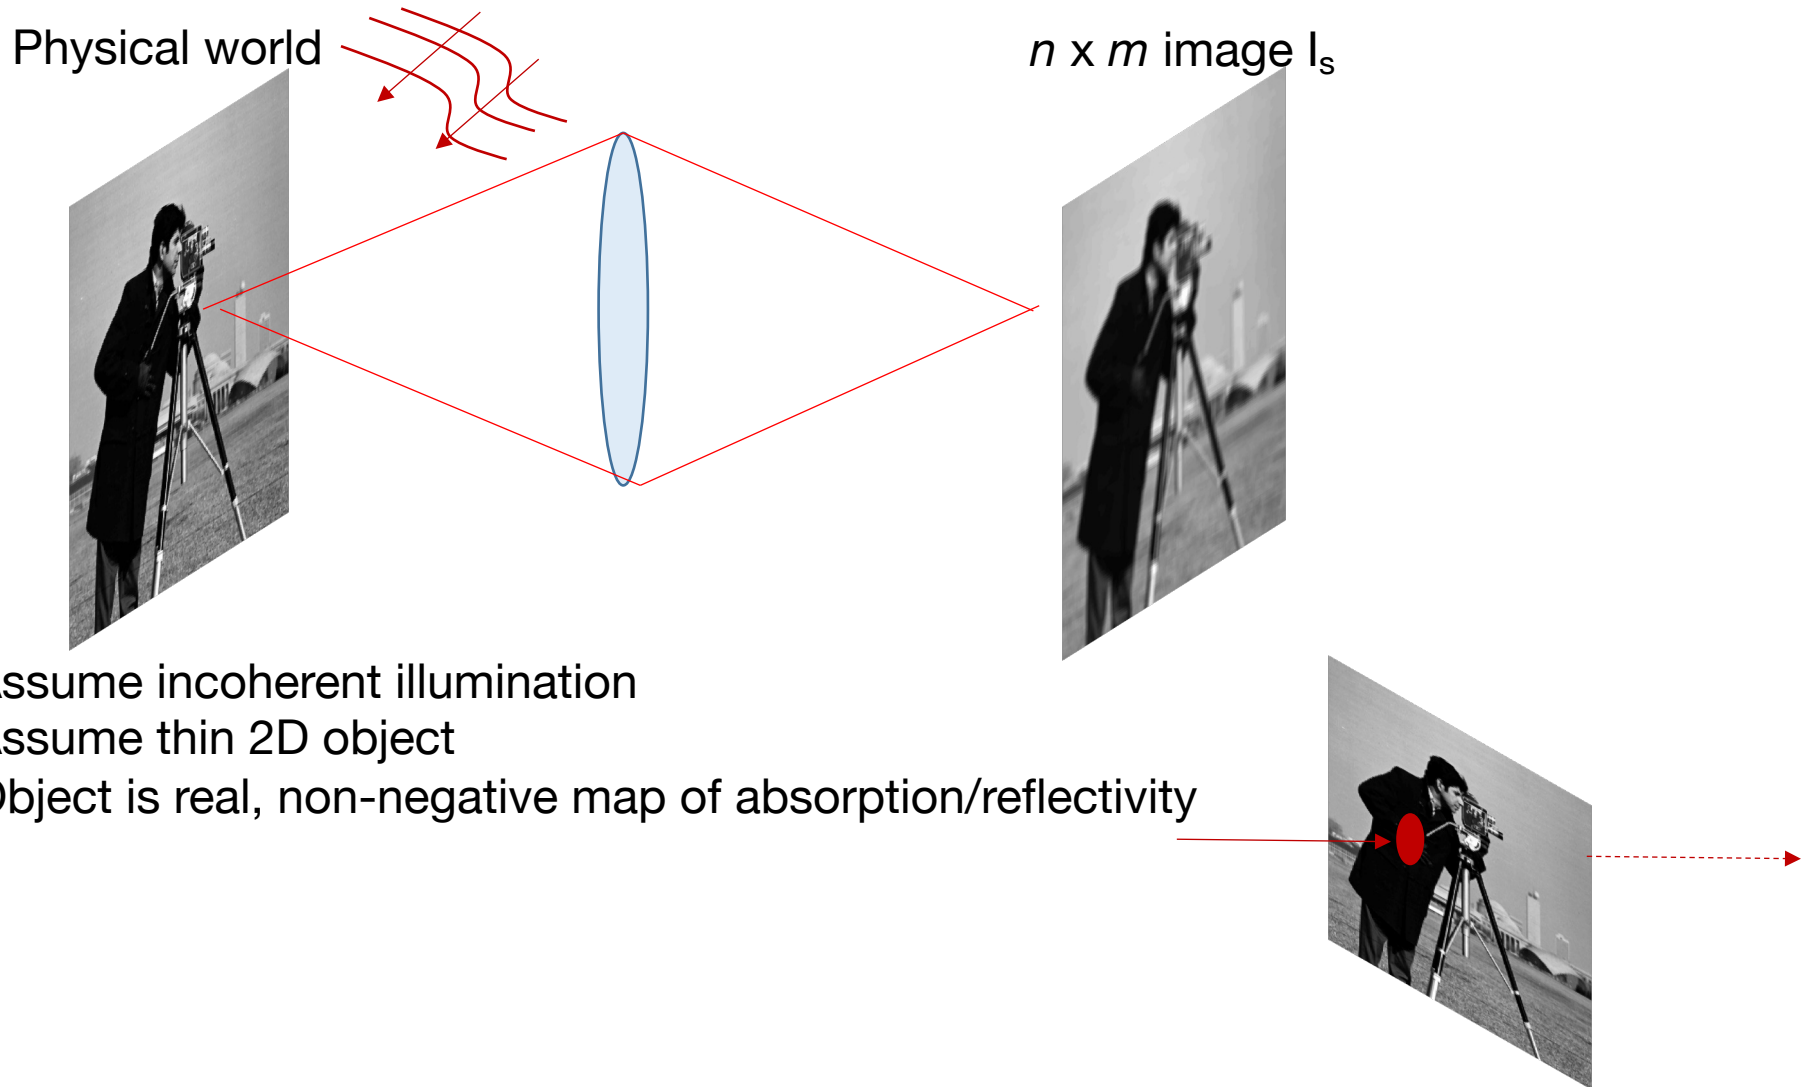
- Assume incoherent illumination
- Assume thin 2D object
- Object is real, non-negative map of absorption/reflectivity

Simple mathematical model of incoherent image formation



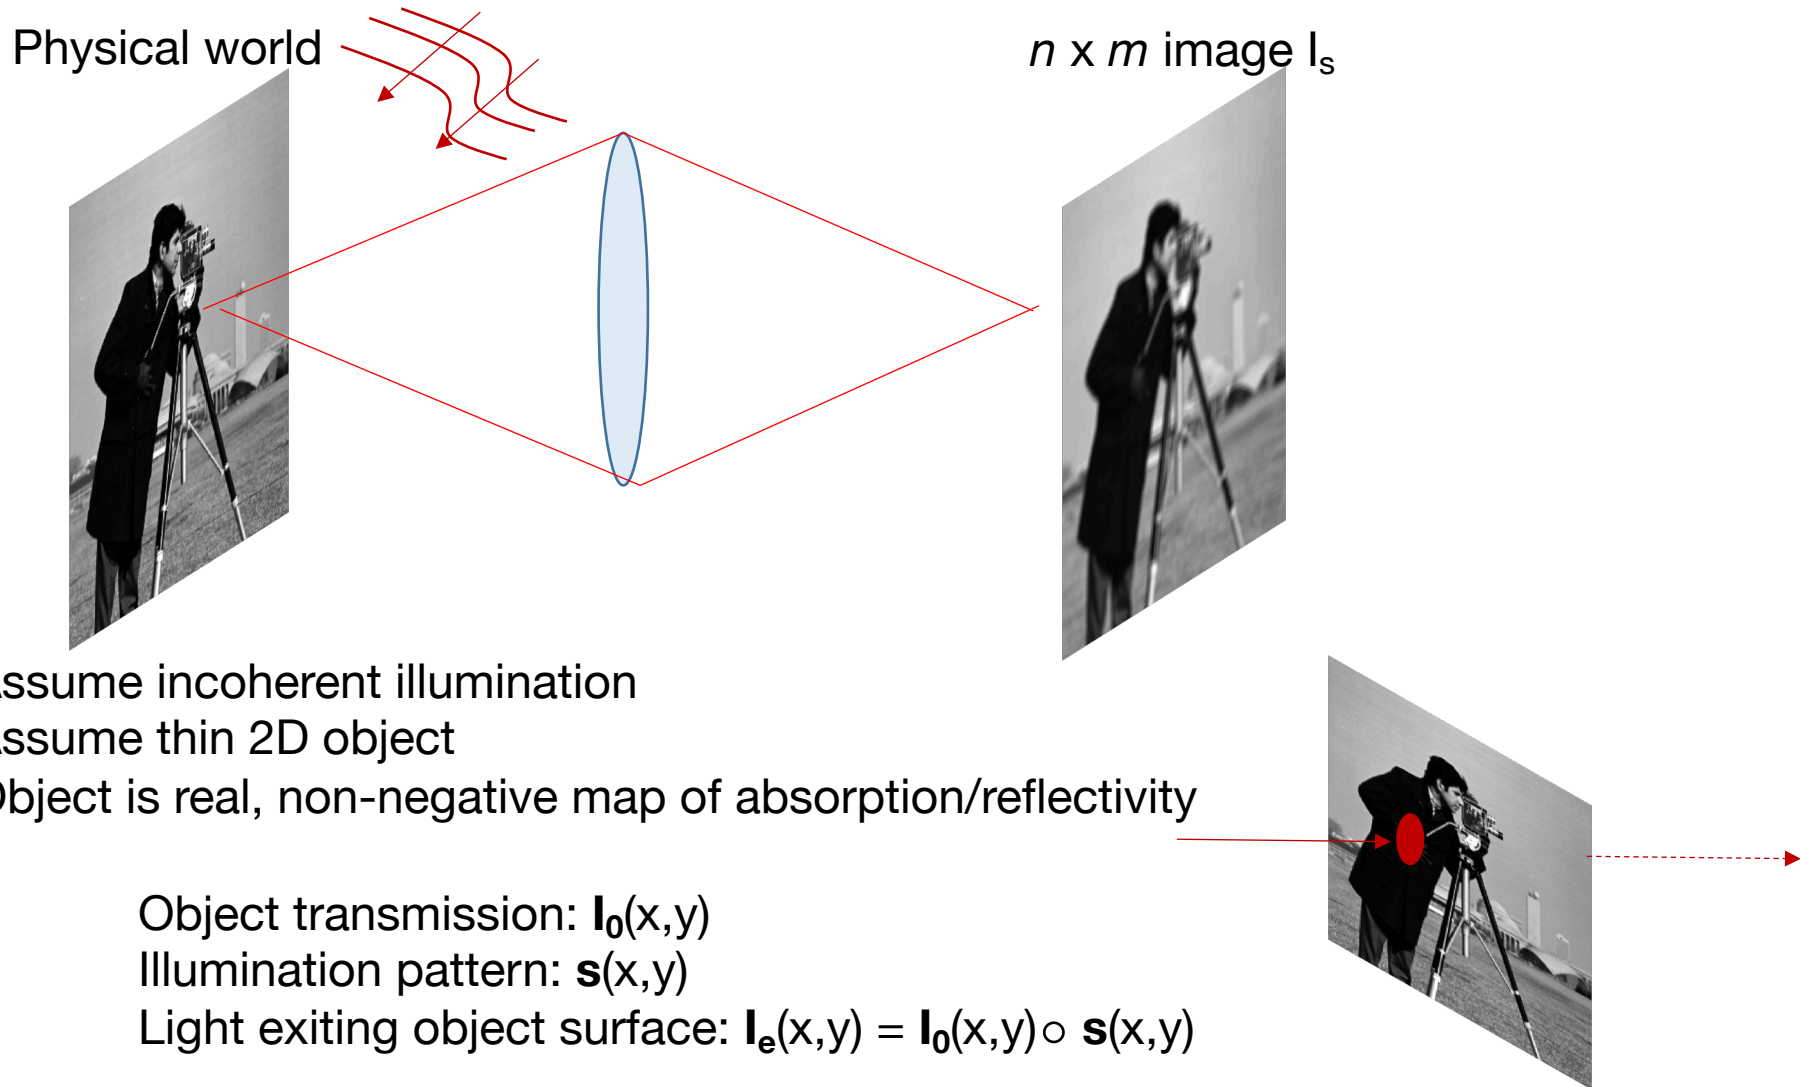
- Assume incoherent illumination
- Assume thin 2D object
- Object is real, non-negative map of absorption/reflectivity

Simple mathematical model of incoherent image formation




- Assume incoherent illumination
- Assume thin 2D object
- Object is real, non-negative map of absorption/reflectivity

Simple mathematical model of incoherent image formation



Simple mathematical model of incoherent image formation




- Assume incoherent illumination
- Assume thin 2D object
- Object is real, non-negative map of absorption/reflectivity


Object *reflectivity*: $I_0(x,y)$

Illumination pattern: $s(x,y)$

Light exiting object surface: $I_e(x,y) = I_0(x,y) \circ s(x,y)$



Simple mathematical model of incoherent image formation



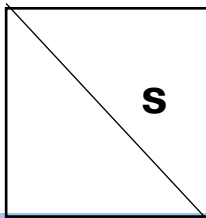
- Assume incoherent illumination
- Assume thin 2D object
- Object is real, non-negative map of absorption/reflectivity

Modeling incoherent illumination


$$I_e(x,y) = I_0(x,y) \circ s(x,y)$$

$$I_e = \mathbf{S} I_0$$

$\text{diag}(\mathbf{S}) = \mathbf{s}$



Example #1: Optimized illumination pattern (one color)



First, assume perfect camera:
intensity at image plane $I_p = I_e = S I_0$

Example #1: Optimized illumination pattern (one color)




Training data:

$[I_0(x, y), y]$

I_0 : 100 x 100

Label y : 1x2

Example #1: Optimized illumination pattern (one color)



Training data:

$[I_0(x, y), y]$

I_0 : 100 x 100

Label y : 1x2




$$I_e(x, y) = \mathbf{S} I_0(x, y)$$


Physical Layer



Example #1: Optimized illumination pattern (one color)



$[I_e(x, y), y]$



Training data:

$[I_0(x, y), y]$

I_0 : 100 x 100

Label y : 1x2




$$I_e(x, y) = S I_0(x, y)$$

Physical Layer




Task = $W_n \dots \text{ReLU}[W_1 \text{ReLU}[W_0 I_e] \dots]$

Example #1: Optimized illumination pattern (one color)



$[I_e(x, y), y]$



Training data:

$[I_0(x, y), y]$

$I_0: 100 \times 100$

Label $y: 1 \times 2$



$$I_e(x, y) = S I_0(x, y)$$

Physical Layer



$$\text{Task} = W_n \dots \text{ReLU}[W_1 \text{ReLU}[W_0 I_e] \dots]$$

Tensorflow
1.0 code

```

training_images = tf.placeholder(tf.float32, [image_size, None])
training_labels = tf.placeholder(tf.float32, [None, 3])
illumination_pattern = tf.truncated_normal([image_size, 1], stddev = 0.1)
illumination_matrix = tf.linalg.diag(illumination_pattern)
illumianted_images = tf.matmul(illumination_matrix, training_images)

#Now, train CNN of your choice using [illumianted_images, training_labels]
#Output of train/test will be CNN weights, classification performance AND illumination_pattern!

```


(Will show CoLab Notebook for implementation in Tensorflow 2.0)

Option 1: `tf.linalg.matmul`

1.0 code


```
training_images = tf.placeholder(tf.float32, [image_size, None])
training_labels = tf.placeholder(tf.float32, [None, 3])
illumination_pattern = tf.truncated_normal([image_size, 1], stddev = 0.1)
illumination_matrix = tf.linalg.diag(illumination_pattern)
illuminated_images = tf.matmul(illumination_matrix, training_images)

#Now, train CNN of your choice using [illuminated_images, training_labels]
#Output of train/test will be CNN weights, classification performance AND illumination_pattern!
```



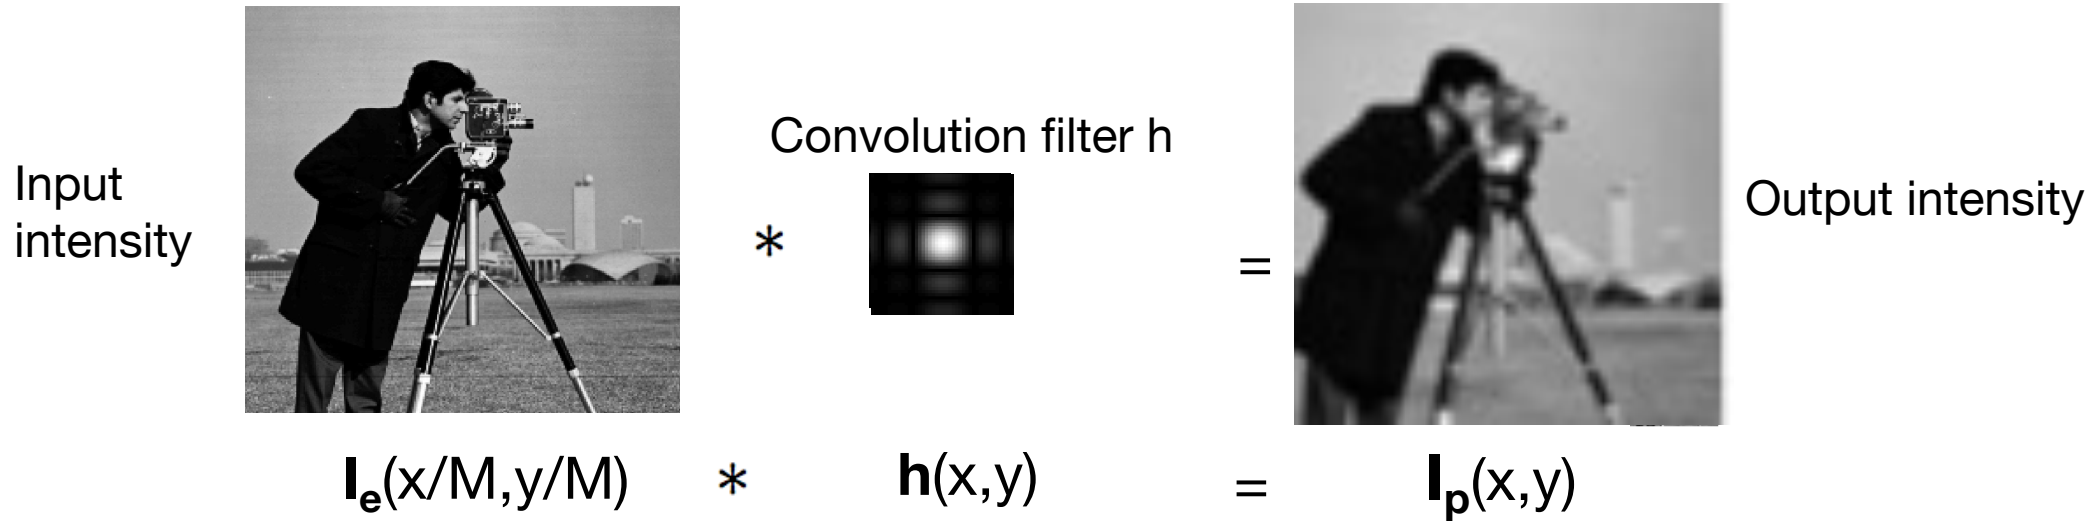
Option 2: `tf.linalg.multiply` (will show in CoLab Notebook)

Simple mathematical model of image formation



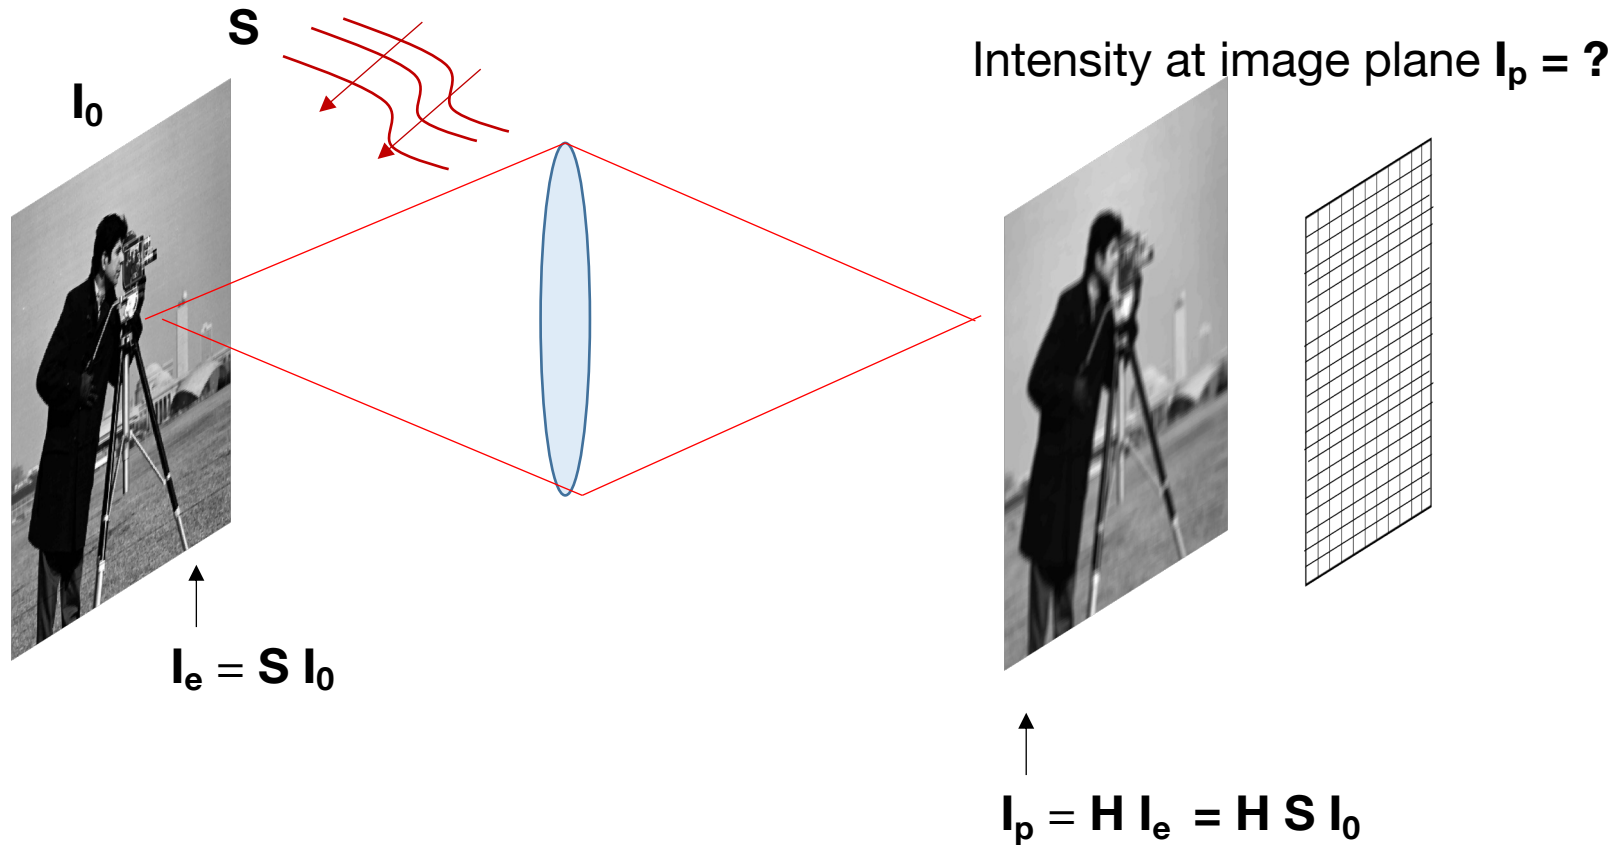
Simple mathematical model of image formation

Lenses blur and rescale images:
(We'll learn how exactly next few weeks)

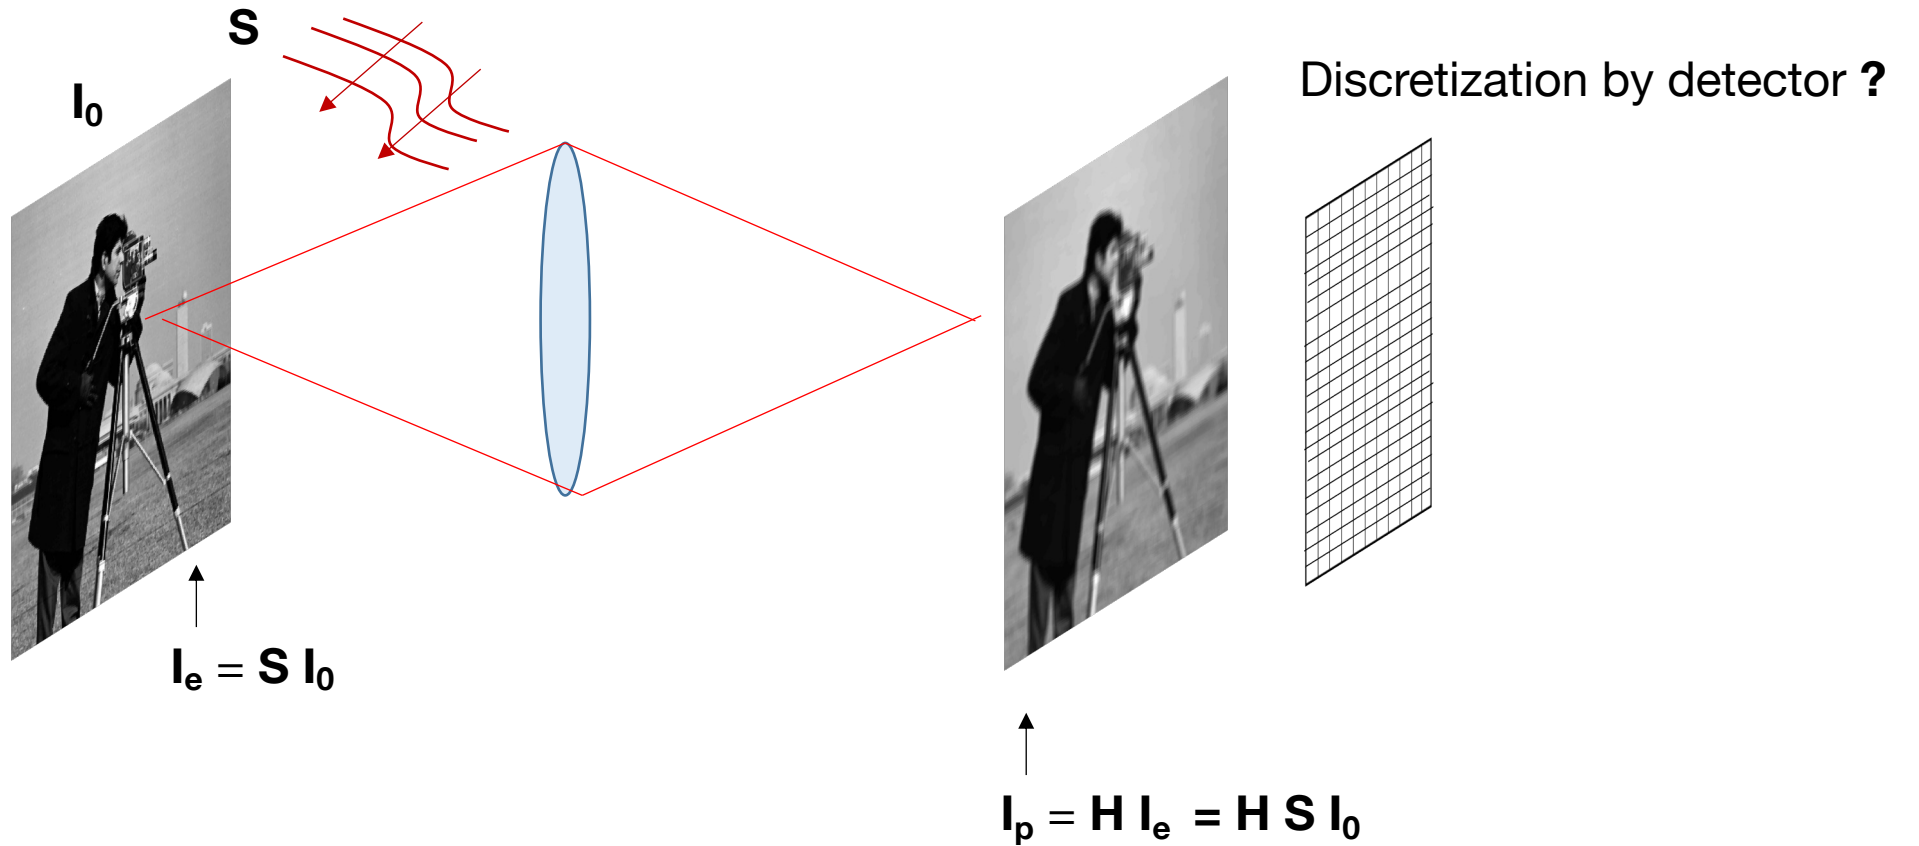


Assuming we've resized by M , $\mathbf{I}_p = \mathbf{I}_e * \mathbf{h} = \mathbf{H} \mathbf{I}_e$

Simple mathematical model of image formation




Simple mathematical model of image formation



Use downsampling matrix
(sum-pooling)

$$D = \begin{bmatrix} 0.5 & 0.5 & 0 & 0 & 0 & \dots \\ & 0.5 & 0.5 & 0 & 0 & 0 & \dots \\ & & & 0.5 & 0.5 & 0 & 0 & 0 & \dots \\ \vdots & & & & & & & & \vdots \end{bmatrix}$$

Simple mathematical model of image formation



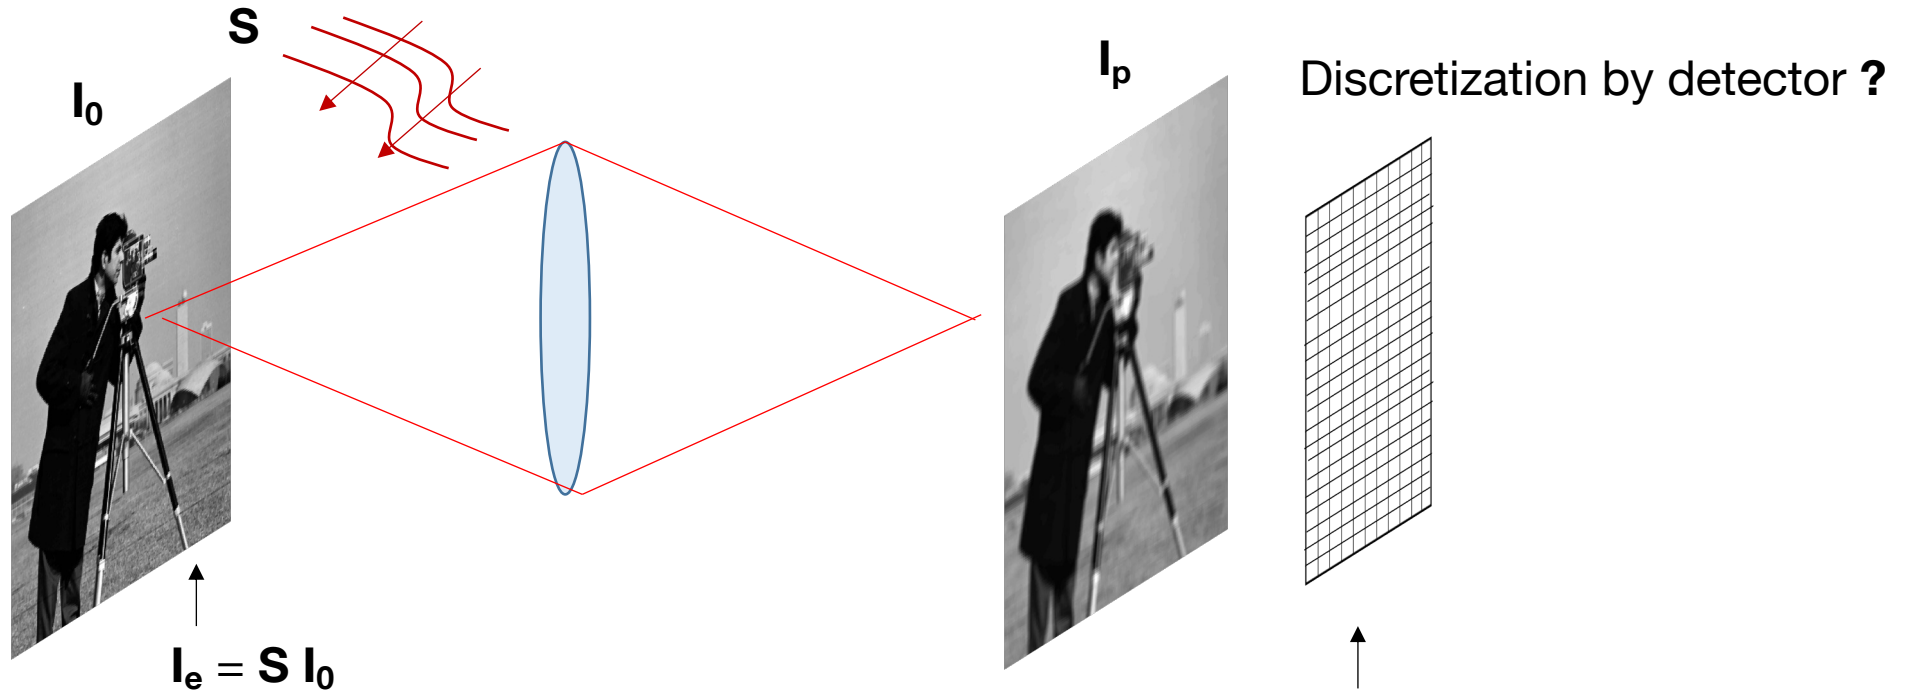
$$I_e = S I_0$$

$$I_s = D I_p = D H S I_0$$

Use downsampling matrix (sum-pooling)

$$D = \begin{bmatrix} 0.5 & 0.5 & 0 & 0 & \dots \\ & 0.5 & 0.5 & 0 & 0 & \dots \\ & & & 0.5 & 0.5 & 0 & 0 & \dots \\ \vdots & & & & & \vdots & & \vdots \end{bmatrix}$$

Simple mathematical model of image formation

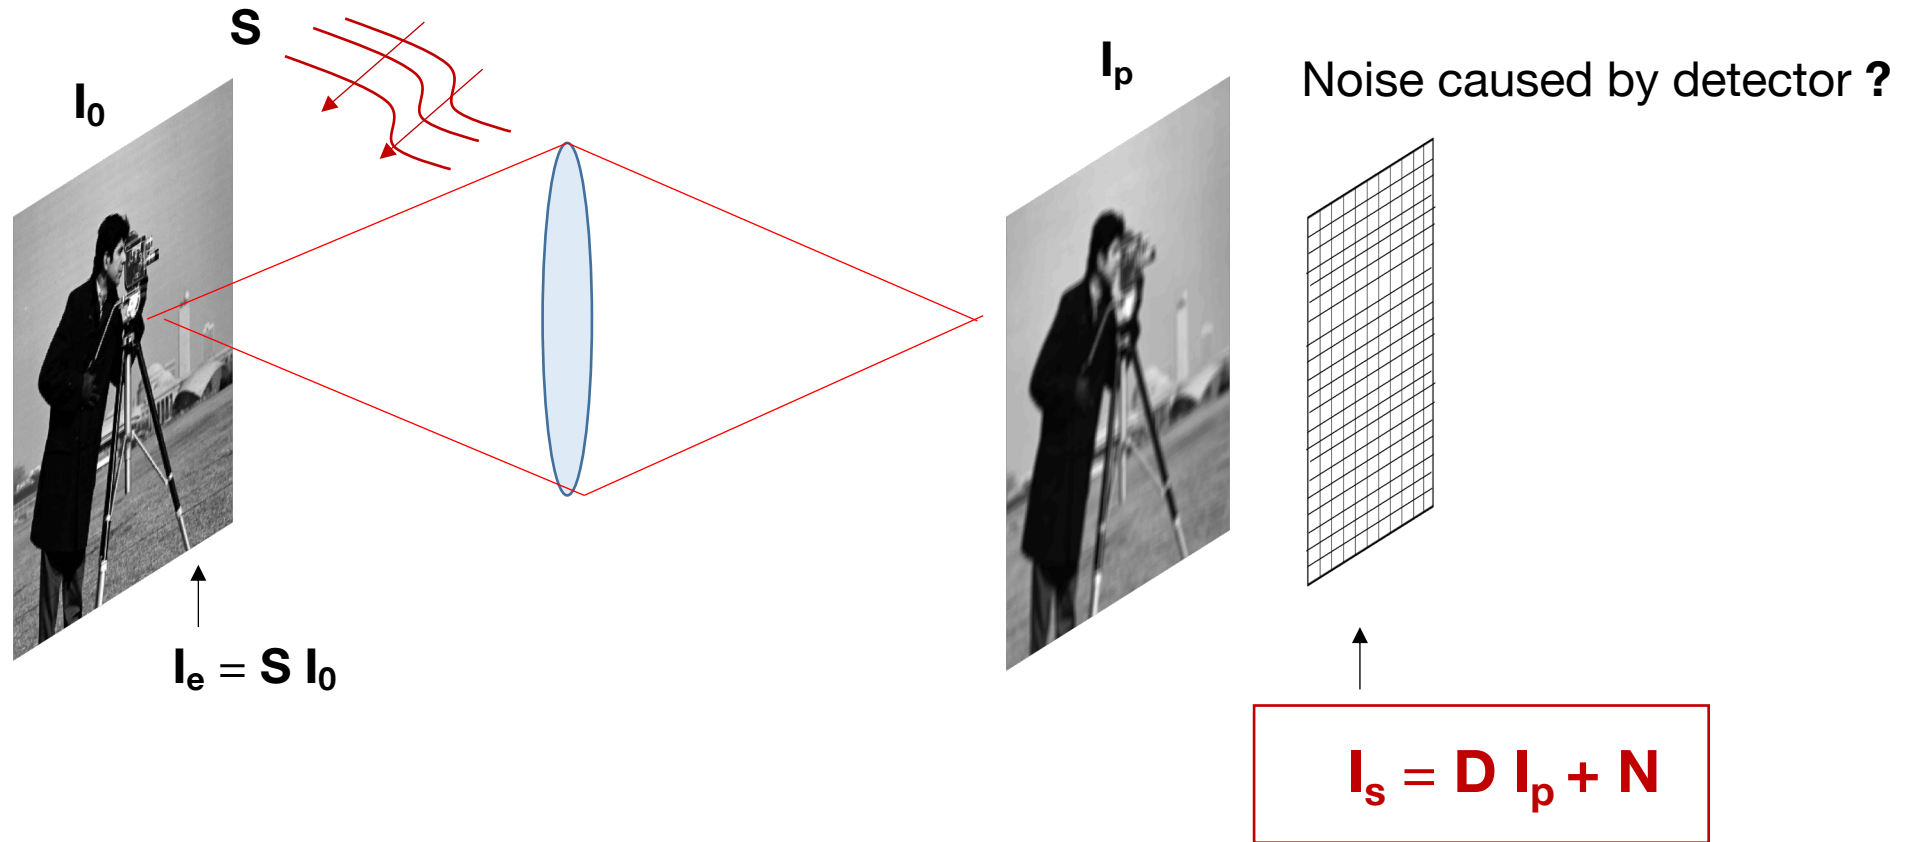


$$I_s = D I_p = D H S I_0$$

$$D = \begin{matrix} w_1 & w_2 & 0 & 0 & 0 & \dots \\ & w_3 & w_4 & 0 & 0 & 0 & \dots \\ & & & w_5 & w_6 & 0 & 0 & 0 & \dots \\ & & & \vdots & & & & \vdots & \end{matrix}$$

*Can make these learnable weights

Simple mathematical model of image formation




Can also add in detector-dependent noise $N = k * \text{np.random.randn}(dx, dy)$

(zero-mean Gaussian noise, for example)

Pause to take a look at:

`physical_layers_example.ipynb`



Task

- CAR
- TRUCK
- VAN
- ...
- BICYCLE

Physical Layers

Digital Layers



Physical world

$$I_0(x_0, y_0)$$



Digital Image

$$I_s(x, y)$$



CONVOLUTION + RELU POOLING CONVOLUTION + RELU POOLING FLATTEN FULLY CONNECTED SOFTMAX

FEATURE LEARNING CLASSIFICATION


Digital layers

Physical layers

$$\text{Task} = \mathbf{W}_n \dots \text{ReLU}[\mathbf{W}_1 \text{ReLU}[\mathbf{W}_0 f[\mathbf{I}_0]] \dots]$$

$$\text{Task} = \mathbf{W}_n \dots \text{ReLU}[\mathbf{W}_1 \text{ReLU}[\mathbf{W}_0 \mathbf{D H S I}_0] \dots]$$

Example #2: Optimized color filter for a grayscale camera





“Ground truth” object:

$$I_0(x, y, \lambda)$$

100 x 100 pix. x 30 spectral channels

Example #2: Optimized color filter for a grayscale camera





Monochromatic
camera sensor

“Ground truth” object:

$$I_0(x, y, \lambda)$$

100 x 100 pix. x 30 spectral channels

Example #2: Optimized color filter for a grayscale camera

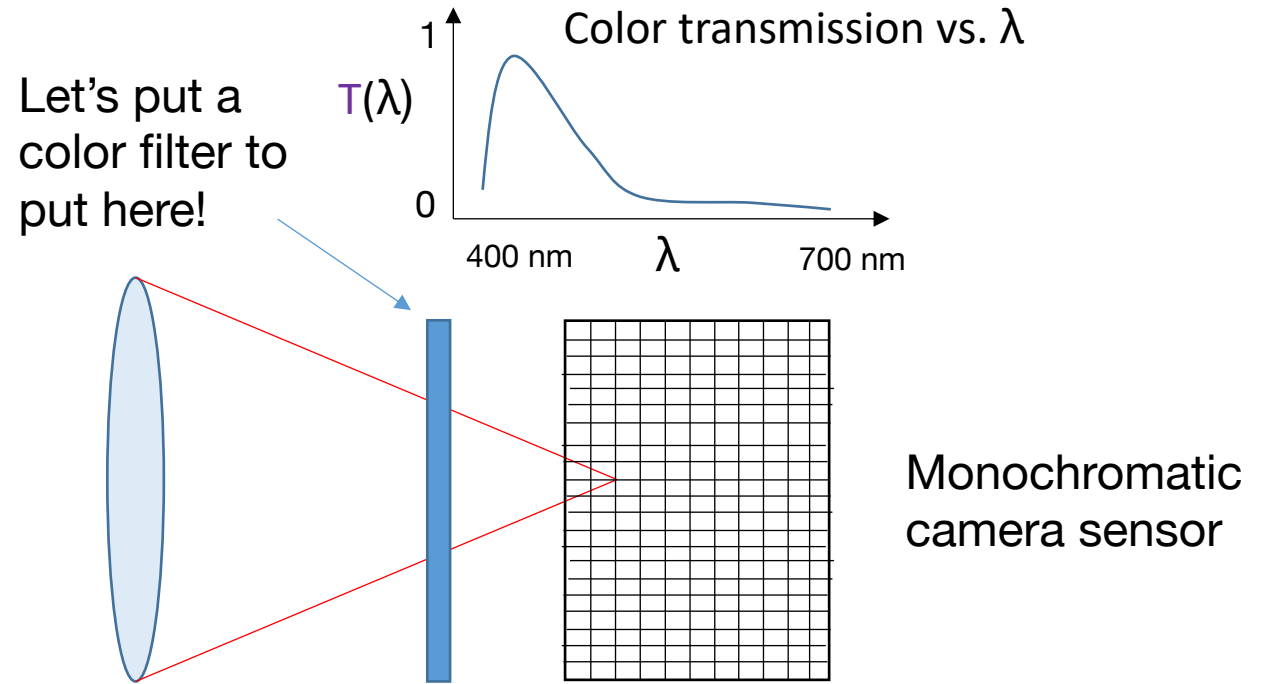



Monochromatic camera sensor

“Ground truth” object:
 $I_o(x, y, \lambda)$
100 x 100 pix. x 30 spectral channels

$$I_s(x, y) = \sum_{\lambda} I_o(x, y, \lambda)$$

Example #2: Optimized color filter for a grayscale camera

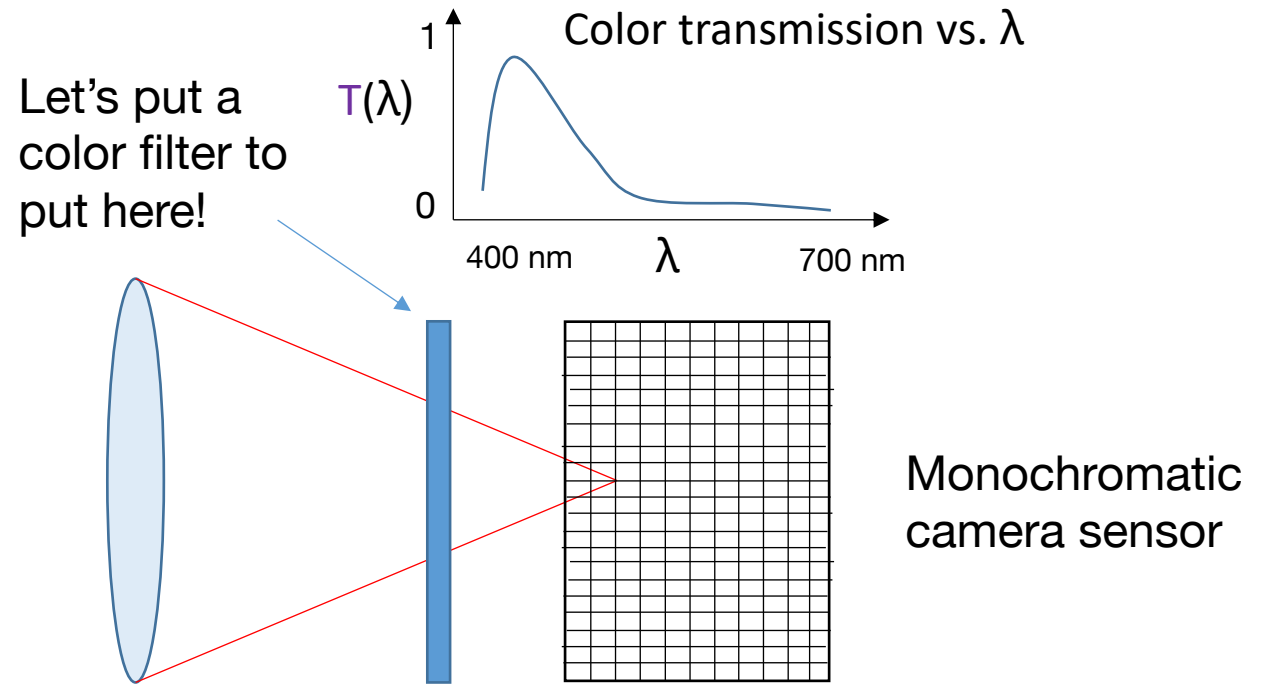



“Ground truth” object:

$$I_0(x, y, \lambda)$$

100 x 100 pix. x 30 spectral channels

Example #2: Optimized color filter for a grayscale camera




“Ground truth” object:

$$I_0(x, y, \lambda)$$



100 x 100 pix. x 30 spectral channels

$$I_s(x, y) = \sum_{\lambda} T(\lambda) I_0(x, y, \lambda)$$

Example #2: Optimized color filter for a grayscale camera



Design optimal color filter for classification:



Training data:

$$[I_0(x, y, \lambda), y]$$


I_0 : 100 x 100 pix. x 30

Label y : 1x3 - pepper, broccoli, green beans

$$I_s(x, y) = \sum_{\lambda} W_0(\lambda) I_0(x, y, \lambda)$$

Physical Layer

Example #2: Optimized color filter for a grayscale camera




Training data:

$[I_0(x, y, \lambda), y]$

I_0 : 100 x 100 x 30

Label y : 1x3

Example #2: Optimized color filter for a grayscale camera



Training data:

$[I_0(x, y, \lambda), y]$

$I_0: 100 \times 100 \times 30$


Label $y: 1 \times 3$



$$I_s(x, y) = \sum_{\lambda} W_0(\lambda) I_0(x, y, \lambda)$$

Physical Layer

Example #2: Optimized color filter for a grayscale camera



Training data:

$$[I_0(x, y, \lambda), y]$$

I_0 : 100 x 100 x 30


Label y : 1x3

$$I_s(x, y) = \sum_{\lambda} W_0(\lambda) I_0(x, y, \lambda)$$

Physical Layer

$$\text{Task} = \mathbf{W}_n \dots \text{ReLU}[\mathbf{W}_1 \text{ReLU}[\mathbf{W}_0 I_s] \dots]$$

Example #2: Optimized color filter for a grayscale camera



Training data:

$I_0(x, y, \lambda), y$
 $I_0: 100 \times 100 \times 30$
 Label $y: 1 \times 3$

$$I_s(x, y) = \sum_{\lambda} W_0(\lambda) I_0(x, y, \lambda)$$

Physical Layer

Task = $W_n \dots \text{ReLU}[W_1 \text{ReLU}[W_0 I_s] \dots]$

```

multispectral_data = tf.placeholder(tf.float32, [None, num_colors, image_size])
veg_labels = tf.placeholder(tf.float32, [None, 3])
filter_weights = tf.truncated_normal([num_colors, 1], stddev = 0.1)
filtered_images = tf.einsum('aij,jk->aik', multispectral_data, filter_weights)

#Now, train CNN of your choice using [filtered_images, veg_labels]
#Output of training/testing will be CNN weights, classification performance AND filter_weights!
  
```

Example implementation with Tensorflow 1.0 code

Tensorflow: operations to sum along 3rd (or higher) dimension

Option 1: Einsum (shown as Tensorflow 1.0 code, and also applicable in Tensorflow 2.0)

```
filtered_images = tf.einsum('aij,jk->aik', multispectral_data, filter_weights)
```

Option 2: tf.reduce_sum

```
filtered_images = tf.reduce_sum(multispectral_data * filter_weights, axis=2)
```

Option 3: Locally connected conv2D with a 1x1 filter size


<https://github.com/keras-team/keras/blob/master/keras/layers/local.py#L183>

Pause to take a look at:

`weighted_image_sum_example.ipynb`

Example 3: learned illumination pattern for improved segmentation

U-Net Architecture




U-Net: Convolutional Networks for Biomedical Image Segmentation

Olaf Ronneberger, Philipp Fischer, and Thomas Brox

Computer Science Department and BIOS Centre for Biological Signalling Studies,
 University of Freiburg, Germany
 ronneber@informatik.uni-freiburg.de,
 WWW home page: <http://lmb.informatik.uni-freiburg.de/>


Example 3: learned illumination pattern for improved segmentation

Variably illuminated images from different LEDs




$$w_1 I_1 + w_2 I_2 + w_3 I_3 + \dots$$

U-Net CNN

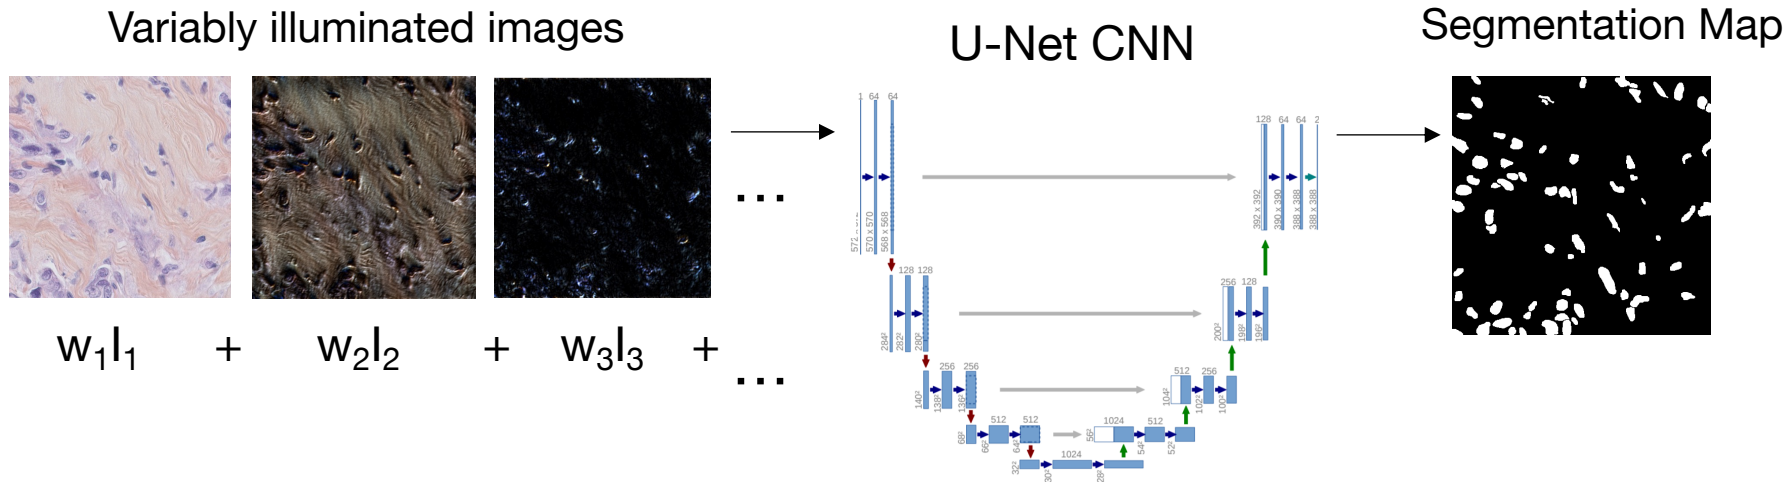


Segmentation Map




*If we allow w 's here to be trainable weights, then we can find ideal brightnesses for different LEDs to illuminate a sample of interest!


Example 3: learned illumination pattern for improved segmentation




Optimized illumination for nuclei segmentation



Standard illumination



Learned illumination




+5-10% accuracy


See C. Cooke et al., "Physics-enhanced machine learning for virtual fluorescence microscopy," ICCV (2021)

Image segmentation –current workflow


Capture: BF images




Capture: Fluorescence



Threshold




Segmentation Mask




(e.g., DAPI-stained nuclei)

Image segmentation –current workflow


Capture: BF images




Capture: Fluorescence



Threshold




Segmentation Mask




(e.g., DAPI-stained nuclei)



Learned illumination



Inference via a trained U-Net



Optimally illuminated



in silico labeling: fluorescence image inference from bright-field data

In Silico Labeling: Predicting Fluorescent Labels in Unlabeled Images

Eric M. Christiansen,^{1,11,*} Samuel J. Yang,¹ D. Michael Ando,^{1,9} Ashkan Javaherian,^{2,9} Gaia Skibinski,^{2,9} Scott Lipnick,^{3,4,8,9} Elliot Mount,^{2,10} Alison O’Neil,^{3,10} Kevan Shah,^{2,10} Alicia K. Lee,^{2,10} Piyush Goyal,^{2,10} William Fedus,^{1,6,10} Ryan Poplin,^{1,10} Andre Esteva,^{1,7} Marc Berndl,¹ Lee L. Rubin,³ Philip Nelson,^{1,*} and Steven Finkbeiner^{2,5,*}

¹Google, Inc., Mountain View, CA 94043, USA

²Taube/Koret Center for Neurodegenerative Disease Research and DaedalusBio, Gladstone Institutes, San Francisco, CA 94158, USA

³Department of Stem Cell and Regenerative Biology, Harvard University, Cambridge, MA 02138, USA


⁴Department of Biomedical Informatics, Harvard Medical School, Boston, MA 02115, USA

⁵Departments of Neurology and Physiology, University of California, San Francisco, 94158, USA

⁶Montreal Institute of Learning Algorithms, University of Montreal, Montreal, QC, Canada

⁷Department of Electrical Engineering, Stanford University, Stanford, CA 94305, USA

⁸Center for Assessment Technology and Continuous Health, Massachusetts General Hospital, Boston, MA 02114, USA




**BF Focal stack
(26+ images)**


Task: bright-field to fluorescence image Inference

“Virtual Staining” – can be used to covert one image type to many others


B1 Autofluorescence to H&E




B2 Phase contrast to multiple FAS (displayed as overlay)




B3 Wide field microscopy to multiple FAS (displayed as overlay)




B4 Transmission microscopy to genetically encoded fluorescence (mCherry)



B5 H&E to IHC : cytokeratin stain



B6 IHC to other IHC stains: Haematoxylin (displayed) and others




L. Kreiss et al., “Digital staining in optical microscopy using deep learning--a review,” Photonix (2024)


B. Bai et al., “Deep learning-enabled virtual histological staining of biological samples,” Nature Light Sci (2023)

Image segmentation versus *in silico* labeling (fluorescence inference)

Capture: Bright field



Capture: Fluorescence




Threshold




Segmentation Mask



Learned illumination for fluorescent image inference?




Directly infer the fluorescence image



Optimally illuminated


Multiple Patterns for Fluorescence image inference




1x Pattern Optimization



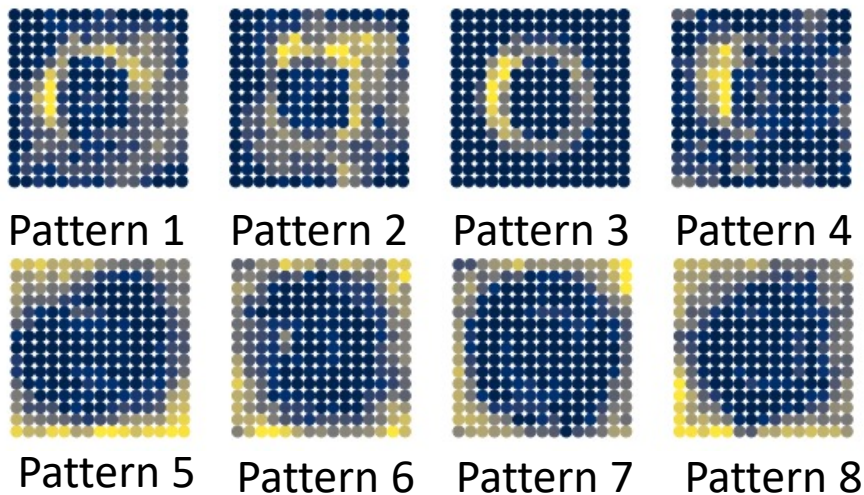
2x Pattern Optimization




4x Pattern Optimization




8x Pattern Optimization




Ground Truth




1 Pattern



4 Patterns



HeLa Task Performance



See C. Cooke et al., "Physics-enhanced machine learning for virtual fluorescence microscopy," ICCV (2021)

Summary: simple physical layers for incoherent imaging

- Deal with sample/image intensities I , real and non-negative

- Effect of illumination is element-wise multiplication λ $I_e(x,y) = \mathbf{S} I_o(x,y)$

- Imaging systems blur the object via point-spread function matrix \mathbf{H}

$$I_b(x,y) = \mathbf{H} I_o(x,y)$$

- Discrete pixels down-sample the object via



$$I_d(x,y) = \mathbf{D} I_o(x,y)$$

- Add noise into measurement $I_N(x,y) = \mathbf{D} I_o(x,y) + \mathbf{N}$

- Different colors add linearly $I_s(x,y) = \sum I_o(x,y,\lambda)$

What is light and how can we model it?



- Interpretation #1: Radiation (*Incoherent*)
- Model: Rays



- Real, non-negative
- Models absorption and brightness

$$I_{\text{tot}} = I_1 + I_2$$

- Interpretation #2: Electromagnetic wave (*Coherent*)
- Model: Waves



- Complex field
- Models Interference

$$E_{\text{tot}} = E_1 + E_2$$

Alternative framework: Modeling coherent radiation as a wave