

Lecture 15: Introduction to Physical Layers in Machine Learning

Machine Learning and Imaging

BME 548L

Roarke Horstmeyer

Other Computer Vision Tasks - We'll pick this back up next class

Semantic Segmentation

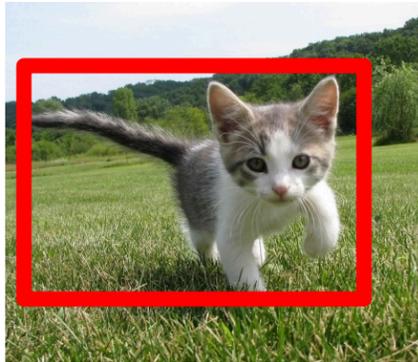


GRASS, CAT,
TREE, SKY

No objects, just pixels

Balanced equation

Classification + Localization

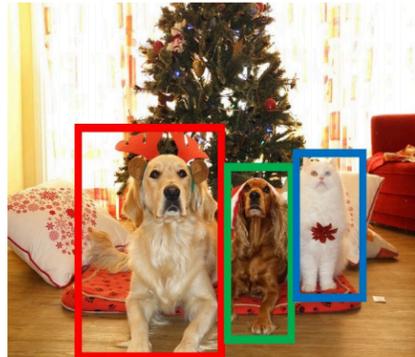


CAT

Single Object

Over-determined

Object Detection



DOG, DOG, CAT

Multiple Object

Over-determined

Instance Segmentation



DOG, DOG, CAT

Over-determined

Super-resolution

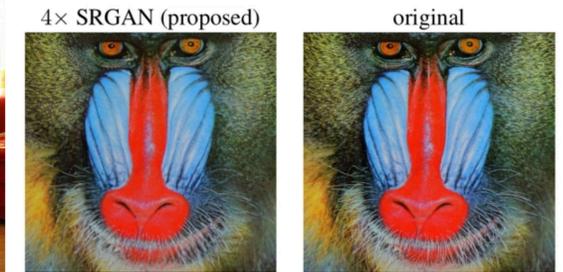
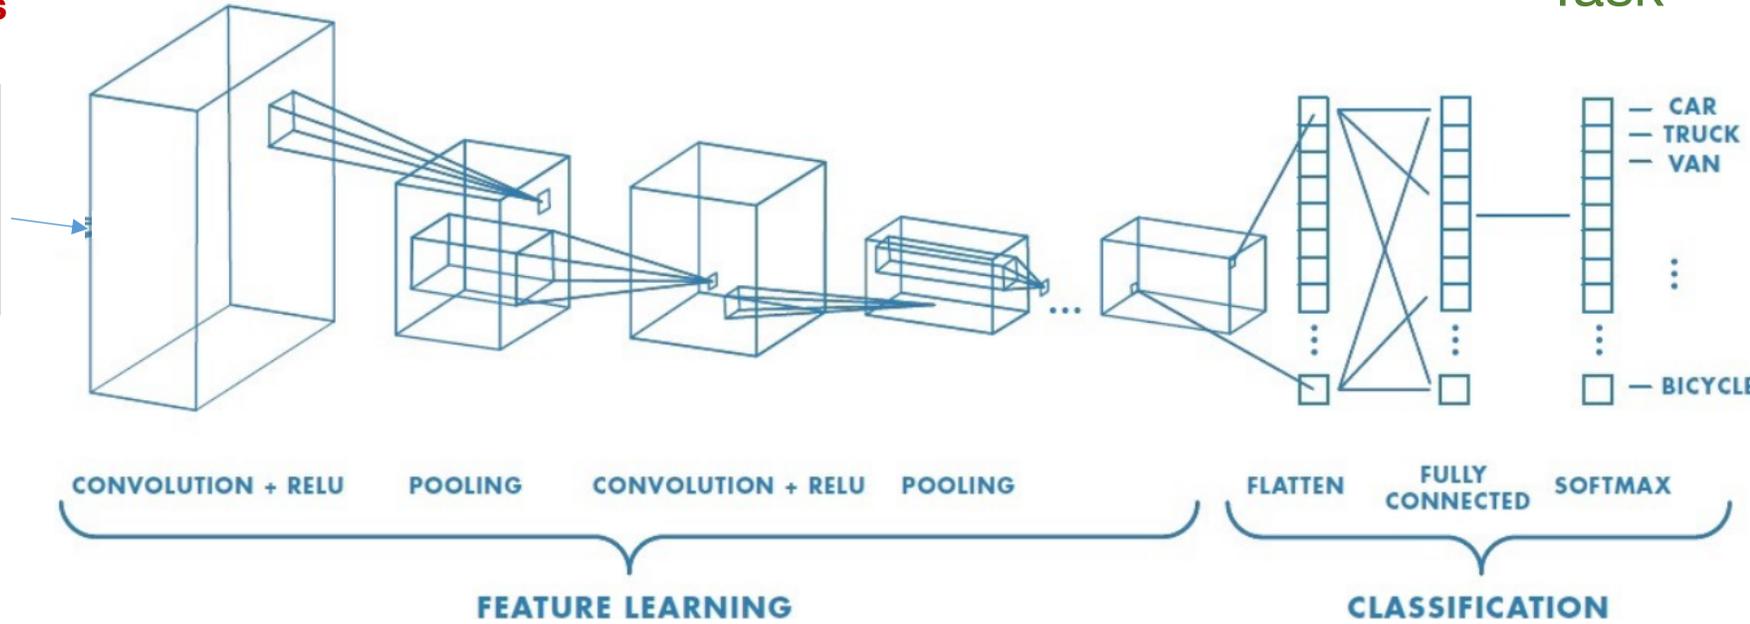
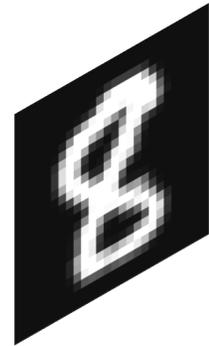


Figure 1: Super-resolved image (left) is almost indistinguishable from original (right). [4x upscaling]

Under-determined

Bringing together physical and digital image representations

Image I_s

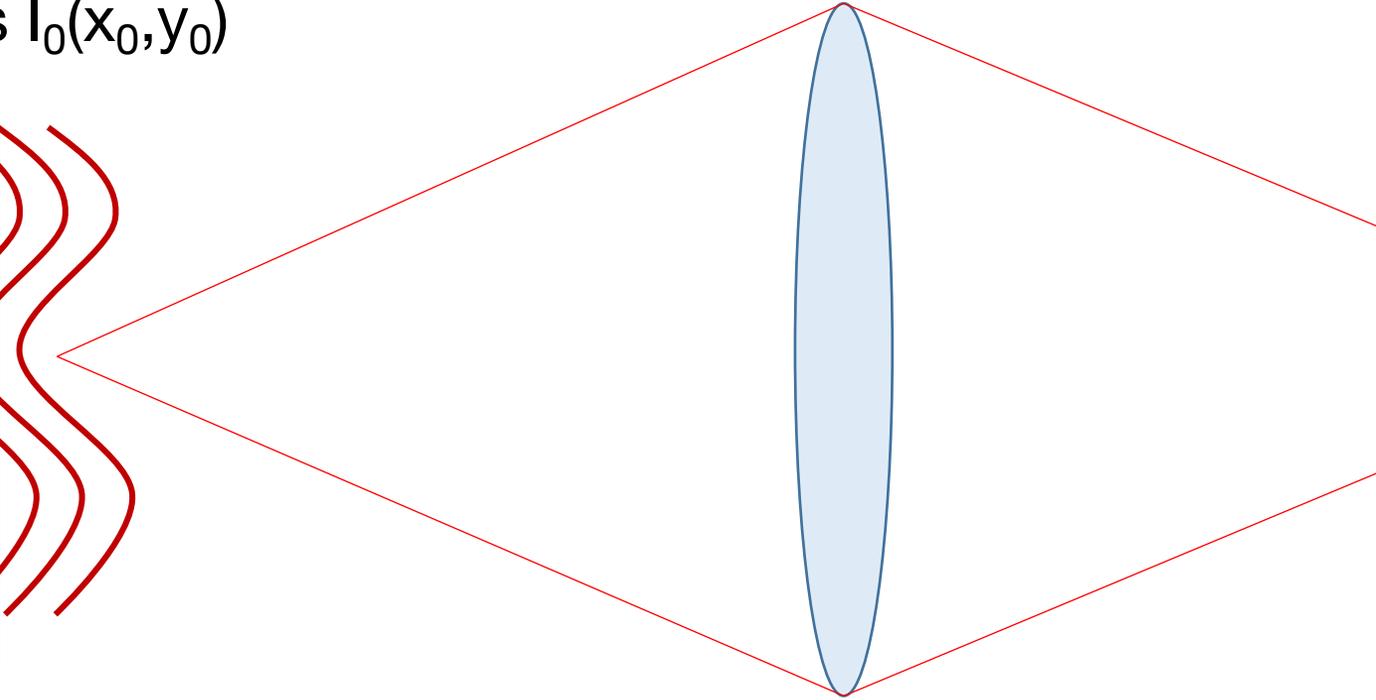
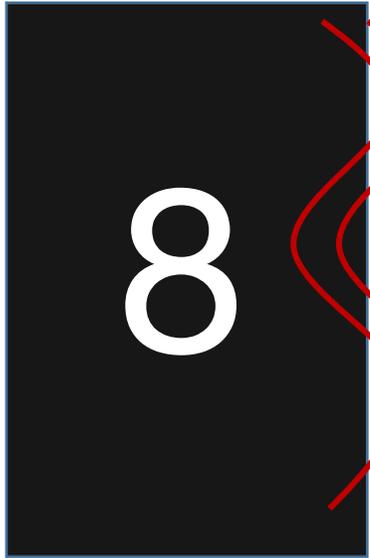


Task

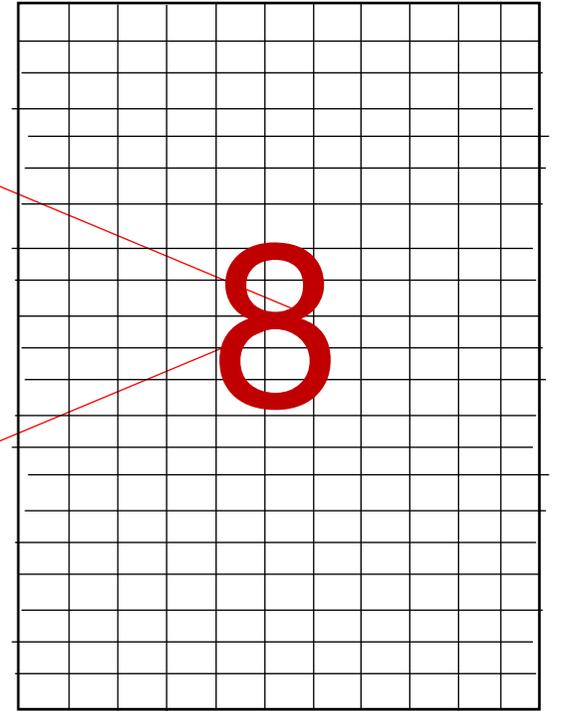
$$\text{Task} = \mathbf{W}_n \dots \text{ReLU}[\mathbf{W}_1 \text{ReLU}[\mathbf{W}_0 I_s] \dots]$$

Simple model of image formation

Continuous $I_0(x_0, y_0)$



$n \times m$ pixel array



Discrete signal

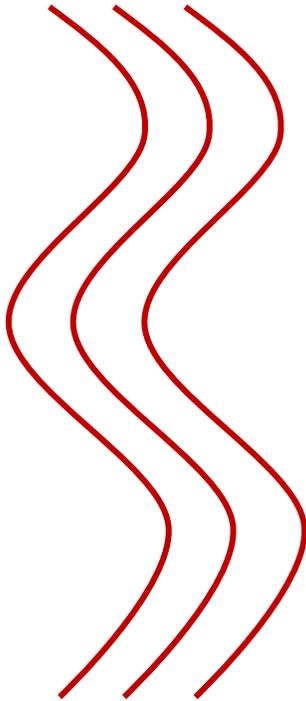
$$I_s(x, y), (x, y) \in \mathbb{Z}^{n \times m}$$

Digitization

Photons to electrons

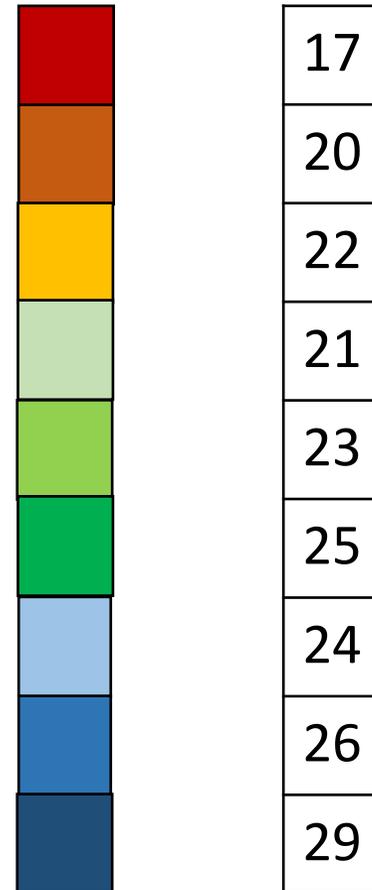
What does the Sampling Theorem mean for us?

Continuous functions



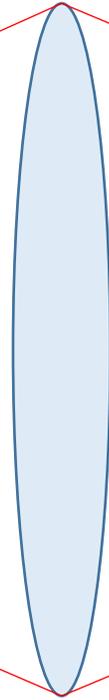
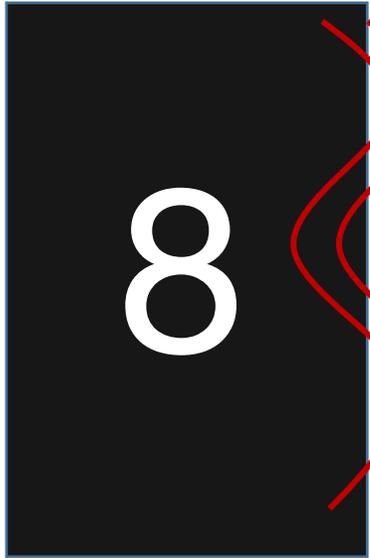
*conditions

Discretize vectors
(and matrices)

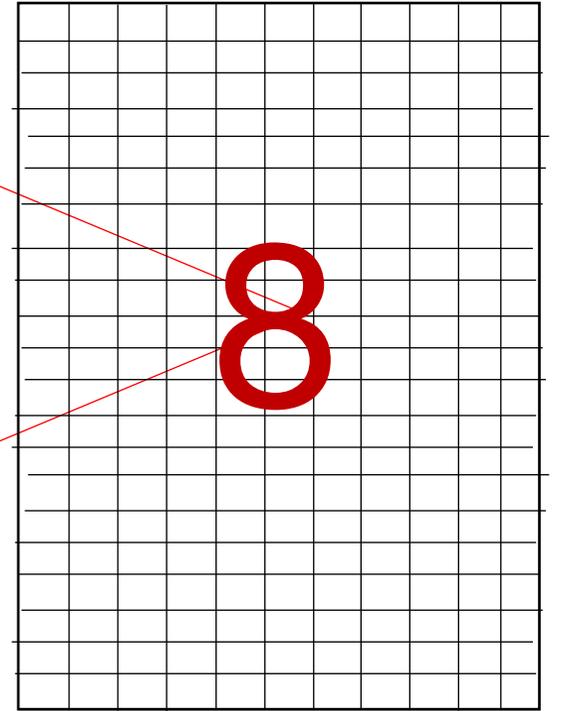


Simple model of image formation

Discrete $I_0(x_0, y_0)$



$n \times m$ pixel array

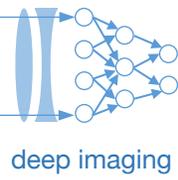


Discrete signal

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

Digitization

Photons to electrons



Bringing together physical and digital image representations

Physical Layers

Digital Layers

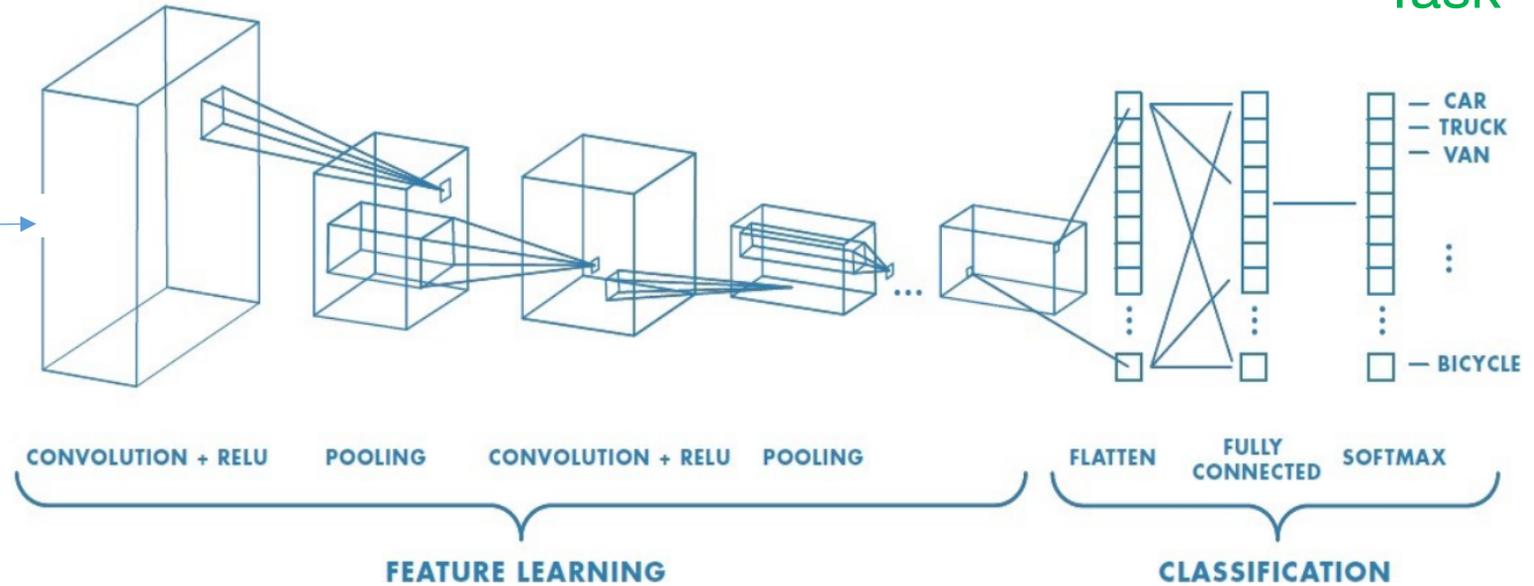
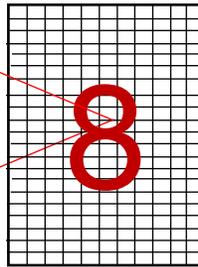
Task

Physical world

Digital Image

$$I_0(x_0, y_0)$$

$$I_s(x, y)$$

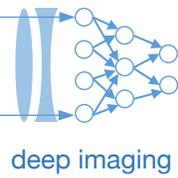


$$I_s = f[I_0]$$

Digital layers

Physical layers

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



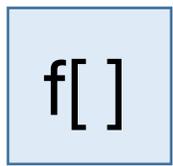
Bringing together physical and digital image representations

Physical Layers

Digital Layers

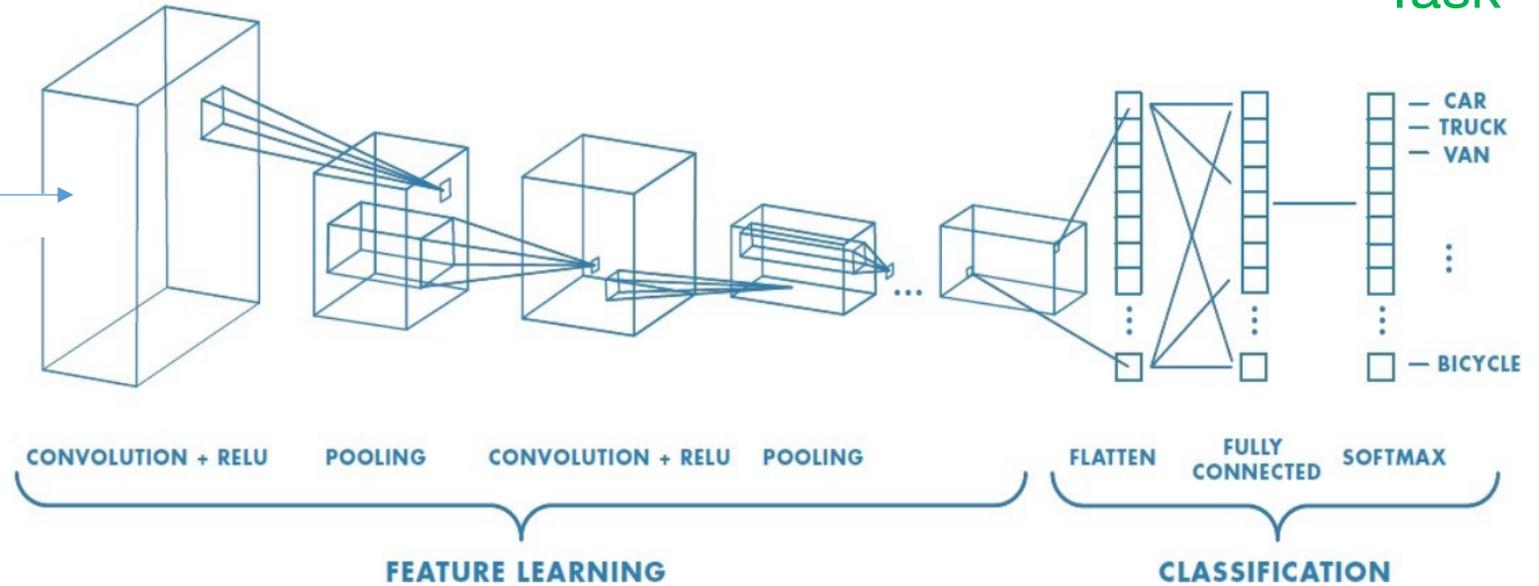
Task

Physical Function I_0



Digitized I_s

$$I_s = f[I_0]$$



Digital layers

Physical layers

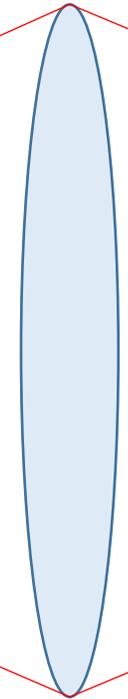
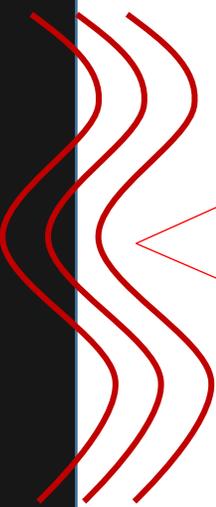
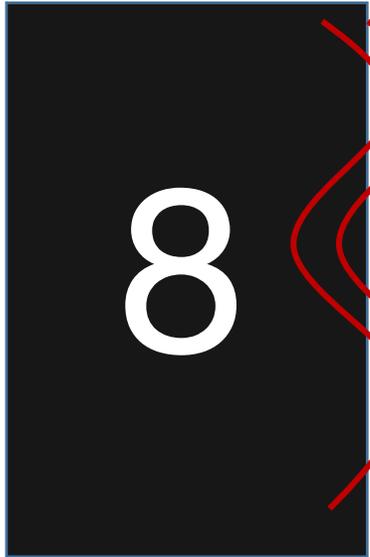
$$\text{Task} = \mathbf{W}_n \dots \text{ReLU}[\mathbf{W}_1 \text{ReLU}[\mathbf{W}_0 f[I_0]] \dots]$$

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

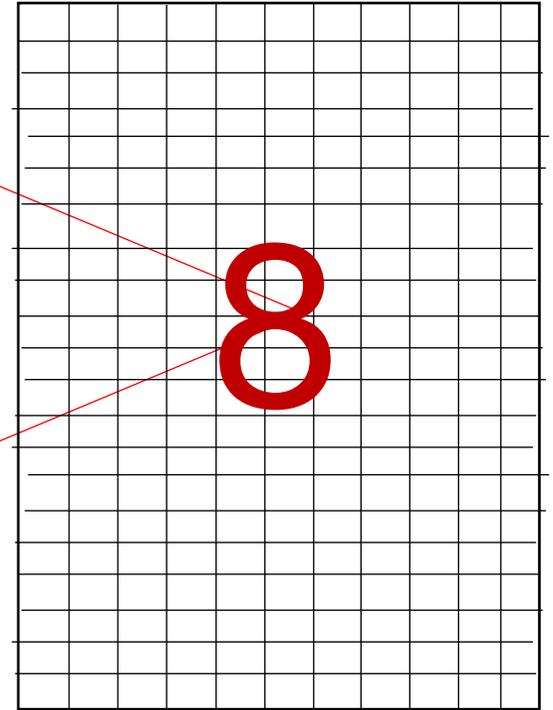
- Finite
- Non-zero gradients
- Differentiable*
- Known structure (for now...)
- Anything else?

What physical parameters effect image formation?

Input: physical object



$n \times m$ pixel array



Output: Detected image

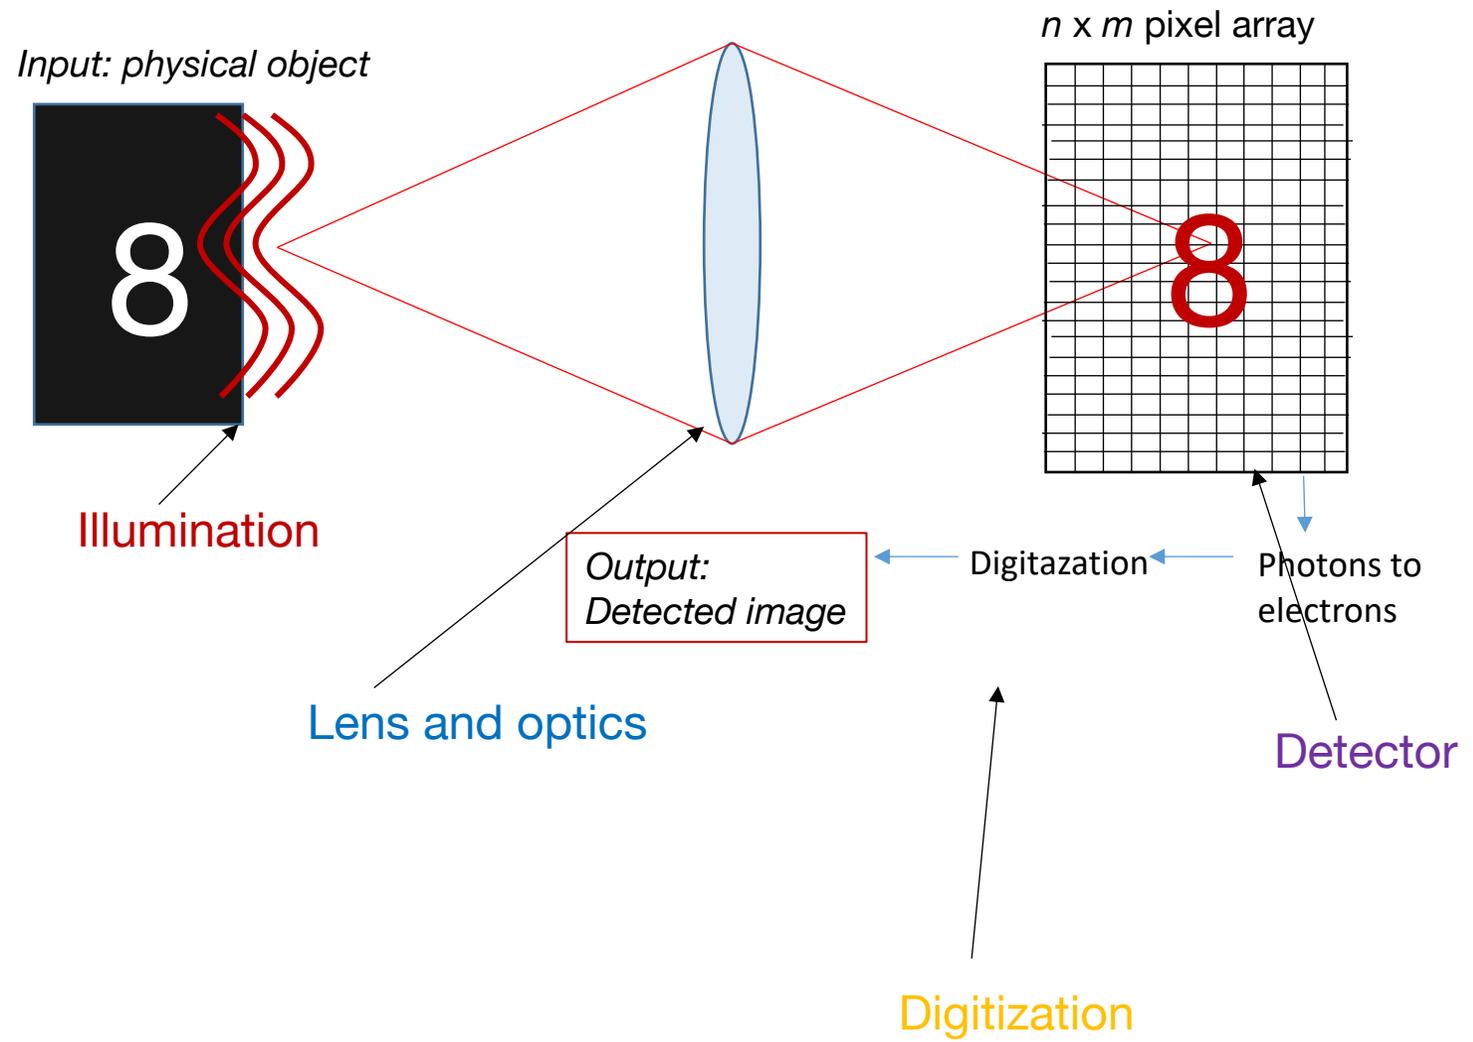
Digitazation

Photons to electrons

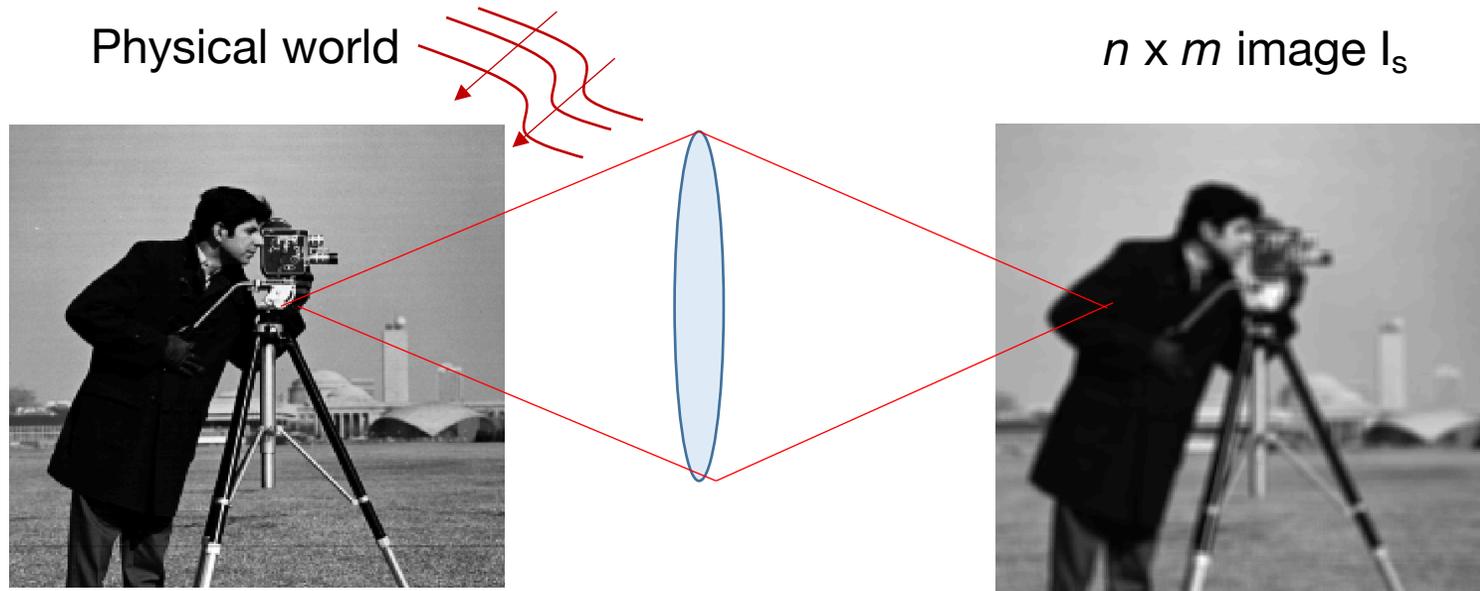


What physical parameters effect image formation?

- **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

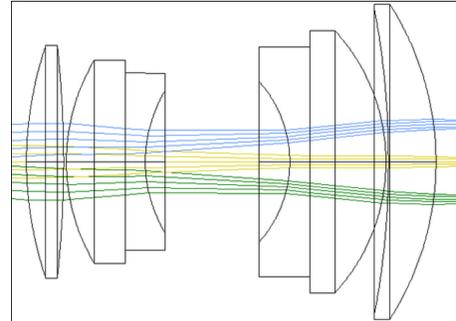
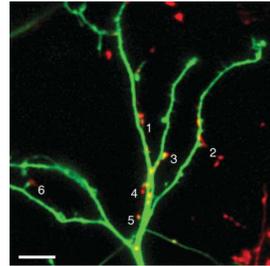


Simple mathematical model of image formation



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

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

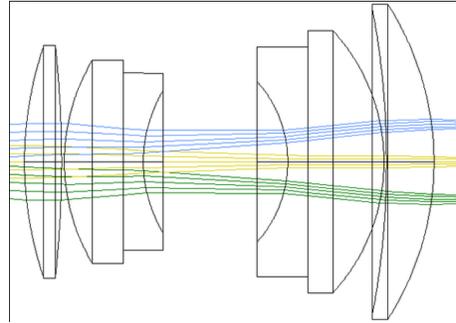
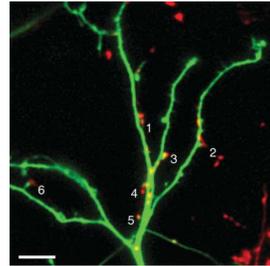


- Real, non-negative
- Models absorption and brightness

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

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

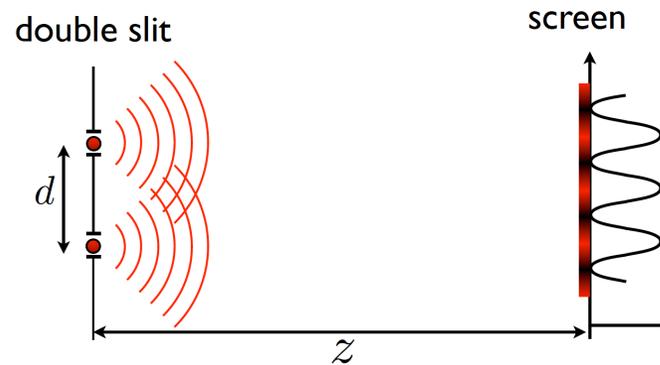
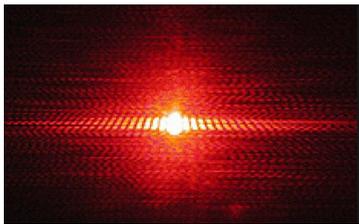
- Interpretation #1: Radiation (*Incoherent*)
- Model: Rays



- Real, non-negative
- Models absorption and brightness

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

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

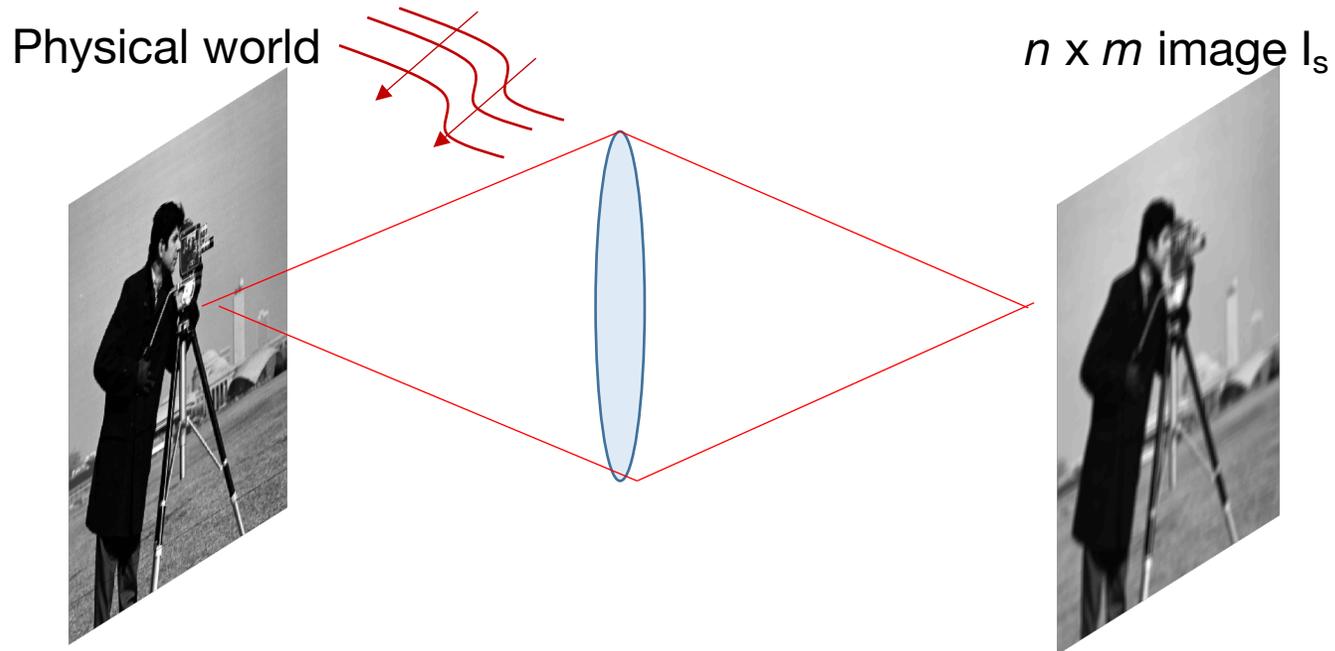


- Complex field
- Models Interference

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

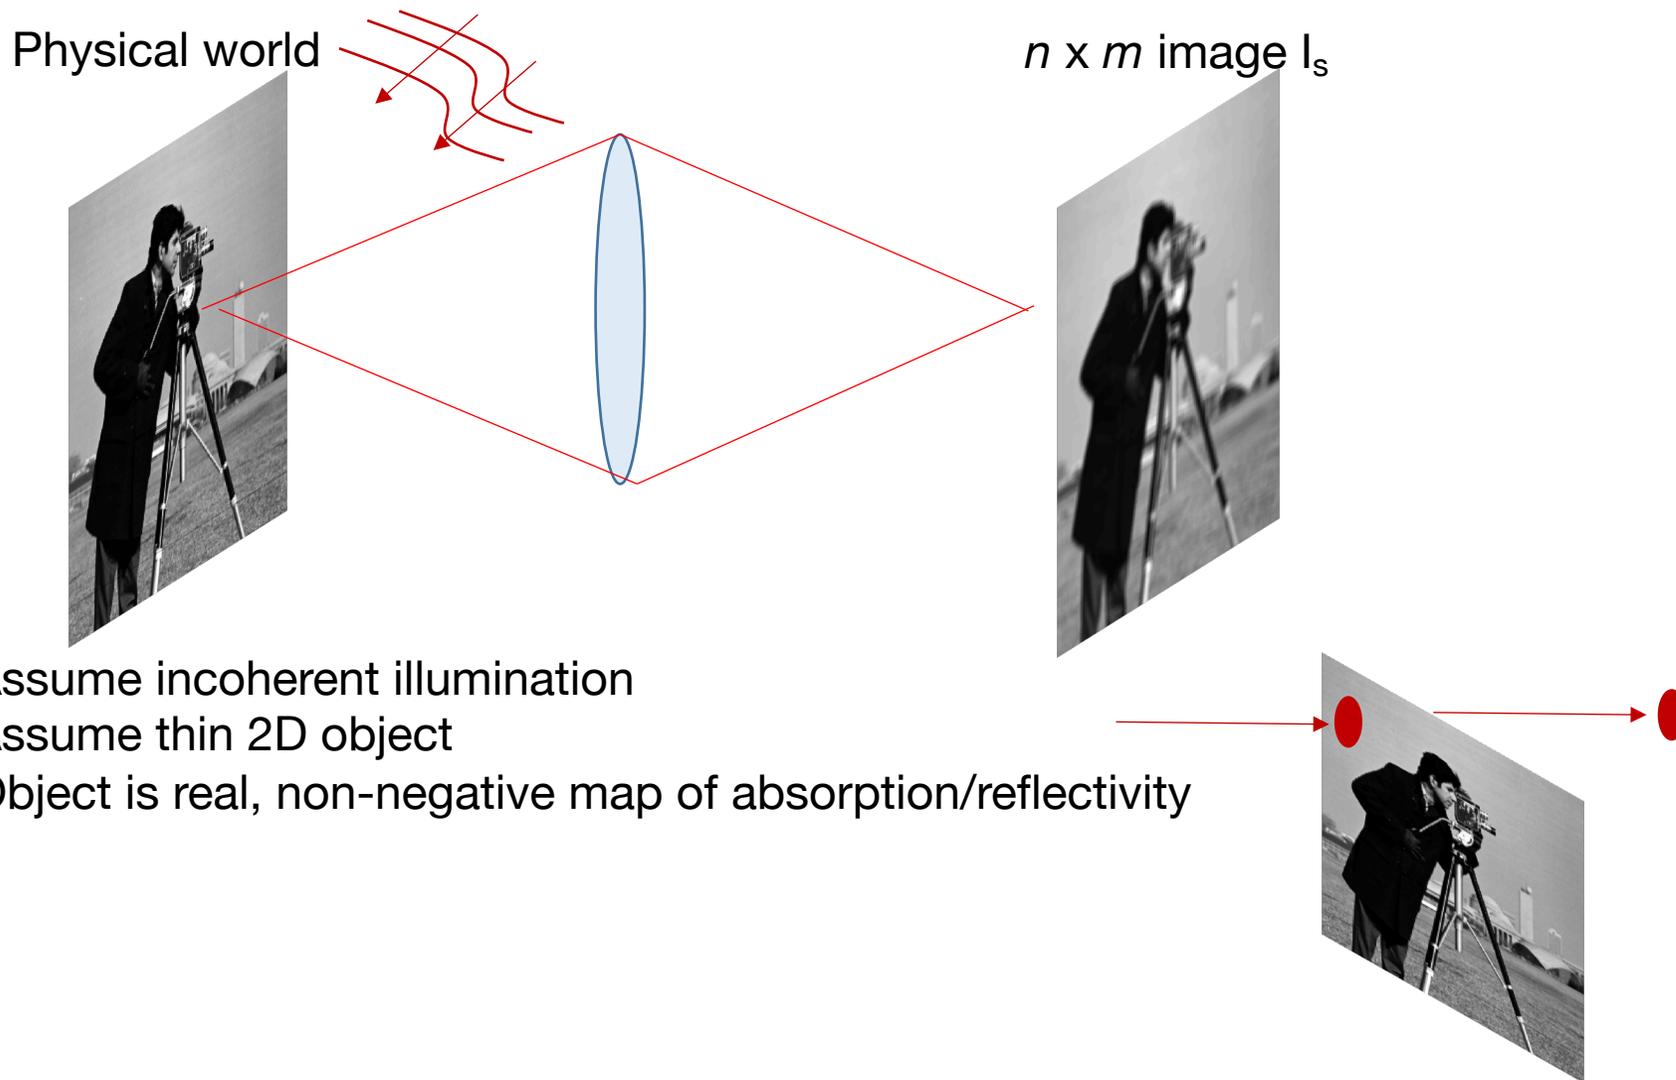
- Interpretation #3: Particle
- Model: Photons

Simple mathematical model of image formation

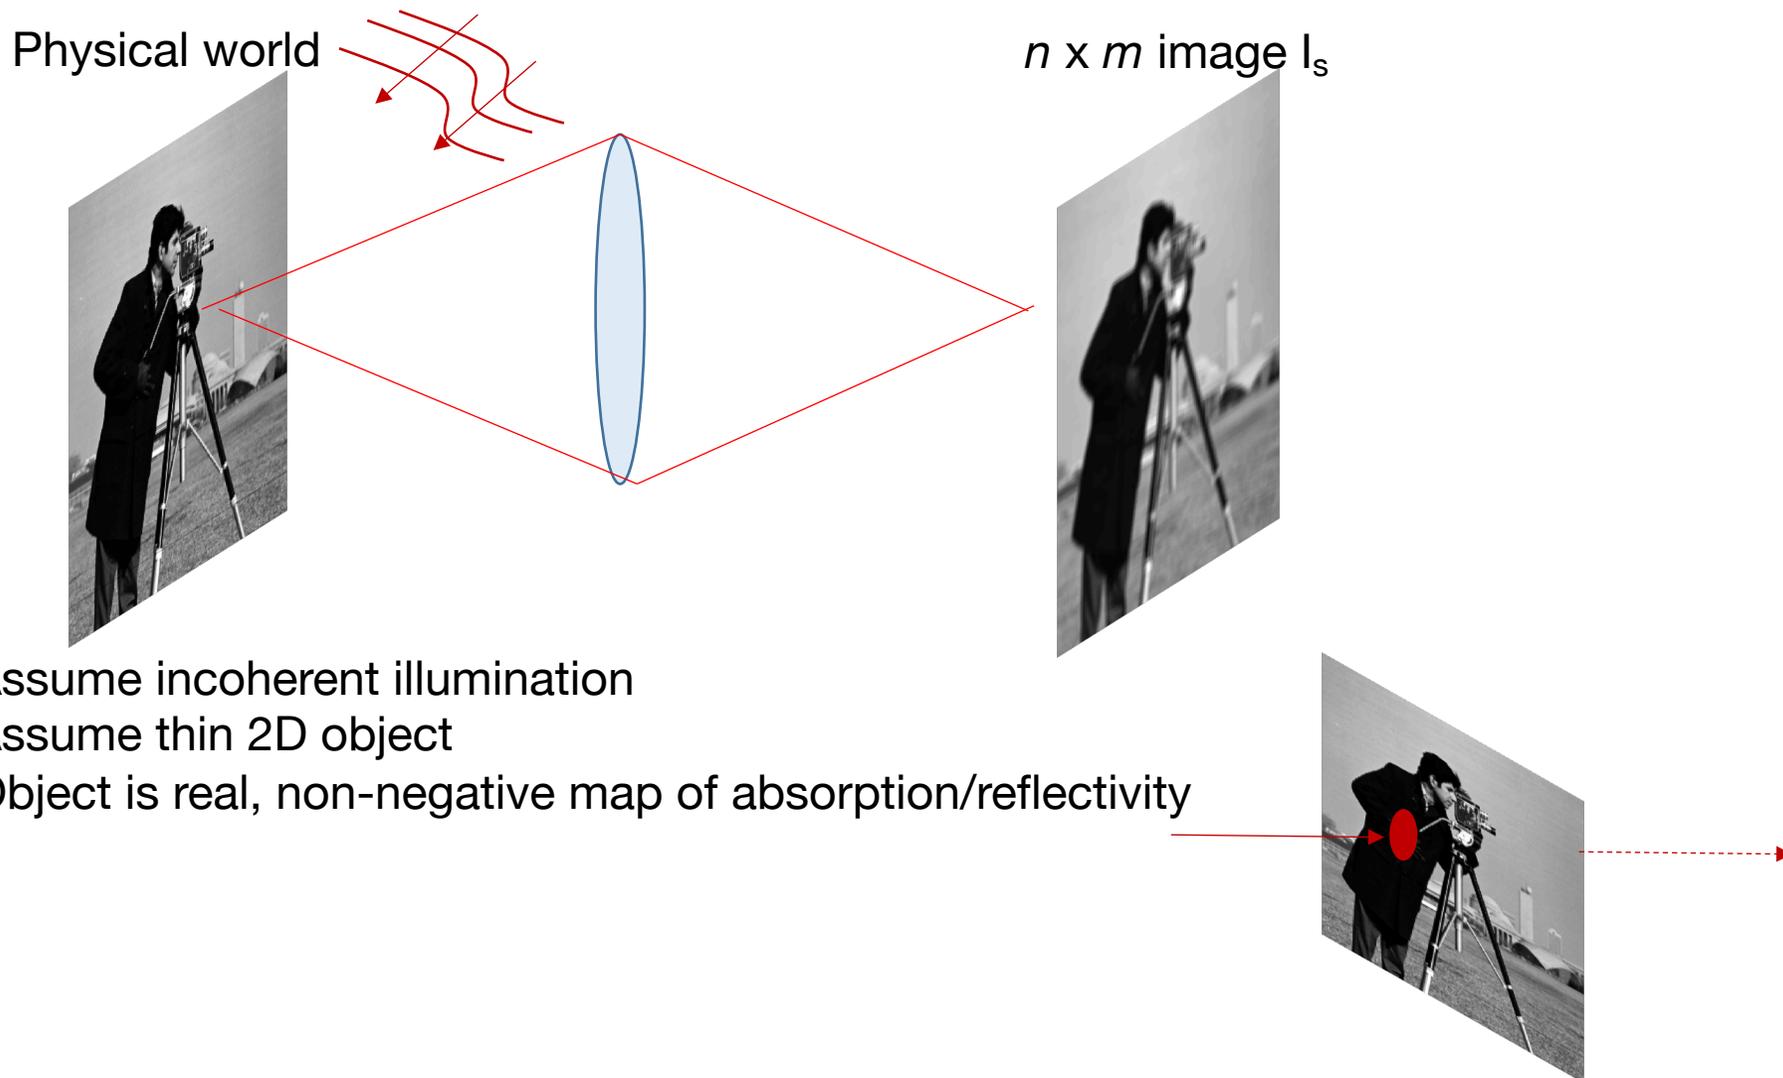


- Assume incoherent illumination
- Assume thin 2D object
- Object is real, non-negative map of absorption/reflectivity

Simple mathematical model of image formation

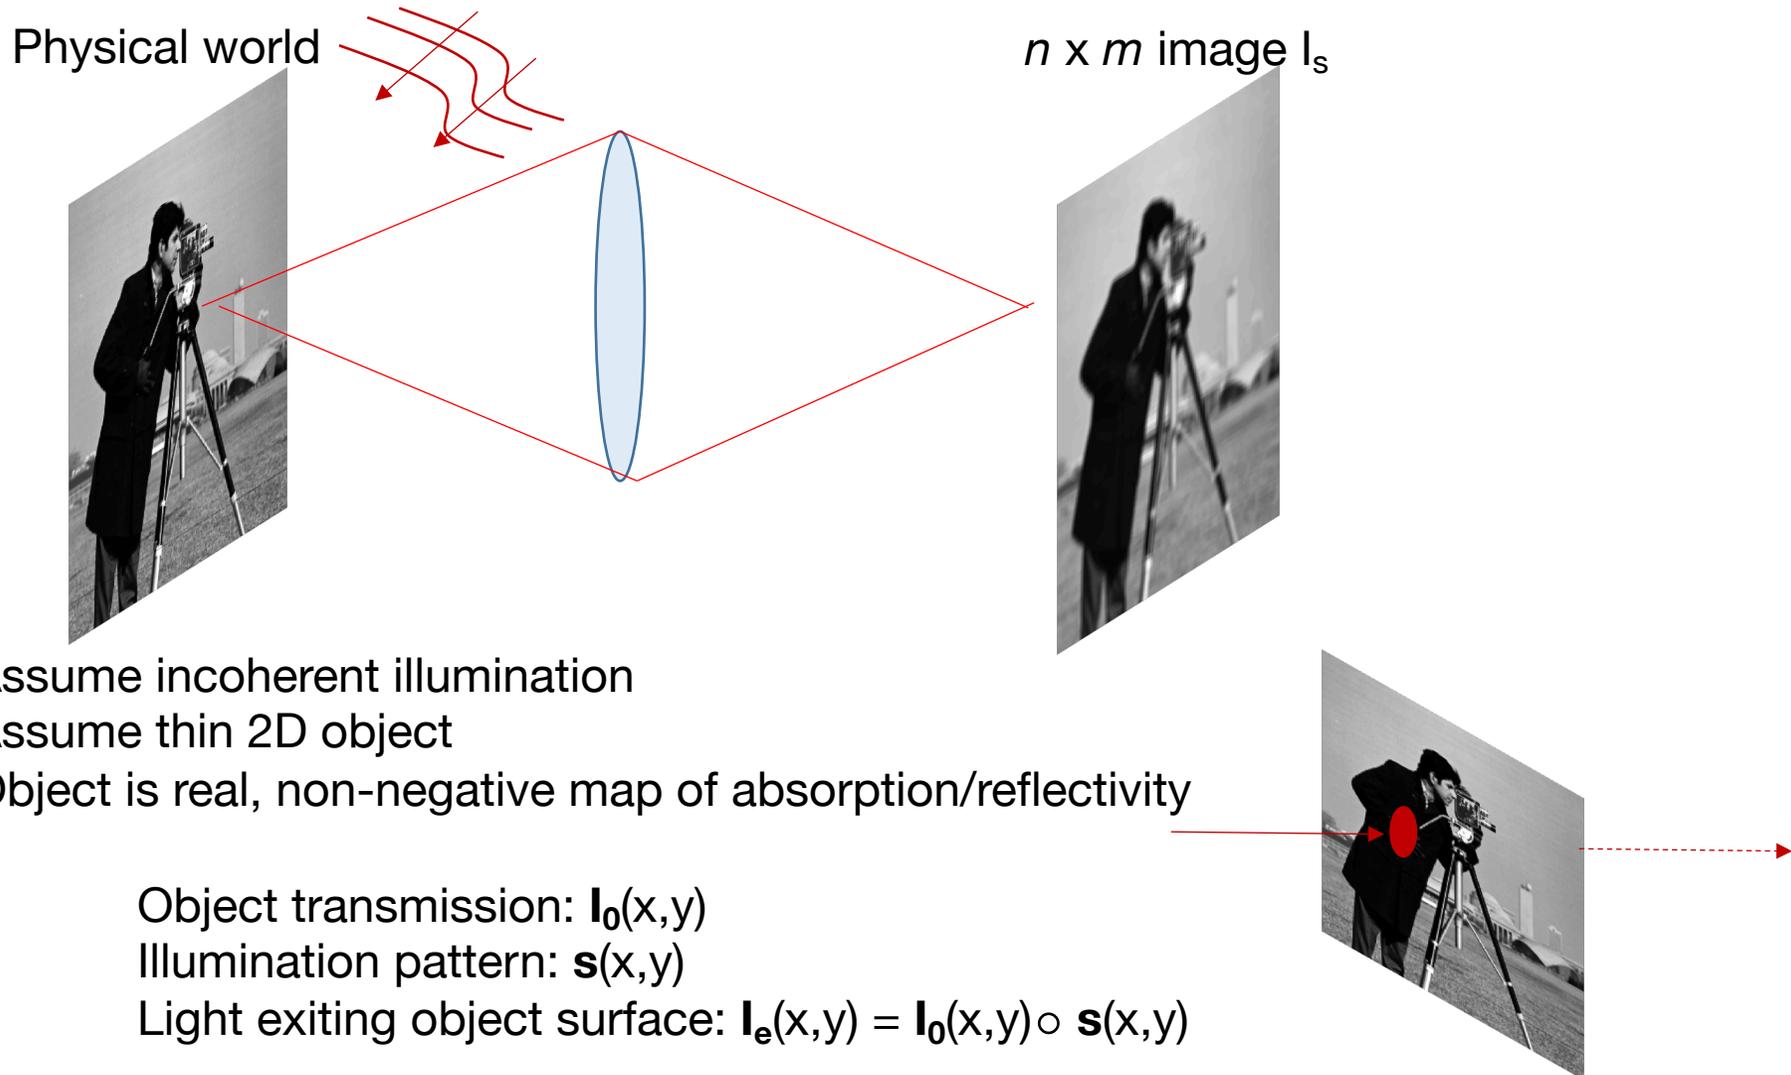


Simple mathematical model of image formation

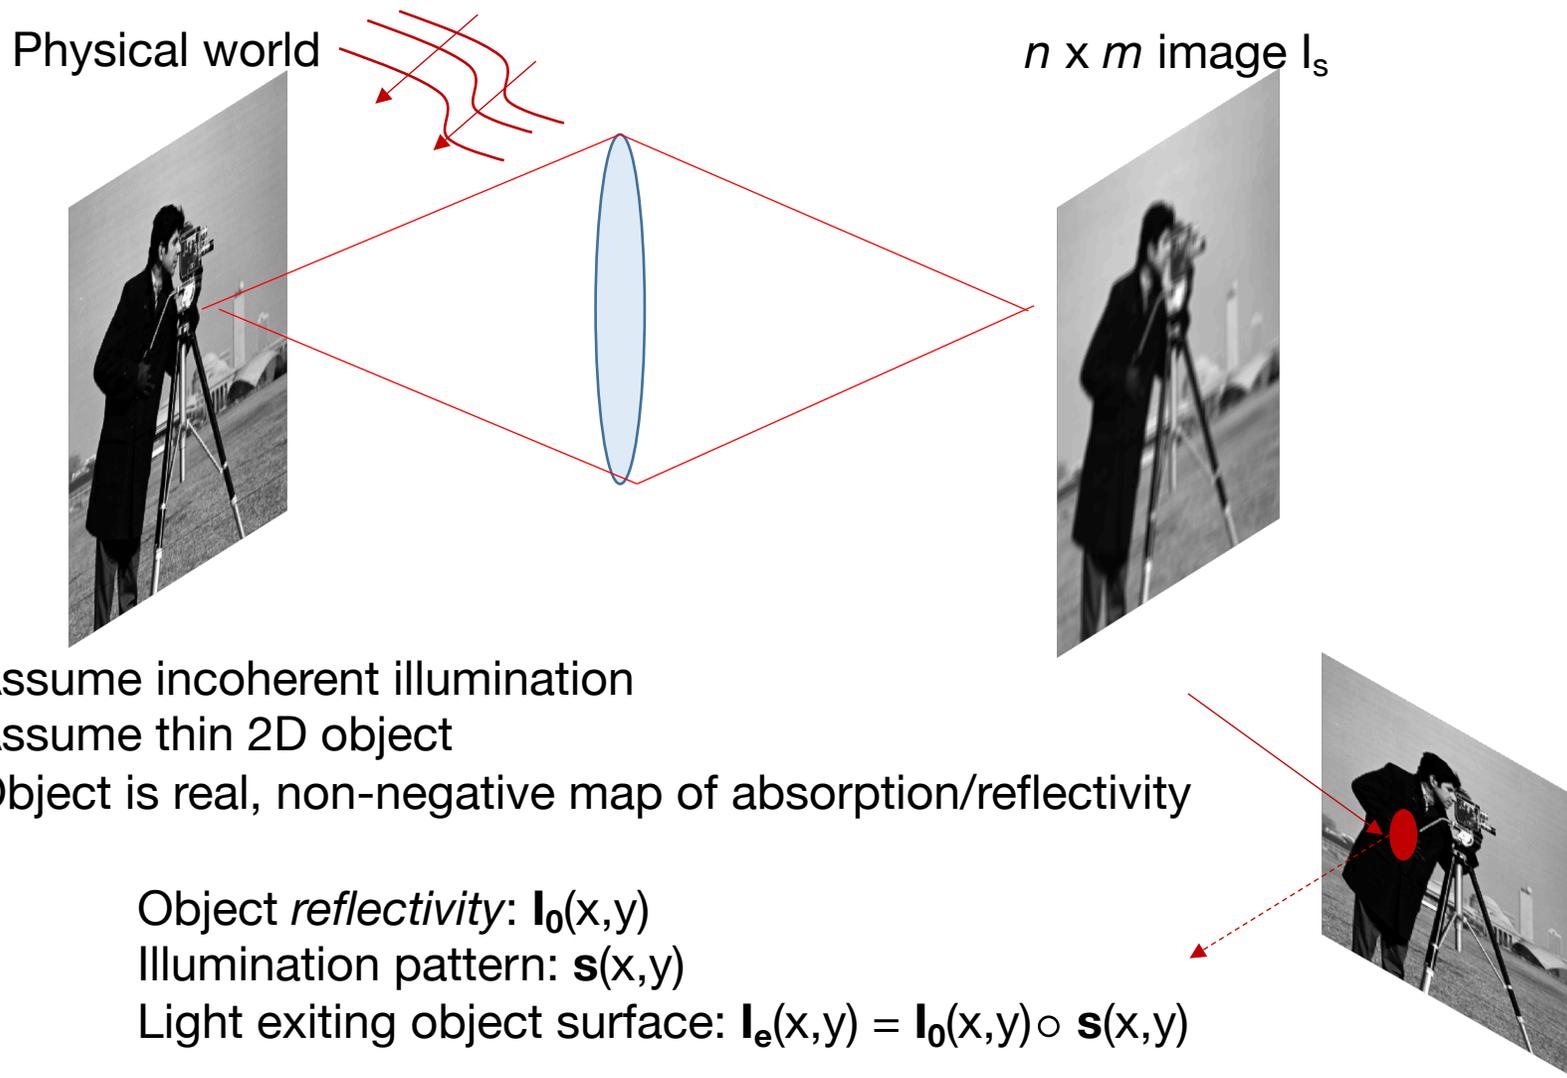


- Assume incoherent illumination
- Assume thin 2D object
- Object is real, non-negative map of absorption/reflectivity

Simple mathematical model of image formation



Simple mathematical model of image formation



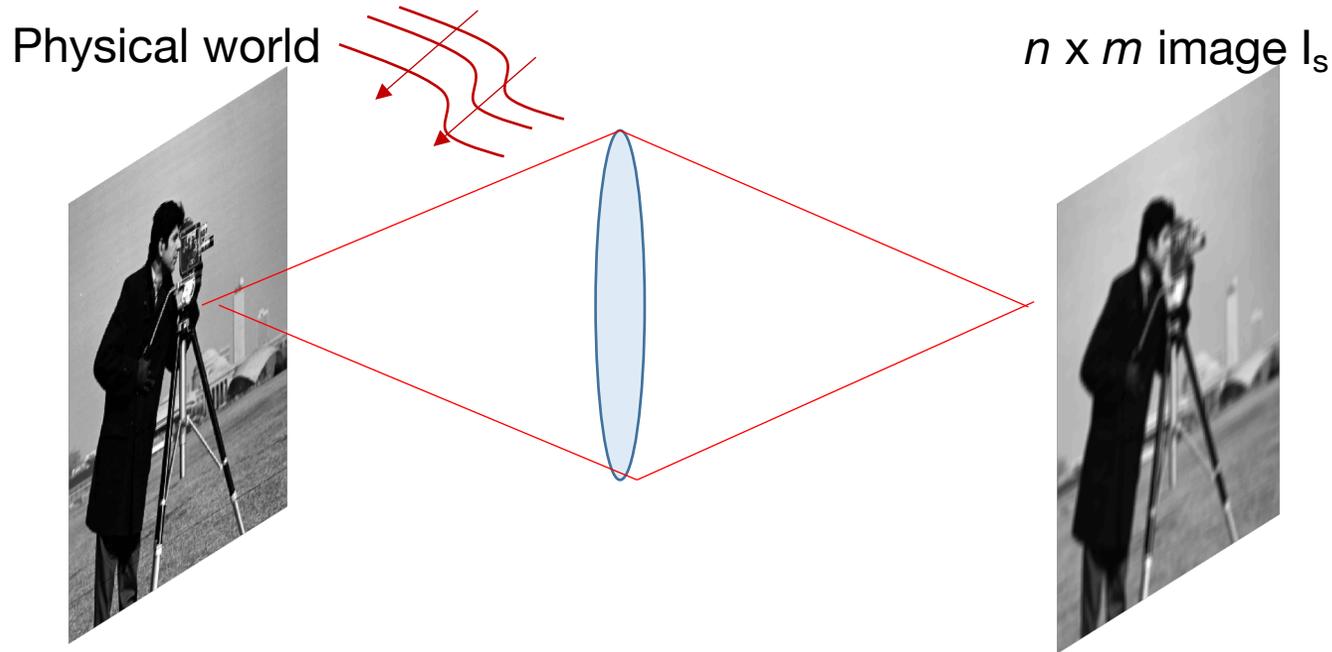
- Assume incoherent illumination
- Assume thin 2D object
- Object is real, non-negative map of absorption/reflectivity

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 image formation



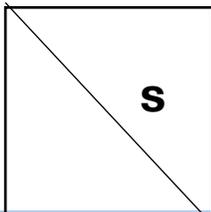
- Assume incoherent illumination
- Assume thin 2D object
- Object is real, non-negative map of absorption/reflectivity

Modeling incoherent illumination

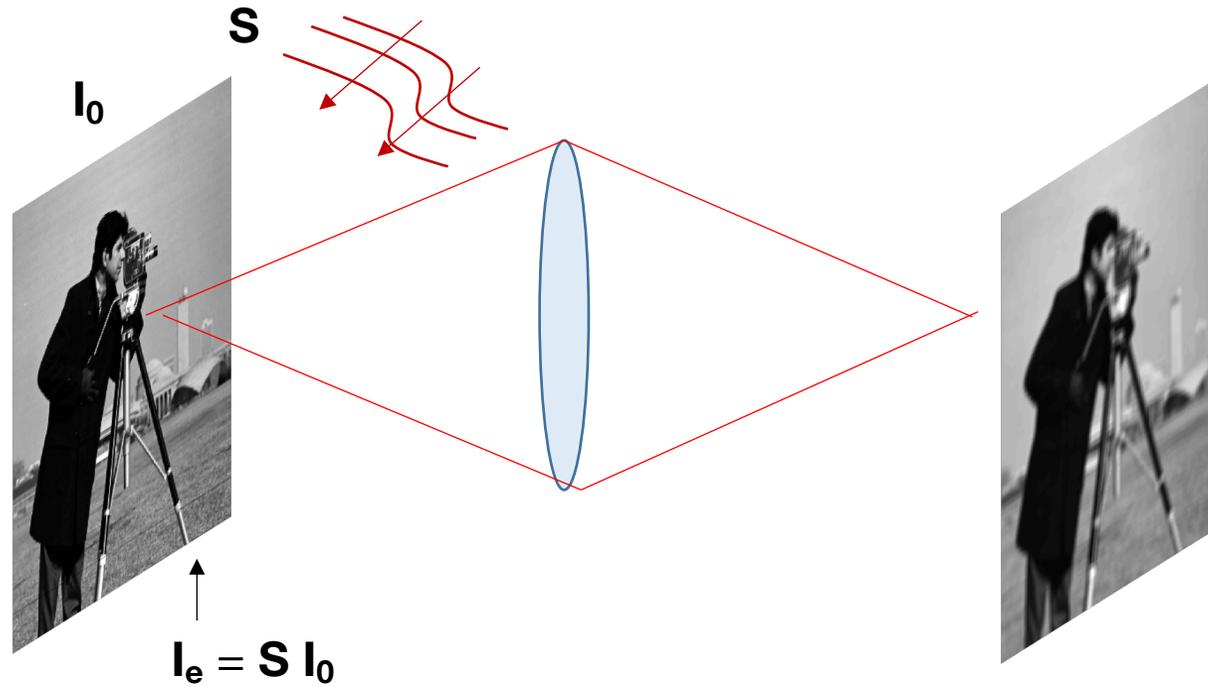
$$I_e(x,y) = I_0(x,y) \circ s(x,y)$$

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

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



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)



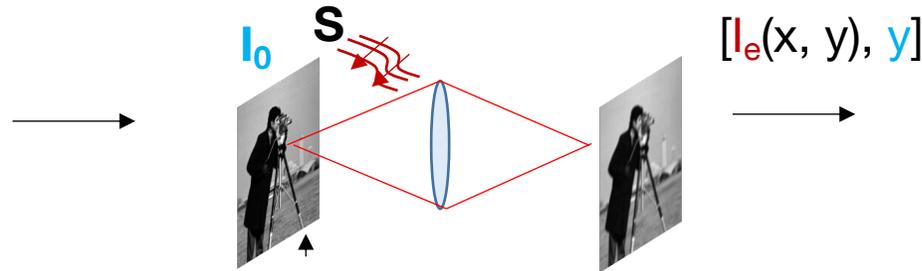
Training data:

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

I_0 : 100 x 100

Label y : 1x2

Example #1: Optimized illumination pattern (one color)



Training data:

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

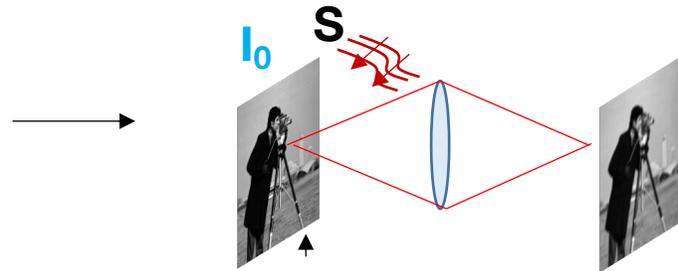
I_0 : 100 x 100

Label y : 1x2

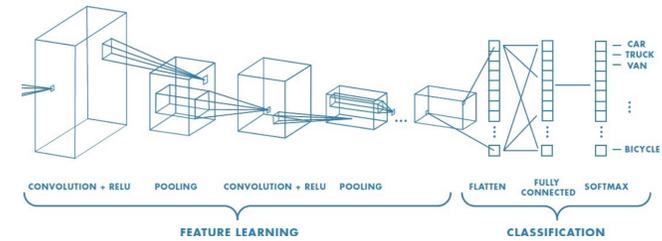
$$I_e(x, y) = \mathbf{S} I_0(x, y)$$

Physical Layer

Example #1: Optimized illumination pattern (one color)



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



Training data:

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

$I_0: 100 \times 100$

Label $y: 1 \times 2$



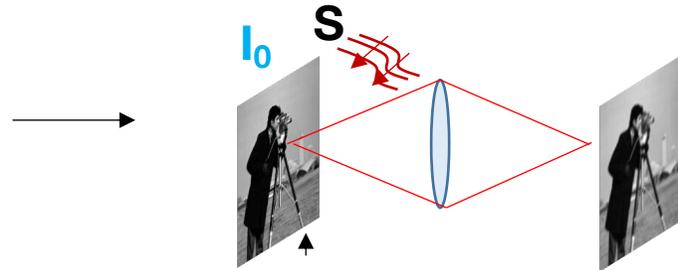
$$I_e(x, y) = S I_0(x, y)$$

Physical Layer

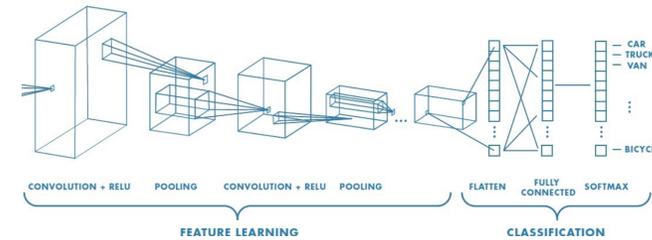


Task = $W_n \dots \text{ReLU}[W_1 \text{ReLU}[W_0 I_e] \dots]$

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!
    
```

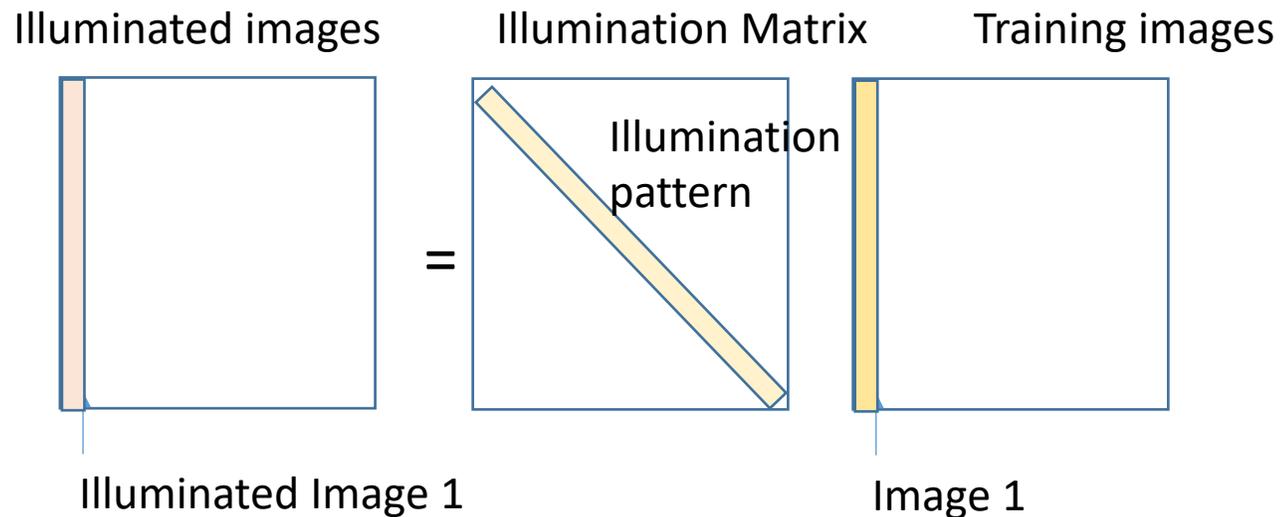
(Will show CoLab Notebook for implementation in Tensorflow 2.0)

Option 1: `tf.linalg.matmul`

1.0 code

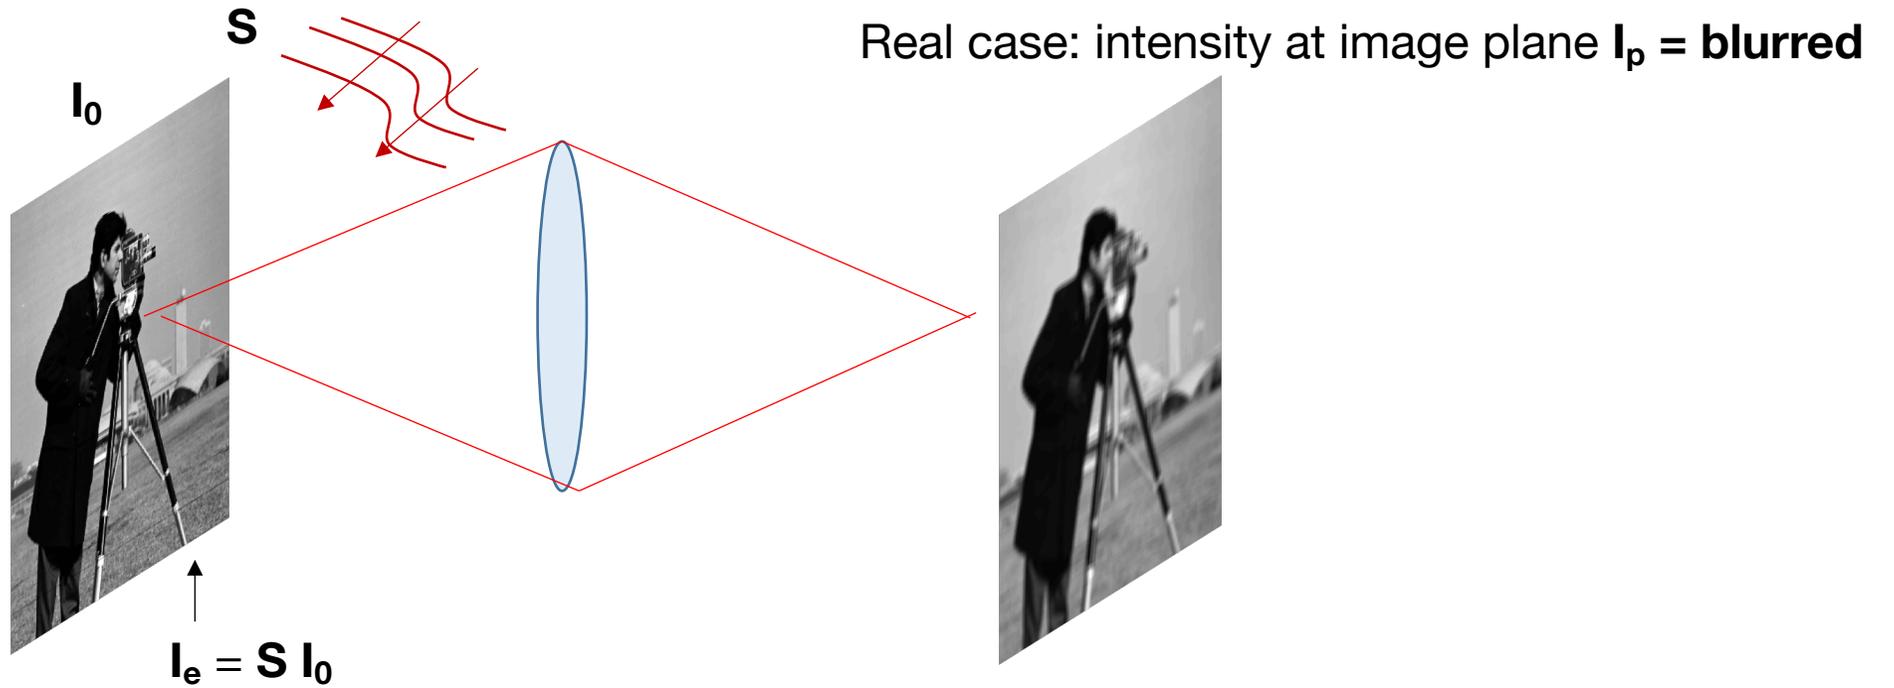
```
training_images = tf.placeholder(tf.float32, [image_size, None])
training_labels = tf.placeholder(tf.float32, [None, 3])
illumination_pattern = tf.truncated_normal([image_size, 1], stddev = 0.1)
illumination_matrix = tf.linalg.diag(illumination_pattern)
illuminated_images = tf.matmul(illumination_matrix,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



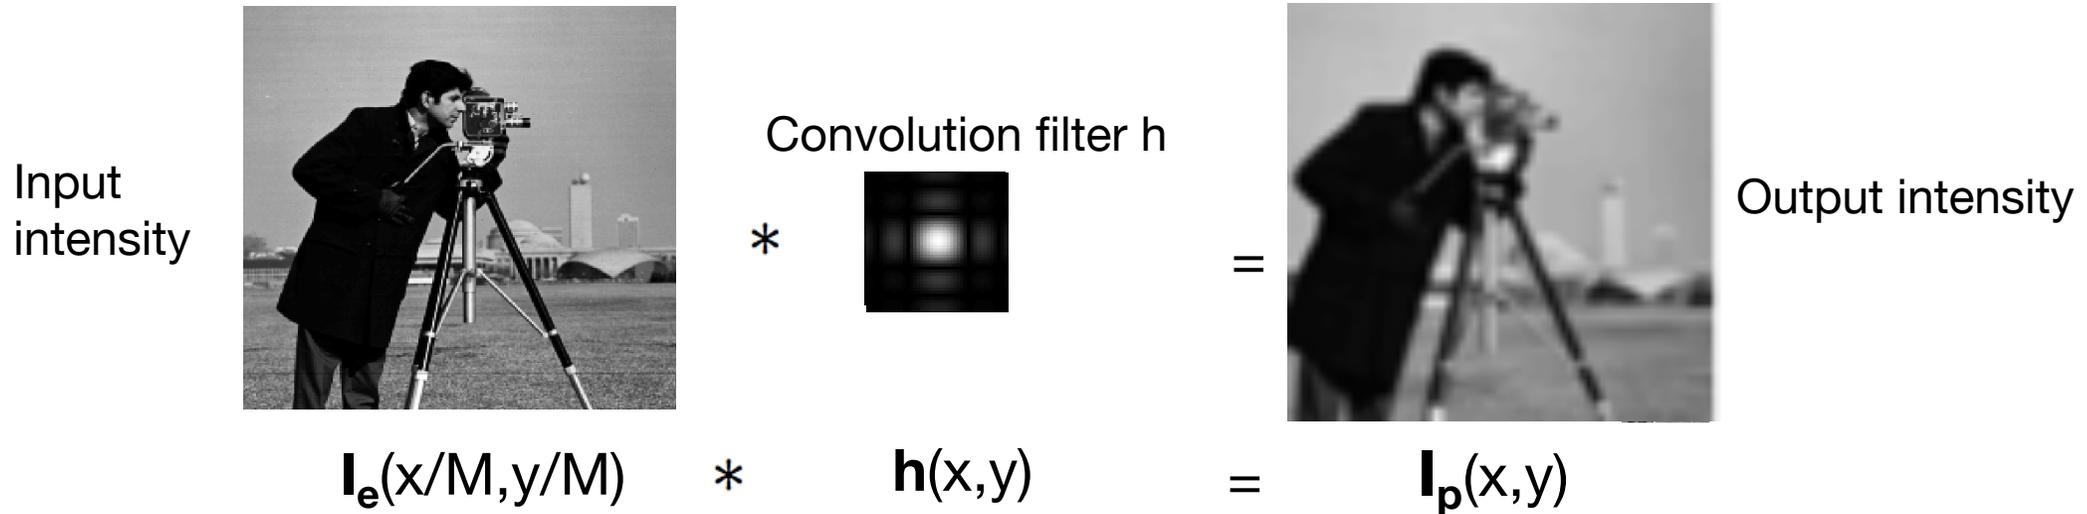
Option 2: `tf.linalg.multiply` (will show in CoLab Notebook)

Simple mathematical model of image formation



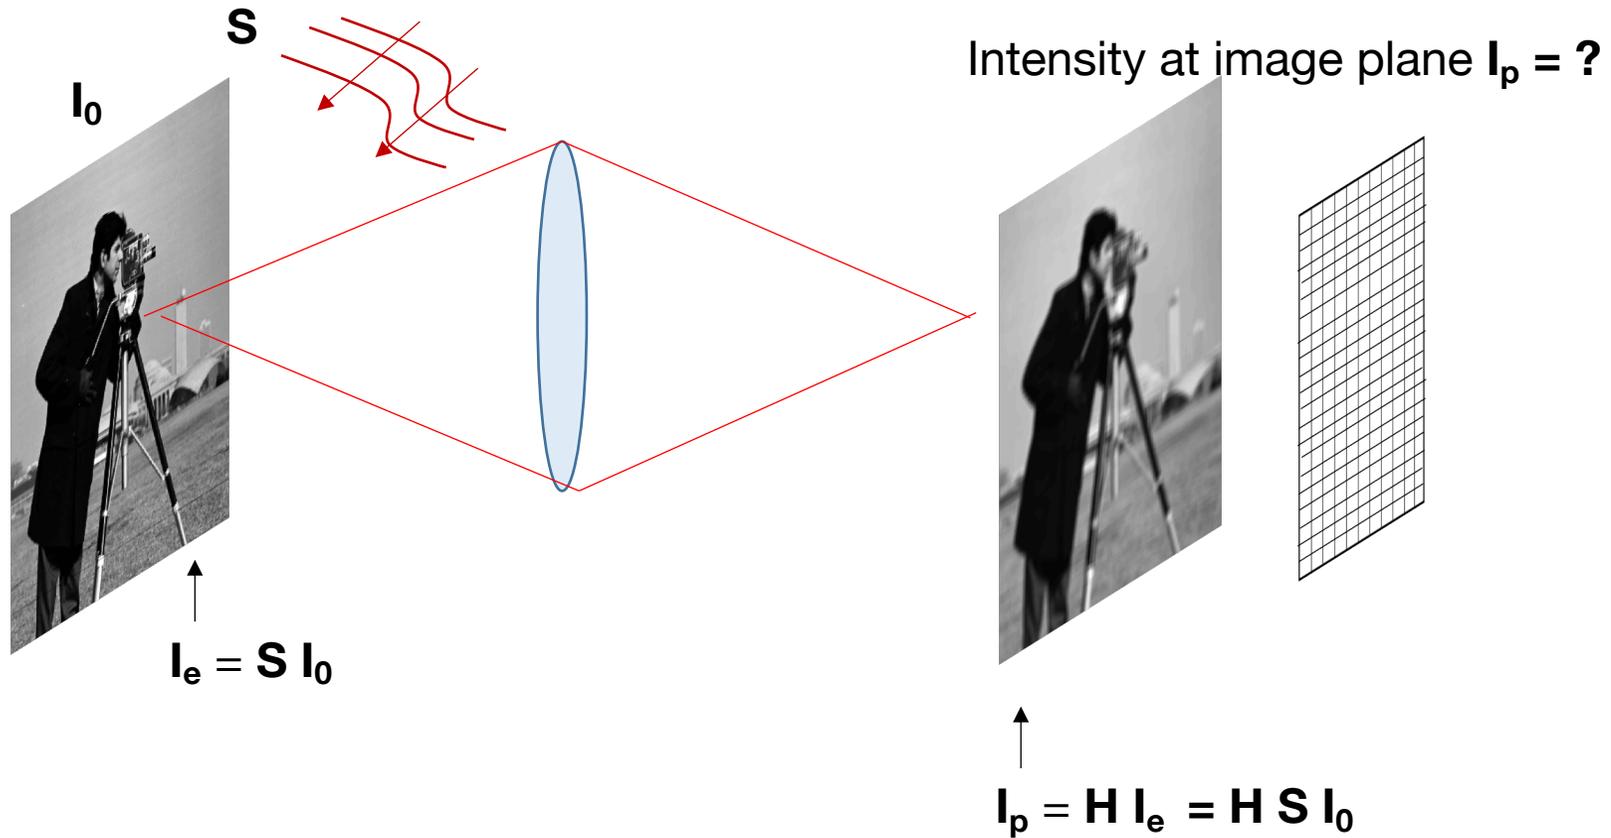
Simple mathematical model of image formation

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

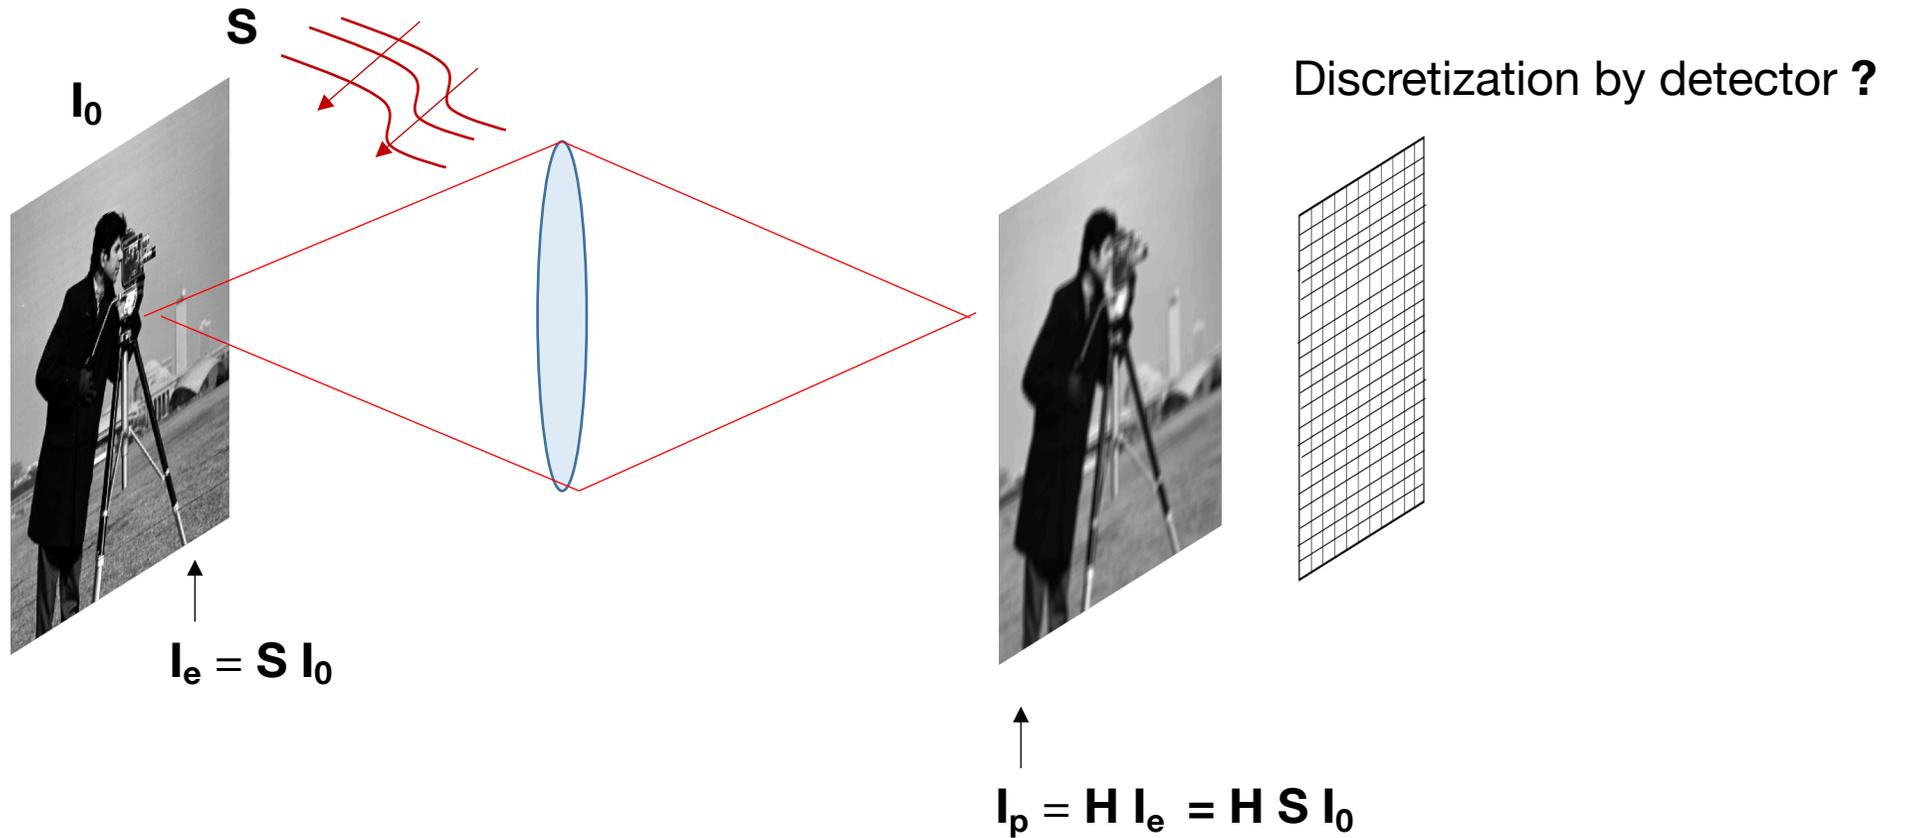


Assuming we've resized by M, $I_p = I_e * h = H I_e$

Simple mathematical model of image formation



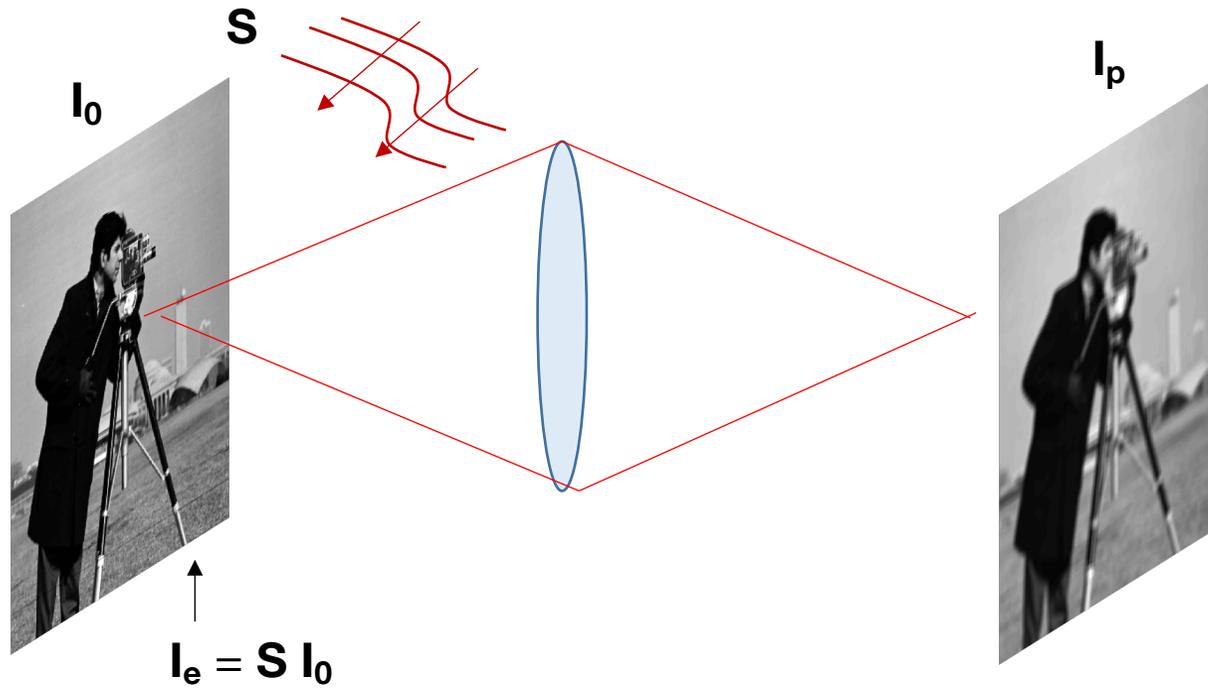
Simple mathematical model of image formation



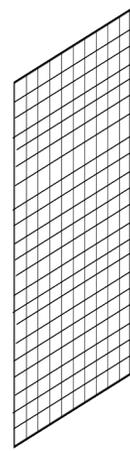
Use downsampling matrix (sum-pooling)

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

Simple mathematical model of image formation



Discretization by detector ?

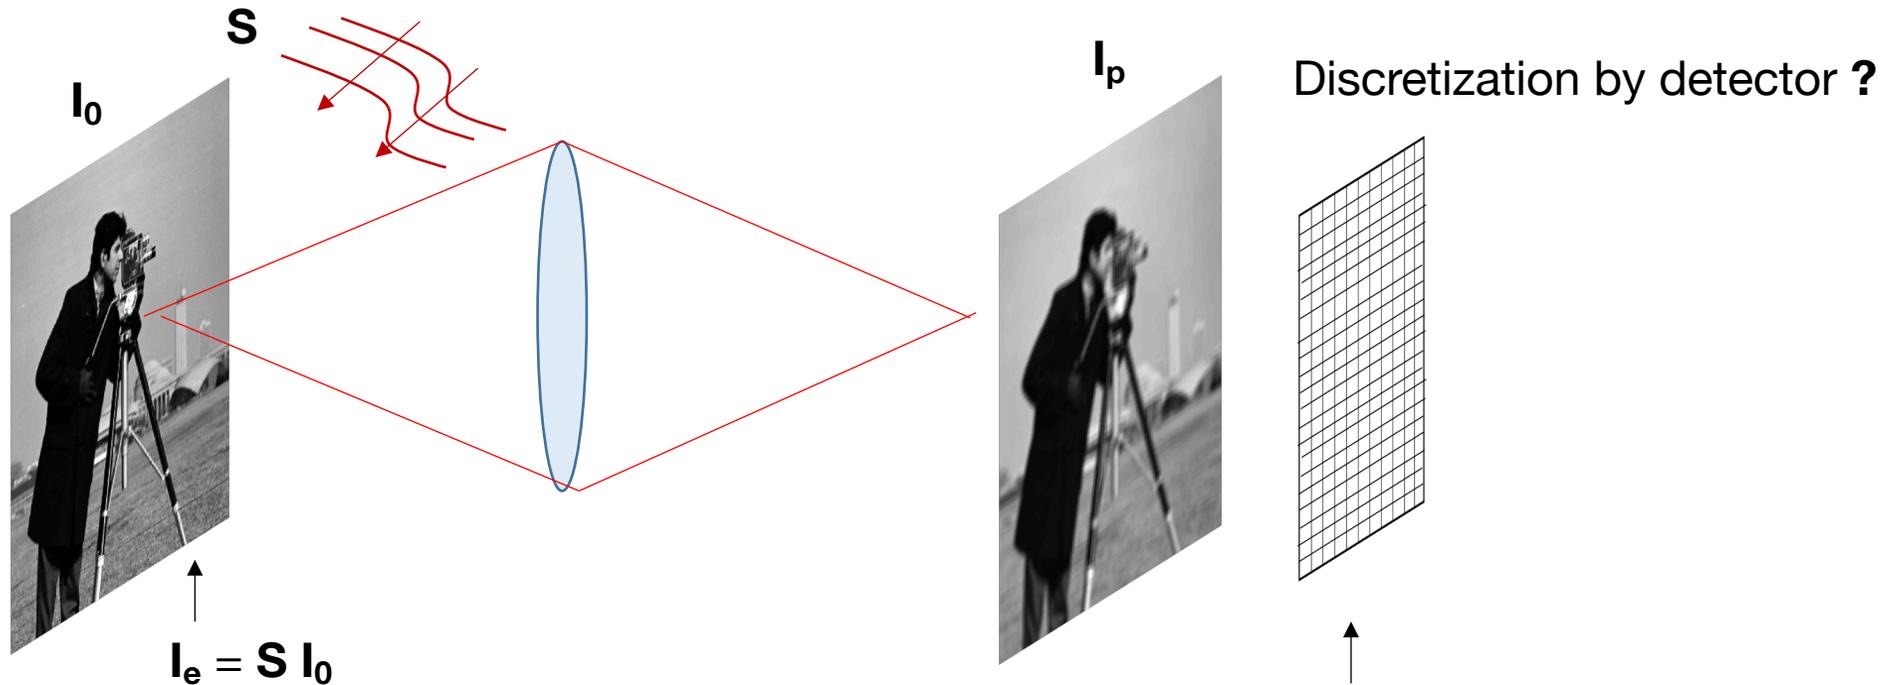


$$I_s = D I_p = D H S I_0$$

Use downsampling matrix (sum-pooling)

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

Simple mathematical model of image formation

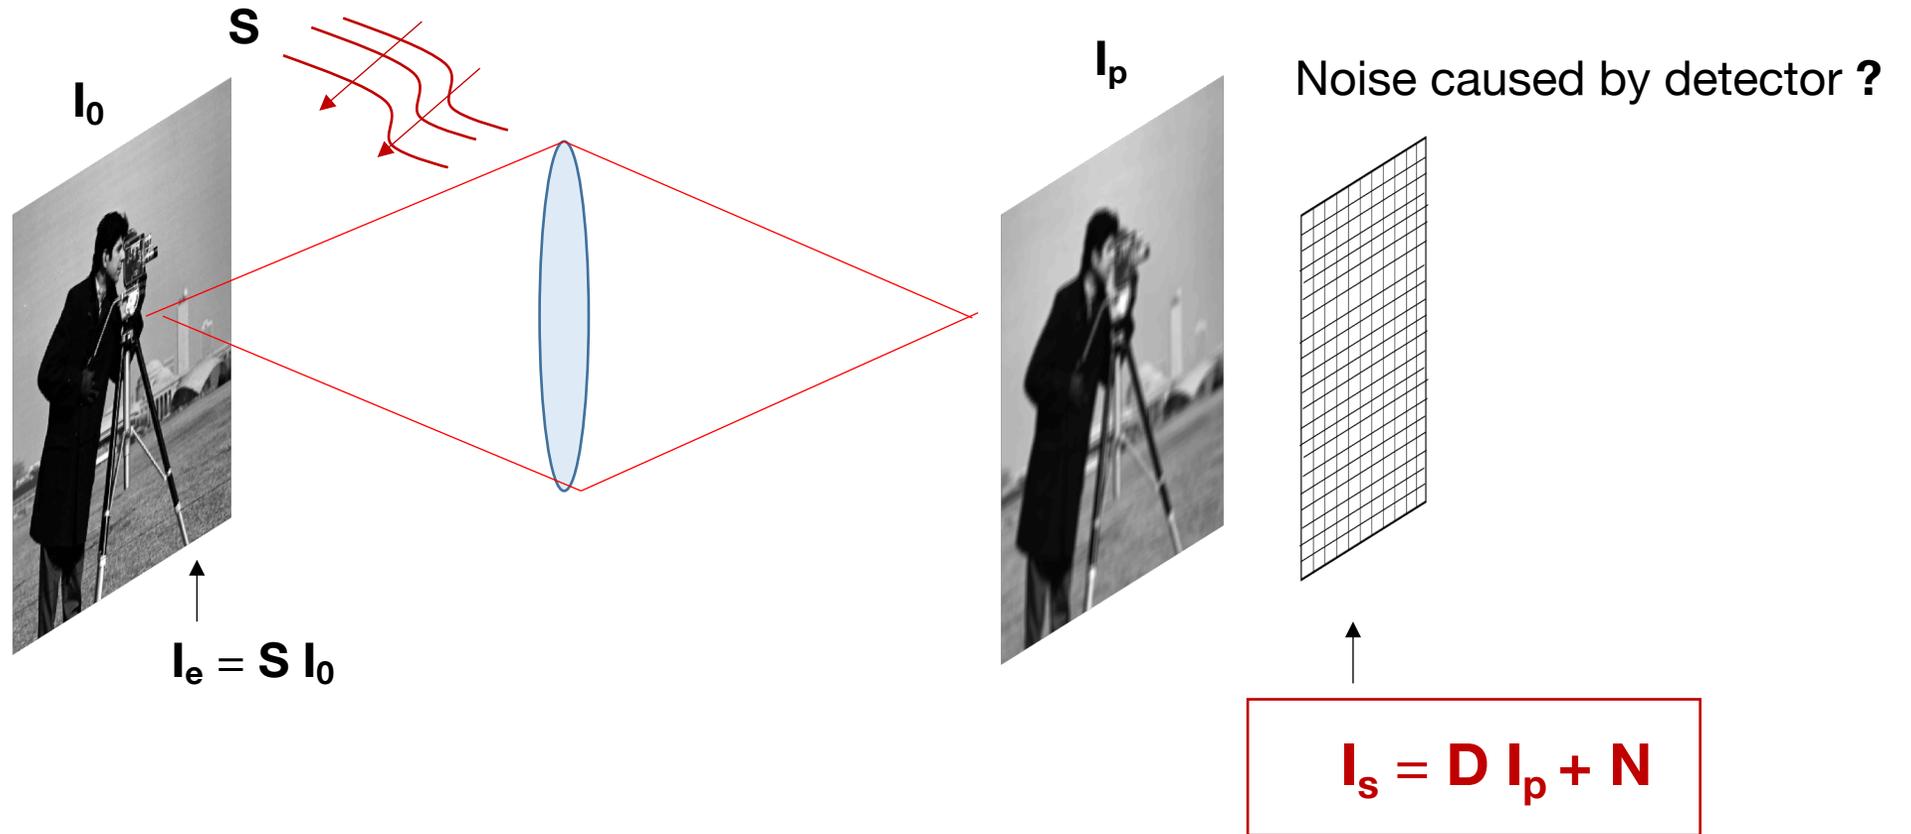


$$I_s = D I_p = D H S I_0$$

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

*Can make these learnable weights

Simple mathematical model of image formation

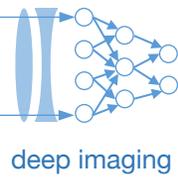


Can also add in detector-dependent noise $N = k * \text{np.random.randn}(dx, dy)$

(zero-mean Gaussian noise, for example)

Pause to take a look at:

`physical_layers_example.ipynb`



Task

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

Digital Layers

Physical Layers

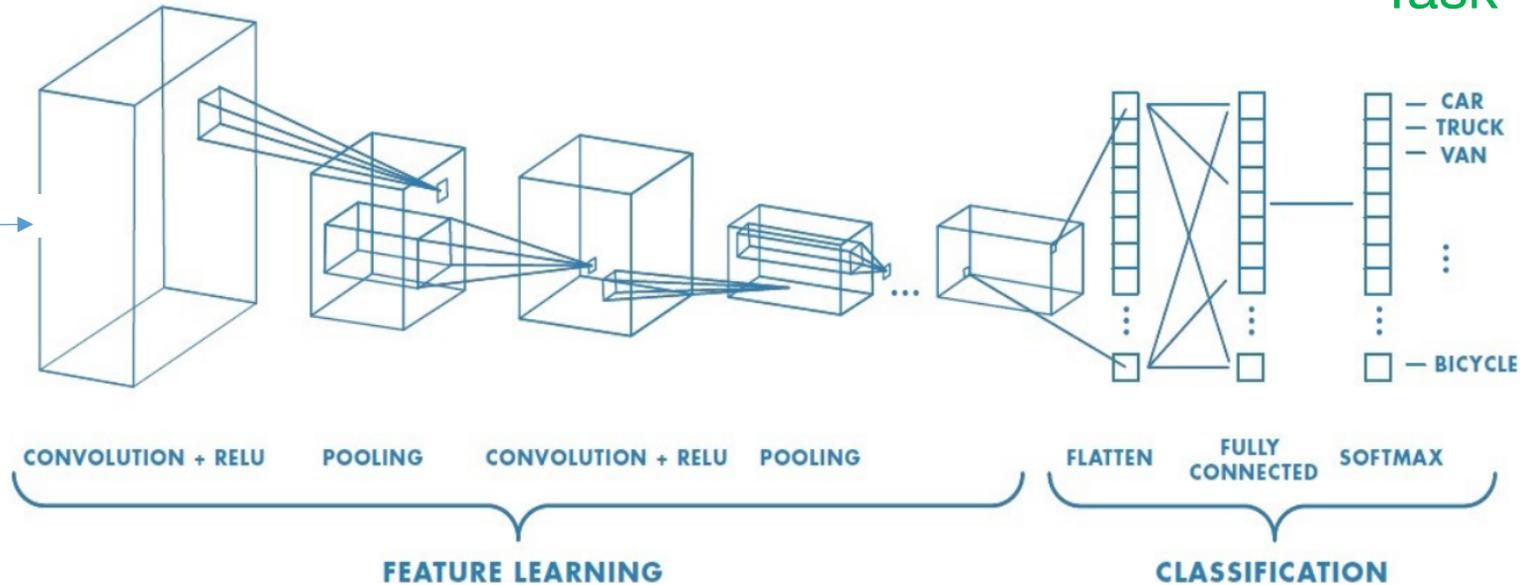
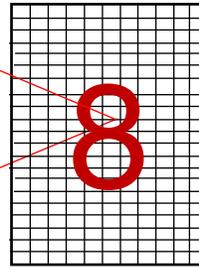
Physical world

$$I_0(x_0, y_0)$$



Digital Image

$$I_s(x, y)$$



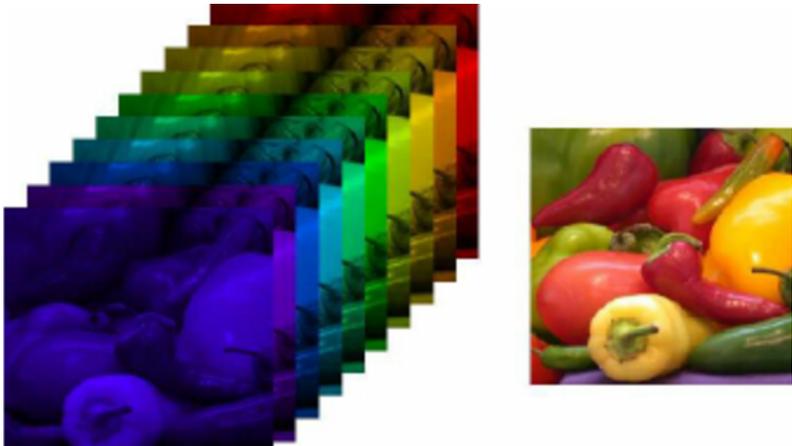
Digital layers

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

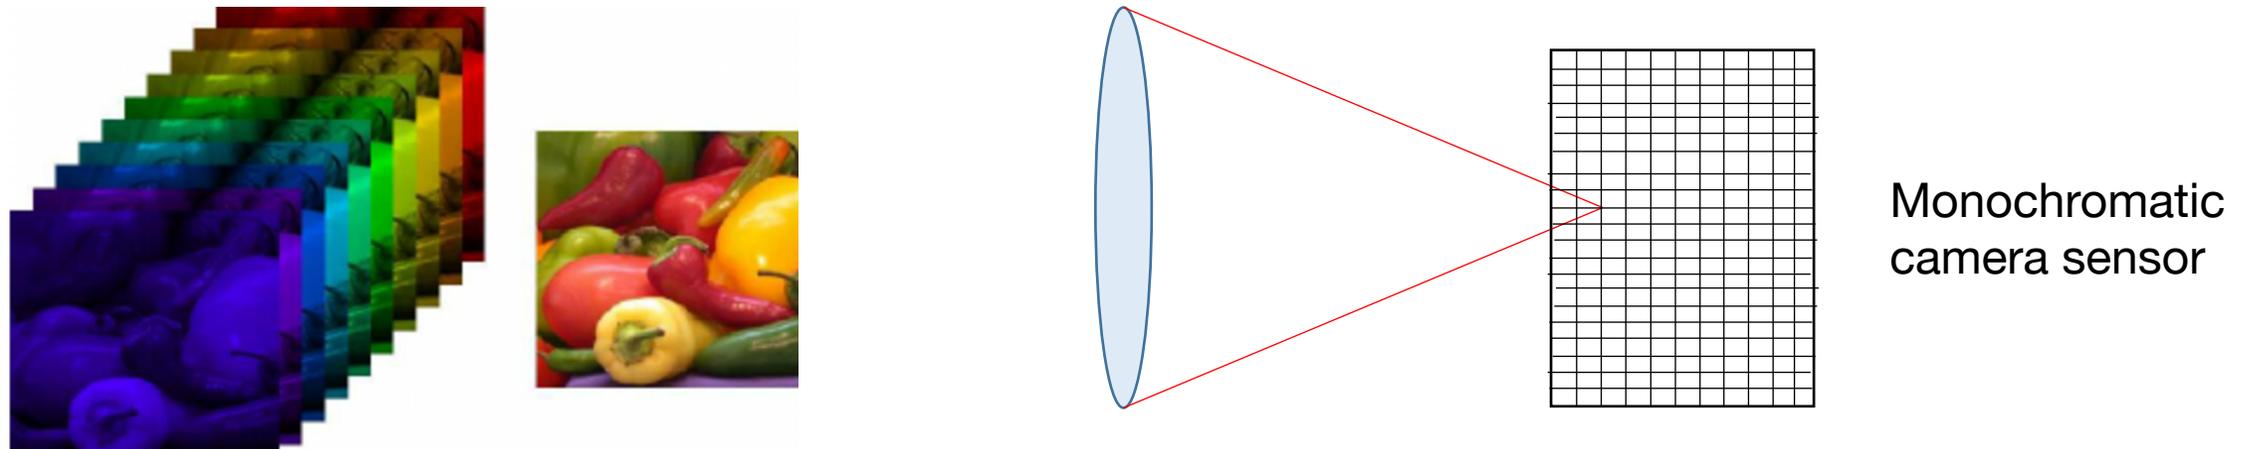


“Ground truth” object:

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

100 x 100 pix. x 30 spectral channels

Example #2: Optimized color filter for a grayscale camera

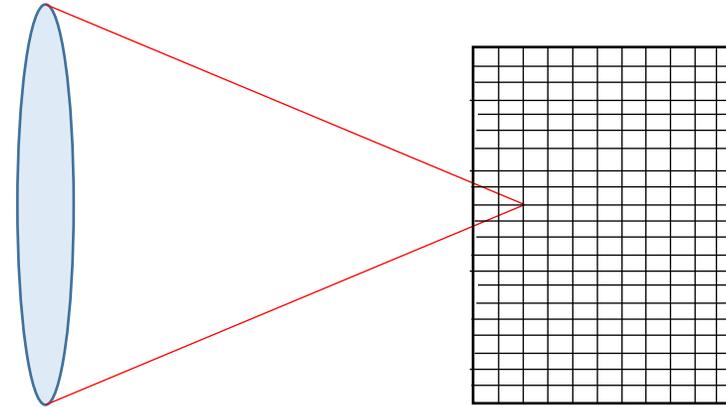
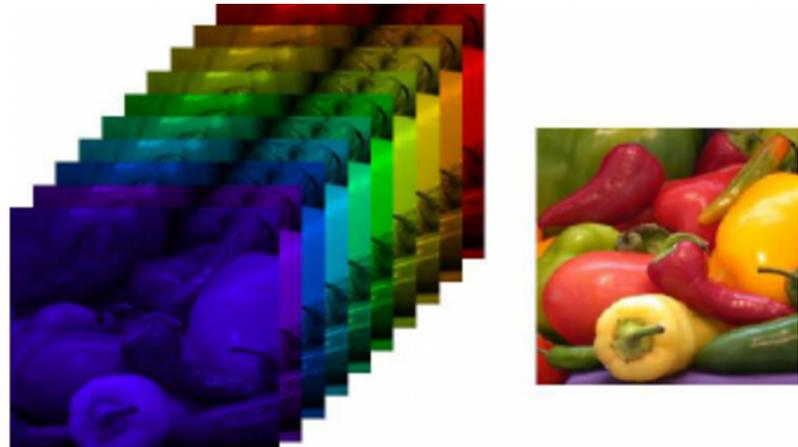


“Ground truth” object:

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

100 x 100 pix. x 30 spectral channels

Example #2: Optimized color filter for a grayscale camera



Monochromatic camera sensor

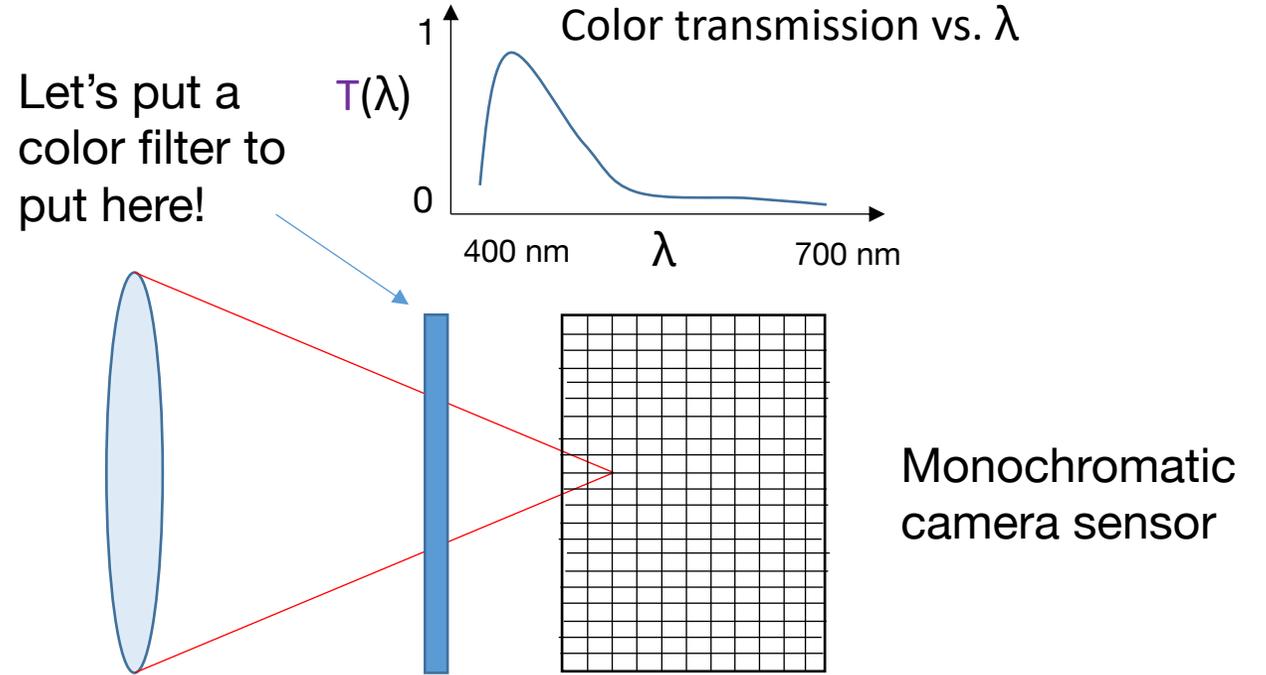
“Ground truth” object:

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

100 x 100 pix. x 30 spectral channels

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

Example #2: Optimized color filter for a grayscale camera

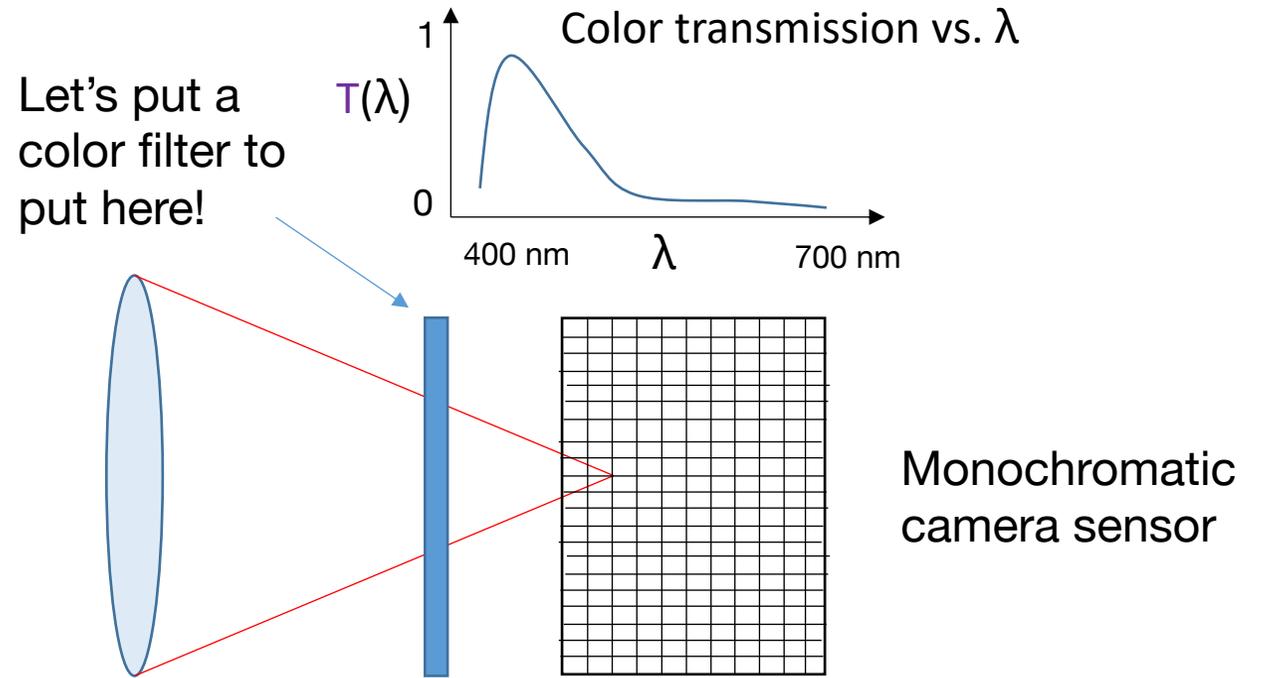
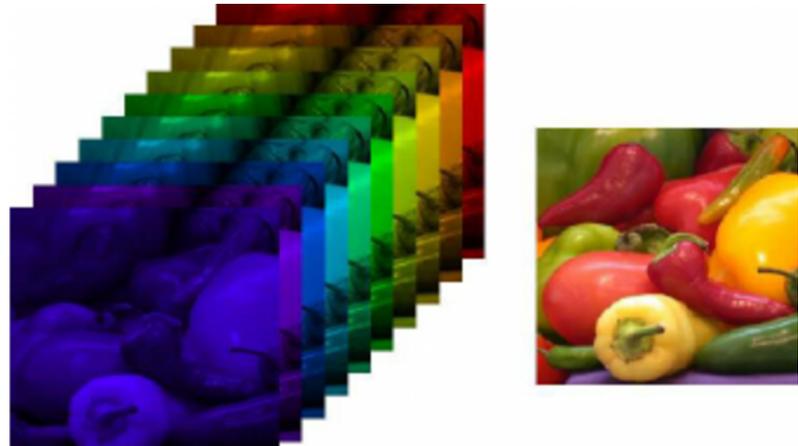


“Ground truth” object:

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

100 x 100 pix. x 30 spectral channels

Example #2: Optimized color filter for a grayscale camera



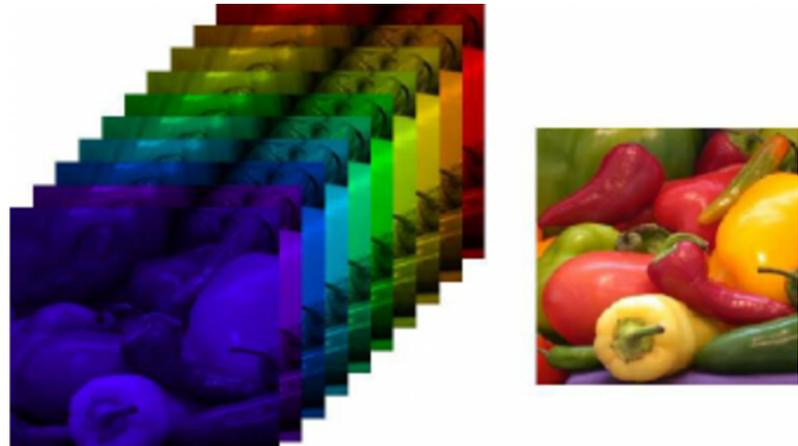
“Ground truth” object:

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

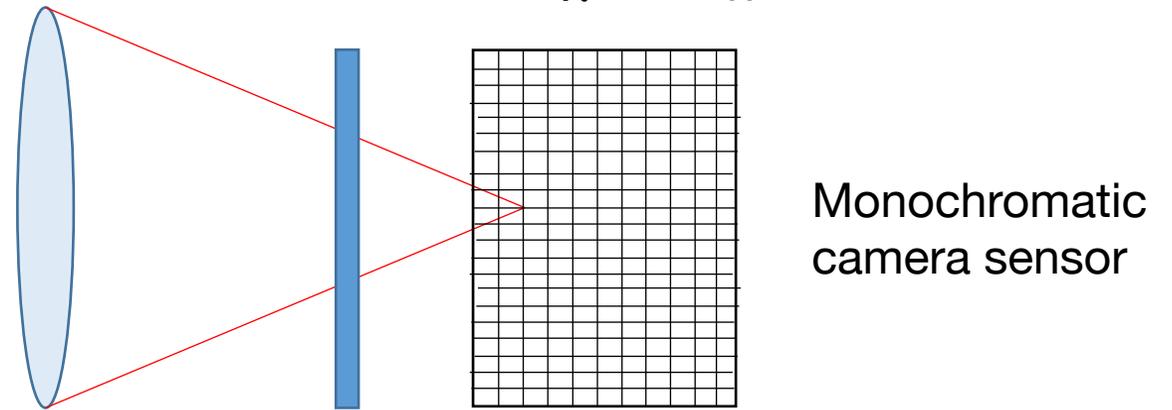
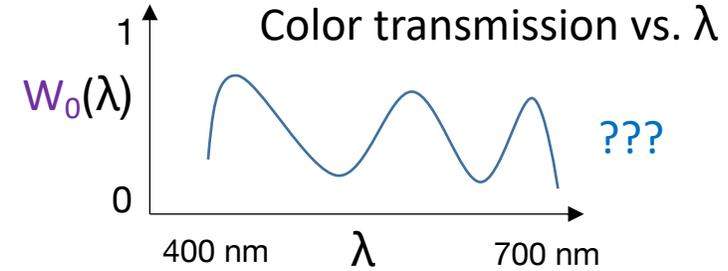
100 x 100 pix. x 30 spectral channels

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

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



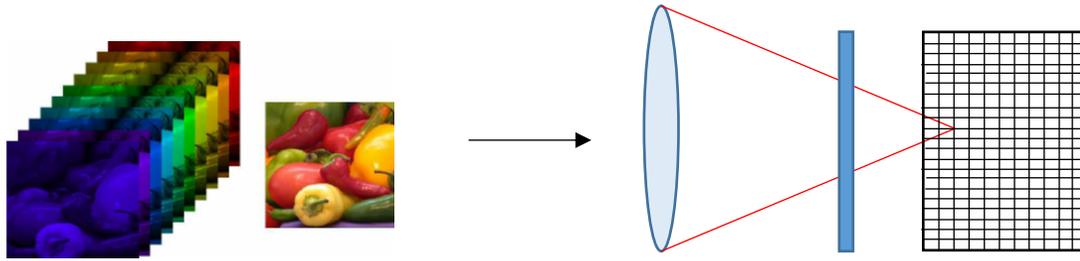
Training data:

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

I_0 : 100 x 100 x 30

Label y : 1x3

Example #2: Optimized color filter for a grayscale camera



Training data:

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

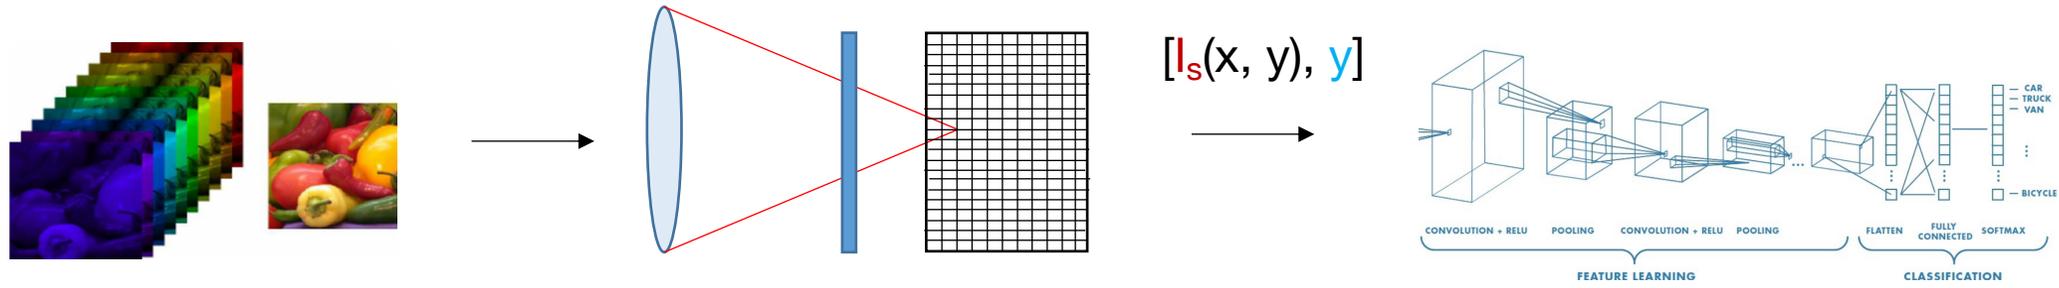
$I_0: 100 \times 100 \times 30$

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

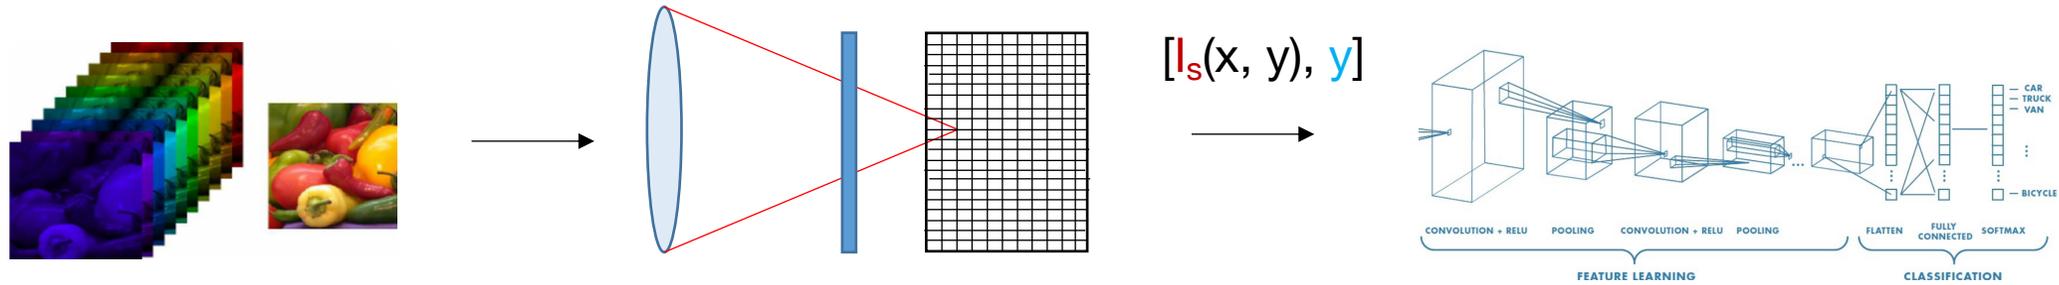
Label y : 1x3

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

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

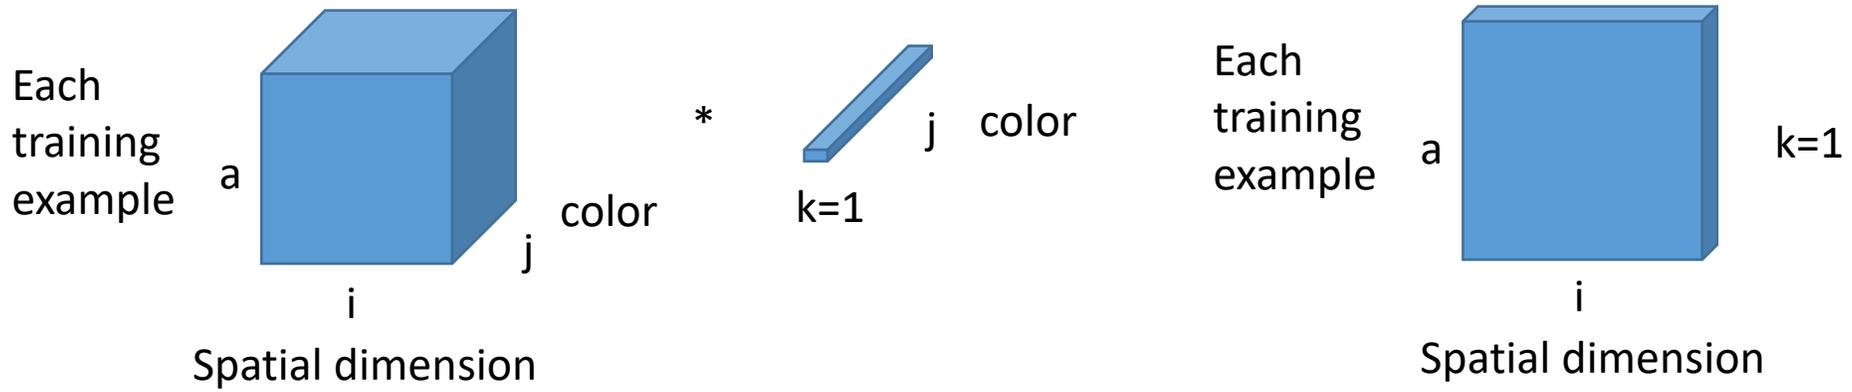
Tensorflow: operations to sum along 3rd (or higher) dimension

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`

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

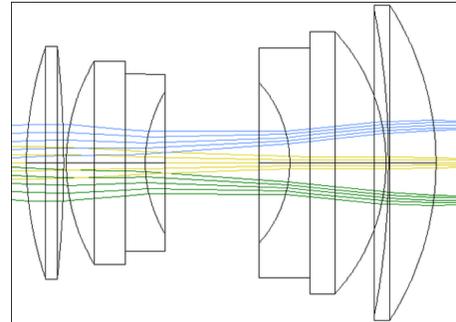
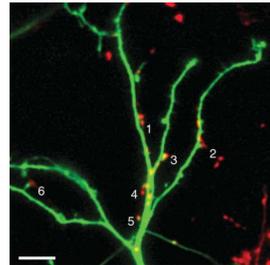
$$I_d(x,y) = \mathbf{D} I_o(x,y)$$

- 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?

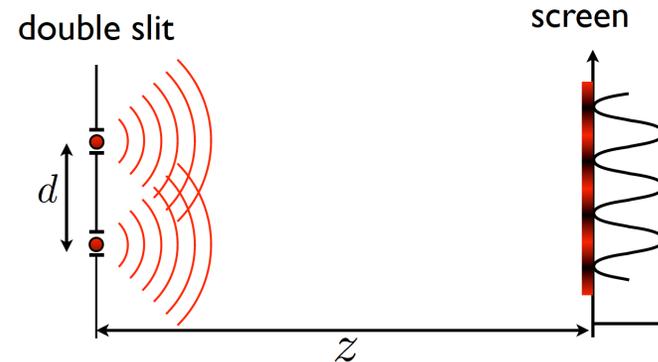
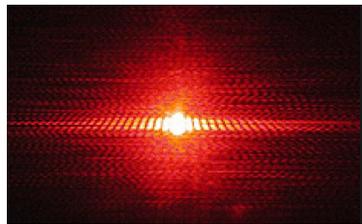
- Interpretation #1: Radiation (*Incoherent*)
- Model: Rays



- 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

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

Next class: Modeling coherent radiation as a wave