

EECS 504: Foundations of Computer Vision

Andrew Owens

Course staff



Haozhu Wang
Graduate student (GSI)



Anthony Liang
Instructional aide (IA)



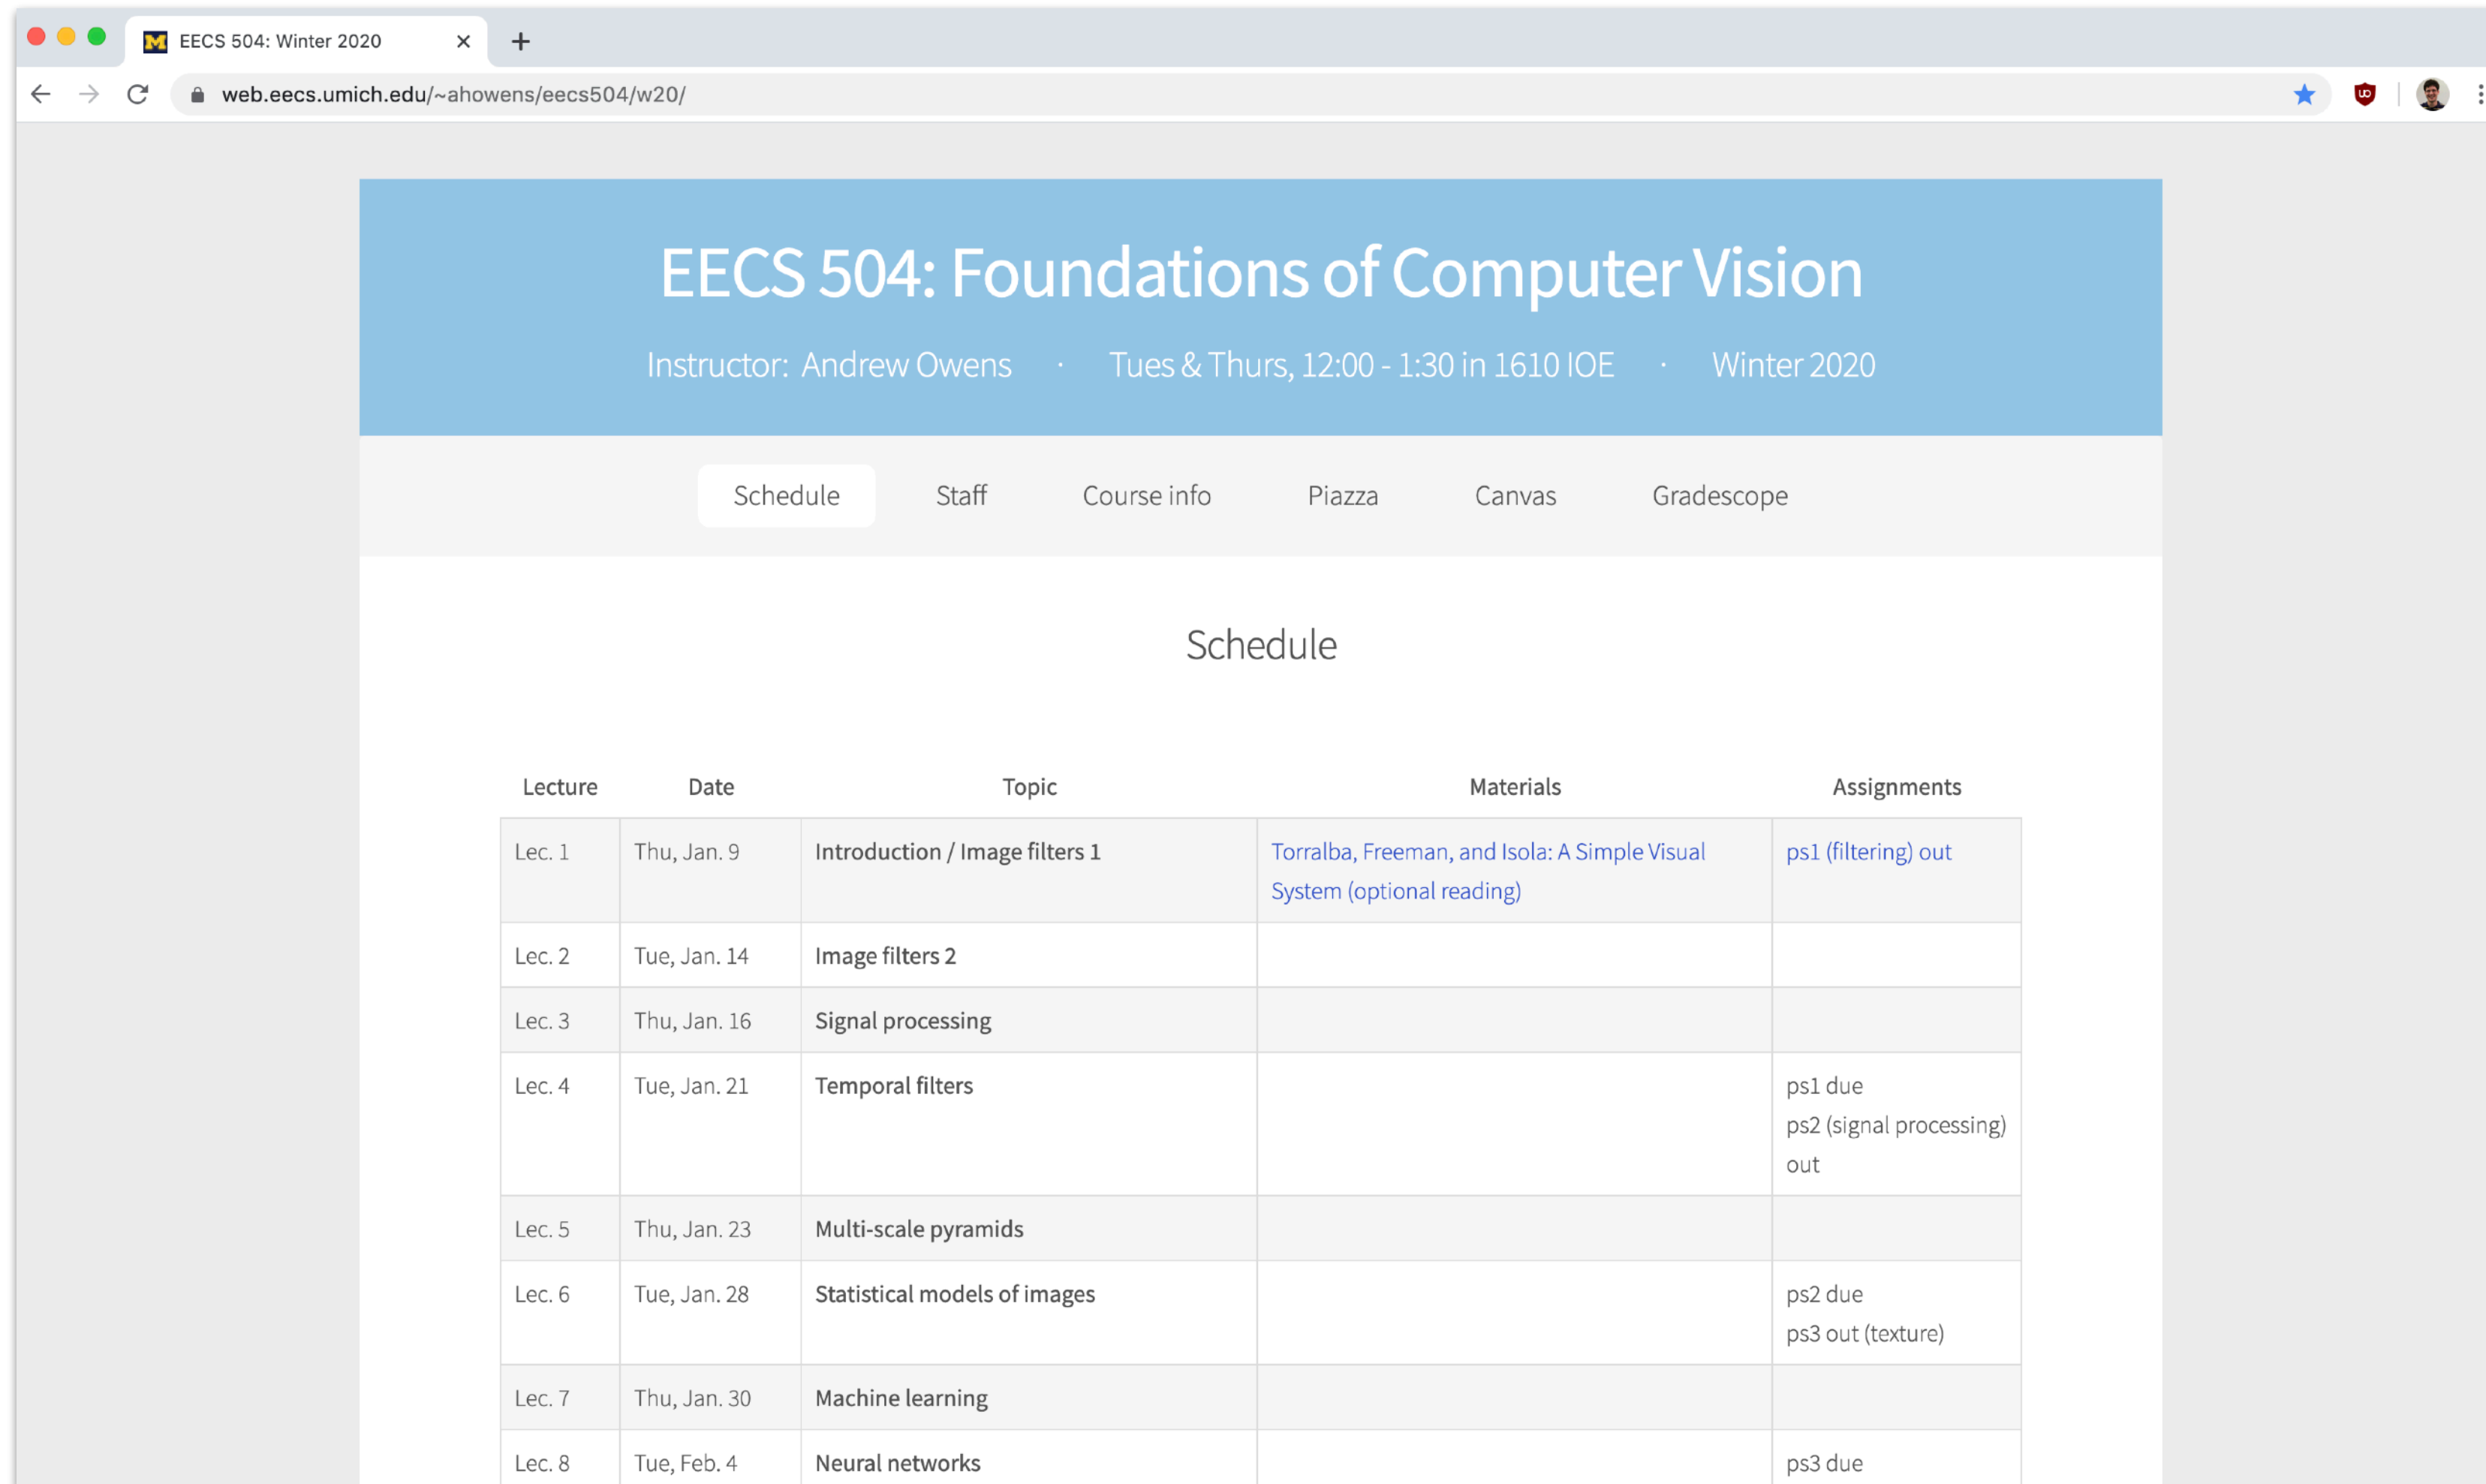
Bingqi Sun
Instructional aide (IA)

Interacting with us

- Ask questions on Piazza.
 - Sign up: <https://bit.ly/36mfYeP>
- Submit written work to Gradescope
- Office hours on website:

Name	Office hour times	Location
Andrew Owens	Fri. 3:00pm - 4:00pm	EECS 4231
Haozhu Wang	Thu. 7:30pm - 8:30pm	EECS 3312
Bingqi Sun	Mon. 12:00 - 1:00	EECS 3312
Anthony Liang	Mon. 1:00 - 2:00pm	EECS 3312

Course website



The screenshot shows a web browser window with the address bar displaying `web.eecs.umich.edu/~ahowens/eecs504/w20/`. The page features a blue header with the course title "EECS 504: Foundations of Computer Vision" and instructor information "Instructor: Andrew Owens · Tues & Thurs, 12:00 - 1:30 in 1610 IOE · Winter 2020". Below the header is a navigation bar with links for "Schedule", "Staff", "Course info", "Piazza", "Canvas", and "Gradescope". The "Schedule" link is highlighted, and the page content displays a table titled "Schedule".

Lecture	Date	Topic	Materials	Assignments
Lec. 1	Thu, Jan. 9	Introduction / Image filters 1	Torralba, Freeman, and Isola: A Simple Visual System (optional reading)	ps1 (filtering) out
Lec. 2	Tue, Jan. 14	Image filters 2		
Lec. 3	Thu, Jan. 16	Signal processing		
Lec. 4	Tue, Jan. 21	Temporal filters		ps1 due ps2 (signal processing) out
Lec. 5	Thu, Jan. 23	Multi-scale pyramids		
Lec. 6	Tue, Jan. 28	Statistical models of images		ps2 due ps3 out (texture)
Lec. 7	Thu, Jan. 30	Machine learning		
Lec. 8	Tue, Feb. 4	Neural networks		ps3 due

<https://web.eecs.umich.edu/~ahowens/eecs504/w20/>

Grading

- Assignments (70%)
- Final project (30%)

Assignments

- Weekly homework assignments (≈ 10 total)
- Due each Tuesday at midnight
- Late submissions penalized 30% per day
 - You have 5 "late days"
- Assignments should be done independently.
 - Encouraged to discuss them
 - Programming and writing should all be yours

Assignments

- Mix of programming and written problems
- Python + numerical computing libraries (numpy, scipy, etc.)
- PyTorch for deep learning
- Linear algebra and multivariable calculus
- Jupyter notebooks and Google Colab for problem sets

Assignments

EECS 504 PS1: Filtering - Colab

colab.research.google.com/drive/1hXN9fx8DaVm4BM4lm2vUj9Ut2_TMdyC5#forceEdit=true&sandboxMode=true&scrollTo=fdVJelRFvw_I

EECS 504 PS1: Filtering

File Edit View Insert Runtime Tools Help Cannot save changes

+ Code + Text Copy to Drive

Table of contents Code snippets Files

EECS 504 PS1: Filtering

Introduction

Starting

Problem 1.2: Pet edge detection (a)

Problem 1.2: Pet edge detection (c)

Problem 1.2: Pet edge detection (d)

Problem 1.2 Pet edge detection (e)

Section

EECS 504 PS1: Filtering

Please provide the following information (e.g. Andrew Owens, ahowens):

[Your first name] [Your last name], [Your UMich uniqname]

Introduction

We'll provide you with starter code, like this, in a Jupyter notebook for most problem sets. Please fill in the code to complete the assignment, and submit your notebook to Canvas as a .ipynb file. You can, of course, initially write your code offline in an editor like Emacs or Vim – we'd just like the final output to be in a notebook format to make grading more consistent.

Please note that we won't run your code. The notebook you submit should already contain all of the results we ask for. In particular, the visualizations of edge responses and blurred images should be computed before you submit. If you'd like to preview what your notebook will look like when we grade it, follow the directions here for converting the .ipynb notebook into an HTML file.

Starting

Run the following code to import the modules you'll need, and to download the images.

```
[4] import numpy as np, matplotlib as mpl, matplotlib.pyplot as plt, urllib, os
import scipy.ndimage # For image filtering
import imageio # For loading images

# Download the images that you'll need
base_url = 'https://web.eecs.umich.edu/~ahowens/eecs504/w20/psets/ps1/ims'
for name in ['dog-1.jpg', 'dog-2.jpg', 'apple.jpg']:
    with open(name, 'wb') as out:
        url = os.path.join(base_url, name)
        out.write(urllib.request.urlopen(url).read())
```

Problem 1.2: Pet edge detection (a)

```
# You can upload images yourself or load them from URLs
im = imageio.imread('dog-1.jpg')
# Convert to grayscale. We'll use floats in [0, 1].
im = im.mean(2)/255.

# Your code here!

# Visualize edge maps using matplotlib
plt.figure()
plt.title('Input image')
plt.axis('off')
plt.imshow(im, cmap = 'gray', vmin = 0, vmax = 1)

# Convolve the image with horizontal and vertical gradient filters
Ix = convolve(im, dx)
Iy = convolve(im, dy)
edges = Ix**2. + Iy**2.

plt.figure()
plt.axis('off')
```

Table of contents Code snippets Files

EECS 504 PS1: Filtering

Introduction

Starting

Problem 1.2: Pet edge detection (a)

Problem 1.2: Pet edge detection (c)

Problem 1.2: Pet edge detection (d)

Problem 1.2 Pet edge detection (e)

Section

Iy = convolve(im, dy)
edges = Ix**2. + Iy**2.

Visualize edge maps using matplotlib
plt.figure()
plt.title('Input image')
plt.axis('off')
plt.imshow(im, cmap = 'gray', vmin = 0, vmax = 1)

plt.figure()
plt.axis('off')
plt.title('Ix')
plt.imshow(Ix)

plt.figure()
plt.title('Iy')
plt.axis('off')
plt.imshow(Iy)

plt.figure()
plt.title('Edges')
plt.axis('off')
Please visualize edge responses using this range of values.
plt.imshow(edges, vmin = 0., vmax = np.percentile(edges, 99))

<matplotlib.image.AxesImage at 0x7f3ab8d60160>


Input image

Ix


Project

- Open-ended! Example projects:
 - Implement and extend a recent computer vision paper
 - Use computer vision in your research
 - We'll also provide a list of project ideas
- Work in small groups (up to 4 people)
- Complete in last month of class.
 - Project proposal (after spring break)
 - Short presentation (finals period)
 - Writeup (finals period)

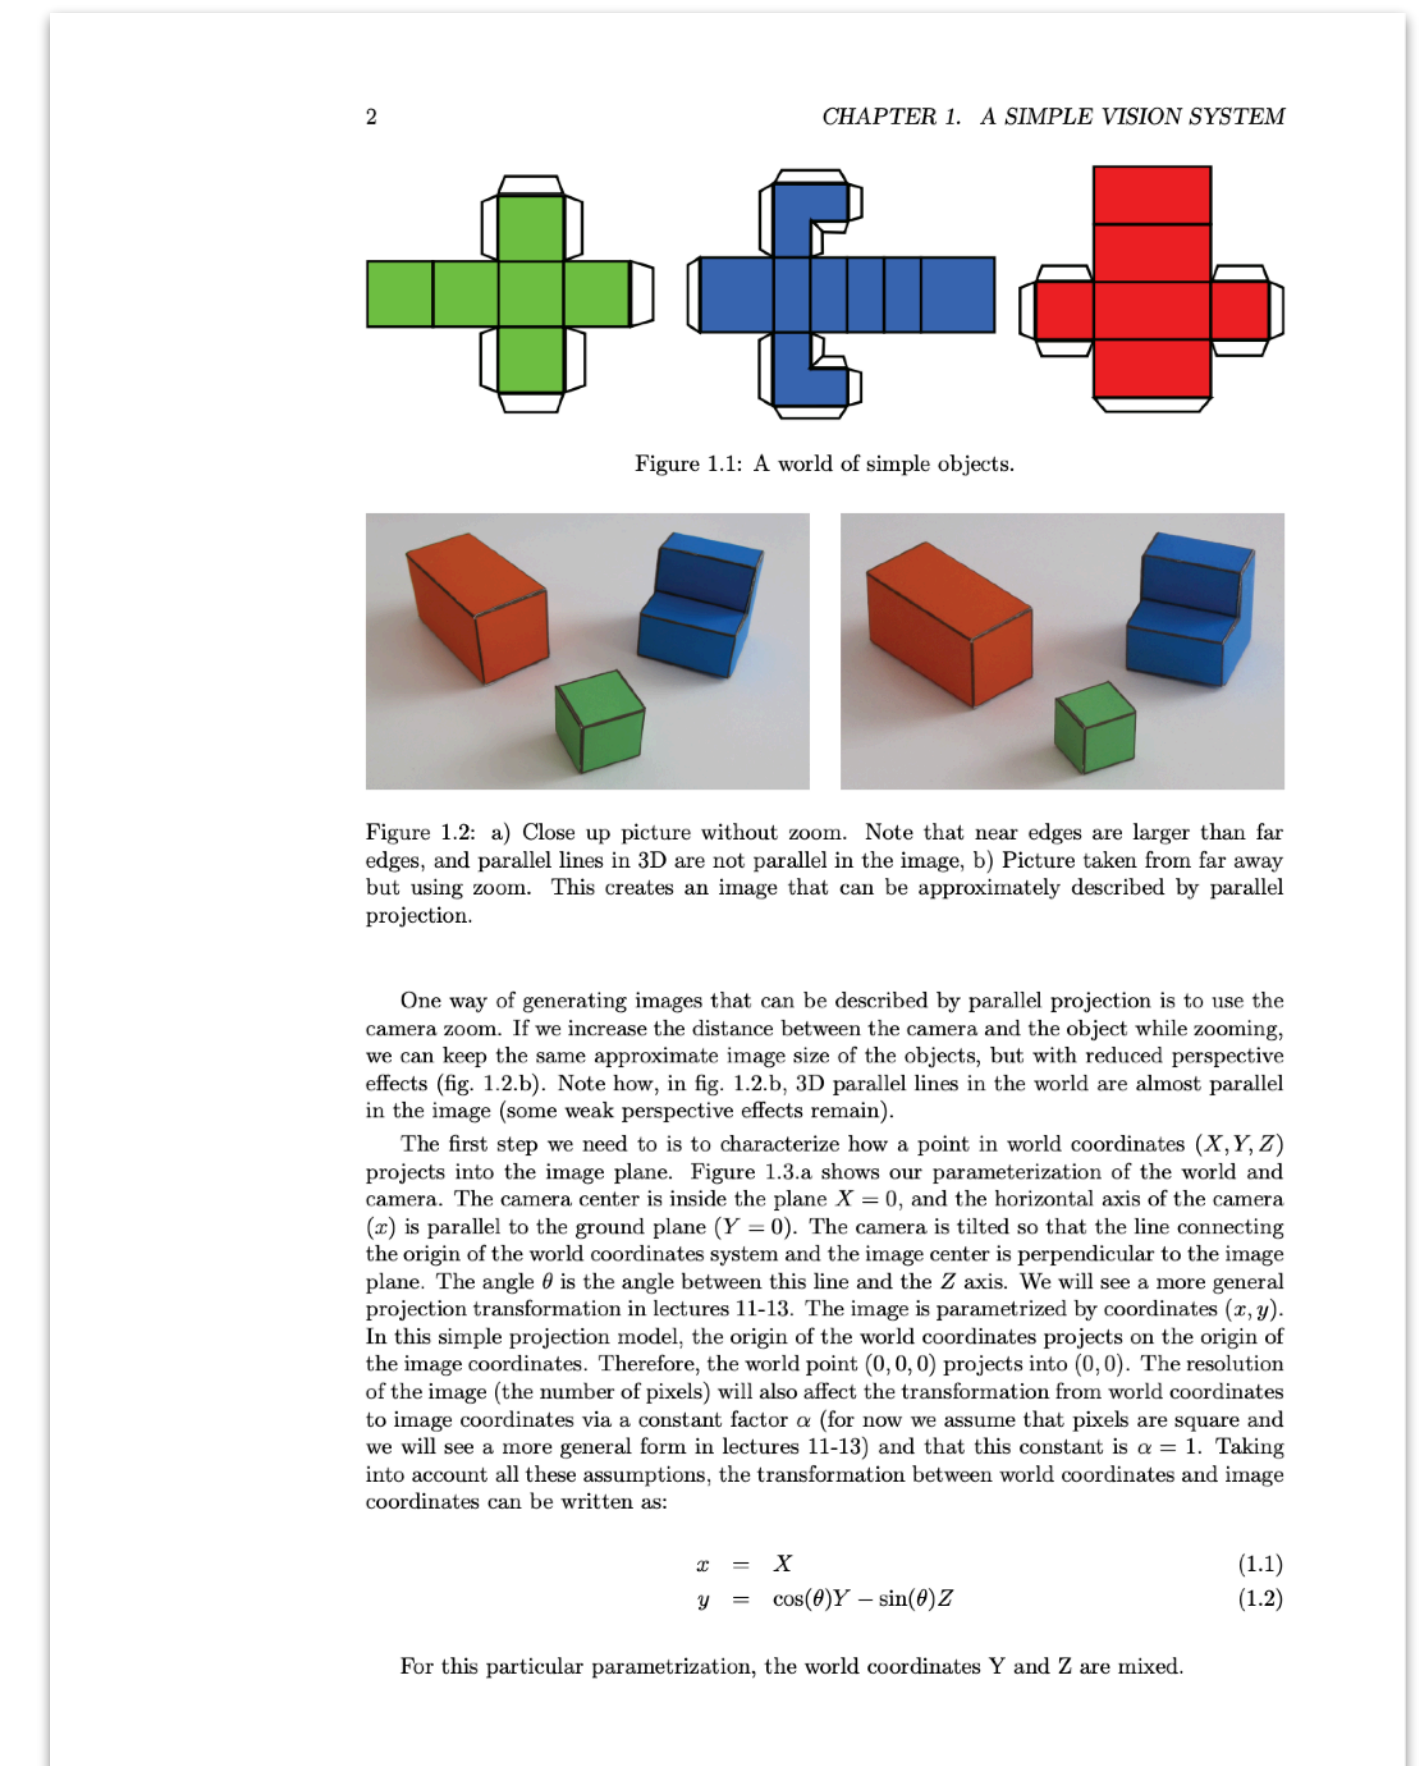
Readings



<http://szeliski.org/Book>



<https://www.deeplearningbook.org>



Manuscript chapters by Torralba, Freeman, and Isola (on course website). Class based on this coursework.

And also occasional paper readings

Class topics

Homework problem:

Lec. 1	Thu, Jan. 9	Introduction / Image filters 1
Lec. 2	Tue, Jan. 14	Image filters 2
Lec. 3	Thu, Jan. 16	Signal processing
Lec. 4	Tue, Jan. 21	Temporal filters
Lec. 5	Thu, Jan. 23	Multi-scale pyramids
Lec. 6	Tue, Jan. 28	Statistical models of images
Lec. 7	Thu, Jan. 30	Machine learning
Lec. 8	Tue, Feb. 4	Neural networks
Lec. 9	Thu, Feb. 6	Optimization
Lec. 10	Tue, Feb. 11	Convolutional networks
Tutorial	Tue, Feb. 11	PyTorch tutorial


■ ■ ■

} Signal



Lec. 1	Thu, Jan. 9	Introduction / Image filters 1
Lec. 2	Tue, Jan. 14	Image filters 2
Lec. 3	Thu, Jan. 16	Signal processing
Lec. 4	Tue, Jan. 21	Temporal filters
Lec. 5	Thu, Jan. 23	Multi-scale pyramids
Lec. 6	Tue, Jan. 28	Statistical models of images
Lec. 7	Thu, Jan. 30	Machine learning
Lec. 8	Tue, Feb. 4	Neural networks
Lec. 9	Thu, Feb. 6	Optimization
Lec. 10	Tue, Feb. 11	Convolutional networks
Tutorial	Tue, Feb. 11	PyTorch tutorial

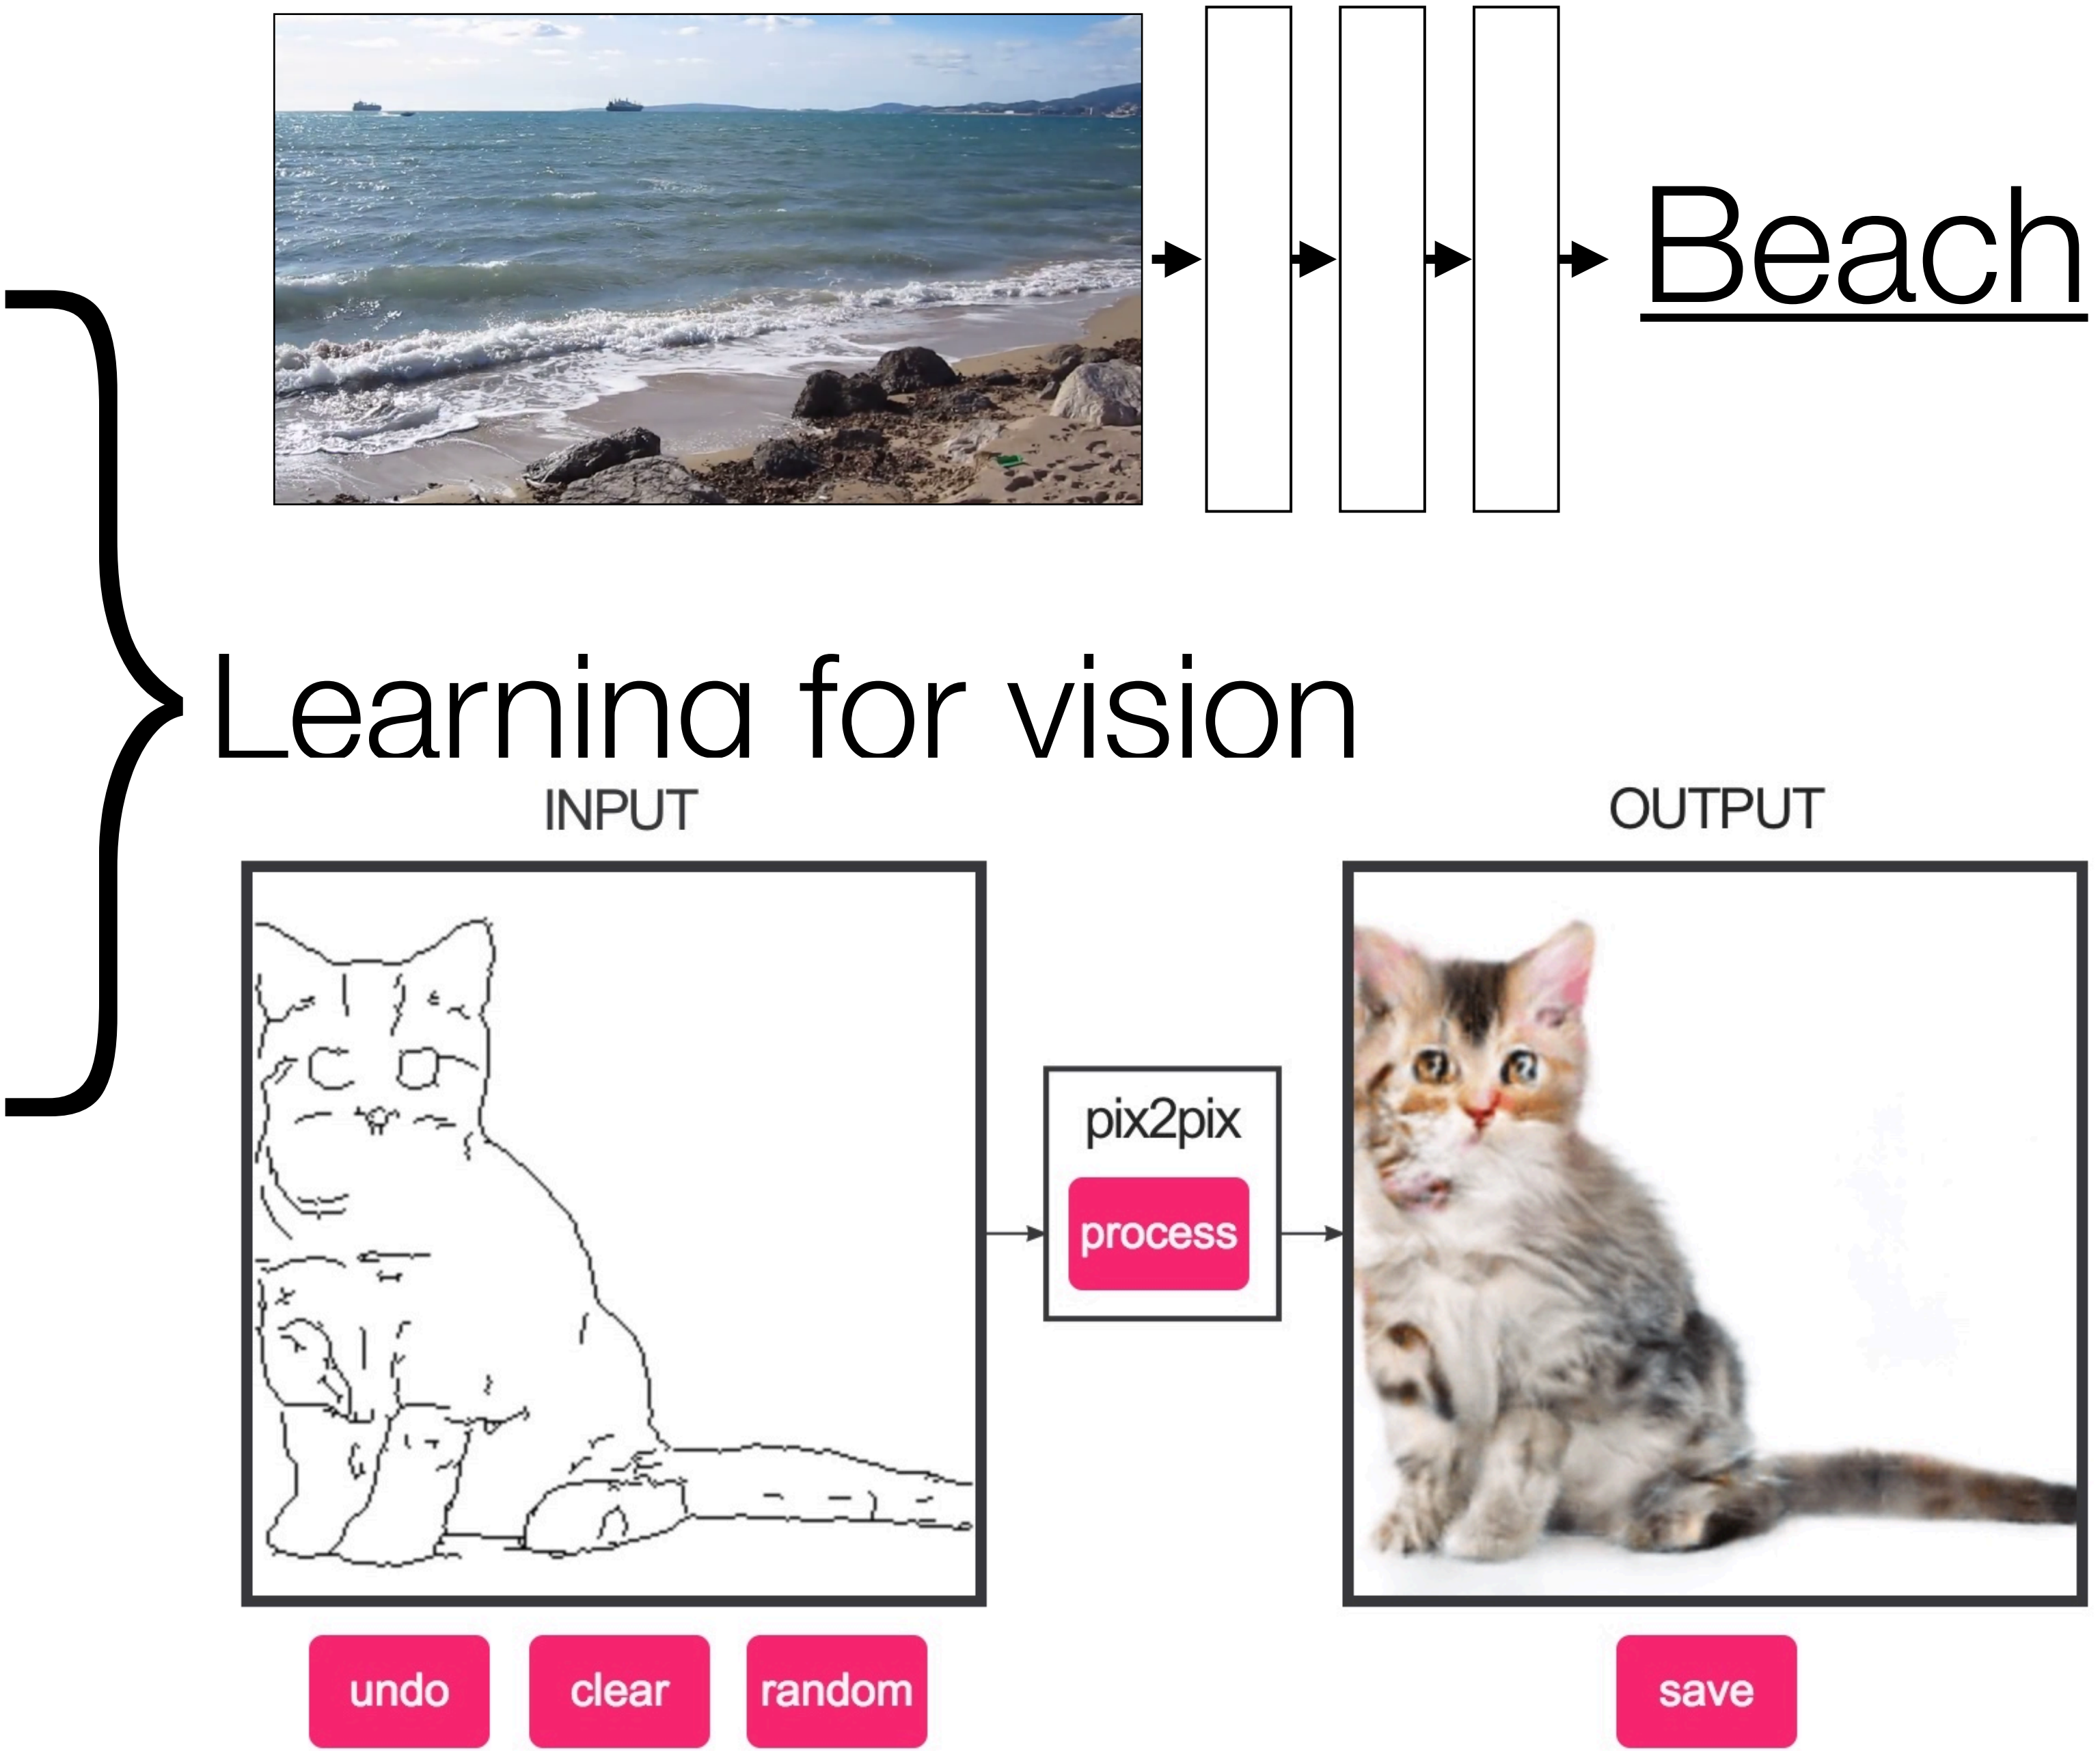
■ ■ ■



Intro to
deep learning

Tutorial	Tue, Feb. 11	PyTorch tutorial
Lec. 11	Thu, Feb. 13	Scene understanding
Lec. 12	Tue, Feb. 18	Object detection
Lec. 13	Thu, Feb. 20	Action recognition
Lec. 14	Tue, Feb. 25	GANs
Lec. 15	Thu, Feb. 27	Image synthesis
Spring break		
Lec. 16	Tue, Mar. 10	Representation learning
Lec. 17	Thu, Mar. 12	Sight, sound, and touch
Lec. 18	Tue, Mar. 17	Optical flow
Lec. 19	Thu, Mar. 19	Multi-view geometry


...



Homework problem:

Lec. 17	Thu, Mar. 12	Sight, sound, and touch
Lec. 18	Tue, Mar. 17	Optical flow
Lec. 19	Thu, Mar. 19	Multi-view geometry
Lec. 20	Tue, Mar. 24	Structure from motion
Lec. 21	Thu, Mar. 26	Depth estimation
Lec. 22	Tue, Mar. 31	Graphical models
Lec. 23	Thu, Apr. 2	Color
Lec. 24	Tue, Apr. 7	Embodied vision
Lec. 25	Thu, Apr. 9	Modern CNNs
Lec. 26	Tue, Apr. 14	Image forensics
Lec. 27	Thu, Apr. 16	Language and vision

■ ■ ■



Can
optical



Lec. 19	Thu, Mar. 19	Multi-view geometry
Lec. 20	Tue, Mar. 24	Structure from motion
Lec. 21	Thu, Mar. 26	Depth estimation
Lec. 22	Tue, Mar. 31	Graphical models
Lec. 23	Thu, Apr. 2	Color
Lec. 24	Tue, Apr. 7	Embodied vision
Lec. 25	Thu, Apr. 9	Modern CNNs
Lec. 26	Tue, Apr. 14	Image forensics
Lec. 27	Thu, Apr. 16	Language and vision
Lec. 28	Tue, Apr. 21	Datasets and bias


} Advanced topics
and applications

Today


1. A bit of vision history
2. Why vision is hard
3. A simple visual system

Exciting times for computer vision

Robotics



Medical applications



3D modeling



Driving



Mobile devices




Accessibility



To see

“What does it mean, to see? The plain man's answer (and Aristotle's, too) would be, to know what is where by looking.”

To discover from images what is present in the world, where things are, what actions are taking place, to predict and anticipate events in the world.



MASSACHUSETTS INSTITUTE OF TECHNOLOGY
PROJECT MAC

Artificial Intelligence Group
Vision Memo. No. 100.


July 7, 1966

THE SUMMER VISION PROJECT



Seymour Papert.


The summer vision project is an attempt to use our summer workers effectively in the construction of a significant part of a visual system. The particular task was chosen partly because it can be segmented into sub-problems which will allow individuals to work independently and yet participate in the construction of a system complex enough to be a real landmark in the development of "pattern recognition".

Just a few years ago...



["HOGgles", Vondrick et al. , ICCV 2013]






["GauGAN", Park et al., CVPR 2019]

Different signals, same methods

Sound




WiFi



(Zhao et al. 2019)

Touch



(Calandra et al. 2018)

Input video



(Owens and Efros 2018)

On-screen audio



(Owens and Efros 2018)


Off-screen audio



(Owens and Efros 2018)

What makes vision hard?


To see: perception vs. measurement



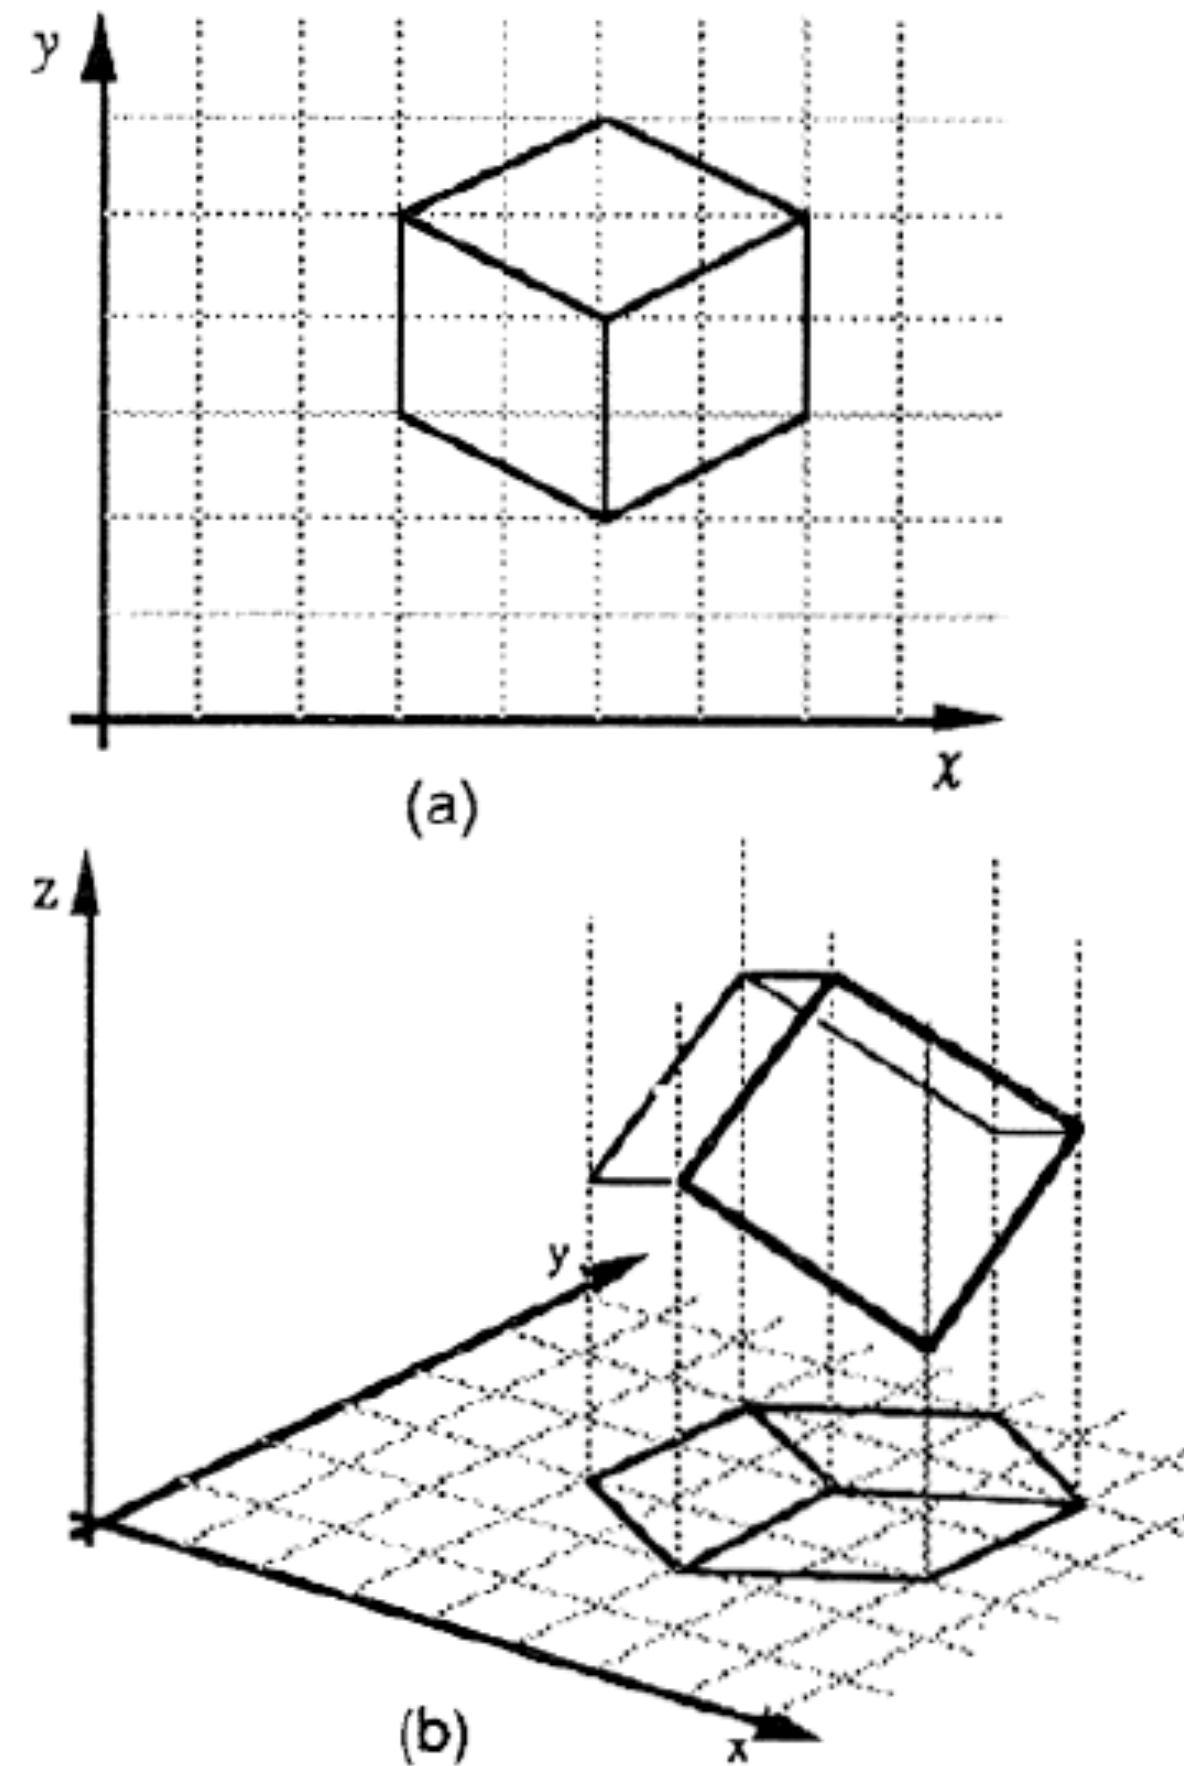
Edward H. Adelson

Slide credit: Torralba, Freeman, Isola

To see: perception vs. measurement



Other ambiguities



Sinha & Adelson 93

Slide credit: Antonio Torralba

Other ambiguities




Figure 1. (a) A line drawing provides information only about the x, y coordinates of points lying along the object contours. (b) The human visual system is usually able to reconstruct an object in three dimensions given only a single 2D projection (c) Any planar line-drawing is geometrically consistent with infinitely many 3D structures.

Sinha & Adelson 93

A simple visual system

- A simple world
- A simple image formation model
- A simple goal

A simple world


MACHINE PERCEPTION OF THREE-DIMENSIONAL SOLIDS

by

LAWRENCE GILMAN ROBERTS


Submitted to the Department of Electrical Engineering
on May 10, 1963, in partial fulfillment of the requirements for the degree of Doctor of Philosophy.

The problem of machine recognition of pictorial data has long been a challenging goal, but has seldom been attempted with anything more complex than alphabetic characters. Many people have felt that research on character recognition would be a first step, leading the way to a more general pattern recognition system. However, the multitudinous attempts at character recognition, including my own, have not led very far. The reason, I feel, is that the study of abstract, two-dimensional forms leads us away from, not toward, the techniques necessary for the recognition of three-dimensional objects. The per-



Complete Convex Polygons. The polygon selection procedure would select the numbered polygons as complete and convex. The number indicates the probable number of sides. A polygon is incomplete if one of its points is a collinear joint of another polygon.


A simple world




A simple image formation model

Simple world rules:

- Surfaces can be horizontal or vertical.
- Objects will be resting on a white horizontal ground plane





A simple image formation model



A simple image formation model

Image and projection of the world coordinate axes into the image plane




World coordinates

$$\begin{aligned} x &= X + x_0 \\ y &= \cos(\theta) Y - \sin(\theta) Z + y_0 \end{aligned}$$

image coordinates


A simple goal

Recover the 3D structure of the world





We want to recover $X(x,y)$, $Y(x,y)$, $Z(x,y)$ using as input $I(x,y)$

Edges



Treating the image as a function



Finding edges in the image




Image gradient:

$$\nabla \mathbf{I} = \left(\frac{\partial \mathbf{I}}{\partial x}, \frac{\partial \mathbf{I}}{\partial y} \right)$$

Approximation image derivative:

$$\frac{\partial \mathbf{I}}{\partial x} \simeq \mathbf{I}(x, y) - \mathbf{I}(x - 1, y)$$

Edge strength

$$E(x, y) = |\nabla \mathbf{I}(x, y)|$$



Edge orientation:

$$\theta(x, y) = \angle \nabla \mathbf{I} = \arctan \frac{\partial \mathbf{I} / \partial y}{\partial \mathbf{I} / \partial x}$$


Edge normal:

$$\mathbf{n} = \frac{\nabla \mathbf{I}}{|\nabla \mathbf{I}|}$$

Finding edges in the image






$$\nabla I = \left(\frac{\partial I}{\partial x}, \frac{\partial I}{\partial y} \right) \quad \mathbf{n} = \frac{\nabla I}{|\nabla I|}$$




$E(x,y)$ and $n(x,y)$

Edge classification

- Figure/ground segmentation
 - Using the fact that objects have color
- Occlusion edges
 - Occlusion edges are owned by the foreground
- Contact edges




From edges to surface constraints




From edges to surface constraints

- Ground



$Y(x,y) = 0$ if (x,y) belongs to a ground pixel

- Contact edge



$Y(x,y) = 0$ if (x,y) belongs to foreground and is a contact edge


- What happens inside the objects?

... now things get a bit more complicated.

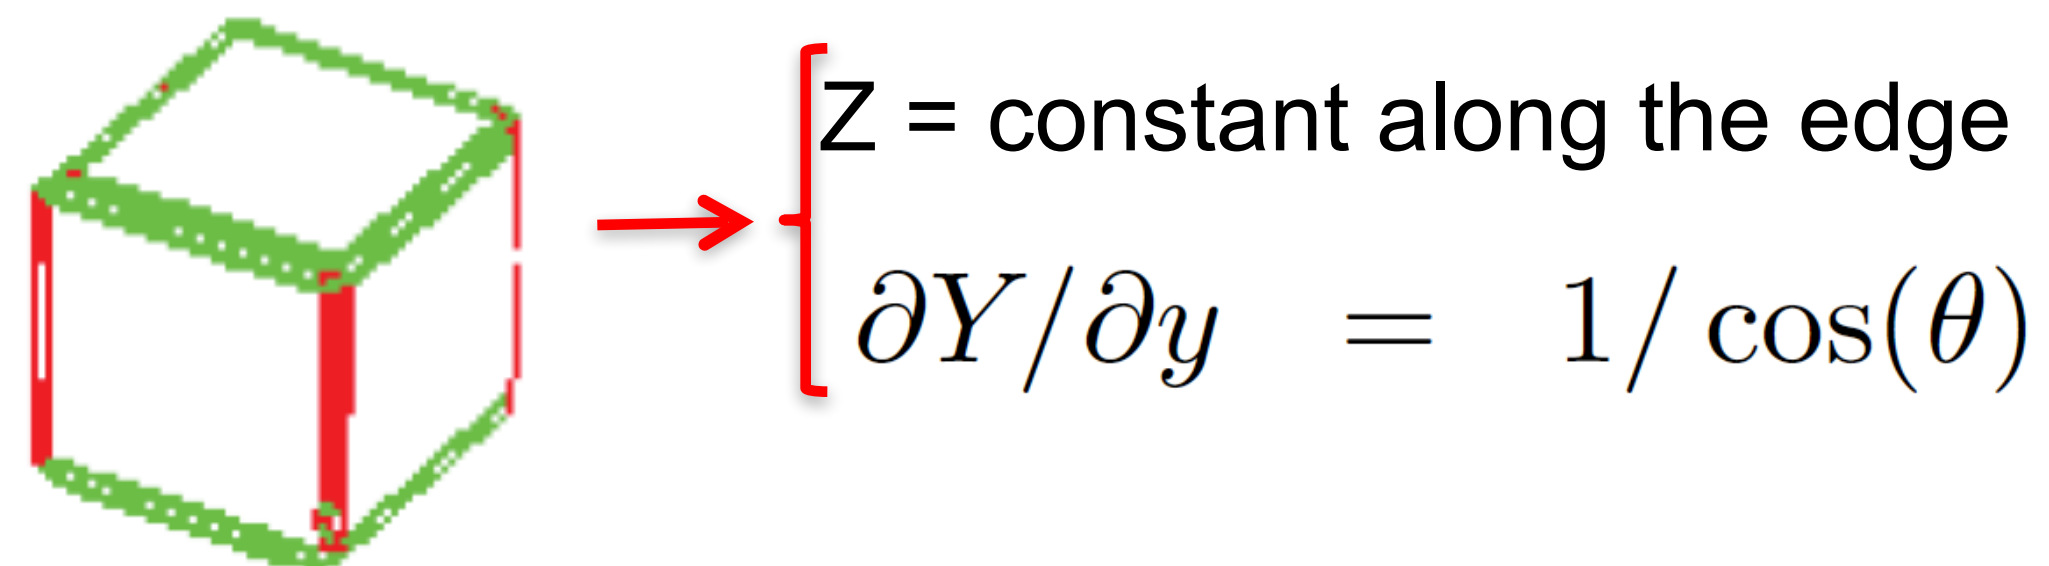
From edges to surface constraints

How can we relate the information in the pixels with 3D surfaces in the world?

Vertical edges




Given the image, what can we say about X , Y and Z in the pixels that belong to a vertical edge?




From edges to surface constraints

- Horizontal edges



Given the image, what can we say about X , Y and Z in the pixels that belong to an horizontal 3D edge?



$Y = \text{constant along the edge}$

$$\partial Y / \partial \mathbf{t} = 0$$


Where \mathbf{t} is the vector parallel to the edge

$$\mathbf{t} = (-n_y, n_x)$$

$$\partial Y / \partial \mathbf{t} = -n_y \partial Y / \partial x + n_x \partial Y / \partial y$$

From edges to surface constraints

- What happens where there are no edges?



Assumption of planar faces:

$$\begin{aligned}\partial^2 Y / \partial x^2 &= 0 \\ \partial^2 Y / \partial y^2 &= 0 \\ \partial^2 Y / \partial y \partial x &= 0\end{aligned}$$

Information has to be propagated from the edges

A simple inference scheme

All the constraints are linear!

$$Y(x,y) = 0$$

if (x,y) belongs to a ground pixel

$$\partial Y / \partial y = 1 / \cos(\theta)$$

if (x,y) belongs to a vertical edge

$$\partial Y / \partial t = 0$$

if (x,y) belongs to an horizontal edge

$$\partial^2 Y / \partial x^2 = 0$$

$$\partial^2 Y / \partial y^2 = 0$$

$$\partial^2 Y / \partial y \partial x = 0$$

if (x,y) is not on an edge

A similar set of constraints could be derived for Z

Discrete approximation

We can transform every differential constraint into a linear constraint on $Y(x,y)$


$Y(x,y)$

111	115	113	111	112	111	112	111
135	138	137	139	145	146	149	147
163	168	188	196	206	202	206	207
180	184	206	219	202	200	195	193
189	193	214	216	104	79	83	77
191	201	217	220	103	59	60	68
195	205	216	222	113	68	69	83
199	203	223	228	108	68	71	77

$$\frac{dY}{dx} \approx Y(x,y) - Y(x-1,y)$$

Discrete approximation


Transform the “image” $Y(x,y)$ into a column vector:



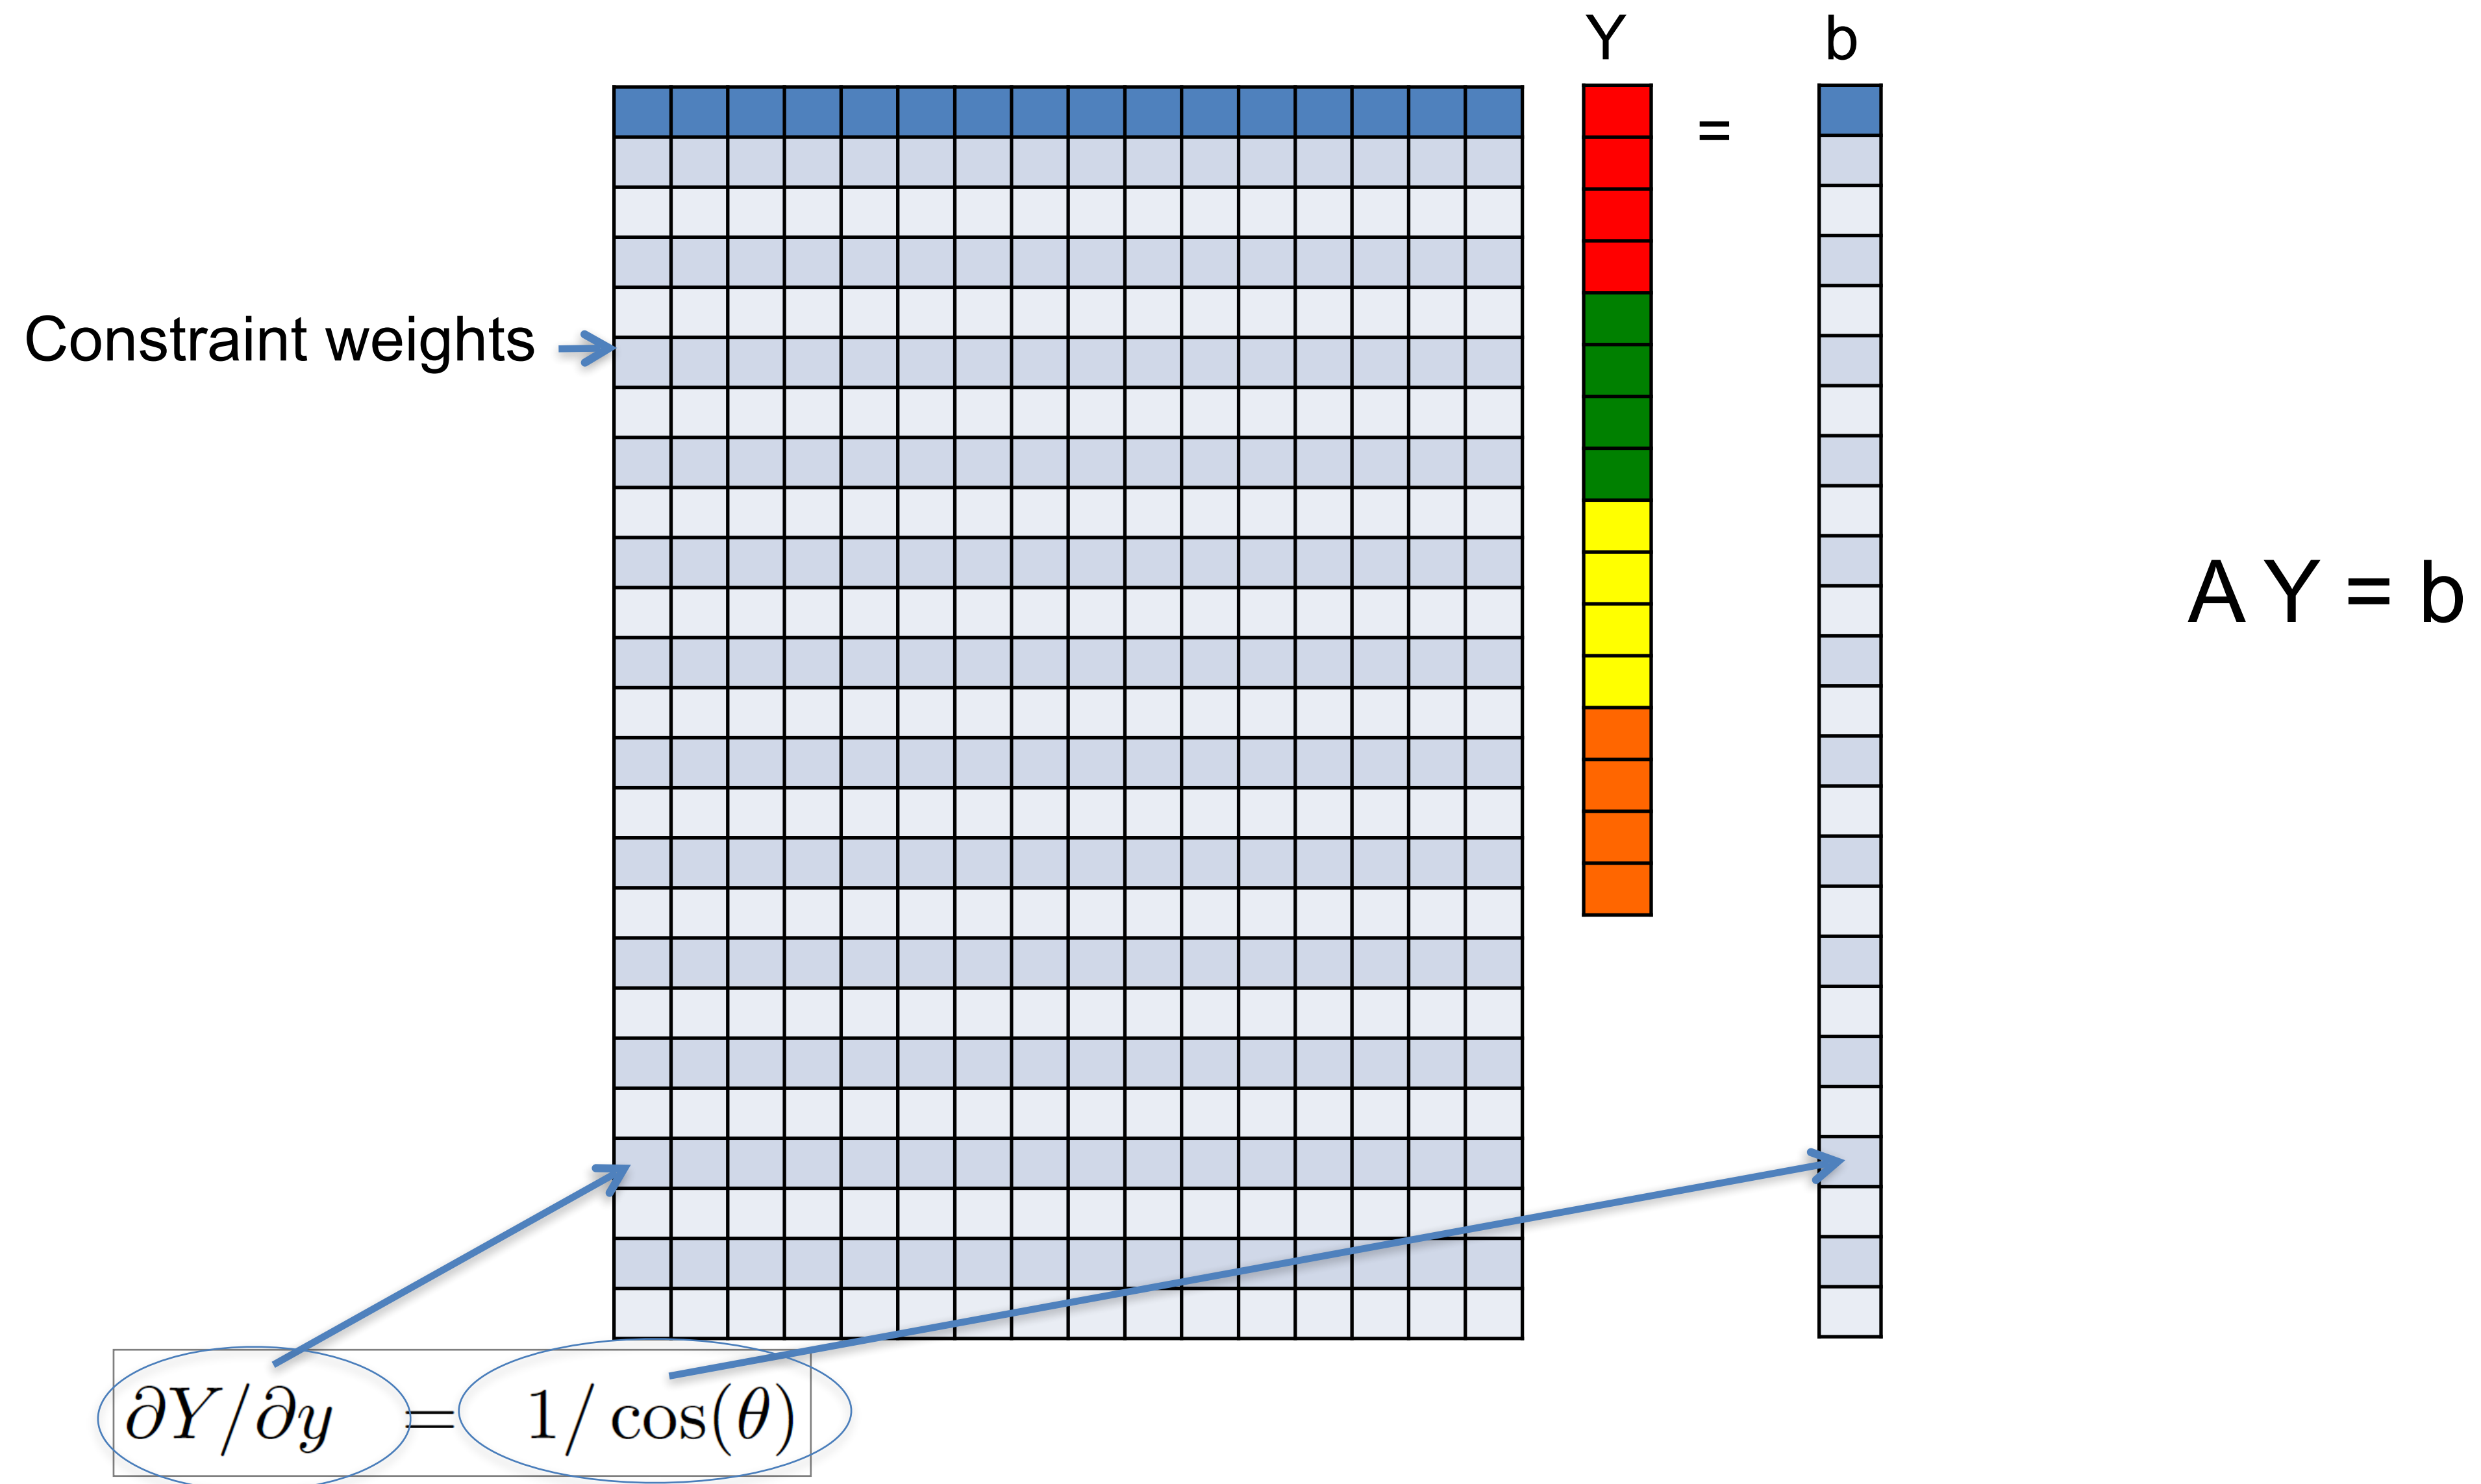
$x=2, y=2$

$$\frac{dY}{dx} \approx Y(x,y) - Y(x-1,y) = Y(2,2) - Y(1,2) =$$


0	0	0	0	0	-1	0	0	0	1	0	0	0	0	0	0
---	---	---	---	---	----	---	---	---	---	---	---	---	---	---	---




A simple inference scheme: solve for Y




Results




Edge normals




Edge strength





3D orientation



Contact edges




Depth discontinuities






Changing view point



Input







New view points:



Failure cases... even in a simple world!



Failure cases... even in a simple world!



Failure cases... even in a simple world!



Questions about the course?