

# Computer Vision I - *Introduction & Filtering*

Carsten Rother

28/10/2015

# Admin Stuff

- Language: German/English; Slides: English  
(all the terminology and books are in English)
- Lecturer: Carsten Rother (Eric Brachmann, Anita Sellent)
- Exercises: Dmitri Schlesinger, Eric Brachmann
- Staff Email: [dmytro.shlezinger@tu-dresden.de](mailto:dmytro.shlezinger@tu-dresden.de)
- Announcements: on our webpage
  
- Course Books:
  - Image Processing/Geometry:  
Computer Vision: Algorithms and Applications  
by Rick Szeliski; Springer 2011. An **earlier** version of the book is online:  
<http://szeliski.org/Book/>
  - Geometry:  
Multiple View Geometry; Hartley and Zisserman;  
Cambridge Press 2004. Second edition. Parts of book are online:  
<http://www.robots.ox.ac.uk/~vgg/hzbook/>
  - Also pointers to conference and journal articles

# Course Overview (total 14 lectures)

Lecture 1 (23.10): Introduction & Filtering

Ex1: Intro to OpenCV

Lecture 2 (30.10): Filtering and Feature detection

Ex2: Intro to exercise: Filtering

Lecture 3 (6.11): Image Matching and Projective Geometry

Ex3: homework

Lecture 4 (13.11): Geometry of One and Two Images

Ex4: homework

Lecture 5 (20.11): Robust Geometry Estimation

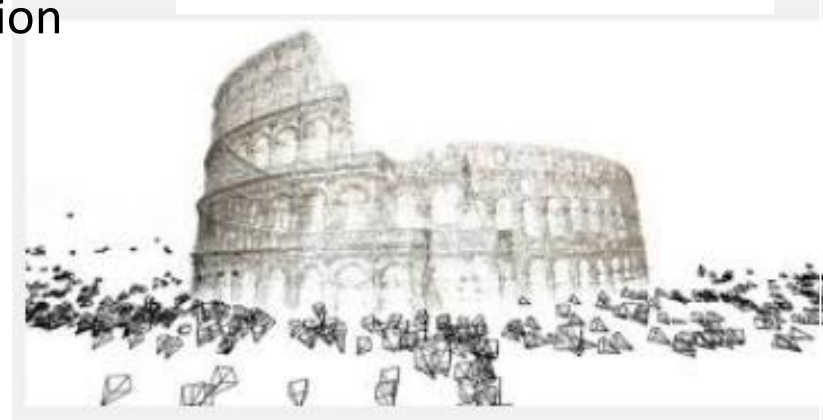
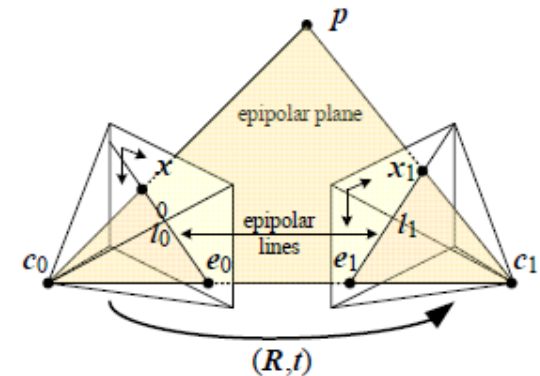
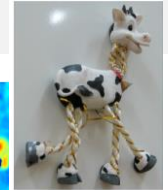
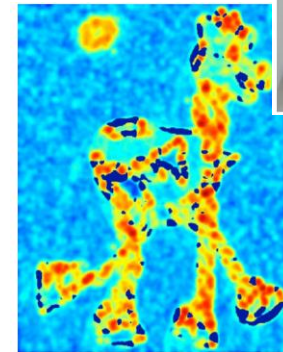
Ex5: Intro to exercise: Panoramic Stitching (Geometry)

Lecture 6 (27.11): Multi-View 3D Reconstruction

Ex6: homework

Lecture 7 (4.12): Object Pose estimation

Ex7: homework



# Course Overview (total 14 lectures)

Lecture 8 (11.12): Tracking – Part 1

Ex8: Intro to exercise: Object Pose Estimation

Lecture 9 (18.12): Tracking – Part 2

Ex9: homework

Lecture 10 (8.1): Dense Matching – Optical Flow

Ex10: homework

Lecture 11 (15.1): Dense Matching – Stereo Part 1

Ex11: homework

Lecture 12 (22.1): Dense Matching – Stereo Part 2

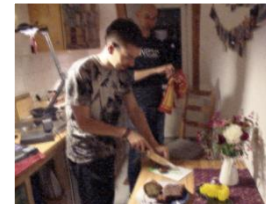
Ex12: homework

Lecture 13 (29.1): Dense Matching – Scene Flow

Ex13: Prepare Poster Session

Lecture 14 (5.2): Poster Session

Ex14: homework



time = 1



time = 2



# Exams and Exercises

- Exam: in Person
- Exercises/homework:
  - There are three blocks
  - In each block you have to do a certain amount of exercises
  - The exercises should to be handed in until end of semester (ideally after each block)
- Exam:
  - Exercises are not mandatory, i.e. you can sit the exam without having done the exercises
  - In the exam I may ask questions about the exact exercise you have done. If you have not done any exercise then this may result in a worse mark

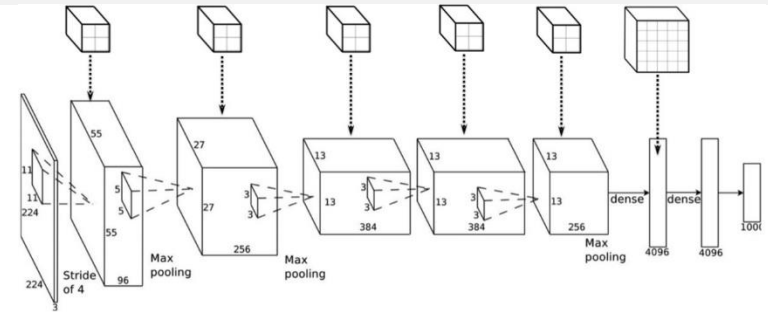
# CVLD Lectures

- WS 15/16

- Computer Vision 1 (2+2)
- Machine Learning 1 (2+2)

- SS 16

- Computer Vision 2 (2+2)
- Machine Learning 2 (2+2)
- Image Processing (1+1)



- For doing a Master/PhD in the CVLD one should do the **computer vision** or **machine learning** track
- Computer graphics (Prof. Gumhold) (Introduction, I, II)  
3D Scanning with structured light; Illumination models; Geometry



# Before we start ... some Advertisement



## Optimization & Learning

with a focus on graphical models, deep models, and large-scale optimization.



## Image Matching

with a focus on 3-9DoF scene flow, as well as jointly recovering multiple physical aspects



## Scene Understanding

with a focus on scene understanding, scene reconstruction, and scene segmentation.



## Bio Imaging

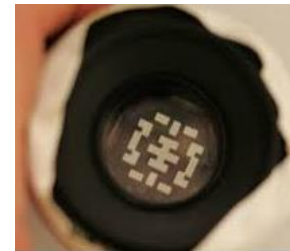
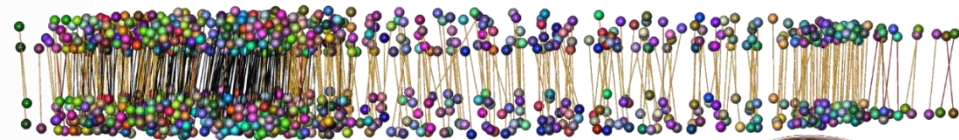
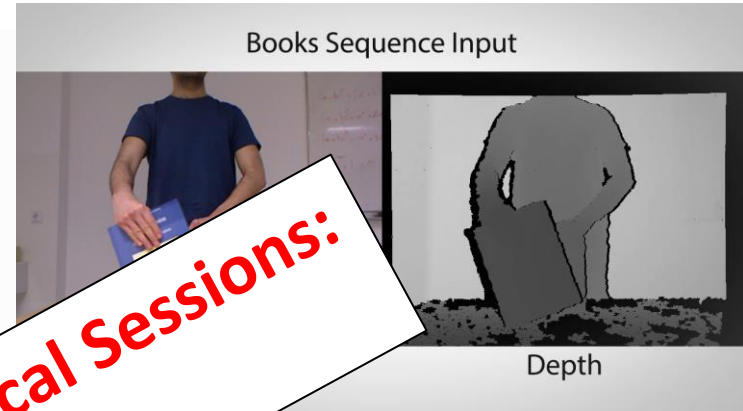
with a focus on segmentation, tracking, and matching – jointly with teams from MPI-CBG.



## Image Analysis

with a focus on segmentation, interactive techniques, as well as camera design.

**Thesis and Practical Sessions:  
Looking for People !**



# Future in Computer Vision

A project work in the CVLD is a good stepping stone if you:

- want to do a PhD in computer vision, graphics, machine learning
- Becoming a researcher or software developer in a research lab (Microsoft Research, Daimler, Google, Adobe, TechniColor, etc)
- If you are interested in doing a start-up
- Other “computer vision related” industry



# Introduction to Computer Vision

What is computer Vision?

(Potential) Definition:

Developing **computational models** and **algorithms** to **interpret digital images and visual data** in order to **understand** the visual world we live in.

# Introduction to Computer Vision

What is computer Vision?

(Potential) Definition:

Developing **computational models** and **algorithms** to **interpret digital images and visual data** in order to **understand** the visual world we live in.

# What does it mean to “understand”?



(Potential) Definition:

Developing **computational models** and **algorithms** to interpret **digital images and visual data** in order to **understand** the visual world we live in.

## Physics-based vision:

Geometry

Segmentation

Camera parameters

Emitted light (sun)

Surface properties: Reflectance, material

## Semantic-based vision:

Objects: class, pose

Scene: outdoor,...

Attributes/Properties:

- old-fashioned train
- A-on-top-of-B

# Image-formation model

Very many  
sources of  
variability

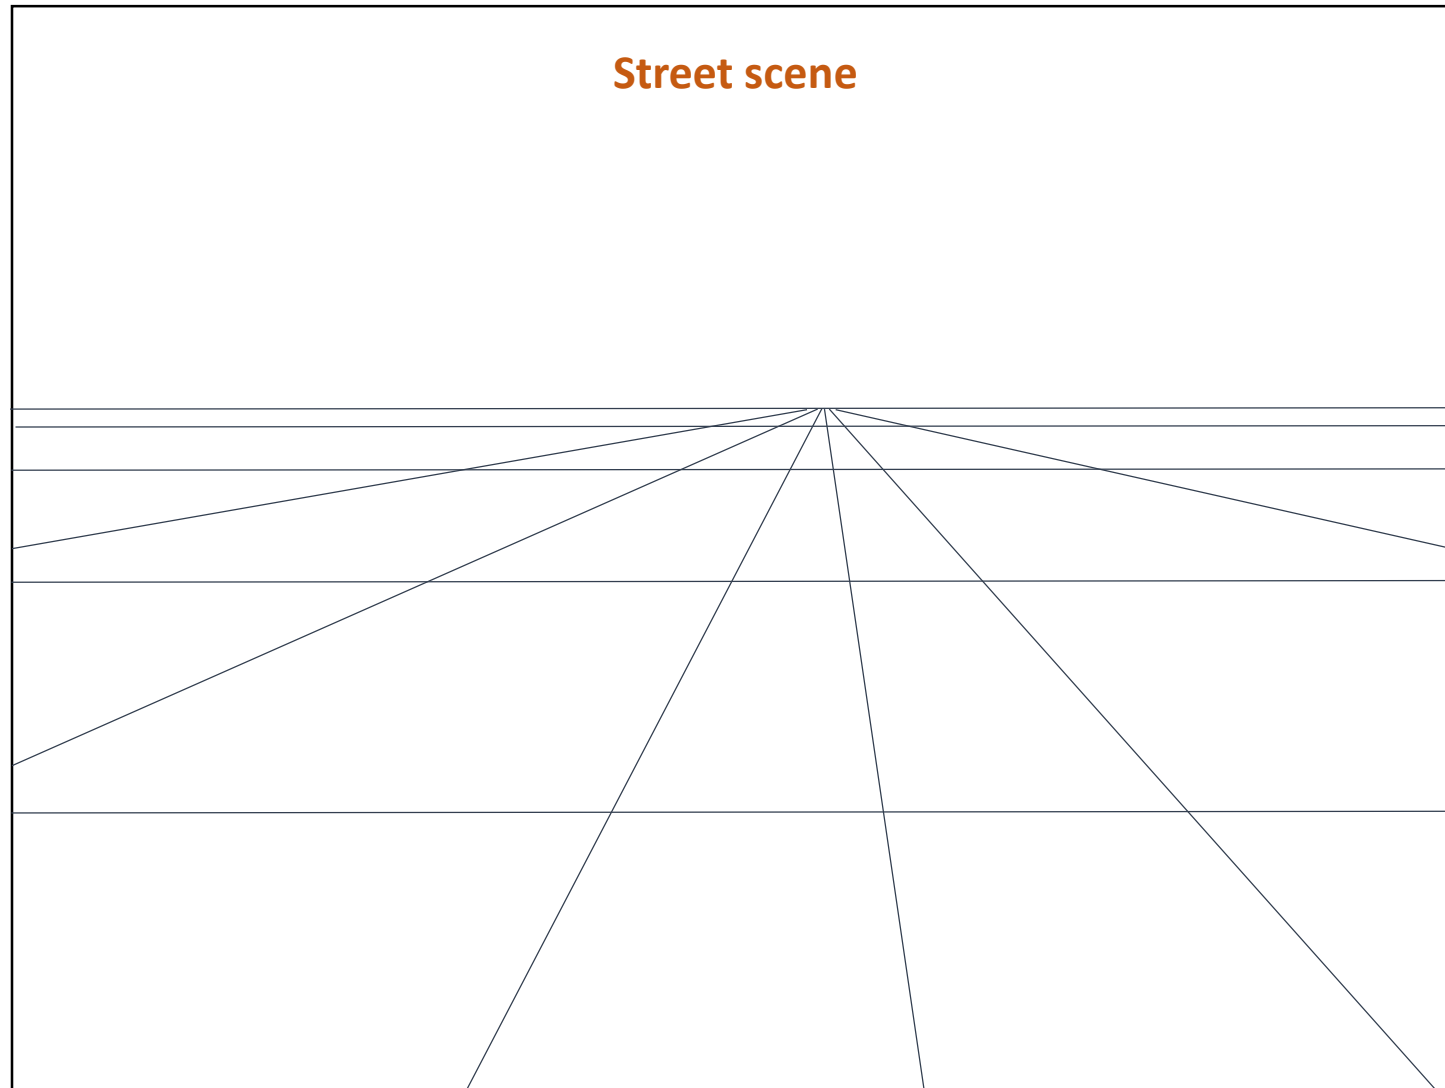


**Image**

[Slide Credits: John Winn, ICML 2008]

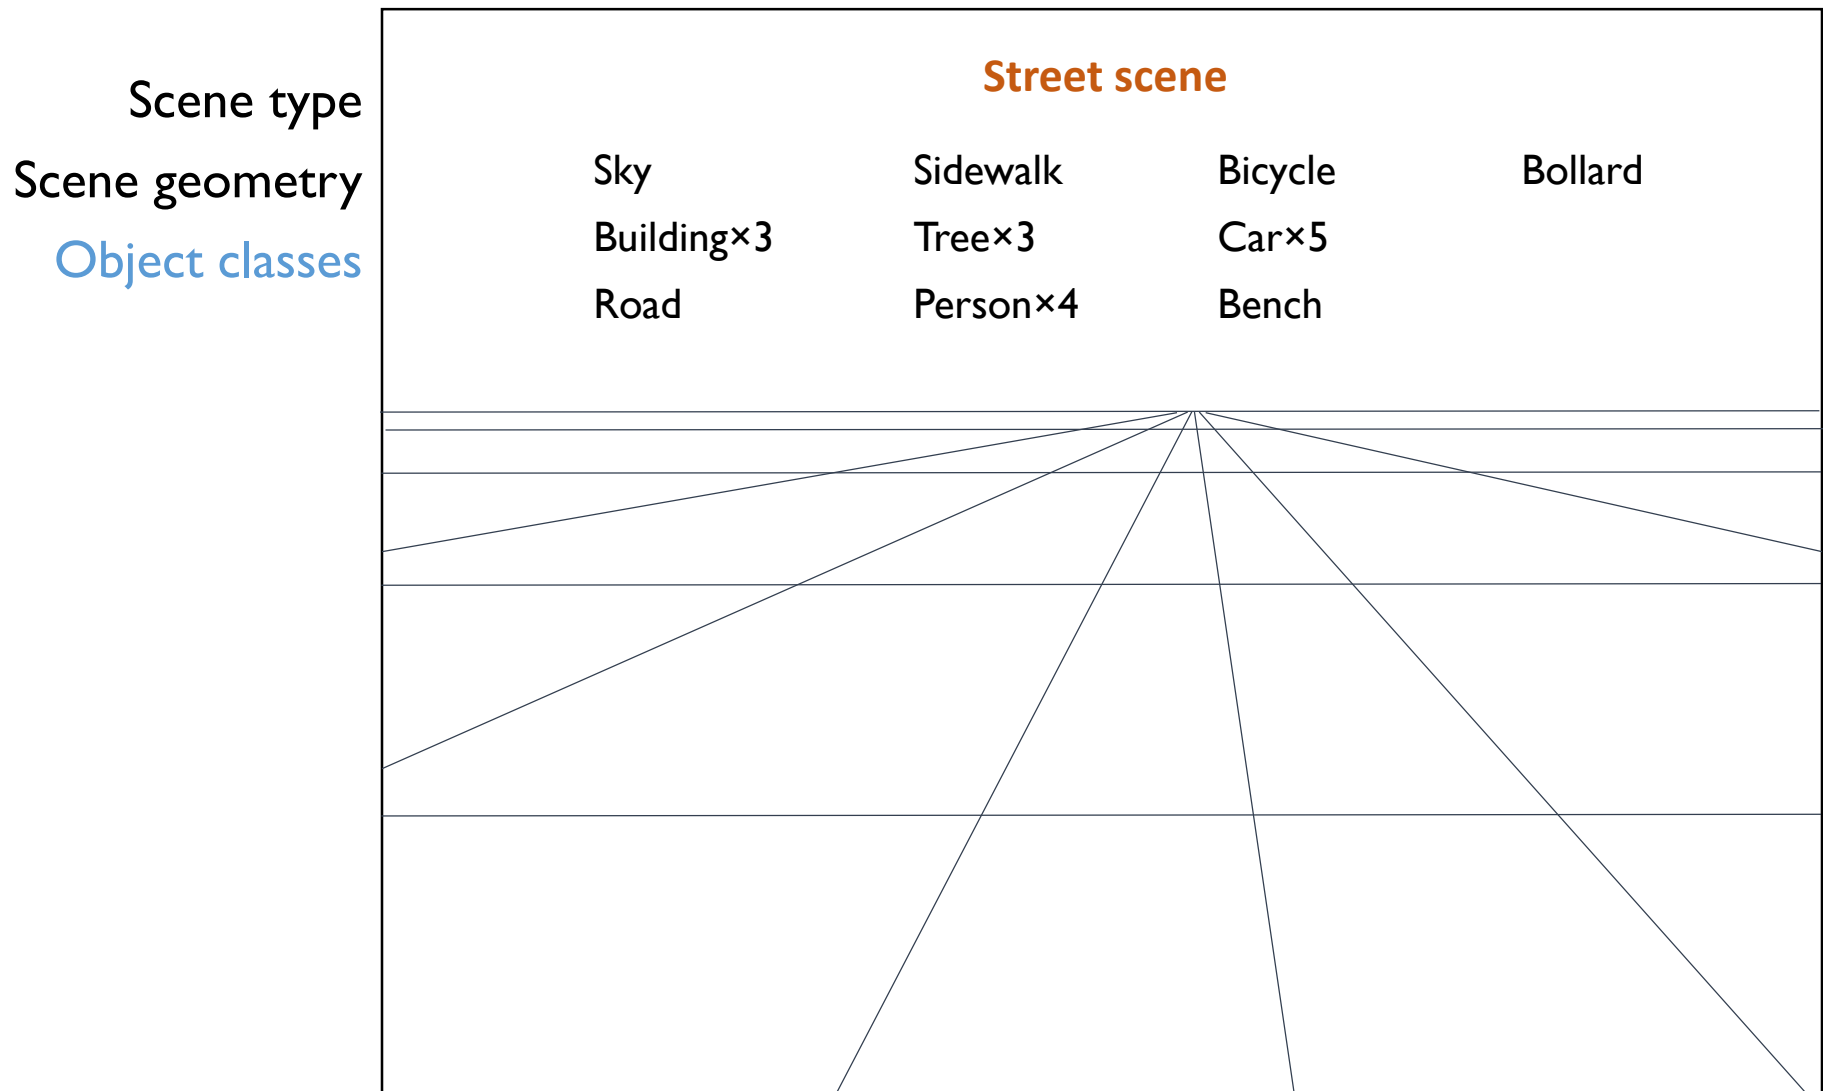
# Image-formation model

Scene type  
Scene geometry



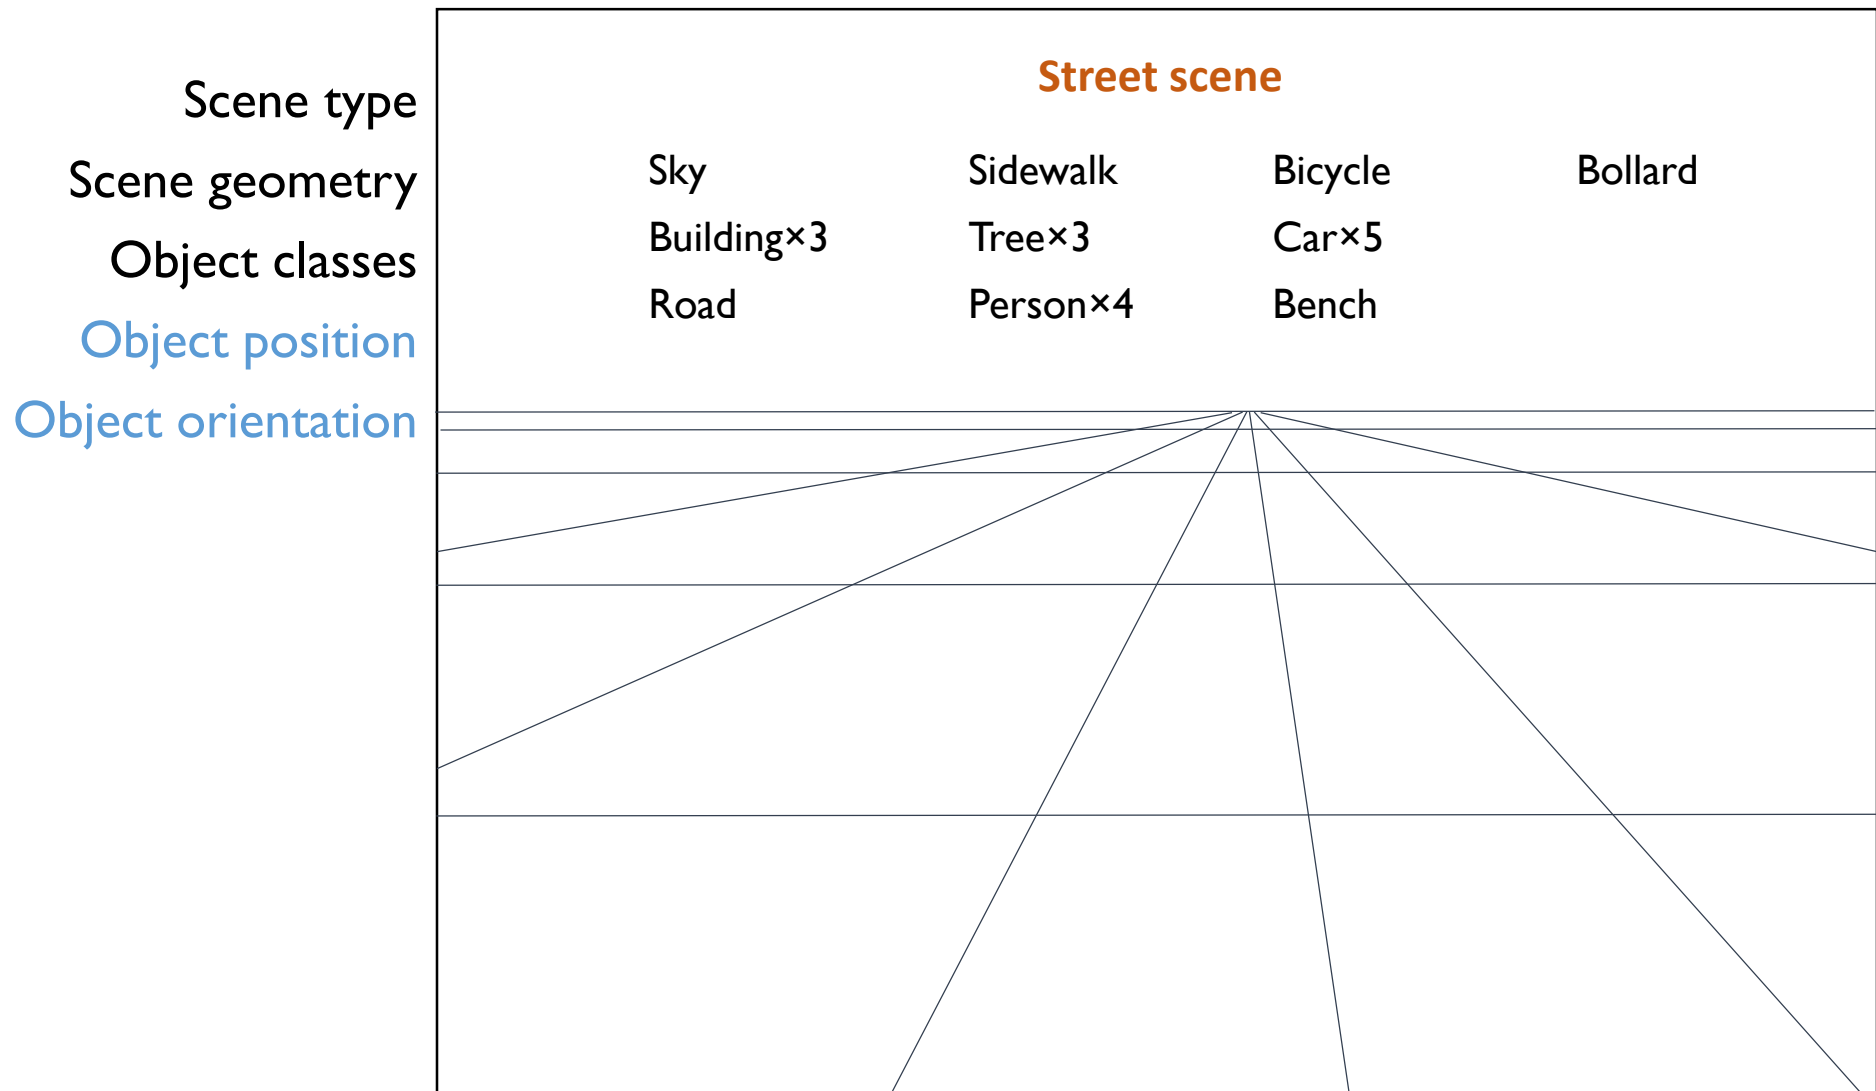
[Slide Credits: John Winn, ICML 2008]

# Image-formation model



[Slide Credits: John Winn, ICML 2008]

# Image-formation model

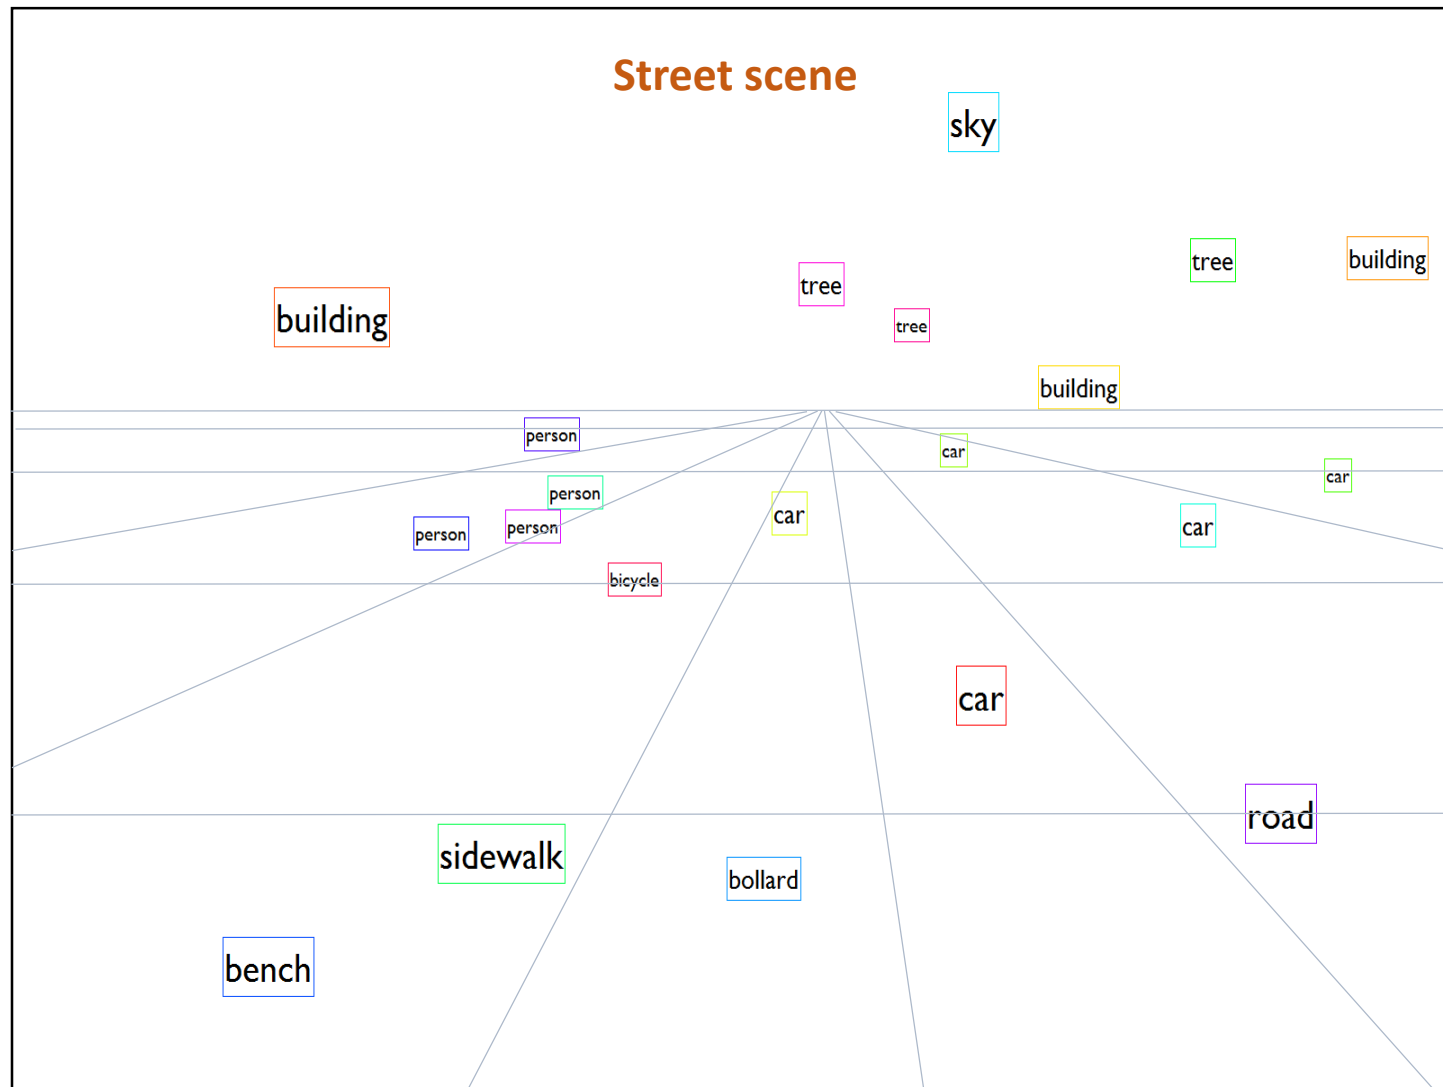


[Slide Credits: John Winn, ICML 2008]



# Image-formation model

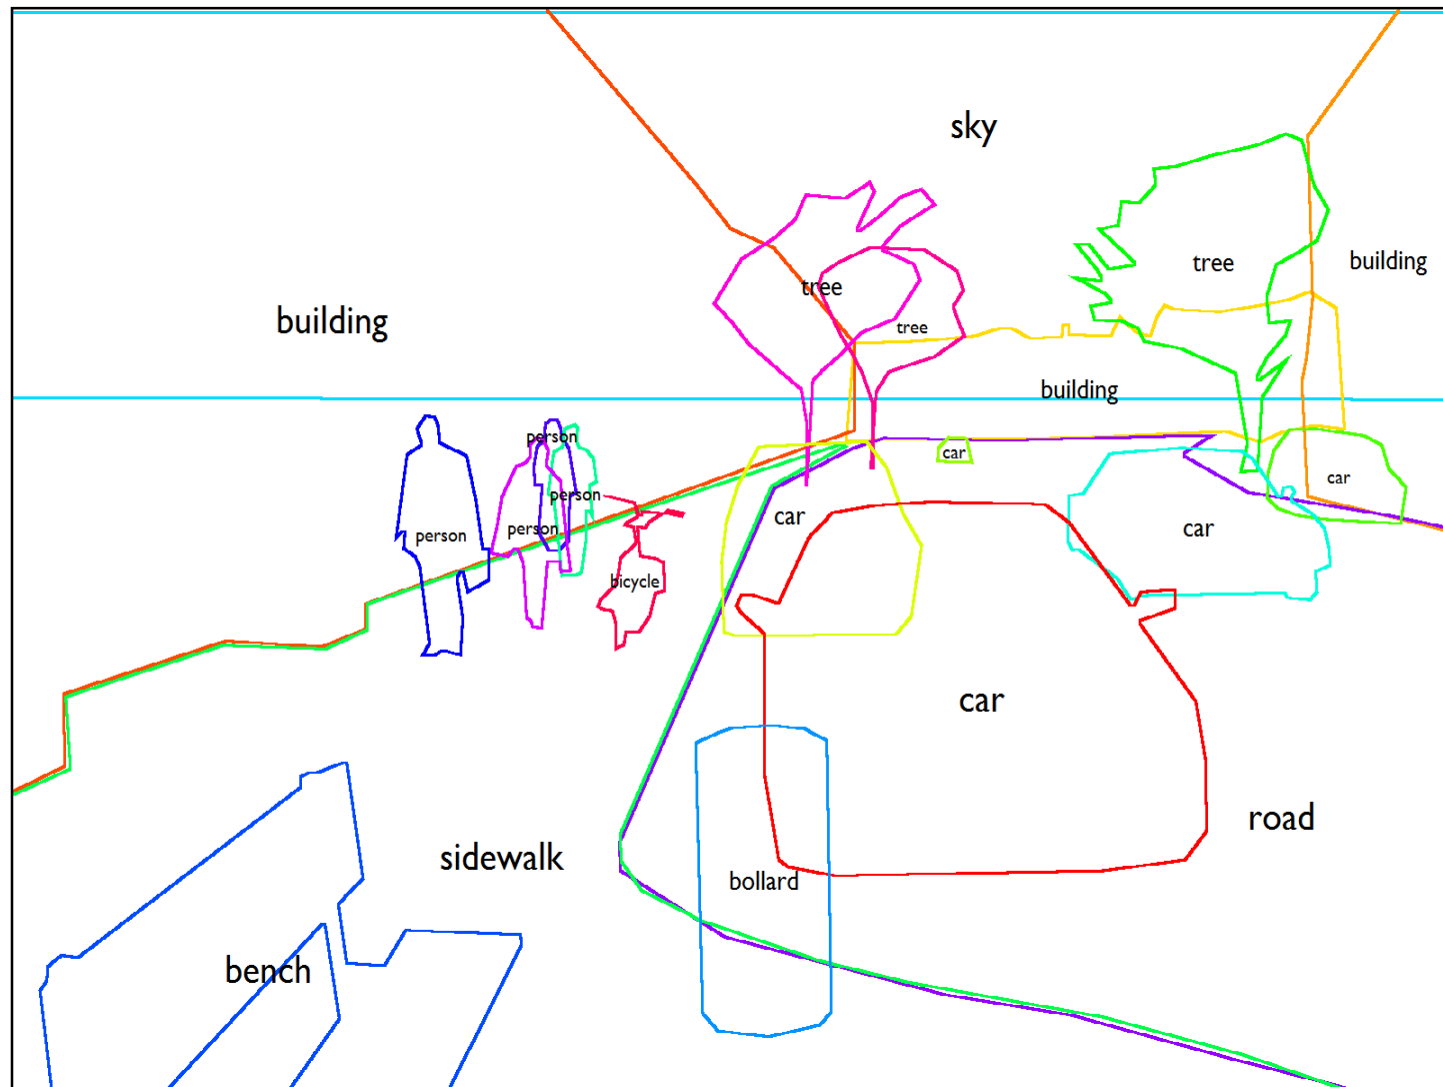
Scene type  
Scene geometry  
Object classes  
Object position  
Object orientation  
Object shape



[Slide Credits: John Winn, ICML 2008]

# Image-formation model

Scene type  
Scene geometry  
Object classes  
Object position  
Object orientation  
Object shape  
Depth/occlusions



[Slide Credits: John Winn, ICML 2008]

# Image-formation model

## Scene type

## Scene geometry

# Object classes

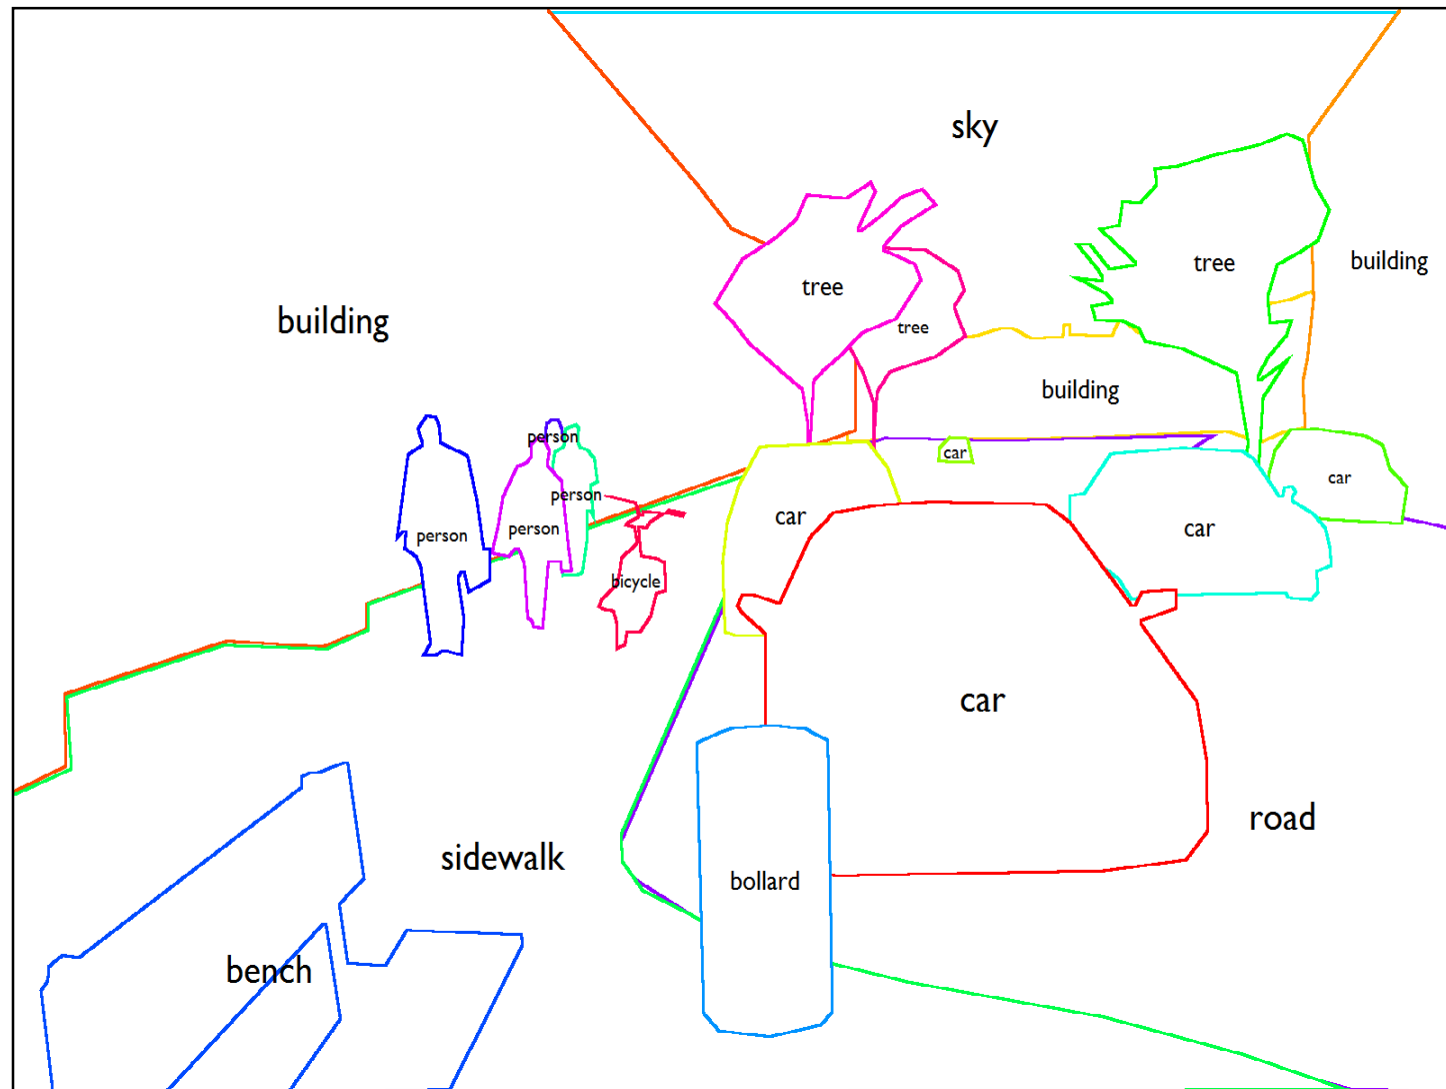
## Object position

## Object orientation

## Object shape

## Depth/occlusions

## Object appearance



[Slide Credits: John Winn, ICML 2008]

# Image-formation model

Scene type

Scene geometry

Object classes

Object position

Object orientation

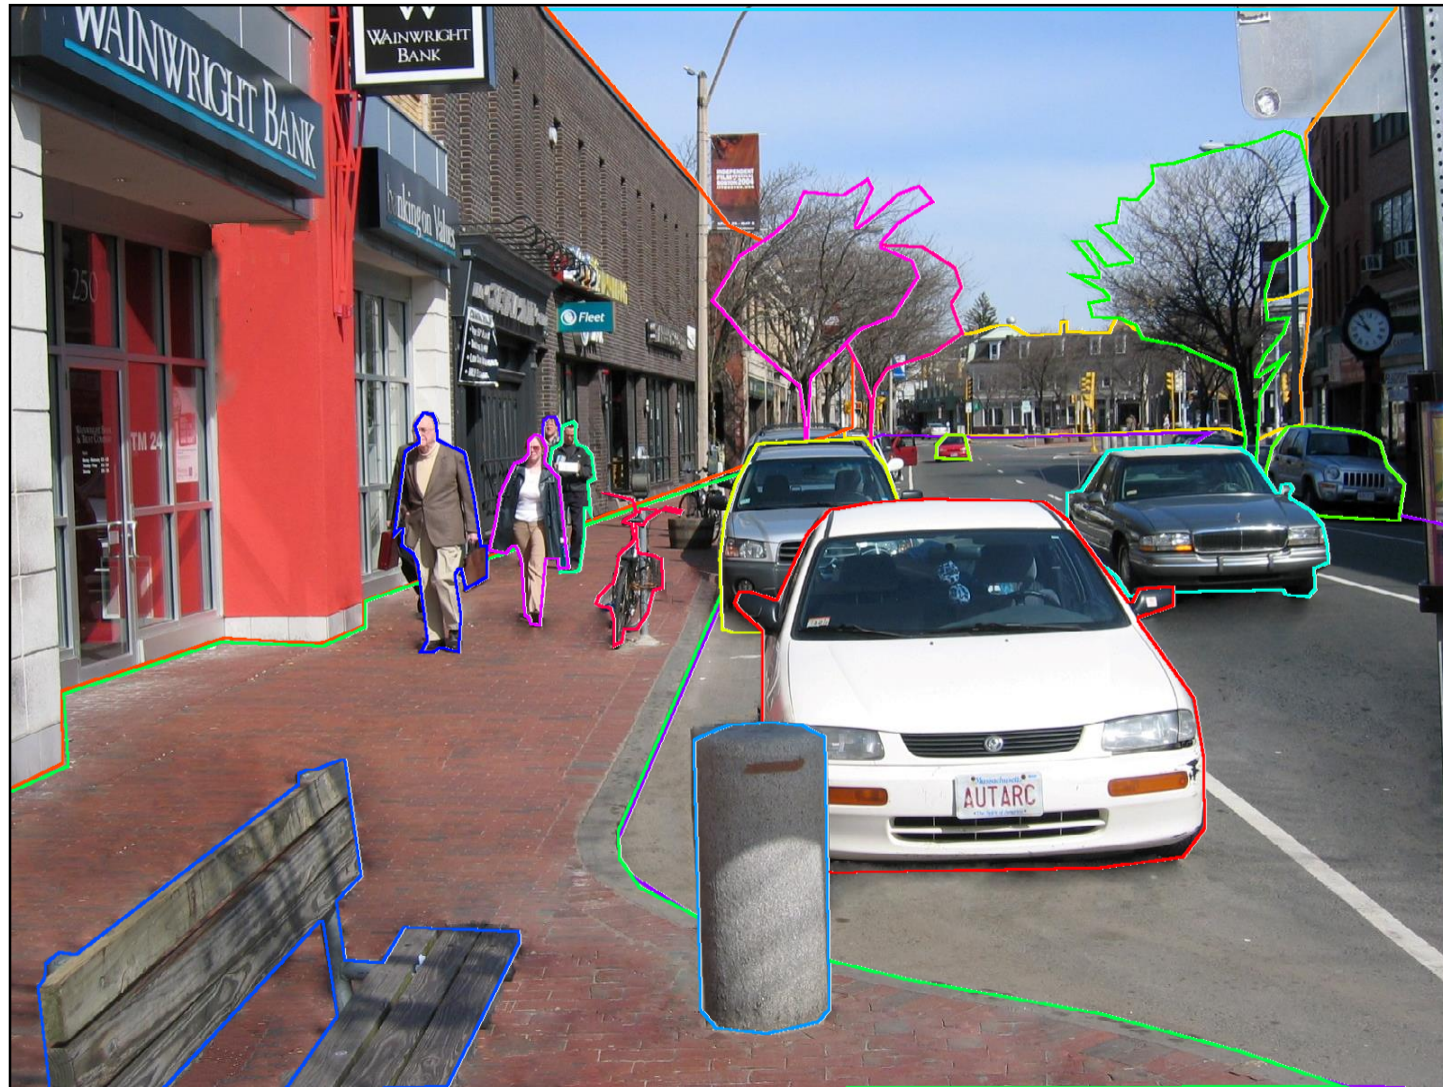
Object shape

Depth/occlusions

Object appearance

Illumination

Shadows



[Slide Credits: John Winn, ICML 2008]



# Image-formation model

Scene type

Scene geometry

Object classes

Object position

Object orientation

Object shape

Depth/occlusions

Object appearance

Illumination

Shadows



[Slide Credits: John Winn, ICML 2008]



# Image-formation model

Scene type

Scene geometry

Object classes

Object position

Object orientation

Object shape

Depth/occlusions

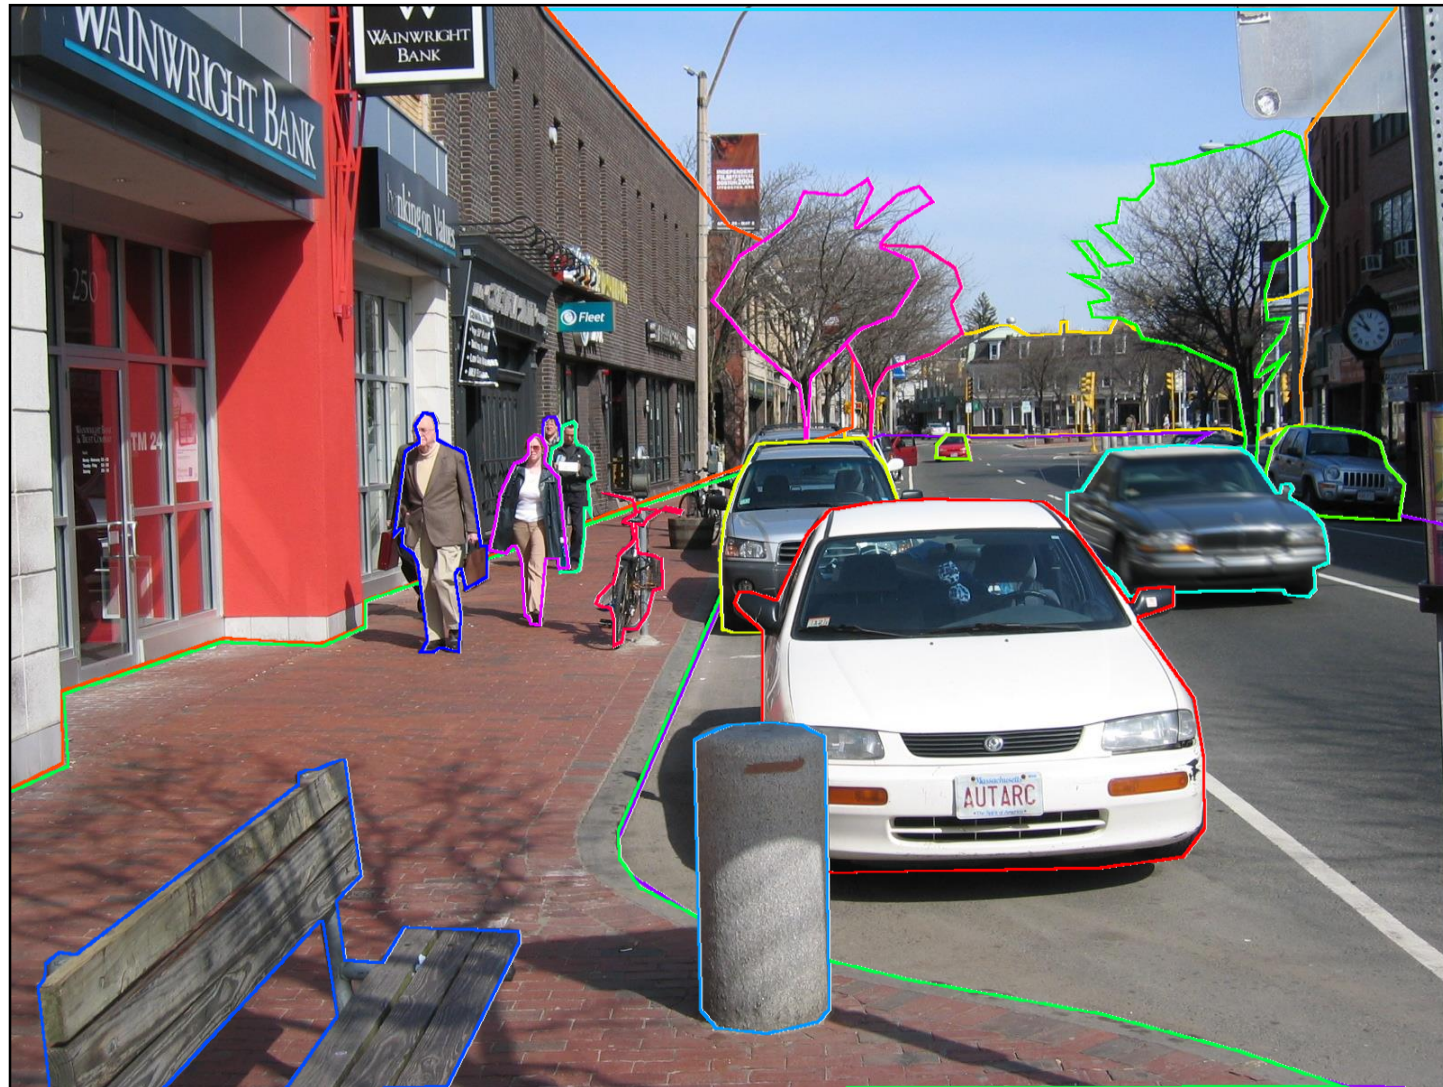
Object appearance

Illumination

Shadows

Motion blur

Camera effects



[Slide Credits: John Winn, ICML 2008]



# Image-formation model

Scene type

Scene geometry

Object classes

Object position

Object orientation

Object shape

Depth/occlusions

Object appearance

Illumination

Shadows

Motion blur

Camera effects



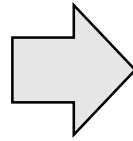
[Slide Credits: John Winn, ICML 2008]



# The “Scene Parsing” challenge --- a “grand challenge” of computer vision



Single image

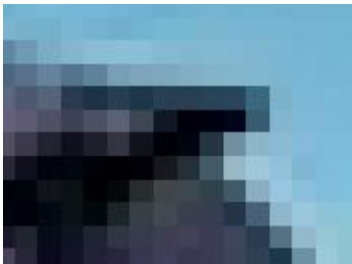


***(Probabilistic) Script*** = {Camera,  
Light, Geometry, Material, Objects,  
Scene, Attributes, Others}

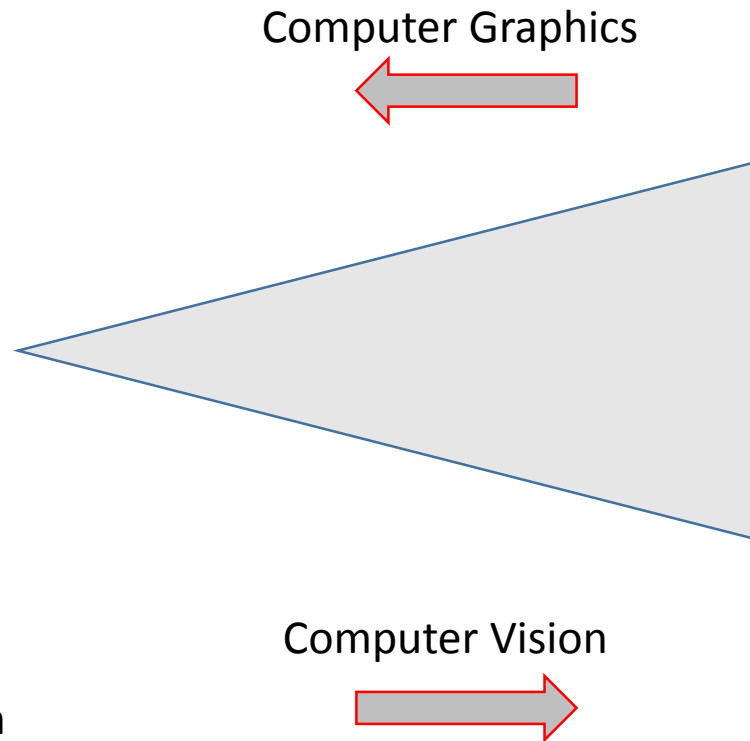
Many applications do not have to extract the full probabilistic script but only a subset, e.g. “does the image contain a car?”

... many examples to come later

# Why is “scene parsing” hard?



2D pixel representation



Computer Graphics

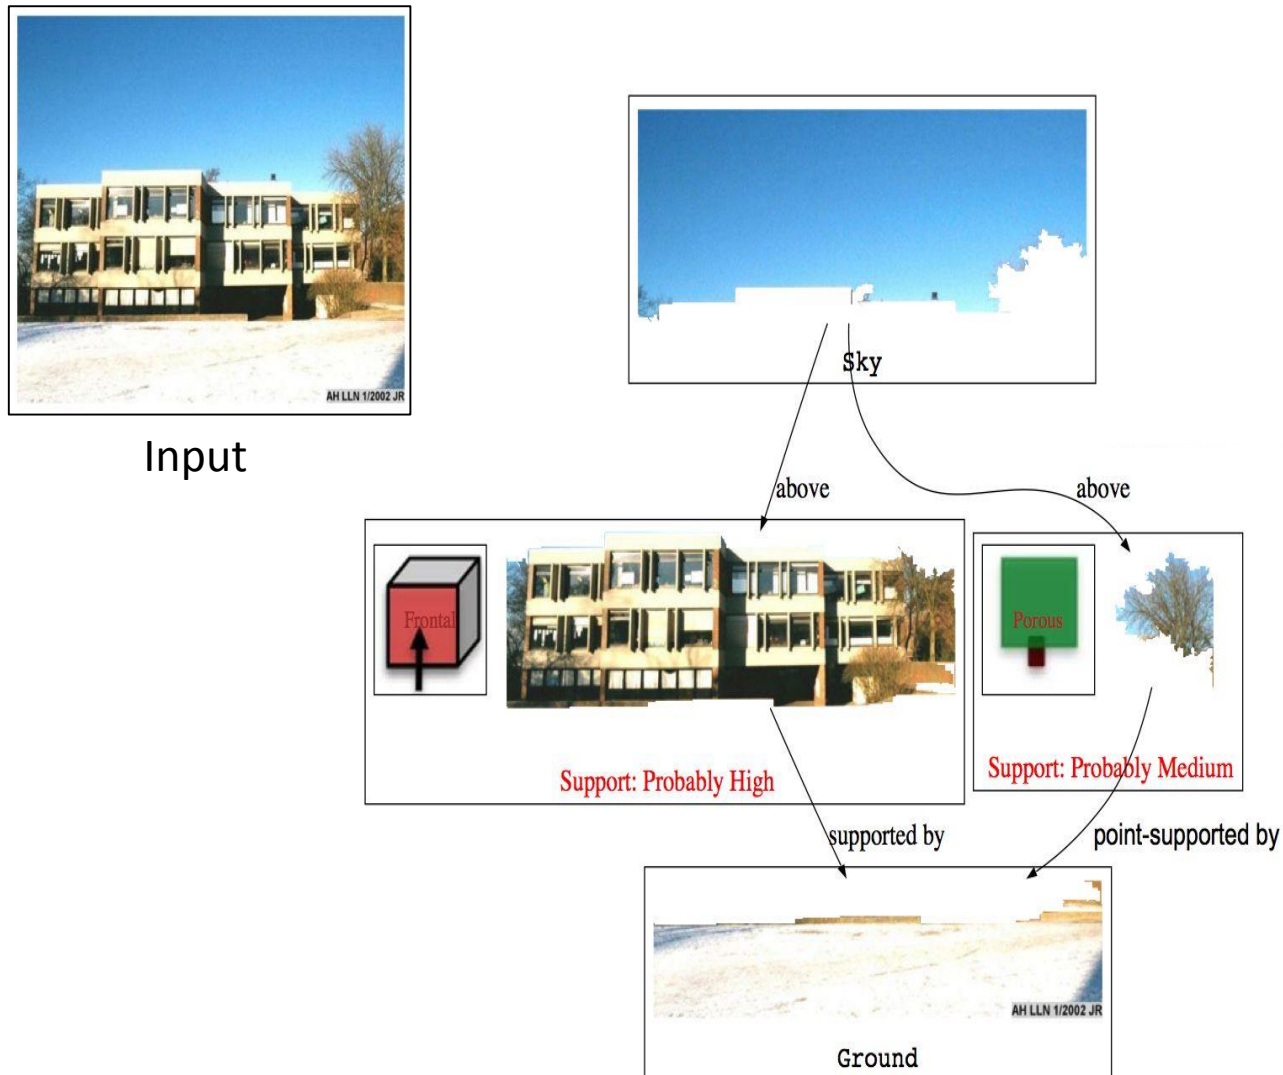
3D Rich Representation,

***Script*** = {Camera, Light, Geometry, Material, Objects, Scene, Attributes, Others}

Computer Vision

Computer Vision can be seen as “**inverse graphics**”

# Example of a recent work



Scene graph

[Gupta, Efros, Herbert, ECCV '10]

# Why is “scene parsing” hard?

## Subgoal for July

Analysis of scenes consisting of non-overlapping objects from the  
balls

bricks with faces of the same or different colors or textures

cylinders.

Each face will be of uniform and distinct color and/or texture.

## Extensions for August

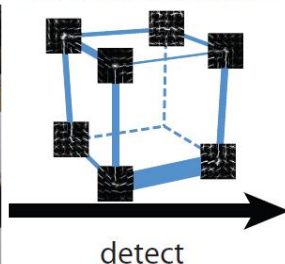
The first priority will be to handle objects of the same sort but  
with complex surfaces and backgrounds, e.g. cigarette pack with writing  
and bands of different color, or a cylindrical battery.

Input Image



[Xiao et al. NIPS 2012]

3D Cuboid Detector



Output Detection Result



Synthesized New Views



[Sussman, Lamport, Guzman 1966]

[Slide credits Andrew Blake]

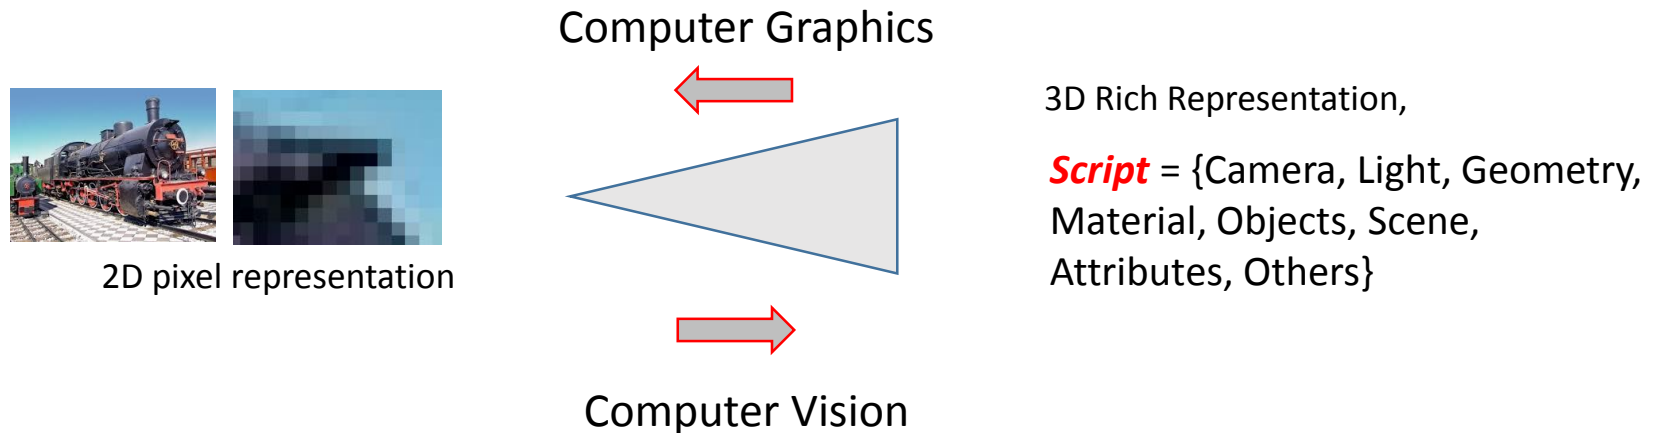
# Introduction to Computer Vision

What is computer Vision?

(Potential) Definition:

Developing **computational models** and **algorithms** to **interpret digital images and visual data** in order to **understand** the visual world we live in.

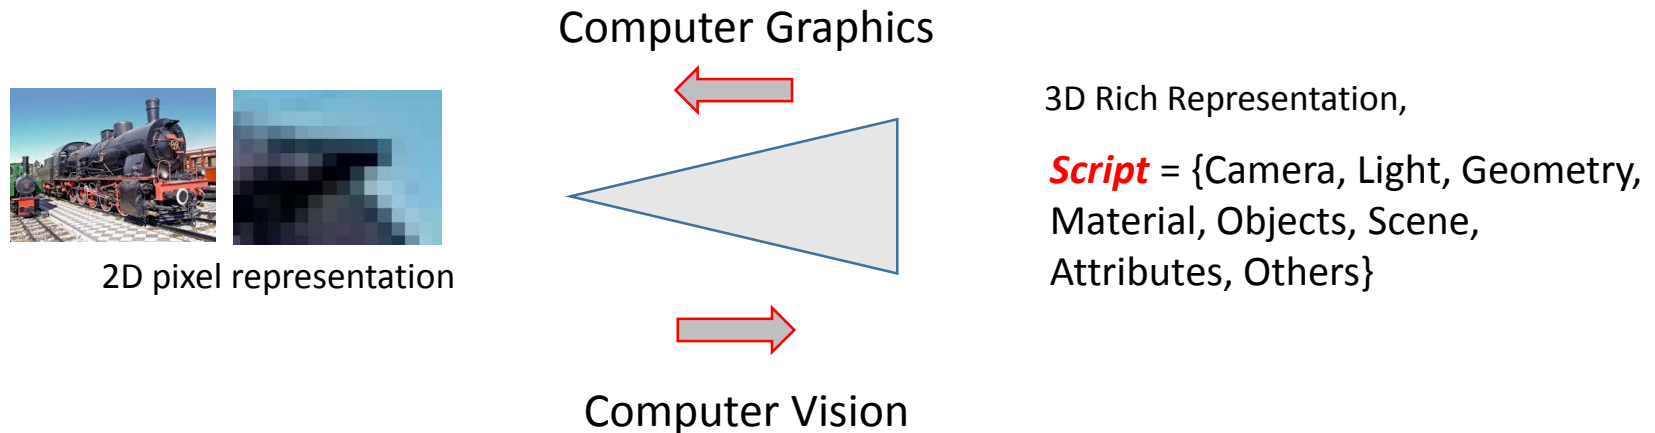
# How can we interpret visual data?



- What general (prior) knowledge of the world (not necessarily visual) can be exploit?
- What properties / cues from the image can be used?

Both aspects are quite well understood (a lot is based on physics) ... but how to use them is efficiently is open challenged (see later)

# How can we interpret visual data?



- What general (prior) knowledge of the world (not necessarily visual) can be exploit?
- What properties / cues from the image can be used?

Both aspects are quite well understood (a lot is based on physics) ... but how to use them is efficiently is open challenged (see later)



# Prior knowledge (examples)

- “Hard” prior knowledge
  - Trains do not fly in the air
  - Objects are connected in 3D
- “Soft” prior knowledge:
  - The camera is more likely 1.70m above ground and not 0.1m.
  - Self-similarity: “all black pixels belong to the same object”

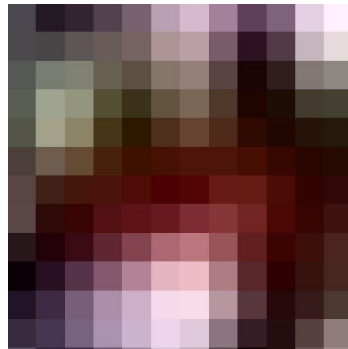


# Prior knowledge – harder to describe

- Describe Image Texture



Real Image



zoom

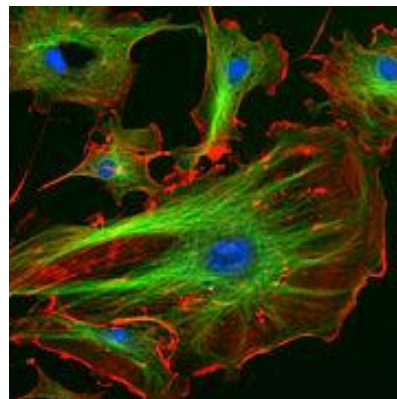


Not a real Image

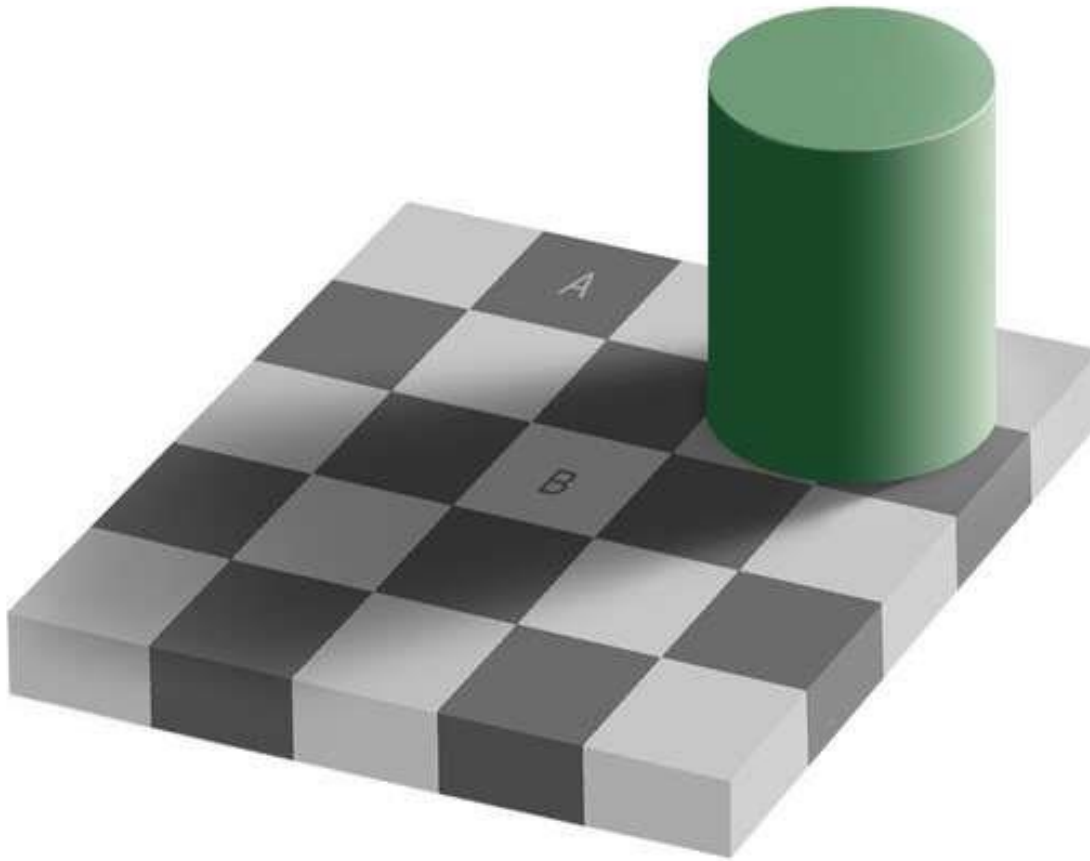


zoom

- Microscopic Images. What is the true shape of these objects



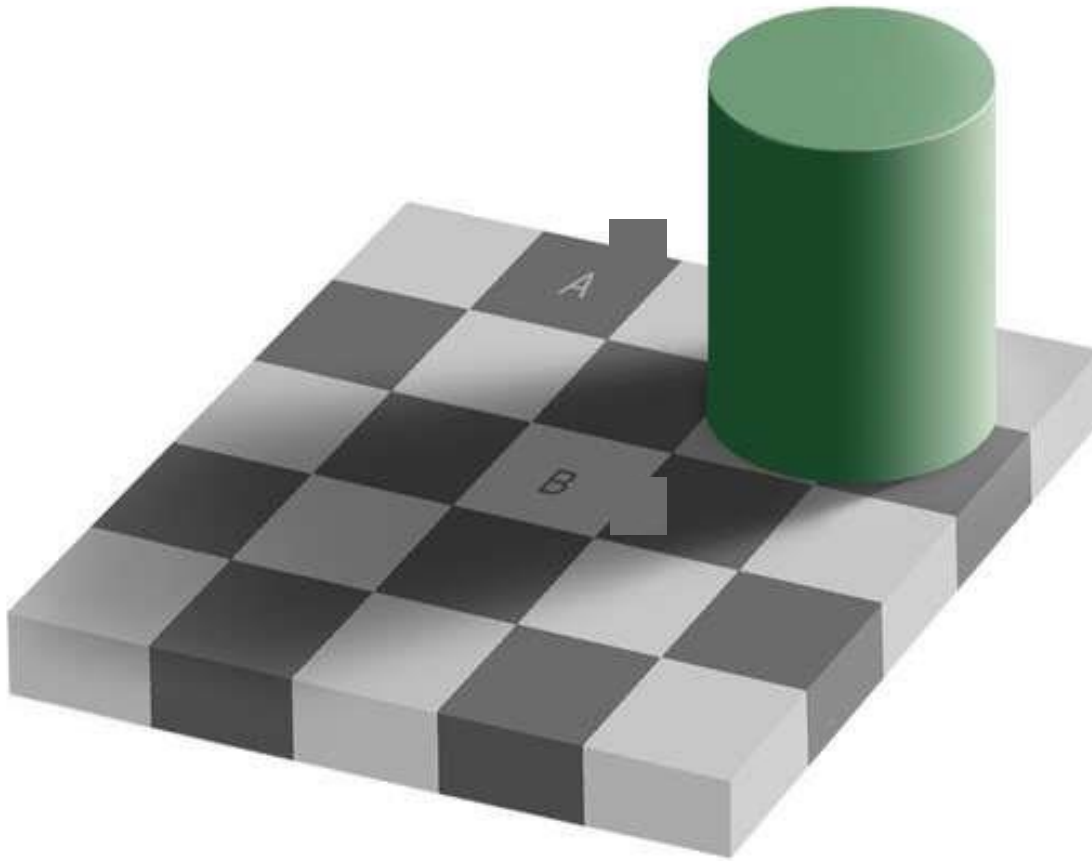
# The importance of Prior knowledge



Which patch is brighter: A or B?

[Edward Adelson]

# The importance of Prior knowledge

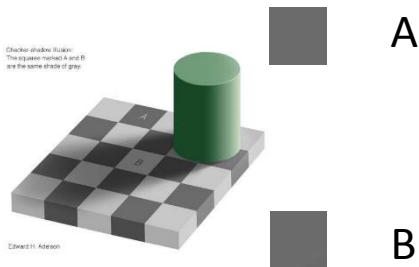


Which patch is brighter: A or B?

[Edward Adelson]

# The importance of Prior knowledge

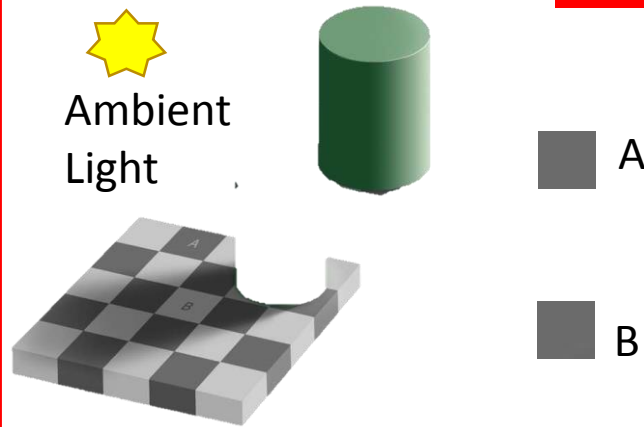
2D



2D Image - local

What the  
computer sees

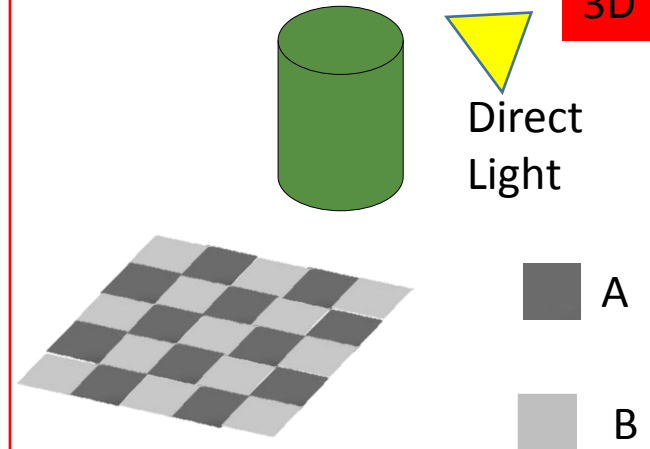
3D



True colors  
in 3D world

An unlikely  
3D representation  
(hard to see for a human)

3D



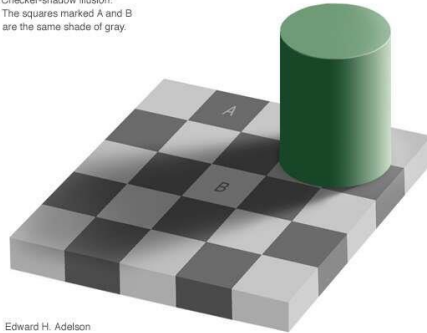
True colours  
In 3D world

The most likely  
3D representation

This is what humans see  
implicitly. Ideally the computer  
sees the same.

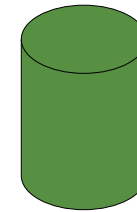
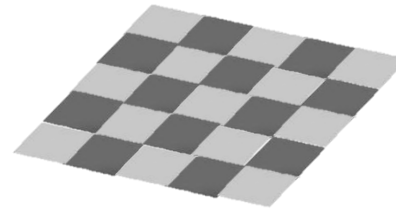
# The importance of Prior knowledge

Checker-shadow illusion:  
The squares marked A and B  
are the same shade of gray.



Edward H. Adelson

2D Image



Light

3D representation

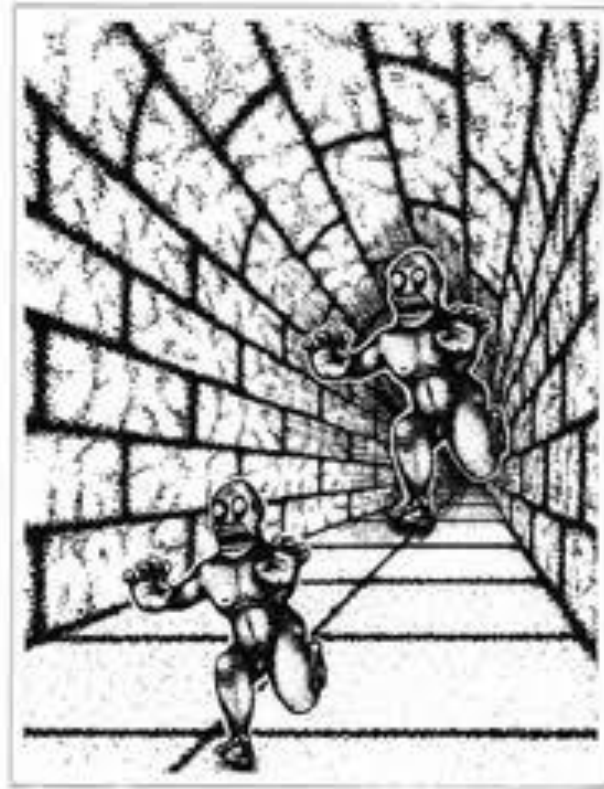
Humans see an image **not** as a set of 2D pixels. They understand an image as a projection of the 3D world we live in.

Humans have the prior knowledge about the world encoded, such as:

- Light cast shadows
- Objects do not fly in the air
- A car is likely to move but a table is unlikely to move

**We have to teach the computer this prior knowledge to understand 2D images as picture of the 3D world**

# The importance of Prior knowledge

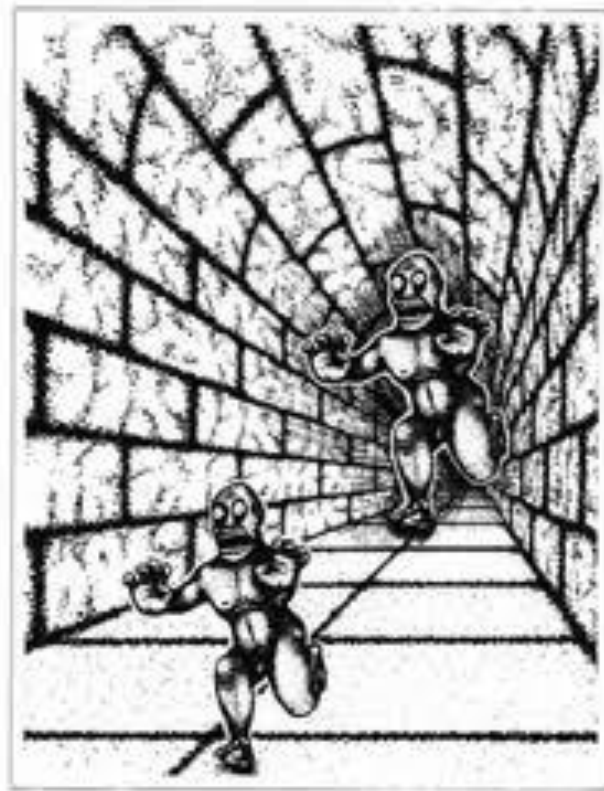


AT TERROR SURTEER

Which monster is bigger?



# The importance of Prior knowledge



AT TERROR SURTEER



In the 2D Image



1meter

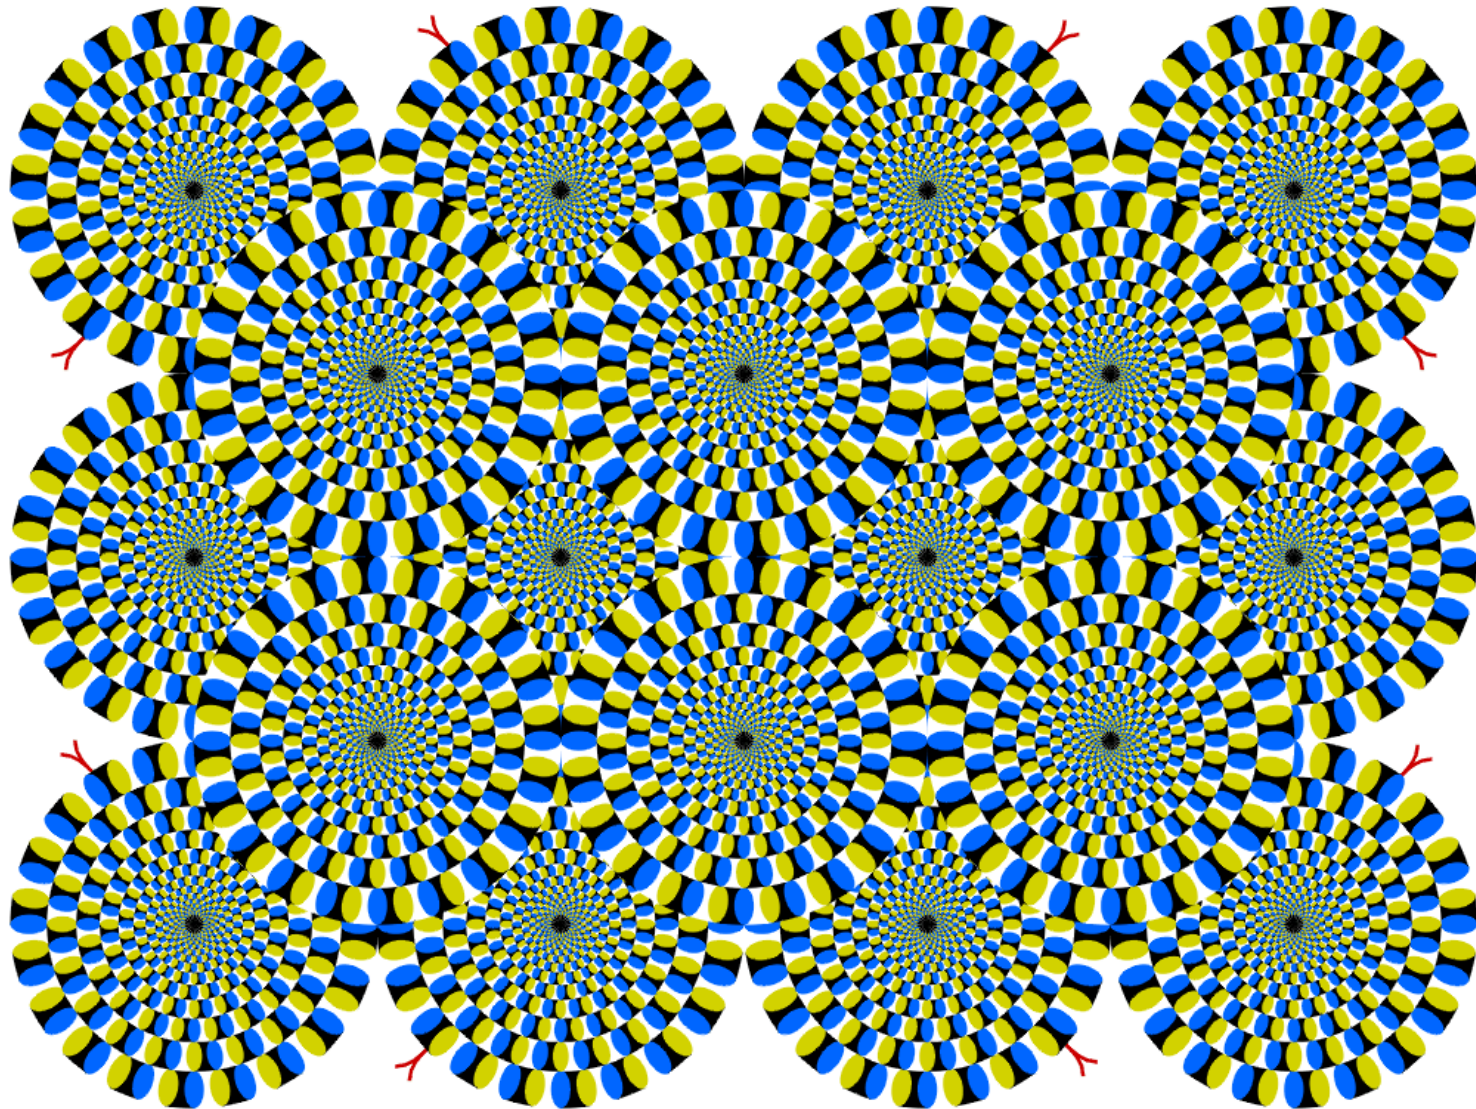


2meter

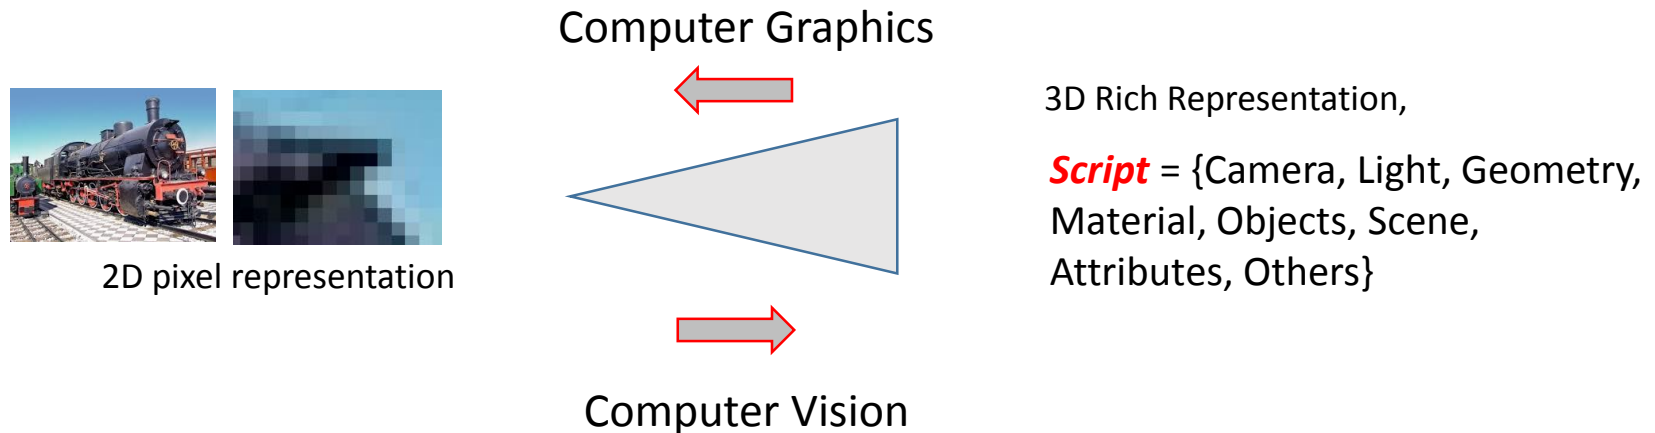
Which monster is bigger?

In the 3D world (true)

# Human Vision can be fooled



# How can we interpret visual data?



- What general (prior) knowledge of the world (not necessarily visual) can be exploit?
- What properties / cues from the image can be used?

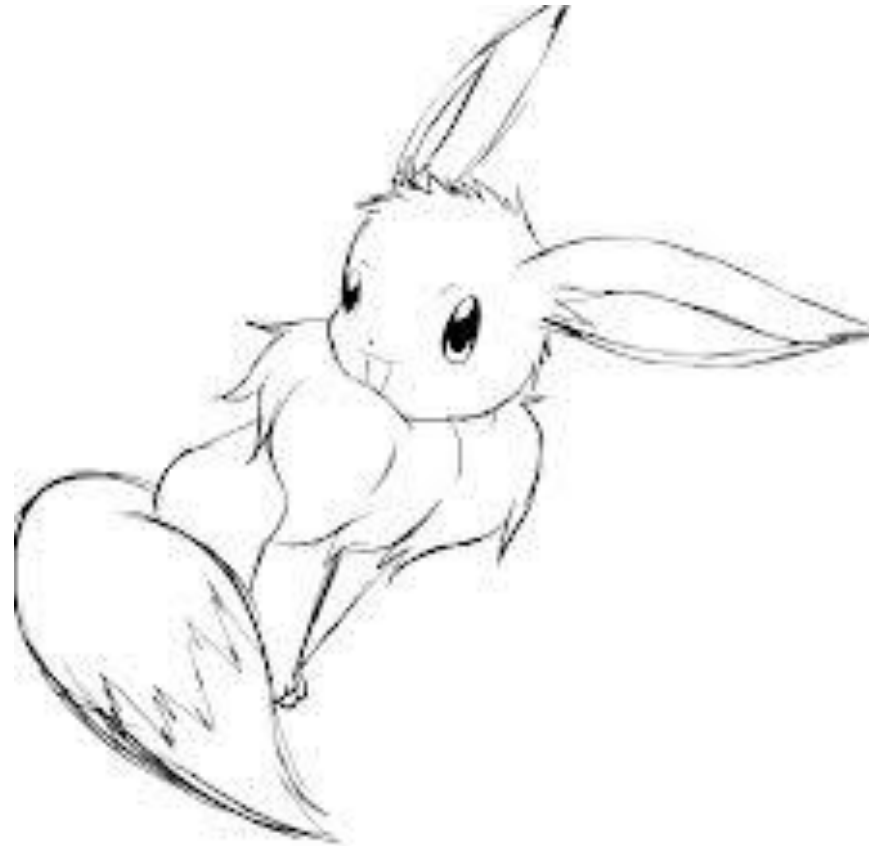
Both aspects are quite well understood (a lot is based on physics) ... but how to use them is efficiently is open challenged (see later)

# Cue: Appearance (Colour, Texture) for object recognition





# Cue: Outlines (shape) for object recognition



# Guess the Object



► Colour



► Texture



► Shape



[from JohnWinn ICML 2008]

# Cue: Context for object recognition

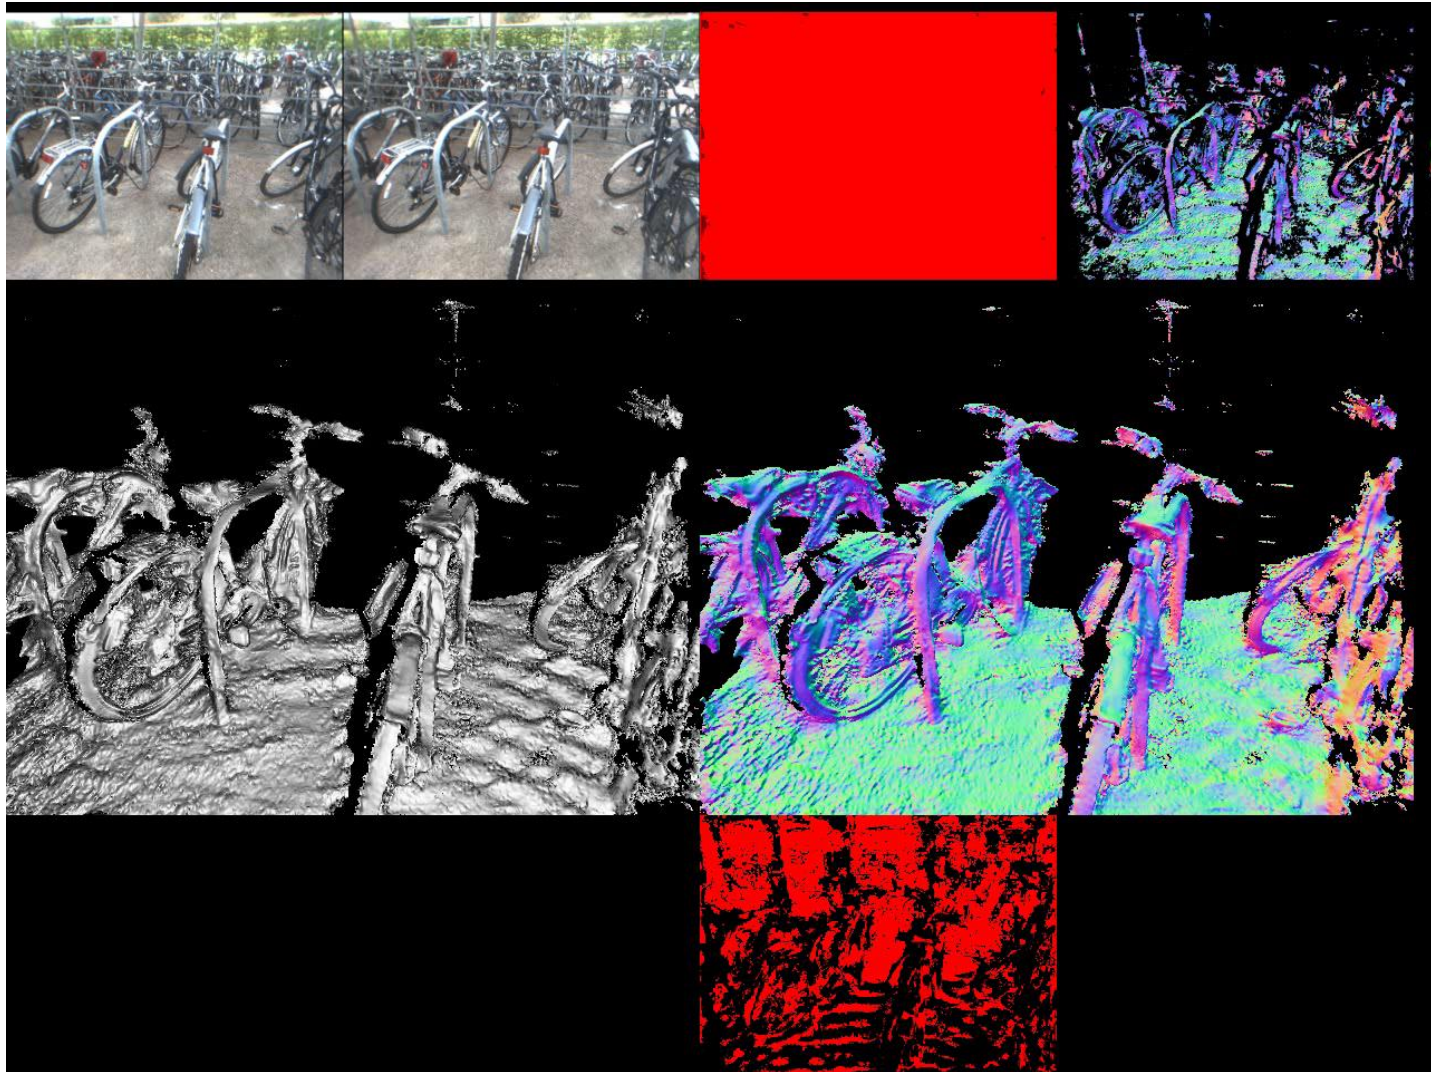


# Cue: Context for object recognition

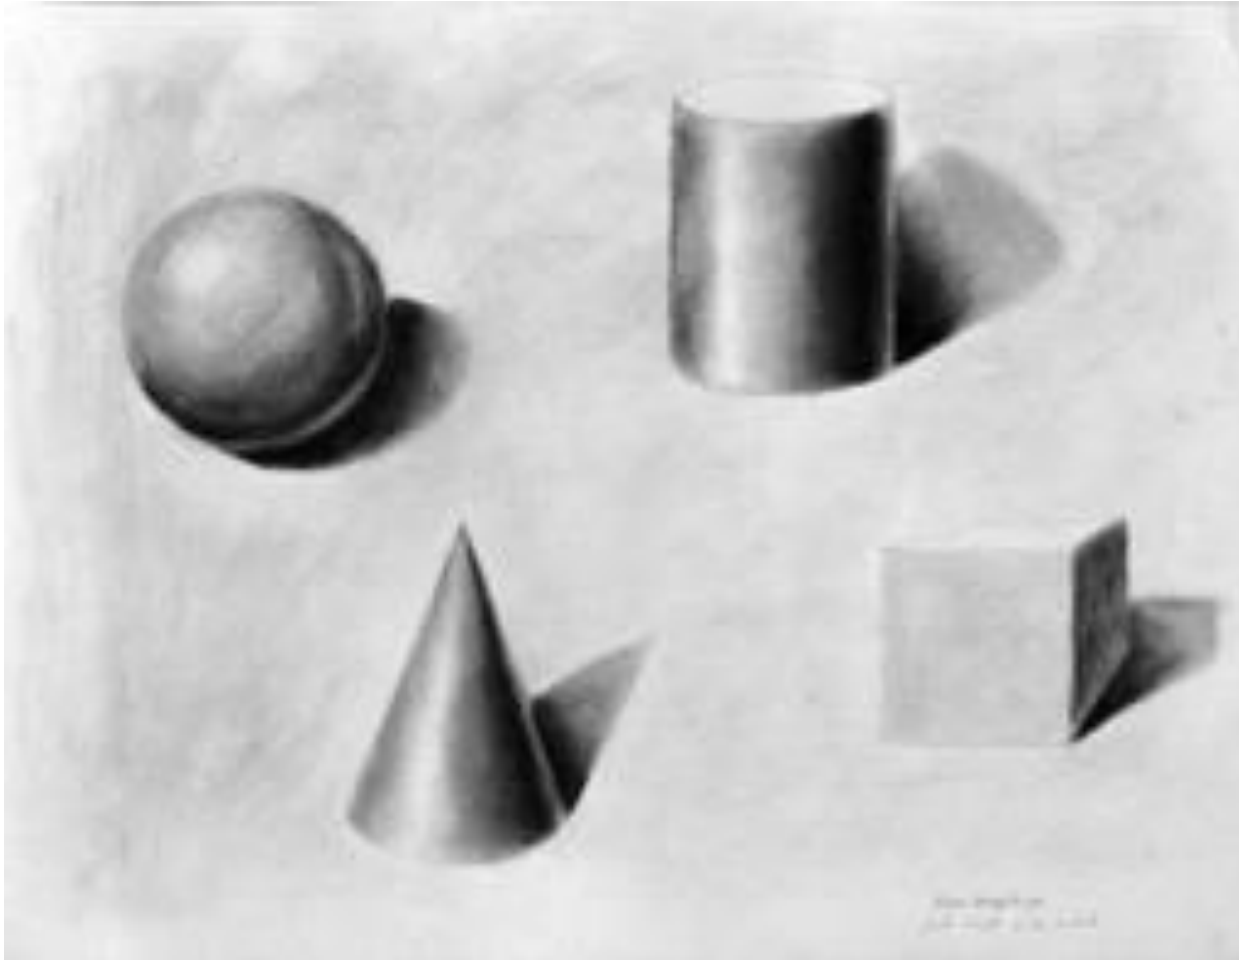




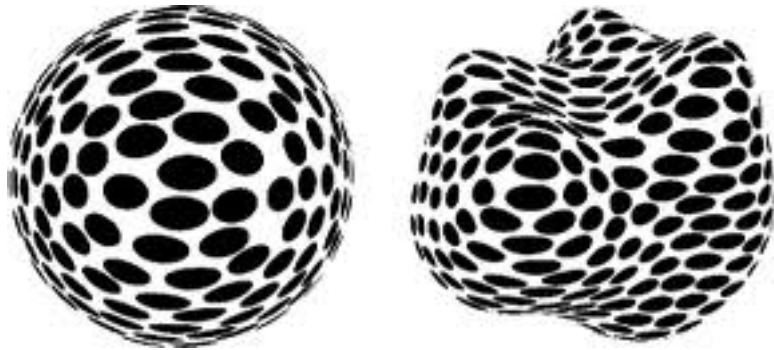
# Cue: Multiple Frames for geometry estimation



# Cue: Shading & shadows for geometry and Light estimation



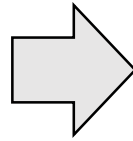
# Texture gradient for geometry estimation



# The “Scene Parsing” challenge --- a “grand challenge” of computer vision



Single image



**(Probabilistic) Script** = {Camera,  
Light, Geometry, Material, Objects,  
Scene, Attributes, Others}

Many applications do not have to extract the full probabilistic script but only a subset, e.g. “does the image contain a car?”

... many examples to come later



# Many application scenarios are in reach

To simplify/tackle the problem:

## 1) **Richer Input:**

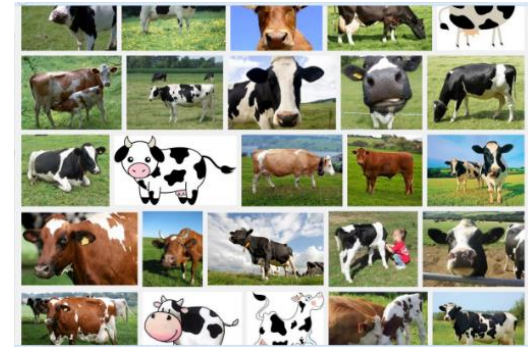
Modern sensing technology; Video;  
Cameras everywhere



## 2) **Rich Models:** Deep Learning

## 3) **Lots of Data** to learn from:

search engines; crowdsourcing;  
graphics engines



## 4) For **many practical** applications:

We do not have to infer the full probabilistic script

# Kinect has simplified computer vision



[Izadi et al. '11]

# Animate the world



[Chen et al. UIST '12]



# Kinect Body tracking and Gesture Recognition

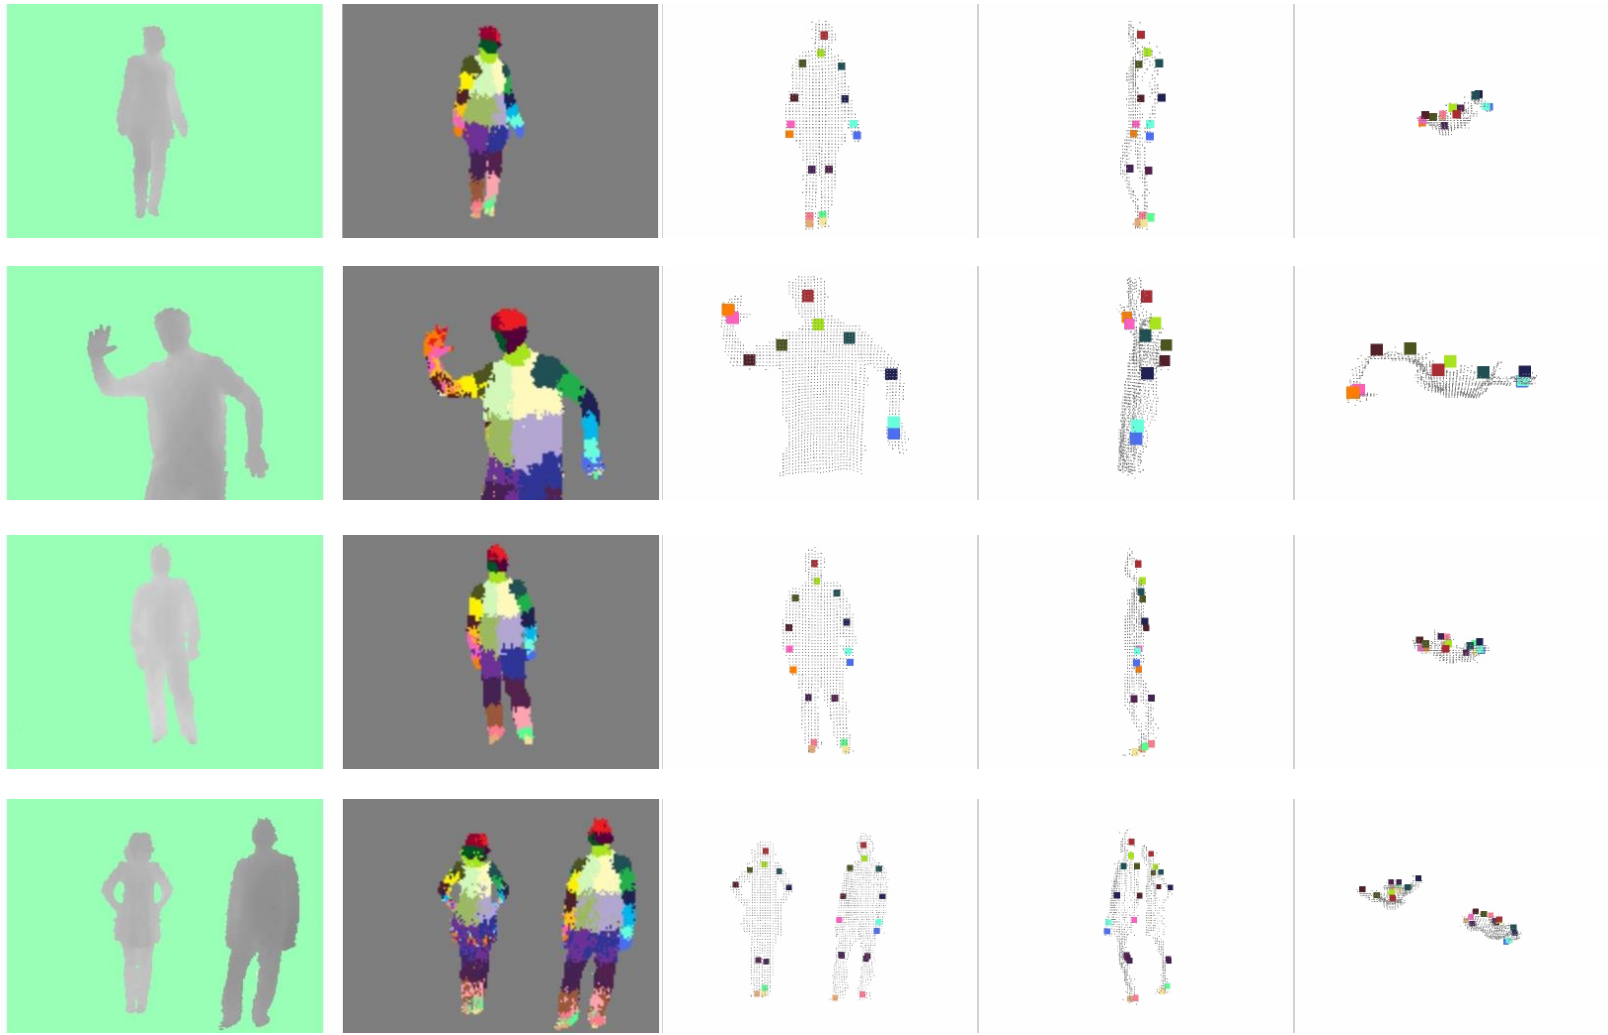


Start-Up 2012: Try Fashion online

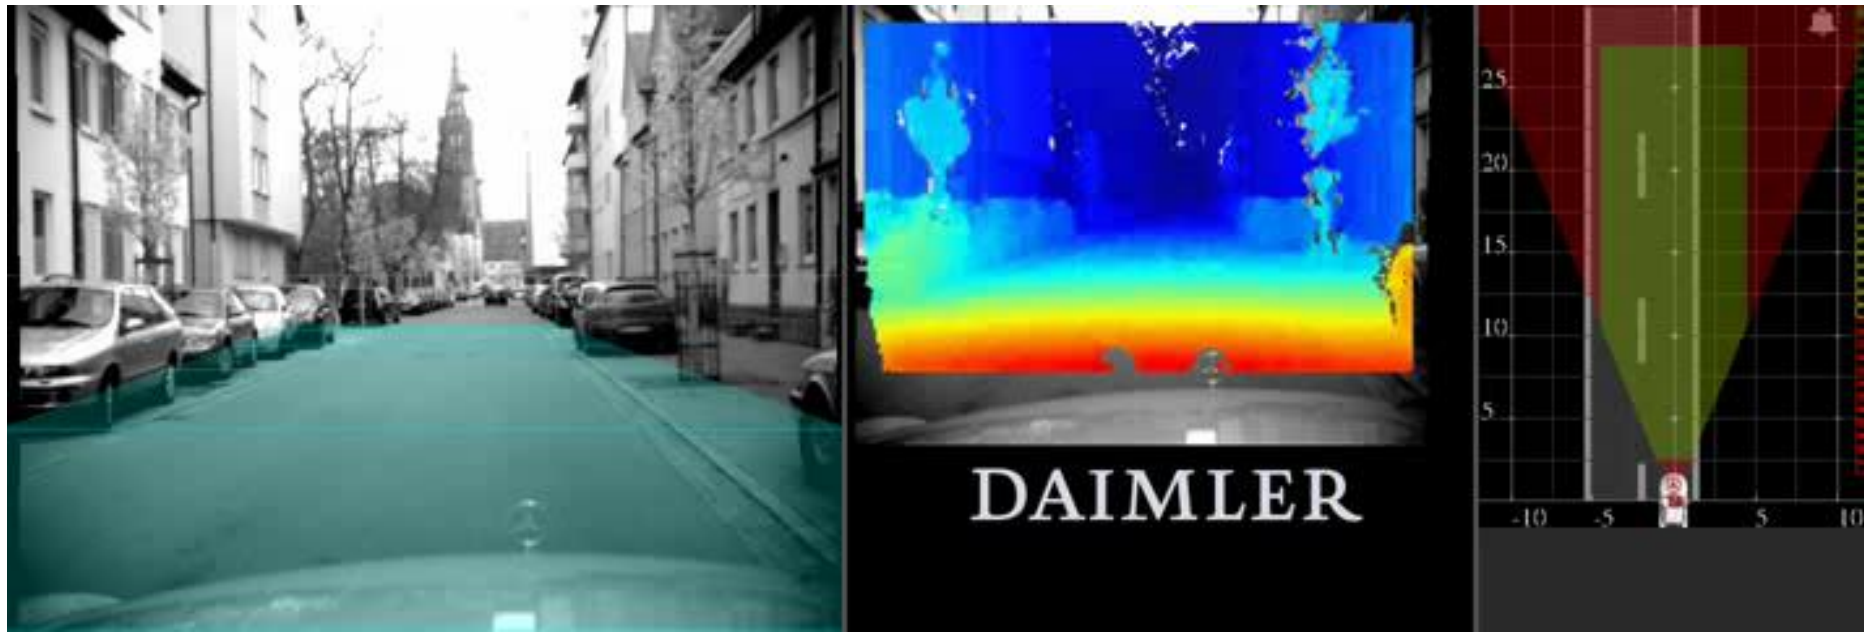
Very large impact in many field: Gaming, Robotics, HCI, Medicine, ...



# Kinect Body Pose estimation and tracking

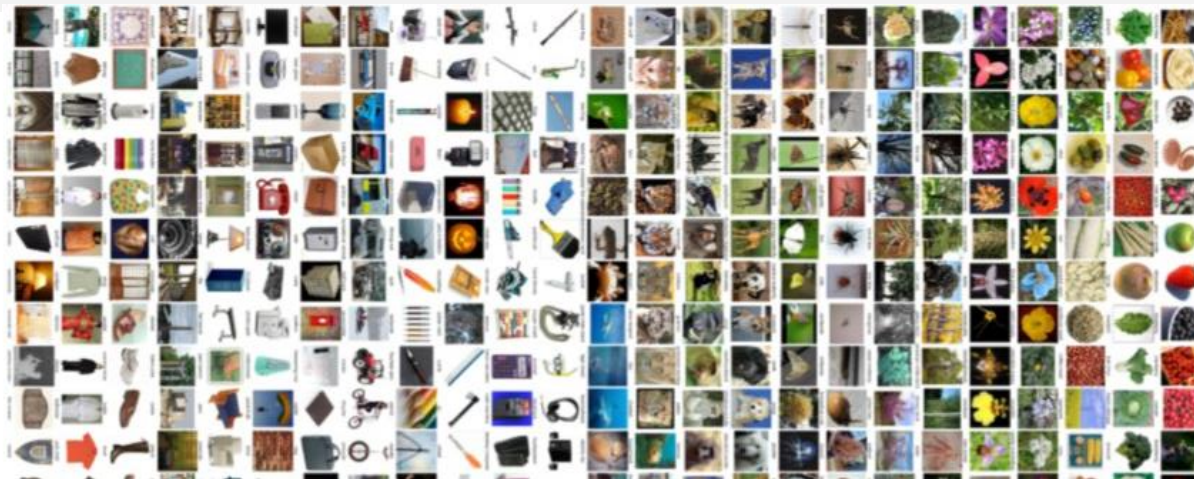


# Real-time pedestrian detection

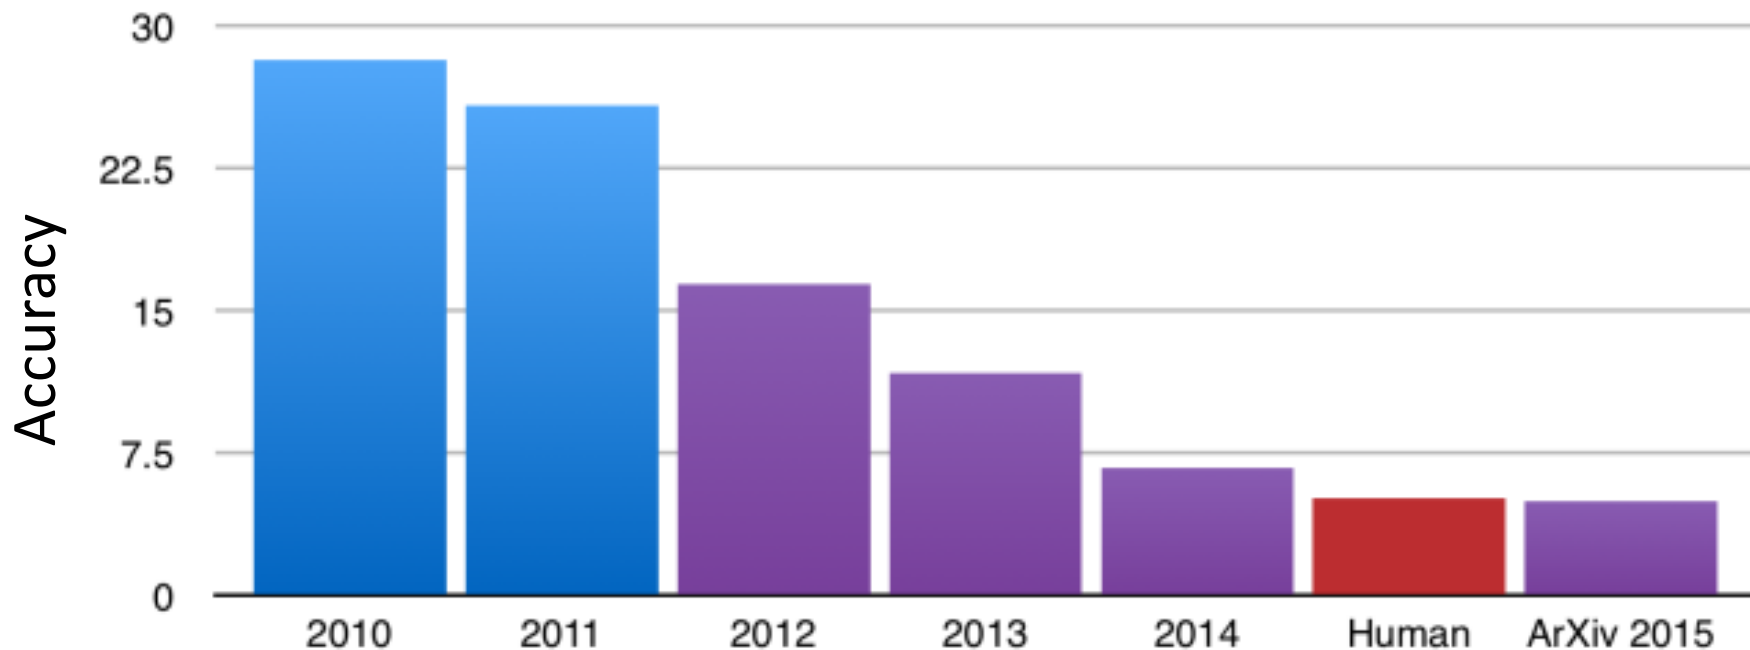


Daimler Research Lab

# Object recognition – ImageNet



What's in the image (1000 classes)



# Start-up Company: Like.com



Vestal

Vestal Alpha Bravo Canvas Watch Blue/Silver/Black

**\$76 online**

Write a review

Add to Shortlist

[Browse Watches »](#)

Go commando with the no-nonsense, military-inspired Vestal Alpha Bravo Canvas Watch. The mineral crystal glass face and stainless-steel housing can take a beating when you're on a mission in hostile territory, and the durable canvas strap provides a comfortable fit for all-day ... [more »](#)

Sponsored ⓘ

**\$76.46**

Free shipping

[Backcountry.com](#)

[Shop »](#)

**\$76.46**

Free shipping

[DogFunk.com](#)

[Shop »](#)

**\$81.00**

[Extreme Supply](#)

★★★★★ (186)

[Shop »](#)

[See all stores from \\$76](#)

## Visually Similar Items



Vestal  
\$75.99



GioiaPura  
\$139.71



Diesel  
\$98.57



Marc Ecko  
\$81.25



Vestal  
\$76.46



Gevril  
\$668.13





# Interactive Image manipulation

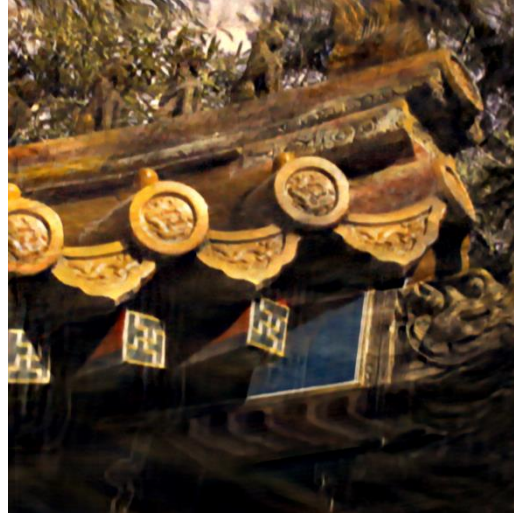


[Agrawal et al '04]

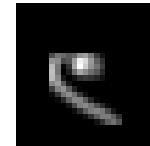
# Image de-convolution



Input



Output



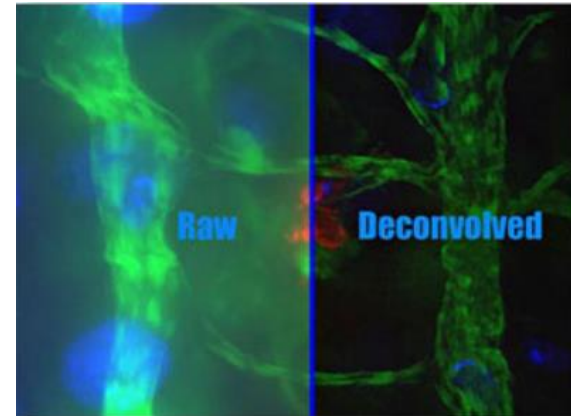
kernel



input



output



[Schmidt, Rother, Nowozin, Jancsary, Roth 2013] Best Student Paper award



[Rav-Acha et al. '08]



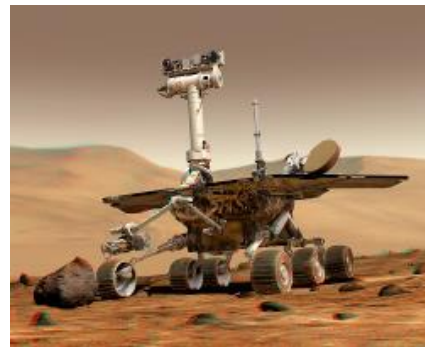
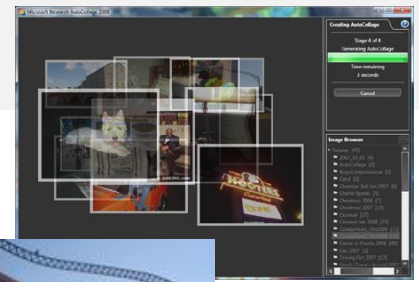
# Industry



Pirates of the Caribbean,  
Industrial Light and Magic



AutoCollage 2008 - Microsoft Research  
[Rother et al. Siggraph 2006]



Robotics



# Introduction to Computer Vision

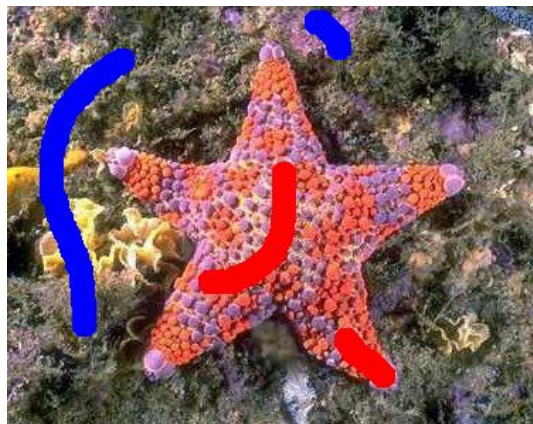
What is computer Vision?

(Potential) Definition:

Developing **computational models** and **algorithms** to **interpret digital images and visual data** in order to **understand** the visual world we live in.

# Model versus Algorithm

Example: Interactive Segmentation



$$\mathbf{z} = (R, G, B)^n$$



Goal



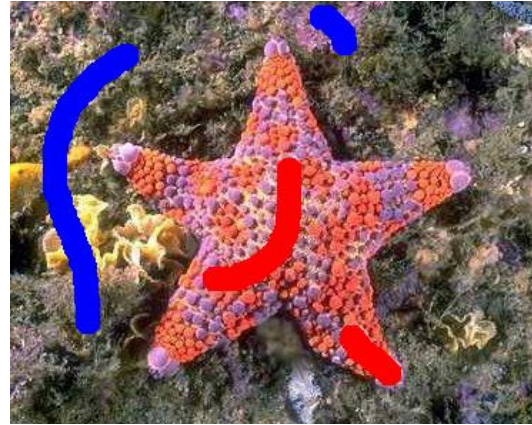
$$\mathbf{x} = \{0,1\}^n$$

Given  $\mathbf{z}$ ; derive binary  $\mathbf{x}$ :

**Model:** Energy function  $E(\mathbf{x})$

**Algorithm** to minimization:  $\mathbf{x}^* = \operatorname{argmin}_{\mathbf{x}} E(\mathbf{x})$

# Model for a starfish



Goal: formulate  $E(\mathbf{x})$  such that



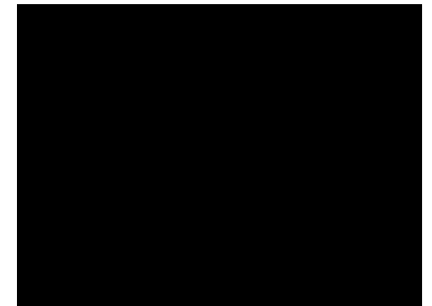
$$E(\mathbf{x}) = 0.01$$



$$E(\mathbf{x}) = 0.05$$



$$E(\mathbf{x}) = 0.05$$



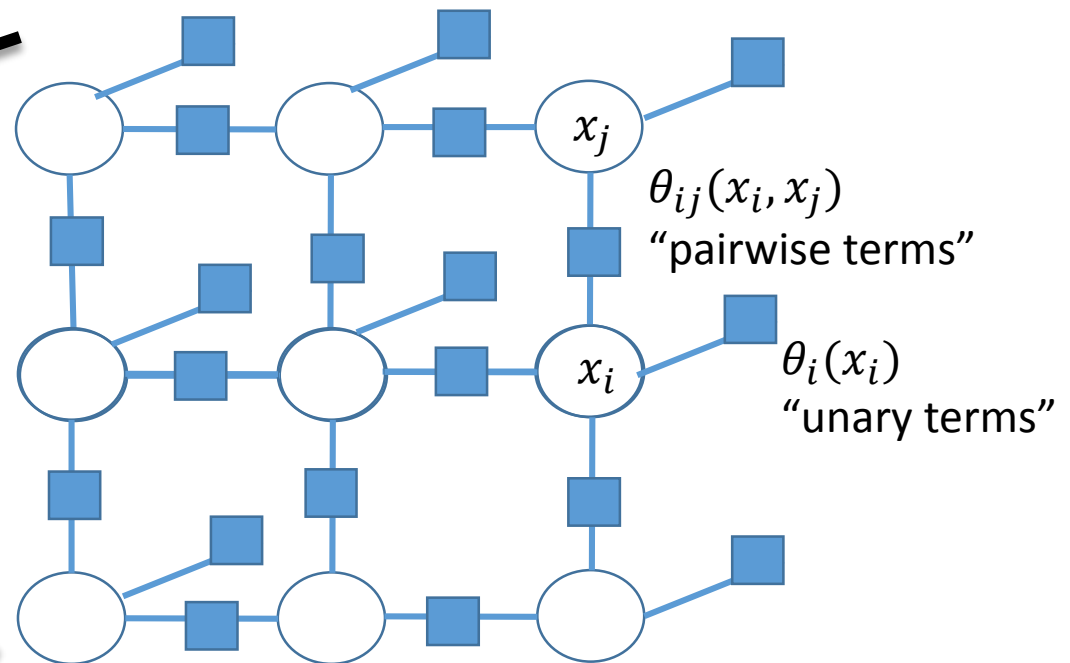
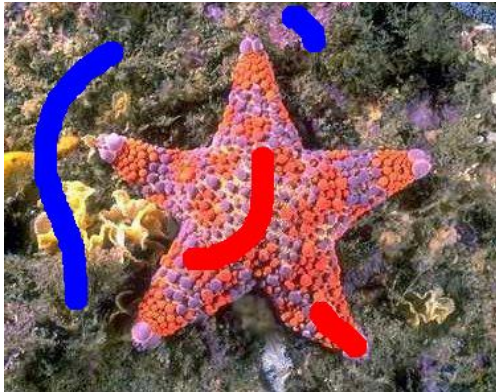
$$E(\mathbf{x}) = 0.1$$

Optimal solution  $\mathbf{x}^* = \operatorname{argmin}_{\mathbf{x}} E(\mathbf{x})$

# How does the energy looks like?

Energy function (sum of local terms):

$$E(\mathbf{x}) = \sum_i \theta_i(x_i) + \sum_{i,j} \theta_{ij}(x_i, x_j)$$



Undirected graphical model

**This is the focus of machine Learning 1 and 2**

# Why is computer vision interesting (to you)?

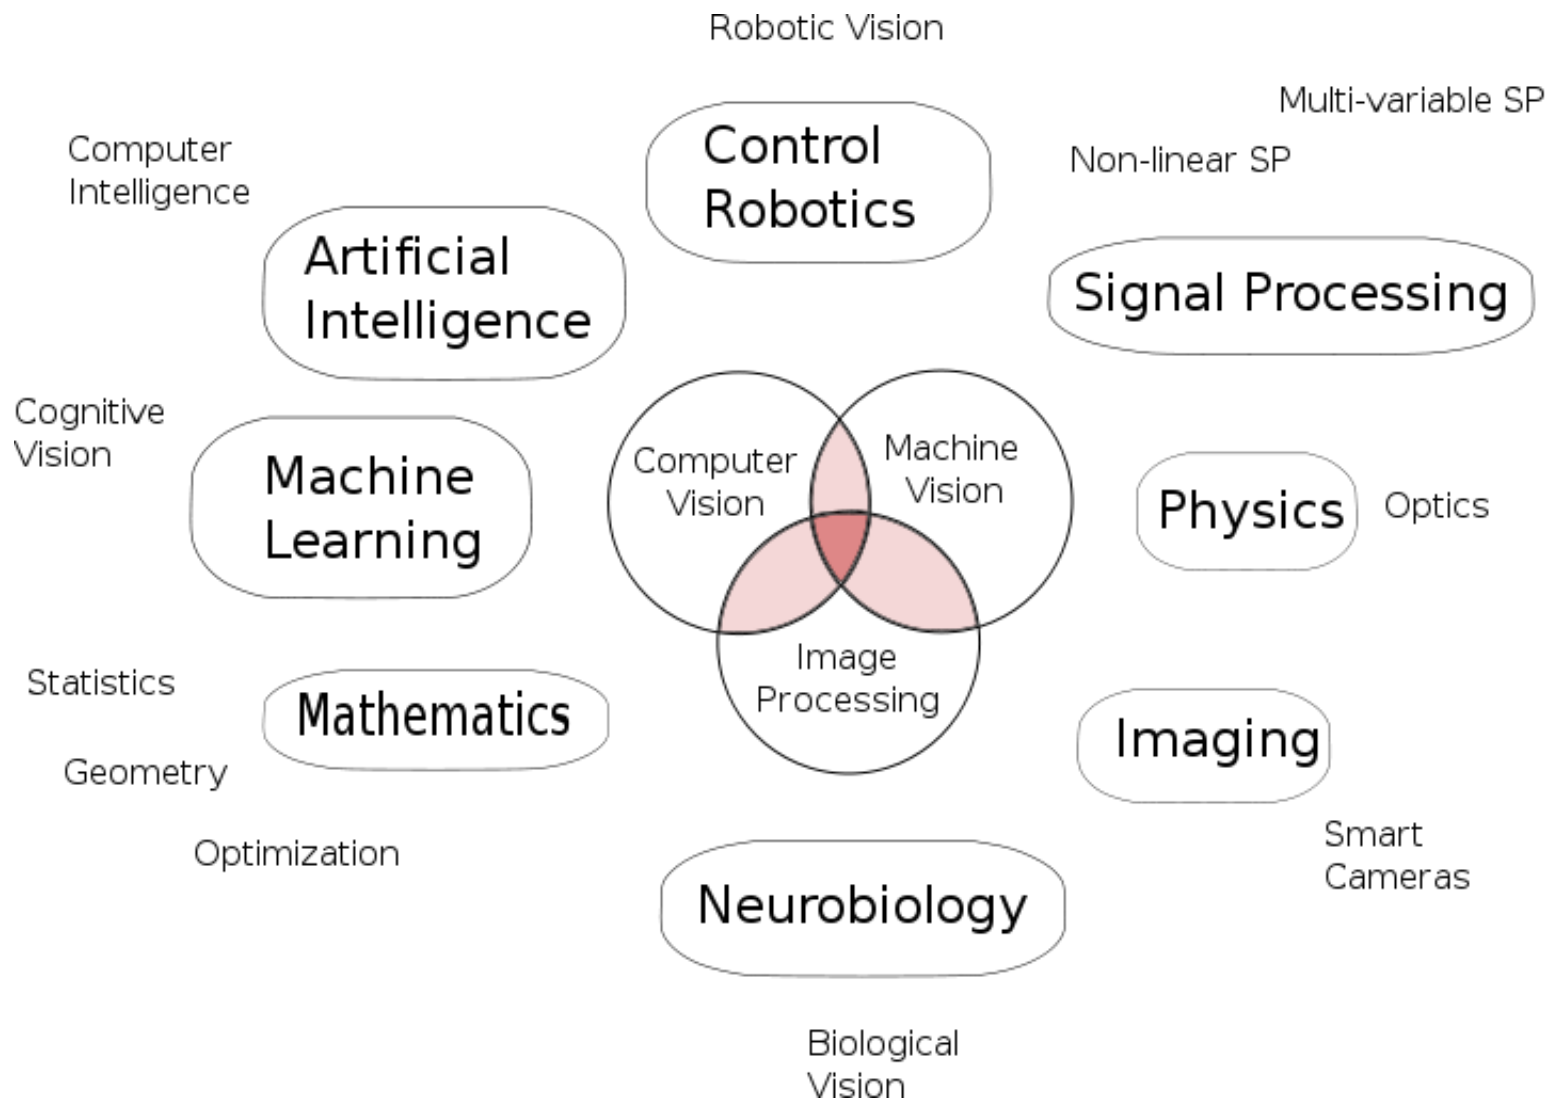
- It is a challenging problem that is far from being solved
- It combines insights and tools from many fields and disciplines:
  - Mathematics and statistics
  - Cognition and perception
  - Engineering (signal processing)
  - Computer science

# Why is computer vision interesting (to you)?

- Allows you to apply theoretical skills
  - ... that you may otherwise only use rarely
- Quite rewarding:
  - Often visually intuitive and encouraging results
- It is a growing field:
  - Cameras are becoming more and more popular
  - There are a lot of companies (big, small, start-up) working in vision
  - Conferences are growing rapidly
  - Deep Learning has revolutionised the field

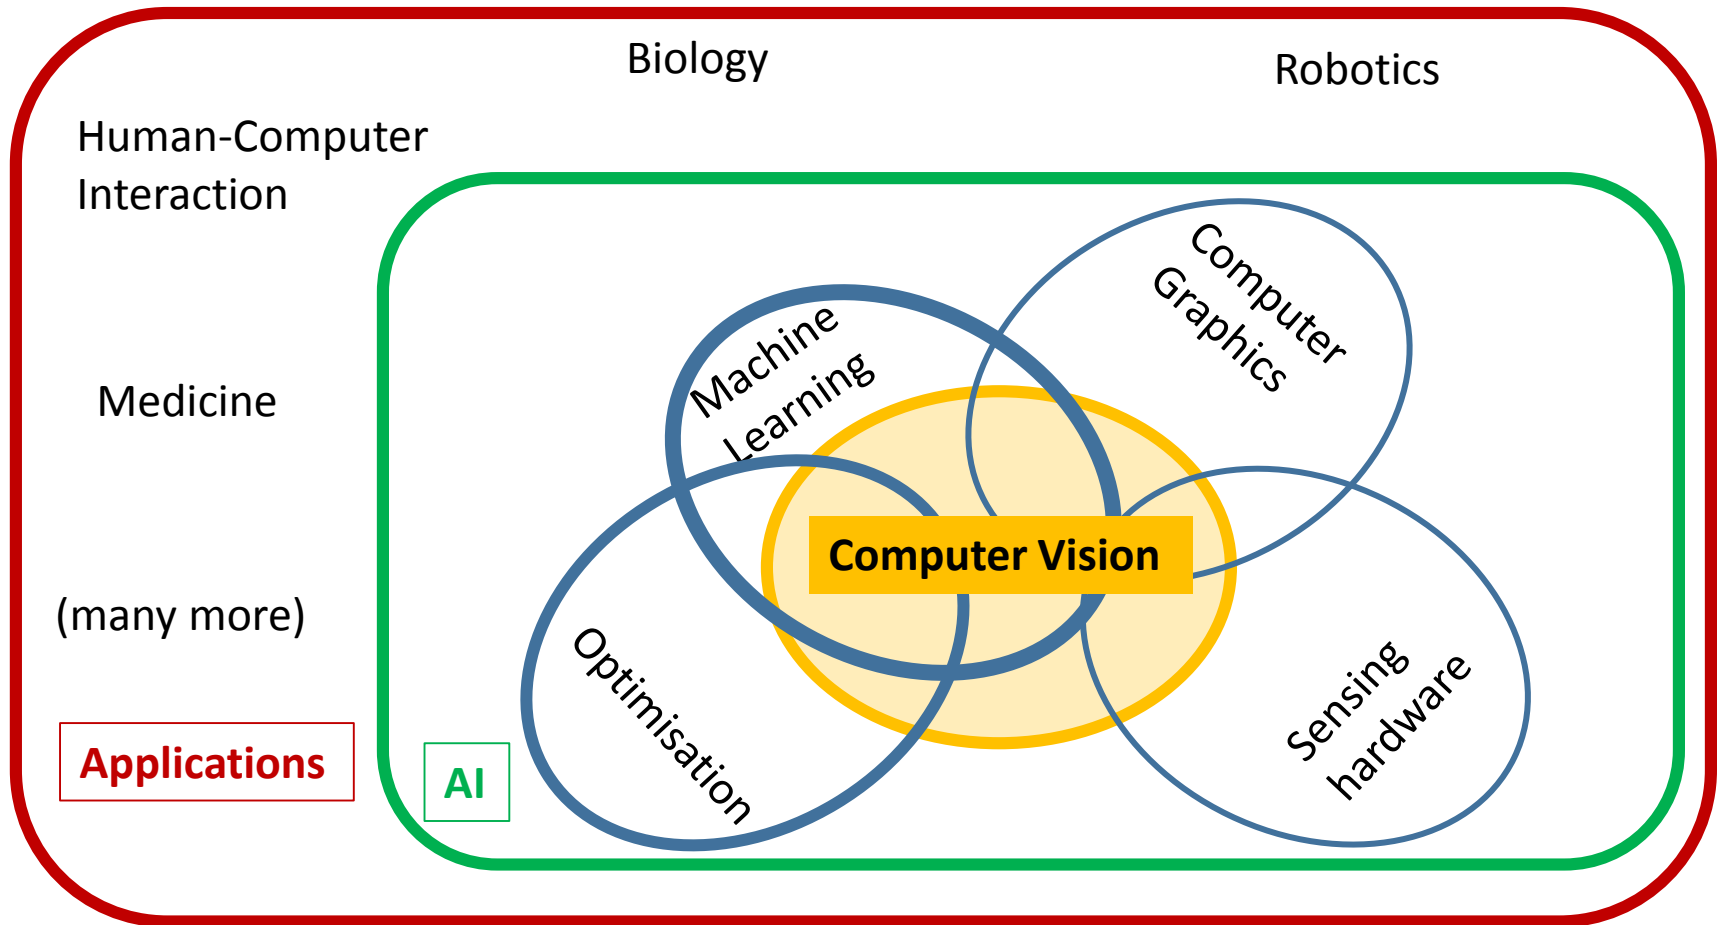


# Relationship to other fields



[Wikipedia]

# Relationship to other fields – my personal view



# Roadmap: Basics of Digital Image Processing

- What is an Image?
- Point operators (ch. 3.1)
- Filtering: (ch. 3.2, ch 3.3, ch. 3.4)
  - Linear filtering
  - Non-linear filtering
- Edges detection (ch. 4.2)
- Interest Point detection (ch. 4.1.1)

# What is an Image

- We can think of the image as a function:

$$I(x, y), \quad I: \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$$

- For every 2D point (pixel) it tells us the amount of light it receives
- The **size** and **range** of the sensor is limited:

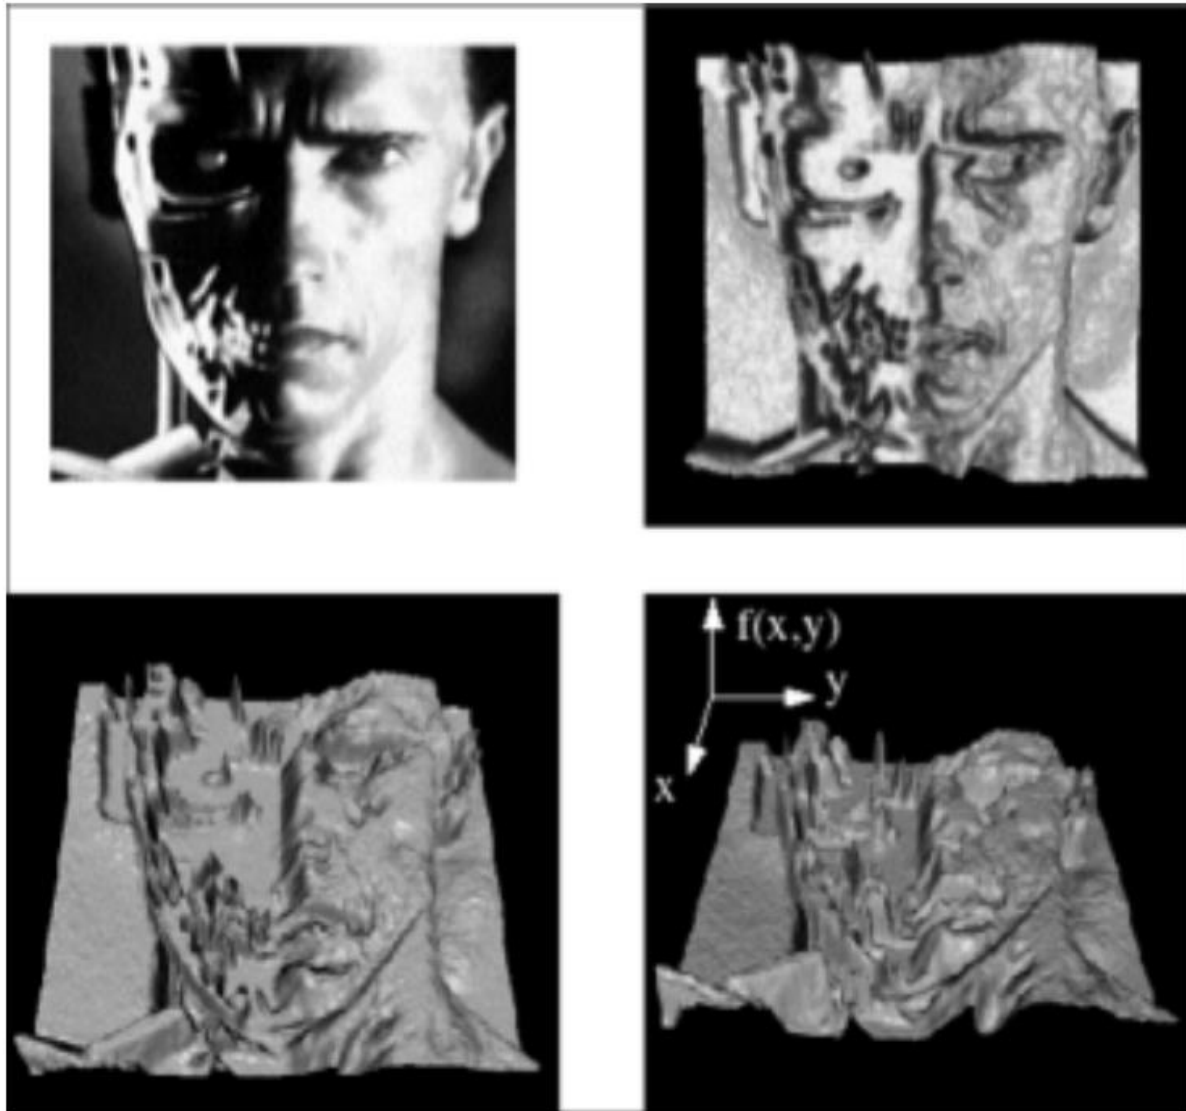
$$I(x, y), \quad I: [a, b] \times [c, d] \rightarrow [0, m]$$

- **Colour image** is then a vector-valued function:

$$I(x, y) = \begin{pmatrix} I_R(x, y) \\ I_G(x, y) \\ I_B(x, y) \end{pmatrix}, \quad I: [a, b] \times [c, d] \rightarrow [0, m]^3$$

- Comment, in most lectures we deal with grey-valued images and extension to colour is “obvious”

# Images as functions

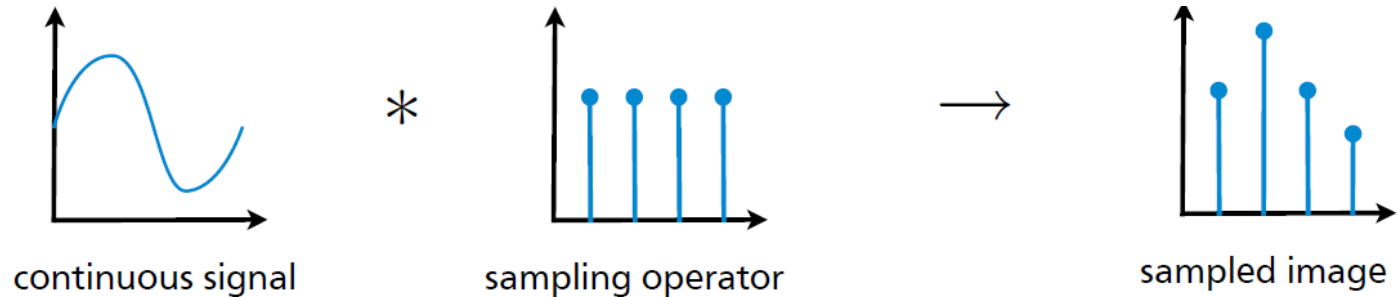


[from Steve Seitz]



# Digital Images

- We usually do not work with spatially continuous functions, since our cameras do not sense in this way.
- Instead we use (spatially) discrete images
- Sample the 2D domain on a regular grid (1D version)

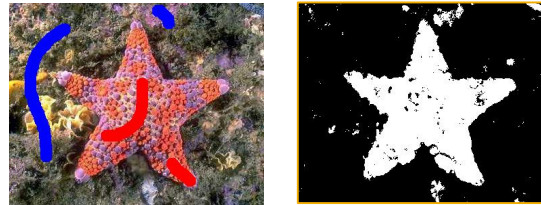


- Intensity/color values usually also discrete.  
Quantize the values per channel  
(e.g. 8 bit per channel)

		x =															
		58	59	60	61	62	63	64	65	66	67	68	69	70	71	72	
y =	41	210	209	204	202	197	247	143	71	64	80	84	54	54	57	58	
	42	206	196	203	197	195	210	207	56	63	58	53	53	61	62	51	
	43	201	207	192	201	198	213	156	69	65	57	55	52	53	60	50	
	44	216	206	211	193	202	207	208	57	69	60	55	77	49	62	61	
	45	221	206	211	194	196	197	220	56	63	60	55	46	97	58	106	
	46	209	214	224	199	194	193	204	173	64	60	59	51	62	56	48	
	47	204	212	213	208	191	190	191	214	60	62	66	76	51	49	55	
	48	214	215	215	207	208	180	172	188	69	72	55	49	56	52	56	
	49	209	205	214	205	204	196	187	196	86	62	66	87	57	60	48	
	50	208	209	205	203	202	186	174	185	149	71	63	55	55	45	56	
	51	207	210	211	199	217	194	183	177	209	90	62	64	52	93	52	
	52	208	205	209	209	197	194	183	187	187	239	58	68	61	51	56	
	53	204	206	203	209	195	203	188	185	183	221	75	61	58	60	60	
	54	200	203	199	236	188	197	183	190	183	196	122	63	58	64	66	
	55	205	210	202	203	199	197	196	181	173	186	105	62	57	64	63	

# Comment on Continuous Domain / Range

- There is a branch of computer vision research (“variational methods”), which operates on continuous domain for input images and output results
- Continuous domain methods are typically used for **physics-based vision**: segmentation, optical flow, etc.



- In this lecture and other lectures we mainly operate in **discrete domain** and **discrete or continuous range** for output results

# Roadmap: Basics of Digital Image Processing

- What is an Image?
- Point operators (ch. 3.1)
- Filtering: (ch. 3.2, ch 3.3, ch. 3.4)
  - Linear filtering
  - Non-linear filtering
- Edges detection (ch. 4.2)
- Interest Point detection (ch. 4.1.1)

# Point operators

- Point operators work on every pixel independently:

$$J(x, y) = h(I(x, y))$$

- Examples for  $h$ :

- Control contrast and brightness;  $h(z) = az^b + c$

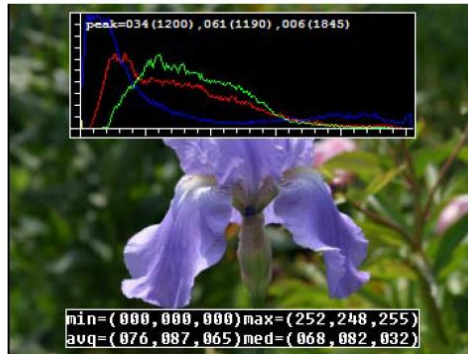


original

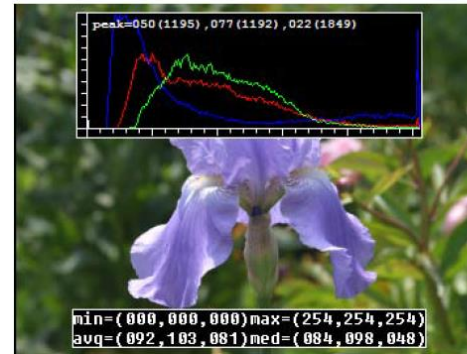


Contrast enhanced

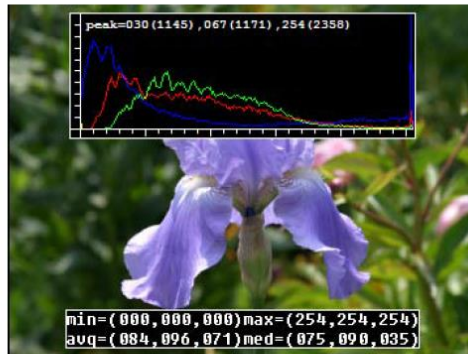
# Example



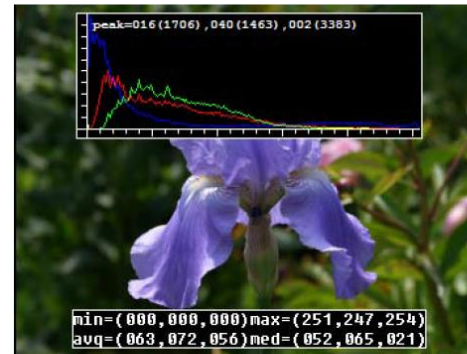
(a)



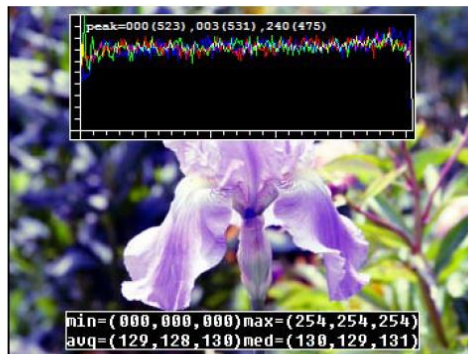
(b)



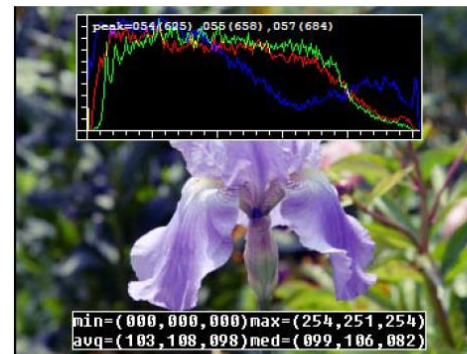
(c)



(d)



(e)



(f)



# Roadmap: Basics of Digital Image Processing

- What is an Image?
- Point operators (ch. 3.1)
- Filtering: (ch. 3.2, ch 3.3, ch. 3.4)
  - Linear filtering
  - Non-linear filtering
- Edges detection (ch. 4.2)
- Interest Point detection (ch. 4.1.1)

# Linear Filters / Operators

- Properties:
  - Homogeneity:  $T[aX] = aT[X]$
  - Additivity:  $T[X + Y] = T[X] + T[Y]$
  - Superposition:  $T[aX + bY] = aT[X] + bT[Y]$
- Example:
  - Convolution
  - Matrix-Vector operations

# Convolution

- Replace each pixel by a linear combination of its neighbours and itself
- 2D convolution (discrete)

$$g = f * h$$

45	60	98	127	132	133	137	133
46	65	98	123	126	128	131	133
47	65	96	115	119	123	135	137
47	63	91	107	113	122	138	134
50	59	80	97	110	123	133	134
49	53	68	83	97	113	128	133
50	50	58	70	84	102	116	126
50	50	52	58	69	86	101	120

image  
 $f(x, y)$

Centred at 0,0

0.1	0.5	0.1
0.1	0.2	0.1
0.1	0.1	0.1

\*

=

69	95	116	125	129	132
68	92	110	120	126	132
66	86	104	114	124	132
62	78	94	108	120	129
57	69	83	98	112	124
53	60	71	85	100	114

filtered image  
 $g(x, y)$

smaller  
output?

$$g(x, y) = \sum_{k,l} f(x - k, y - l)h(k, l) \quad \text{“the image } f \text{ is implicitly mirrored”}$$

# Convolution - Properties

- Linear  $h * (f_0 + f_1) = h * f_0 + h * f_1$
- Associative  $(f * g) * h = f * (g * h)$
- Commutative  $f * h = h * f$
- Can be written in Matrix form:  $g = H f$
- Correlation (not mirrored filter):

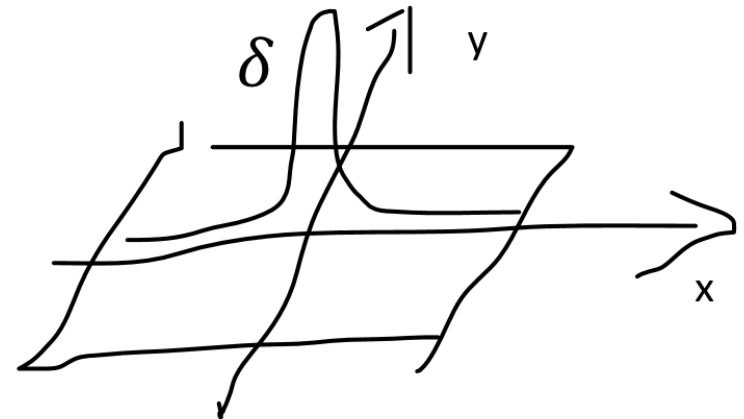
$$g(x, y) = \sum_{k,l} f(x + k, y + l)h(k, l)$$

$$\begin{bmatrix} 72 & 88 & 62 & 52 & 37 \end{bmatrix} * \begin{bmatrix} 1/4 & 1/2 & 1/4 \end{bmatrix} \Leftrightarrow$$

$$\frac{1}{4} \begin{bmatrix} 2 & 1 & . & . & . \\ 1 & 2 & 1 & . & . \\ . & 1 & 2 & 1 & . \\ . & . & 1 & 2 & 1 \\ . & . & . & 1 & 2 \end{bmatrix} \begin{bmatrix} 72 \\ 88 \\ 62 \\ 52 \\ 37 \end{bmatrix}$$

# Examples

- Impulse function:  $f = f * \delta$



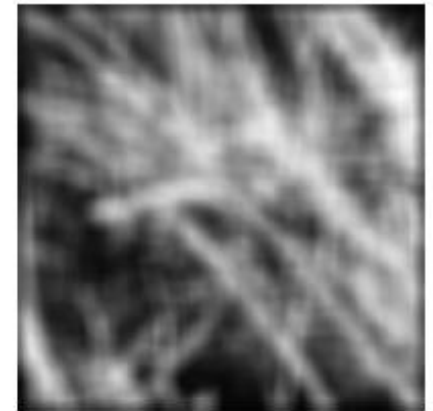
- Box Filter:

$$\frac{1}{9} \cdot \begin{bmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{bmatrix}$$

Original Image



Box-filtered image





# Application: Noise removal

- Noise is what we are not interested in:  
sensor noise (Gaussian, shot noise), quantisation artefacts, etc
- Typical assumption is that the noise is not correlated between pixels
- Basic Idea: neighbouring pixel contain information about intensity

2	3	3
3	20	2
3	2	3

 → 

2	3	3
3	3	2
3	2	3

# The box filter does noise removal

- Box filter takes the mean in a neighbourhood



Image



Noise



Pixel-independent  
Gaussian noise added

$$\frac{1}{9} \cdot \begin{array}{|c|c|c|} \hline 1 & 1 & 1 \\ \hline 1 & 1 & 1 \\ \hline 1 & 1 & 1 \\ \hline \end{array}$$

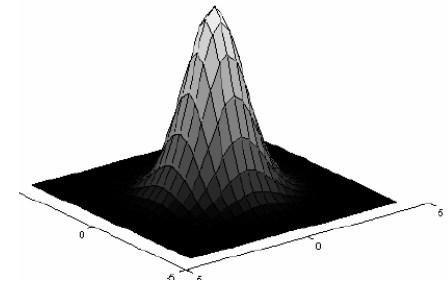
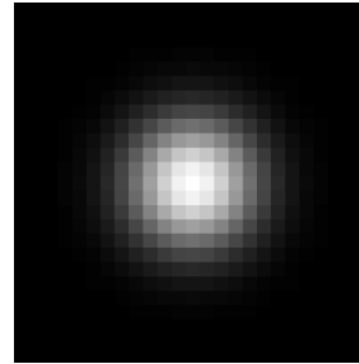


Filtered Image

# Gaussian (Smoothing) Filters

- Nearby pixels are weighted more than distant pixels
- Isotropic Gaussian (rotational symmetric)

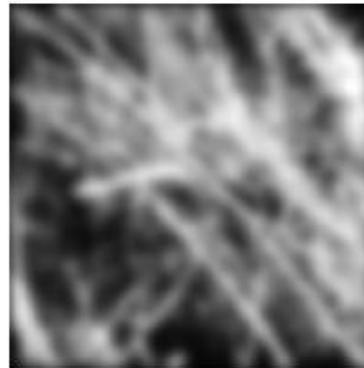
$$G_{\sigma}(x, y) = \frac{1}{2\pi\sigma^2} \exp\left(-\frac{x^2 + y^2}{2\sigma^2}\right)$$



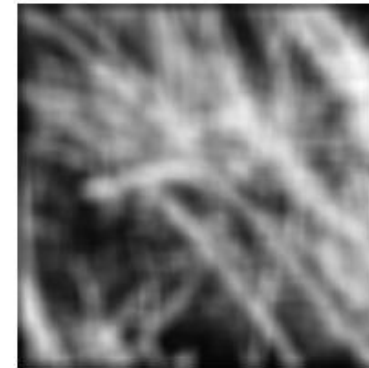
Original Image



Gaussian-filtered image

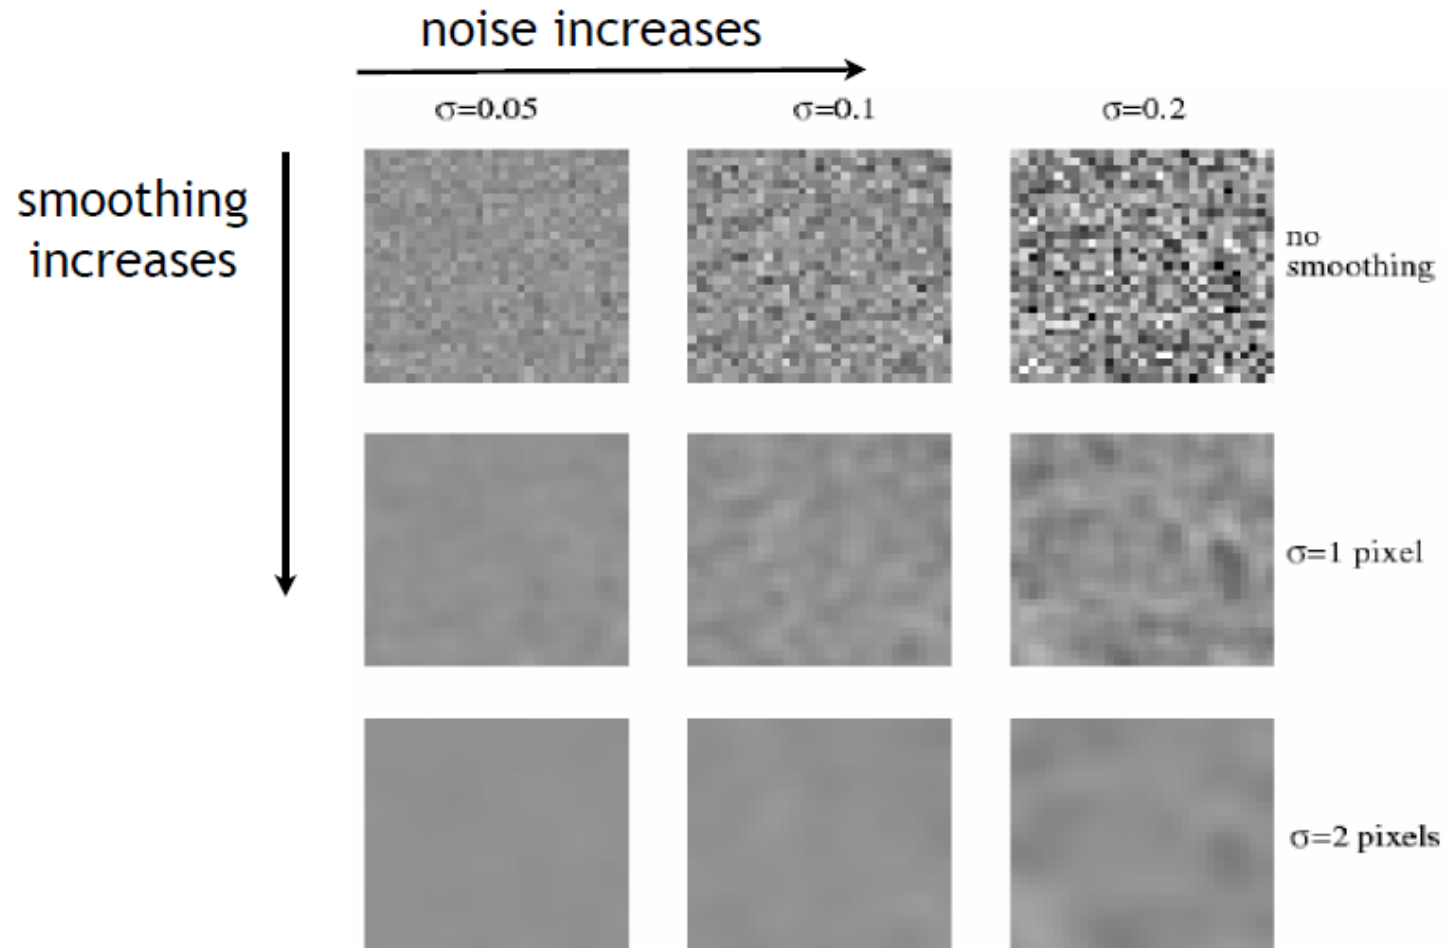


Box-filtered image



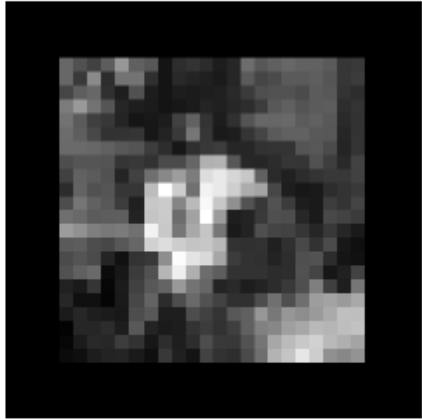
# Gaussian Filter

Input: constant grey-value image

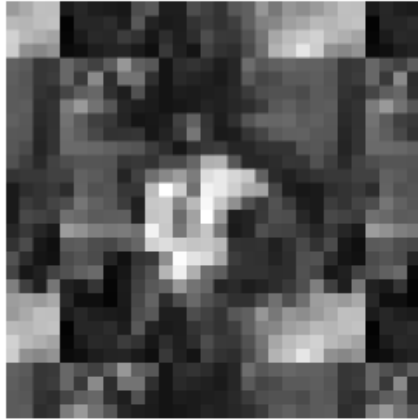


More noise needs larger sigma

# Handling the Boundary (Padding)



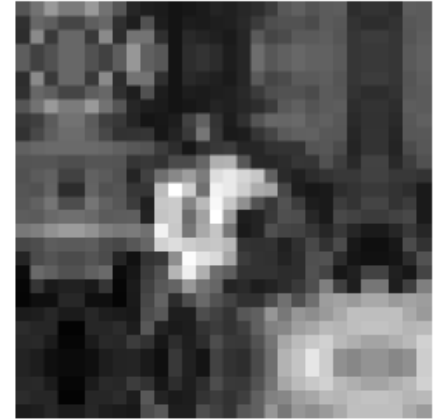
zero



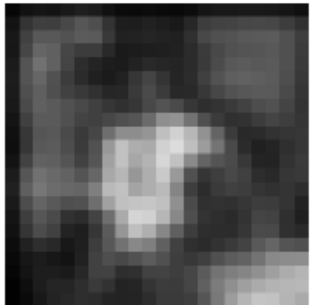
wrap



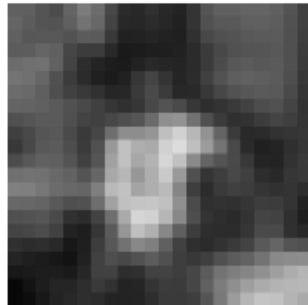
clamp



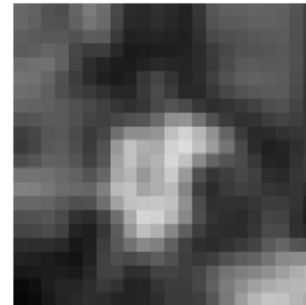
mirror



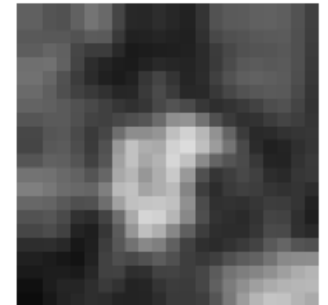
blurred zero



normalized zero



blurred clamp



blurred mirror

# How to compute convolution efficiently?

- Separable filters (next)
- Fourier transformation
- Integral Image trick (see exercise)

## Important for later (integral Image trick):

- Naive implementation would be  $O(Nw)$   
where  $w$  is the number of elements in box filter
- The Box filter (and a few others) can be computed in  $O(N)$

45	60	98	127	132	133	137	133
46	65	98	123	126	128	131	133
47	65	96	115	119	123	135	137
47	63	91	107	113	122	138	134
50	59	80	97	110	123	133	134
49	53	68	83	97	113	128	133
50	50	58	70	84	102	116	126
50	50	52	58	69	86	101	120

image

\*

0.1	0.1	0.1
0.1	0.2	0.1
0.1	0.1	0.1

filter (kernel)

=

69	95	116	125	129	132
68	92	110	120	126	132
66	86	104	114	124	132
62	78	94	108	120	129
57	69	83	98	112	124
53	60	71	85	100	114

filtered image

$$\frac{1}{9} \cdot \begin{bmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{bmatrix}$$

$$g(x, y) = \sum_{k,l} f(x - k, y - l)h(k, l)$$



# Separable filters

For some filters we have:  $f * h = f * (h_x * h_y)$

Where  $h_x, h_y$  are 1D filters.

Example Box filter:

$$\frac{1}{9} \cdot \begin{array}{|c|c|c|} \hline 1 & 1 & 1 \\ \hline 1 & 1 & 1 \\ \hline 1 & 1 & 1 \\ \hline \end{array} = \frac{1}{3} \cdot \begin{array}{|c|c|c|} \hline 1 & 1 & 1 \\ \hline \end{array} * \frac{1}{3} \cdot \begin{array}{|c|} \hline 1 \\ \hline 1 \\ \hline 1 \\ \hline \end{array}$$

$h_x * h_y$                        $h_x$                        $h_y$

Now we can do two 1D convolutions:

$$f * h = f * (h_x * h_y) = (f * h_x) * h_y$$

Naïve implementation for 3x3 filter: 9N operations versus 3N+3N operations

# Can any filter be made separable?

Note: 
$$\frac{1}{9} \cdot \begin{matrix} & h_x * h_y \\ \begin{matrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{matrix} \end{matrix} = \frac{1}{3} \cdot \begin{matrix} & h_x \\ \begin{matrix} 1 & 1 & 1 \end{matrix} \end{matrix} * \frac{1}{3} \cdot \begin{matrix} & h_y \\ \begin{matrix} 1 \\ 1 \\ 1 \end{matrix} \end{matrix} = \frac{1}{3} \cdot \begin{matrix} & h_x \\ \begin{matrix} 1 \\ 1 \\ 1 \end{matrix} \end{matrix} \cdot \frac{1}{3} \cdot \begin{matrix} & h_y \\ \begin{matrix} 1 & 1 & 1 \end{matrix} \end{matrix}$$

Apply SVD to the kernel matrix:

$$\begin{aligned} \mathbf{A} &= \left[ \begin{array}{c|c|c} \mathbf{u}_0 & \cdots & \mathbf{u}_{p-1} \end{array} \right] \left[ \begin{array}{ccc} \sigma_0 & & \\ & \ddots & \\ & & \sigma_{p-1} \end{array} \right] \left[ \begin{array}{c} \mathbf{v}_0^T \\ \cdots \\ \mathbf{v}_{p-1}^T \end{array} \right] \\ &= \sum_{j=0}^{p-1} \sigma_j \mathbf{u}_j \mathbf{v}_j^T, \end{aligned}$$

If all  $\sigma_i$  are 0 (apart from  $\sigma_0$ ) then it is separable.

# Example of separable filters

