CMP717
Image Processing
Feb. 15, 2017

Introduction

Erkut Erdem
Hacettepe University
Computer Vision Lab (HUCVL)

Today

- About me
- About you
- Course outline and logistics
- Introduction to Image Processing

Today

- About me
- About you
- Course outline and logistics
- Introduction to Image Processing

Erkut Erdem

Computer Vision Lab

http://web.cs.hacettepe.edu.tr/~erkut

My research interests concern computer vision and machine learning. I specifically investigate the use of spatial, temporal and cross-modal context in visual processing. My recent research activities cover topics such as saliency prediction, integrated vision and language, image editing and HDR image processing.

2001
BSc, Comp. Eng.
METU, Ankara

2003
MSc, Comp. Eng.
METU, Ankara

2004
Visiting Research Scholar
Virginia Tech, Blacksburg

2006
Visiting Student
UCLA, Los Angeles

2008
PhD, Comp. Eng.
METU, Ankara

2009
Post-doc Researcher
Telecom ParisTech (ENST), Paris

2012
Founder
Computer Vision Lab

2014
Assistant Professor
Hacettepe, Ankara

Structure Preserving Image Smoothing via Region Covariances

Levent Karacan,
Erkut Erdem,
Aykut Erdem

2013

ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia 2013)

Image Matting with KL-Divergence Based Sparse Sampling

Levent Karacan,
Aykut Erdem,
Erkut Erdem

2015

IEEE International Conference on Computer Vision (ICCV 2015)



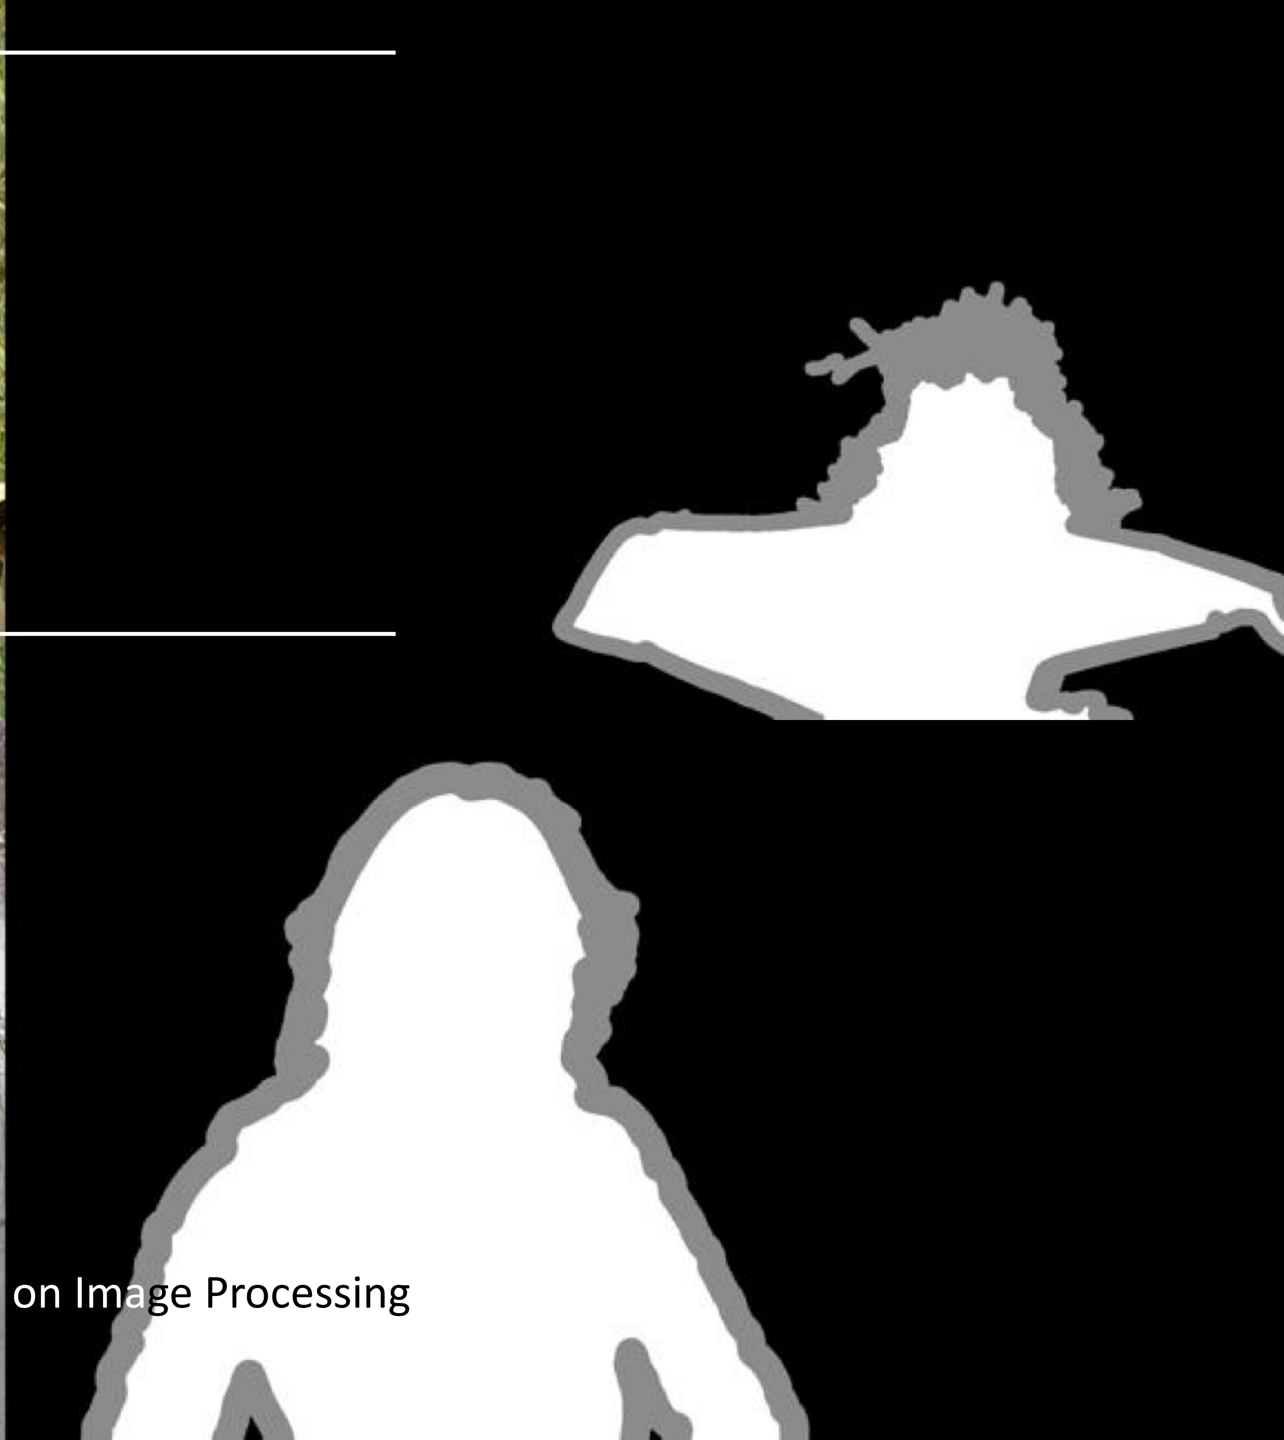

Alpha Matting with KL-Divergence Based Sparse Sampling

Levent Karacan,
Aykut Erdem
Erkut Erdem

2017

IEEE Transactions on Image Processing

Visual saliency estimation by nonlinearly integrating features using region covariances

Erkut Erdem,
Aykut Erdem

2013 Journal of Vision

Top down saliency estimation via superpixel based discriminative dictionaries

Aysun Kocak,
Kemal Cizmeciler,
Aykut Erdem,
Erkut Erdem

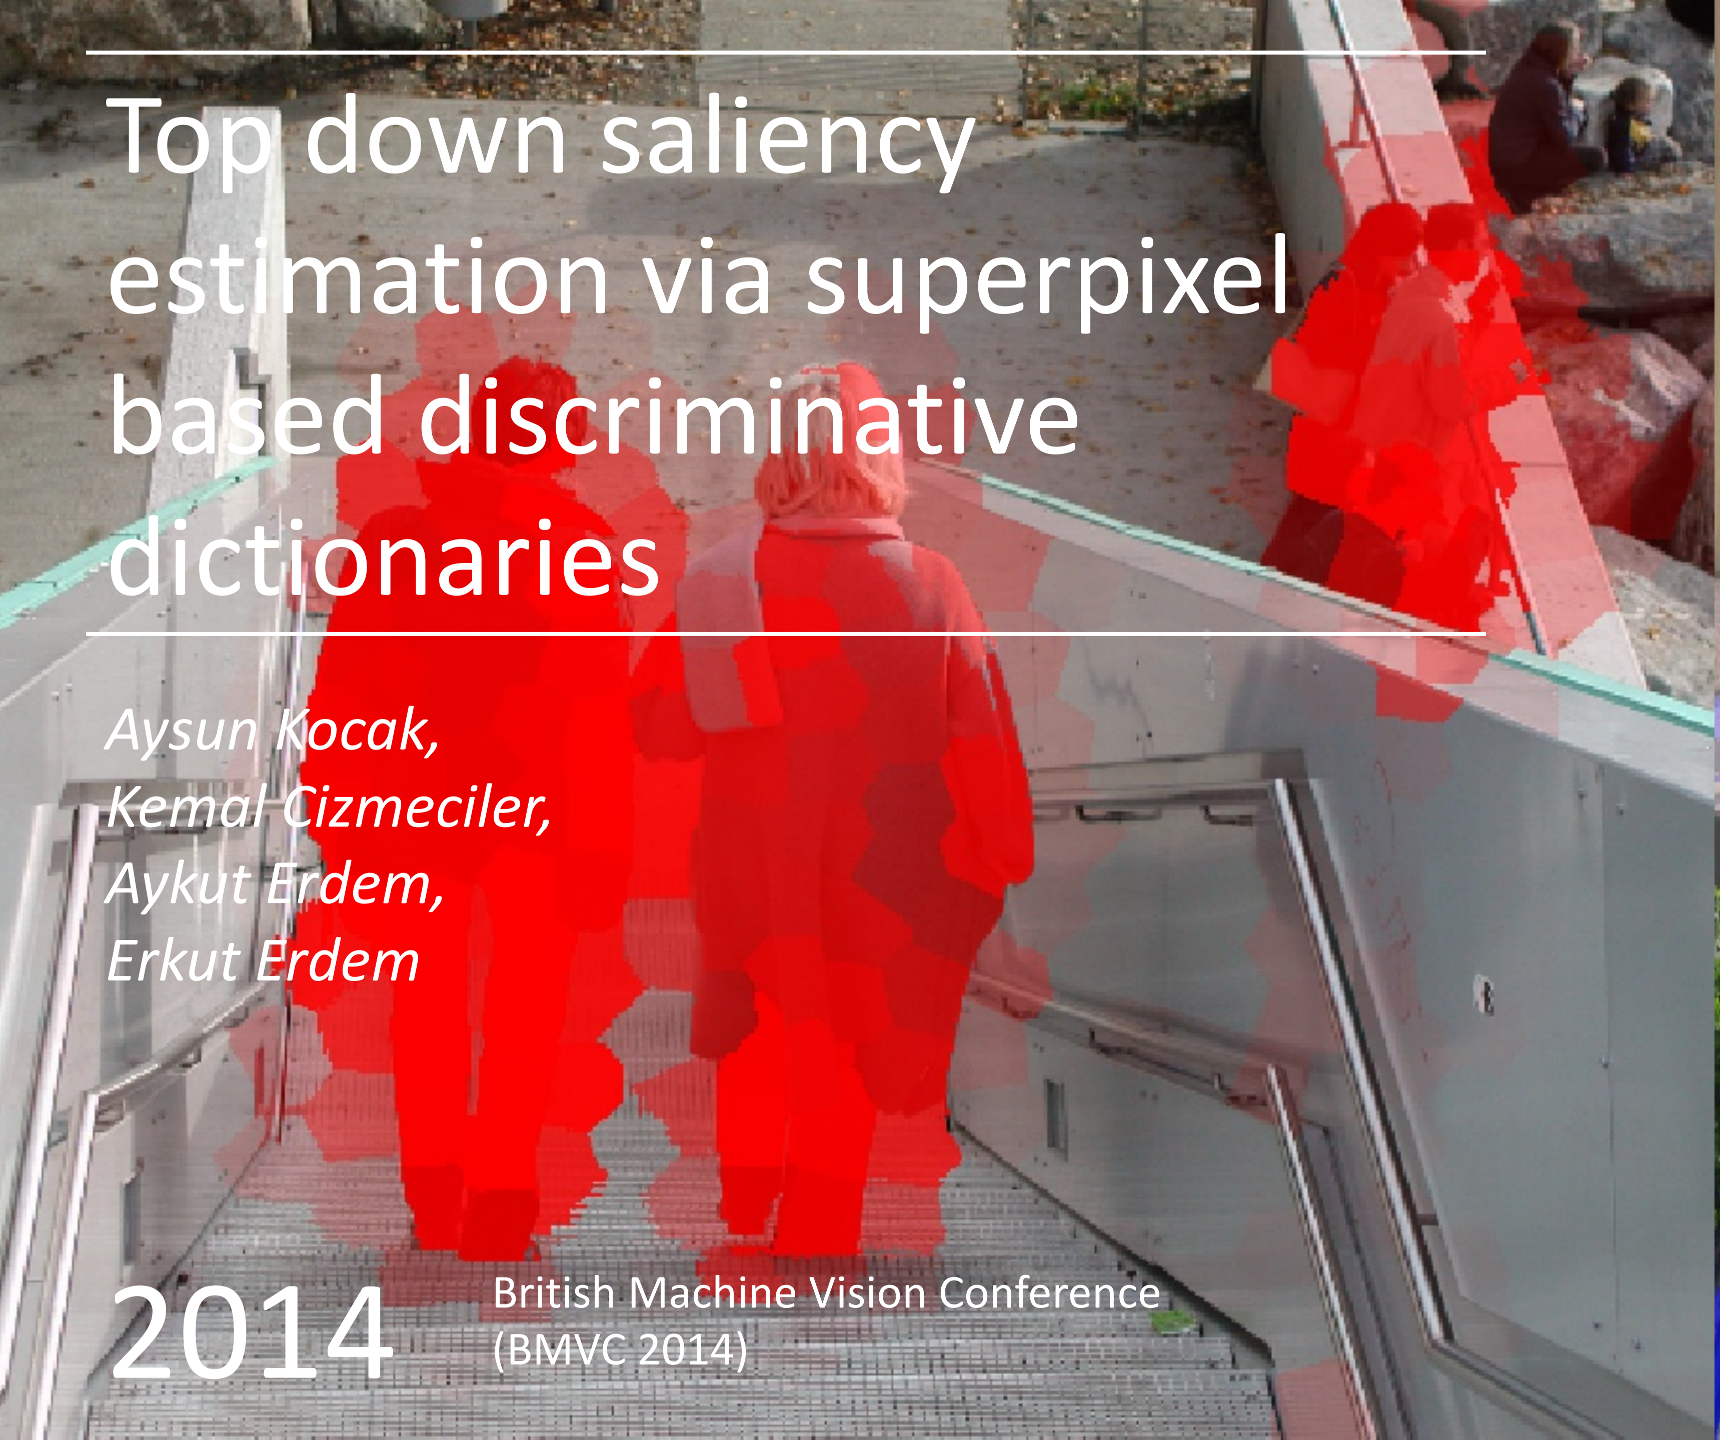

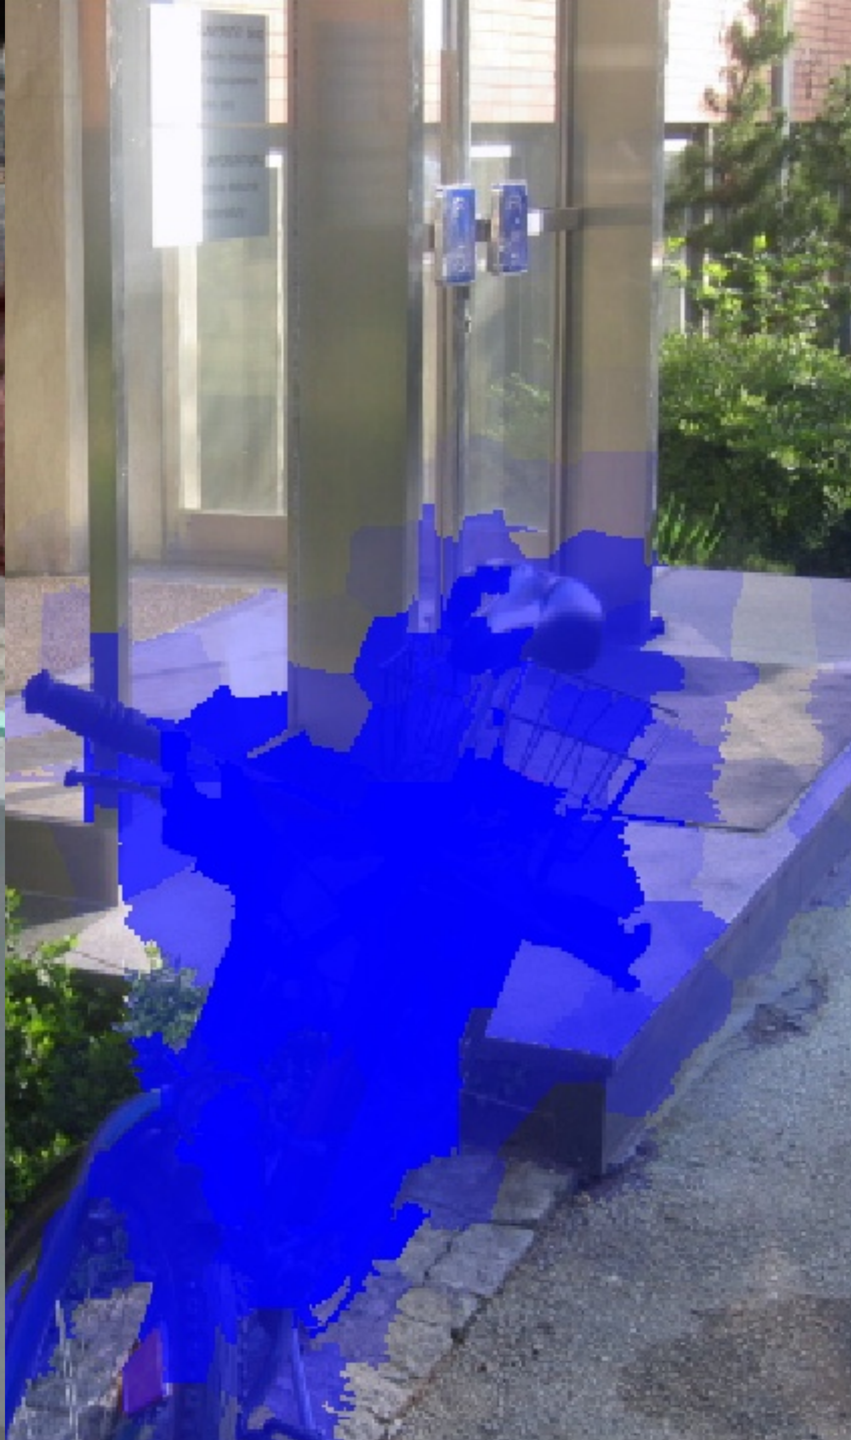

2014

British Machine Vision Conference
(BMVC 2014)

Two-Stream Convolutional Networks for Dynamic Saliency Prediction

Cagdas Bak,
Aykut Erdem,
Erkut Erdem

2018 IEEE Transactions on Multimedia

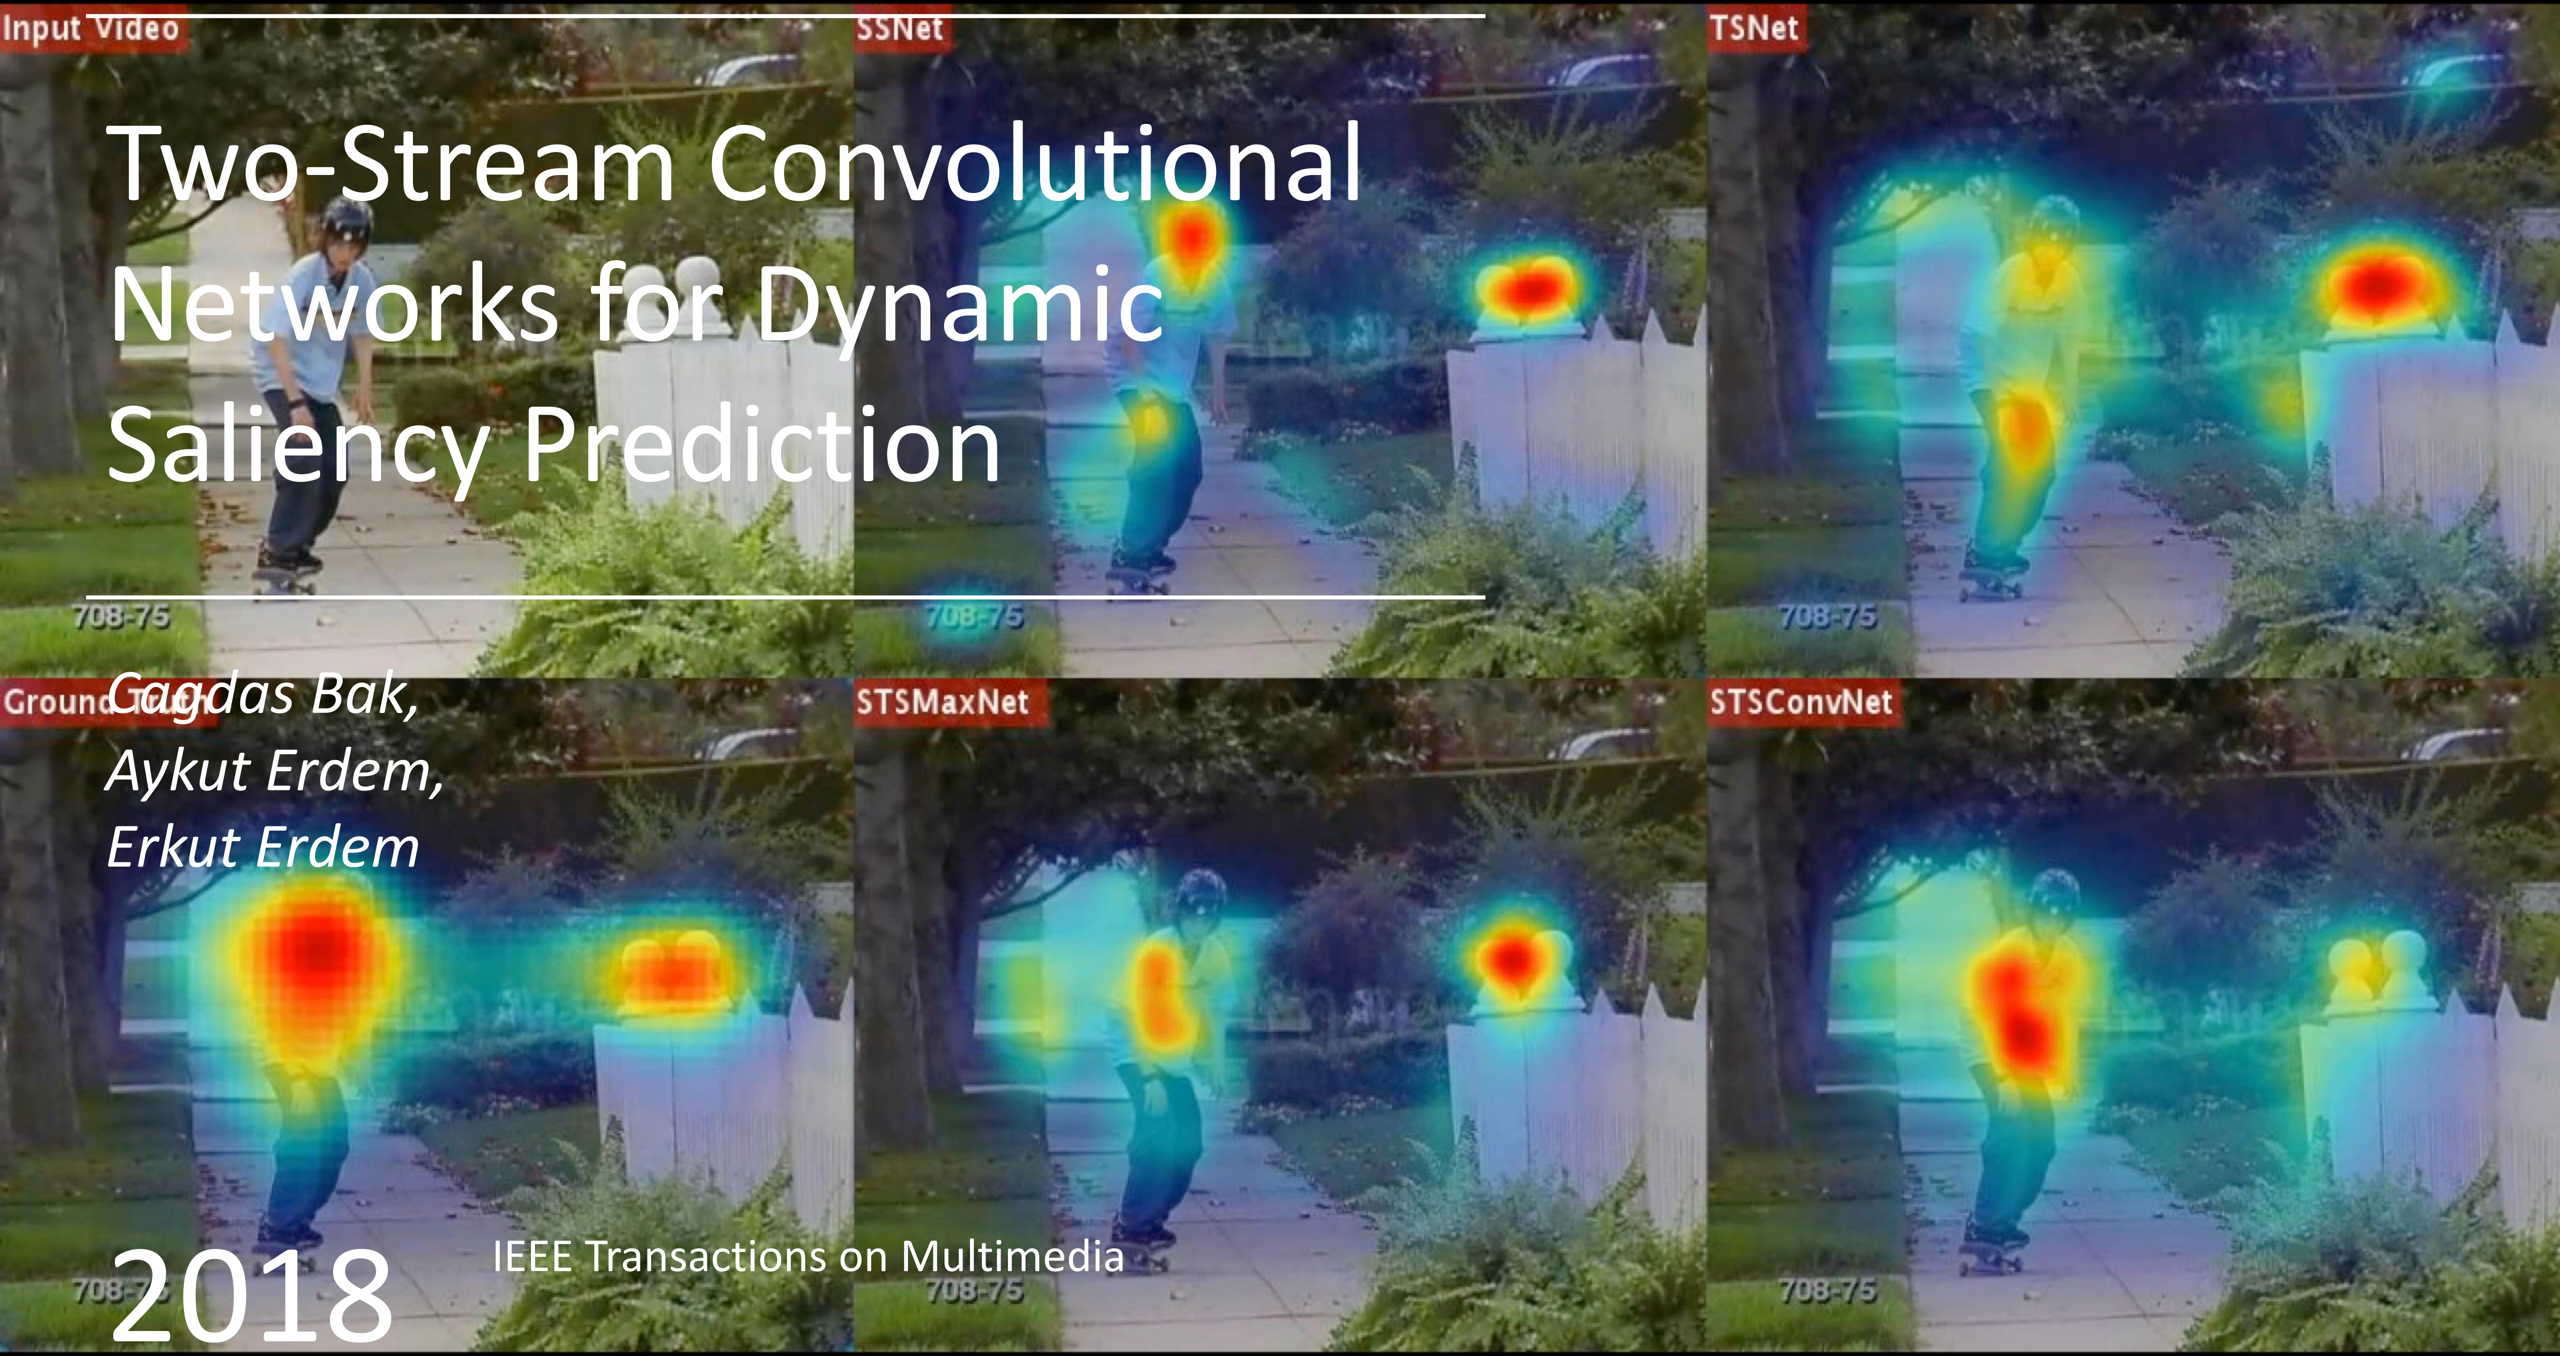

A Comparative Study for Feature Integration Strategies in Dynamic Saliency Estimation

Yasin Kavak,
Erkut Erdem,
Aykut Erdem

2017

Signal Processing: Image Communication

The State of the Art in HDR Deghosting: A Survey and Evaluation

Okan Tarhan Tursun,
Ahmet Oguz Akyuz,
Aykut Erdem,
Erkut Erdem

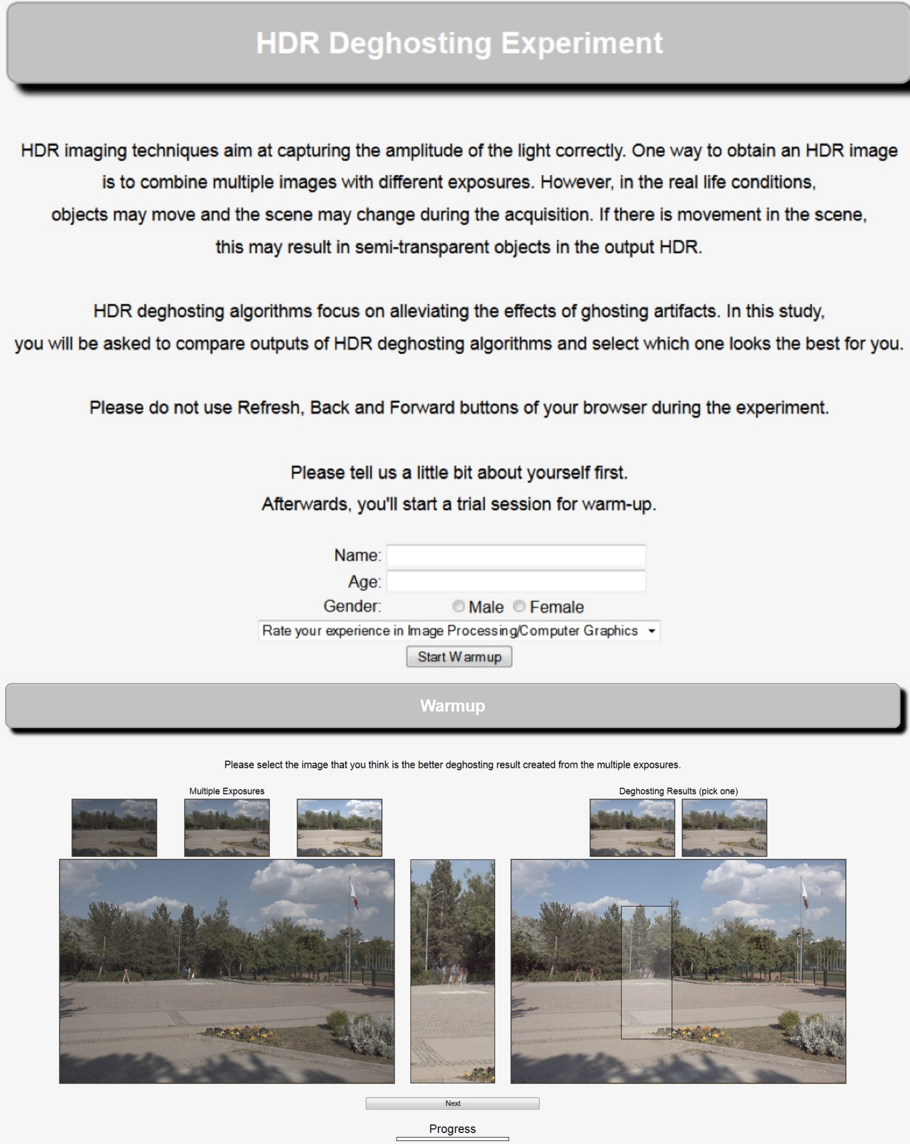

2015 Computer Graphics Forum (Eurographics STAR 2015)

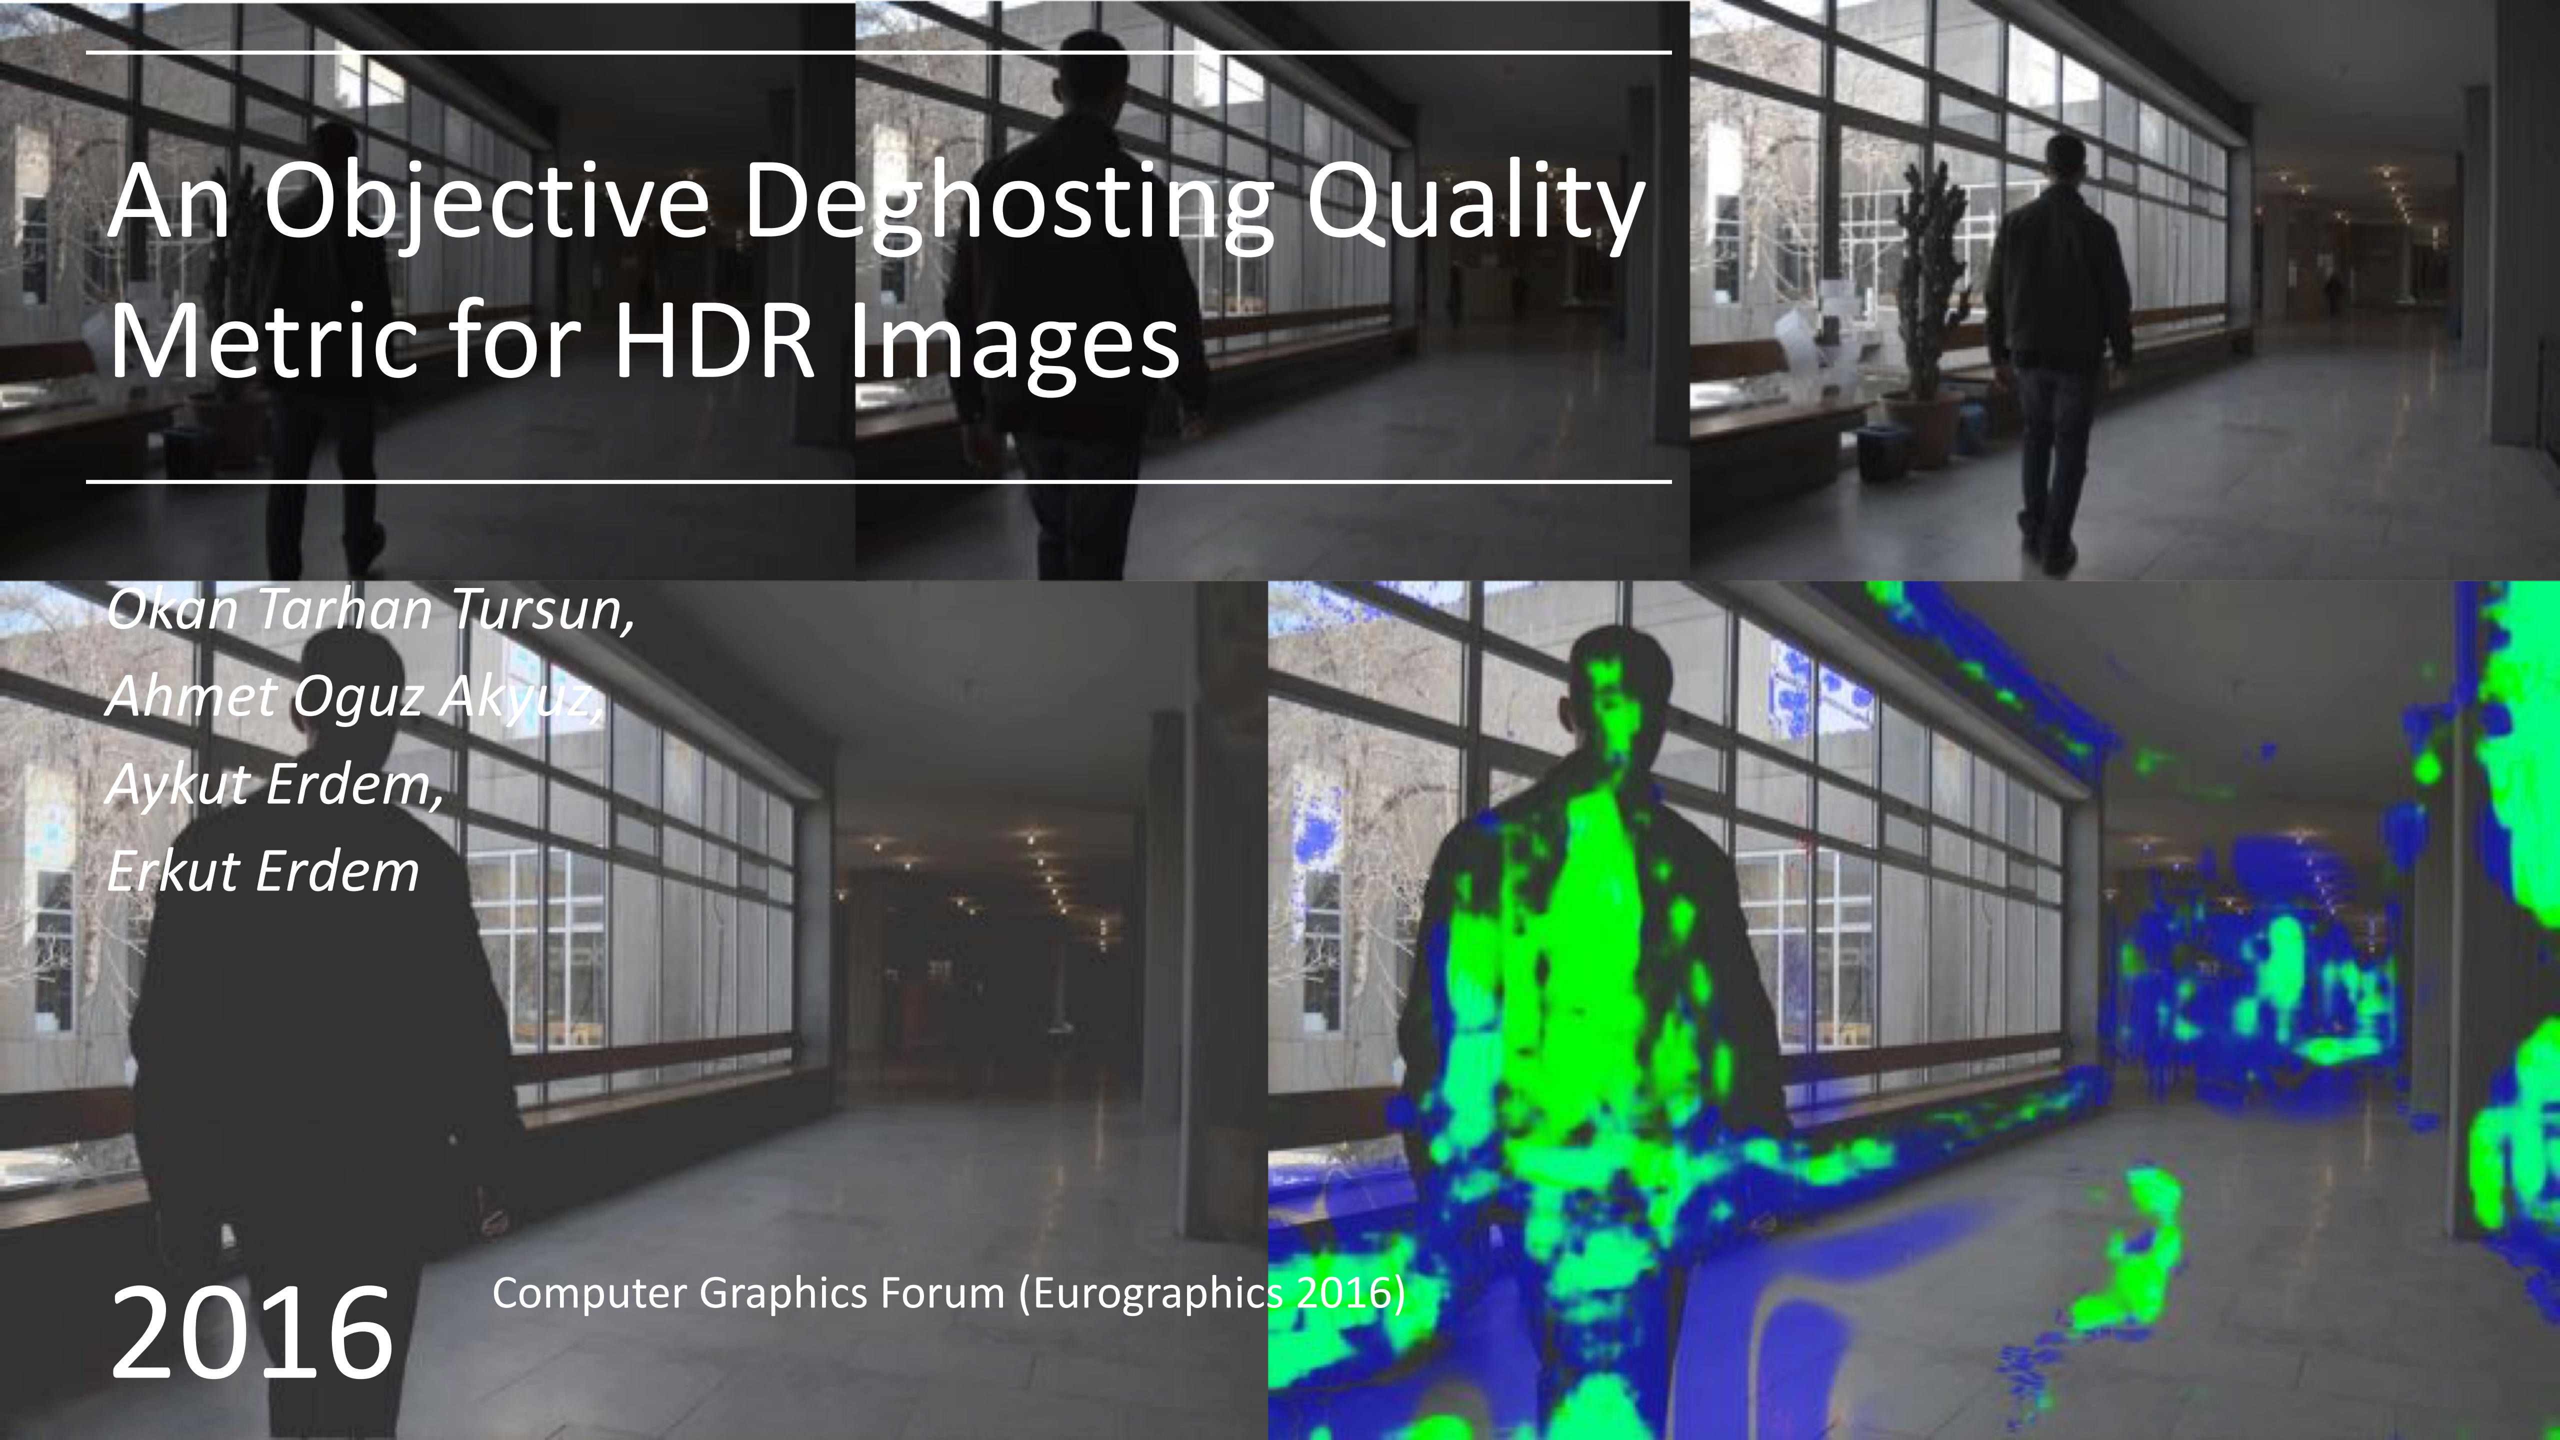

An Objective Deghosting Quality Metric for HDR Images

Okan Tarhan Tursun,
Ahmet Oguz Akyuz,
Aykut Erdem,
Erkut Erdem

2016

Computer Graphics Forum (Eurographics 2016)

Today

- About me
- About you
- Course outline and logistics
- Introduction to Image Processing

Introduce yourselves

- Who are you?
- What do you know about image processing?
- Why you want to take CMP717?

- Send me a short e-mail including your answers to these questions.

Today

- About me
- About you
- Course outline and logistics
- Introduction to Image Processing

Logistics

- Asst. Prof. Erkut ERDEM
- erkut@cs.hacettepe.edu.tr
- Office: 112

- Lectures: Thursday, 09:30-12:30
- Office Hour: *By appointment.*

About CMP717

- This course provides a comprehensive overview of fundamental topics in image processing for graduate students.
- The goal of this course is to provide a deeper understanding of the state-of-the-art methods in image processing literature and to study their connections.
- The course makes the students gain knowledge and skills in key topics and provides them the ability to employ them in their advanced-level studies.

Communication

- The course webpage will be updated regularly throughout the semester with lecture notes, programming and reading assignments and important deadlines.
  http://web.cs.hacettepe.edu.tr/~erkut/cmp717.s18

- All other communications will be carried out through Piazza. Please enroll it by following the link
  https://piazza.com/hacettepe.edu.tr/spring2018/cmp717

Prerequisites

- Good programming skills (for practicals and the course project)
- Calculus (differentiation, chain rule) and linear algebra (vectors, matrices, eigenvalues/vectors)
- Basic probability and statistics (random variables, expectations, multivariate Gaussians, Bayes rule, conditional probabilities)
- Undergraduate level image processing (e.g. BBM413)
- Machine learning (e.g. BBM406 and CMP712)
- Optimization (cost functions, taking gradients, regularization)

Reference Books

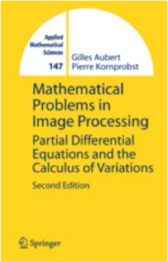

- Mathematical Problems in Image Processing: Partial Differential Equations and the Calculus of Variations, G. Aubert and P. Kornprobst, 2nd Edition, Springer-Verlag, 2006

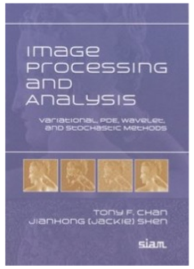

- Image Processing And Analysis: Variational, PDE, Wavelet, And Stochastic Methods, T. Chan and J. Shen, Society for Industrial and Applied Mathematics, 2005

- Markov Random Fields For Vision And Image Processing, Edited by A. Blake, P. Kohli and C. Rother, MIT Press, 2011

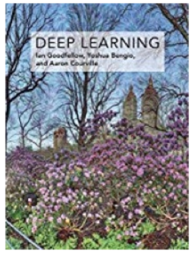

- Deep Learning, Ian Goodfellow, Aaron Courville, and Yoshua Bengio, preparation for MIT Press,

Reading Material

- Lecture notes and handouts
- Papers and journal articles

Grading Policy

- Midterm Exam (20%)
- Paper Presentations (14%)
- Weekly Quizzes (10%)
- Practicals (24%) (3 practicals x 8% each)
- Course Project (presentations and reports) (32%)

Paper presentations and Quizzes

- The students will be required to present at least one research paper either of their choice or from the suggested reading list.
- These papers should be read by every student as the quizzes about the presented papers will be given on the weeks of the presentations.
- <u>The schedule for the presentations will be determined shortly.</u>

Structures of paper presentations

- High-level overview with contributions
- Main motivation
- Clear statement of the problem
- Overview of the technical approach
- Strengths/weaknesses of the approach
- Overview of the experimental evaluation
- Strengths/weaknesses of evaluation
- Discussion: future direction, links to other work

Practicals

- 3 practicals (8% each)
- Involves implementation, analysis, reporting
- Should be done individually

- Late policy: You have 5 skip days in the semester.

- Tentative Dates
  - Practical 1 Out: February 22nd, Due: March 8th
  - Practical 2 Out: March 15th, Due: Mar 29th
  - Practical 3 Out: April 5th, Due: April 19th

Project

- Aim: To give the students some experience on conducting research.
- Students should work individually or groups in two.
- This project may involve
  - Design of a novel approach and its experimental analysis,
  - an extension to a recent study of non-trivial complexity and its experimental analysis
- Deliverables

- Proposals	March 22, 2018
- Project progress reports	April 12, 2018
- Final project presentations	May 17, 2018
- Final reports	June 8, 2018

Tentative Outline

- Overview of Image Processing
- Linear Filtering, Edge/Boundary Detection, Image Segmentation
- Nonlinear Filtering, Snakes, Variational Segmentation Models
- Modern Image Filtering
- Sparse Coding

Tentative Outline

- Graphical Models
- Deep Learning Basics
- Convolutional Neural Networks
- Semantic Segmentation
- Image Deblurring
- Visual Saliency
- Deep Generative Networks

Today

- About me
- About you
- Course outline and logistics
- Introduction to Image Processing

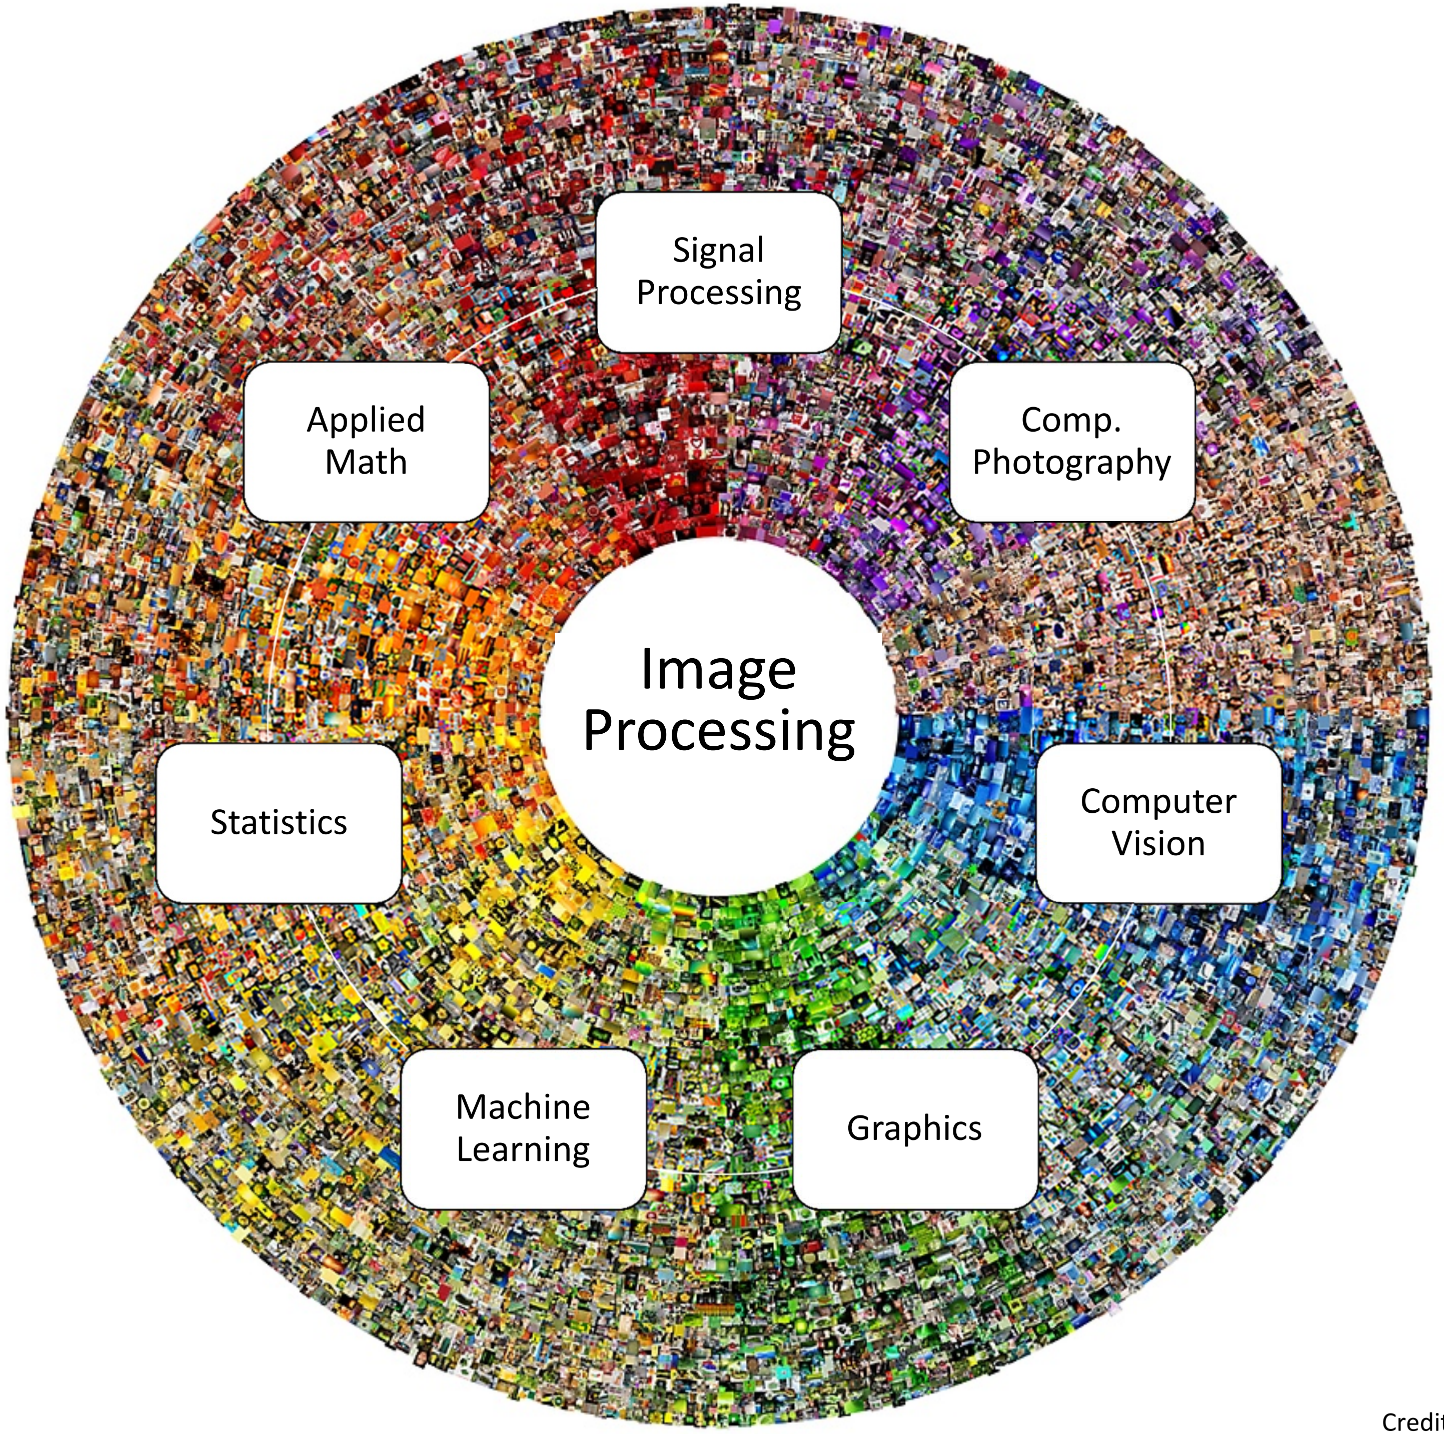

Credit: P. Milanfar

What does it mean, to see?

- *"The plain man's answer (and Aristotle's, too) would be, to know what is where by looking. In other words, vision is the process of discovering from images what is present in the world, and where it is."* David Marr, Vision, 1982

- Our brain is able to use an image as an input, and interpret it in terms of objects and scene structures.

What does Salvador Dali’s *Study for the Dream Sequence in Spellbound (1945)* say about our visual perception?

What does Paul Signac's *Place des Lices* (1893) say about our visual perception?

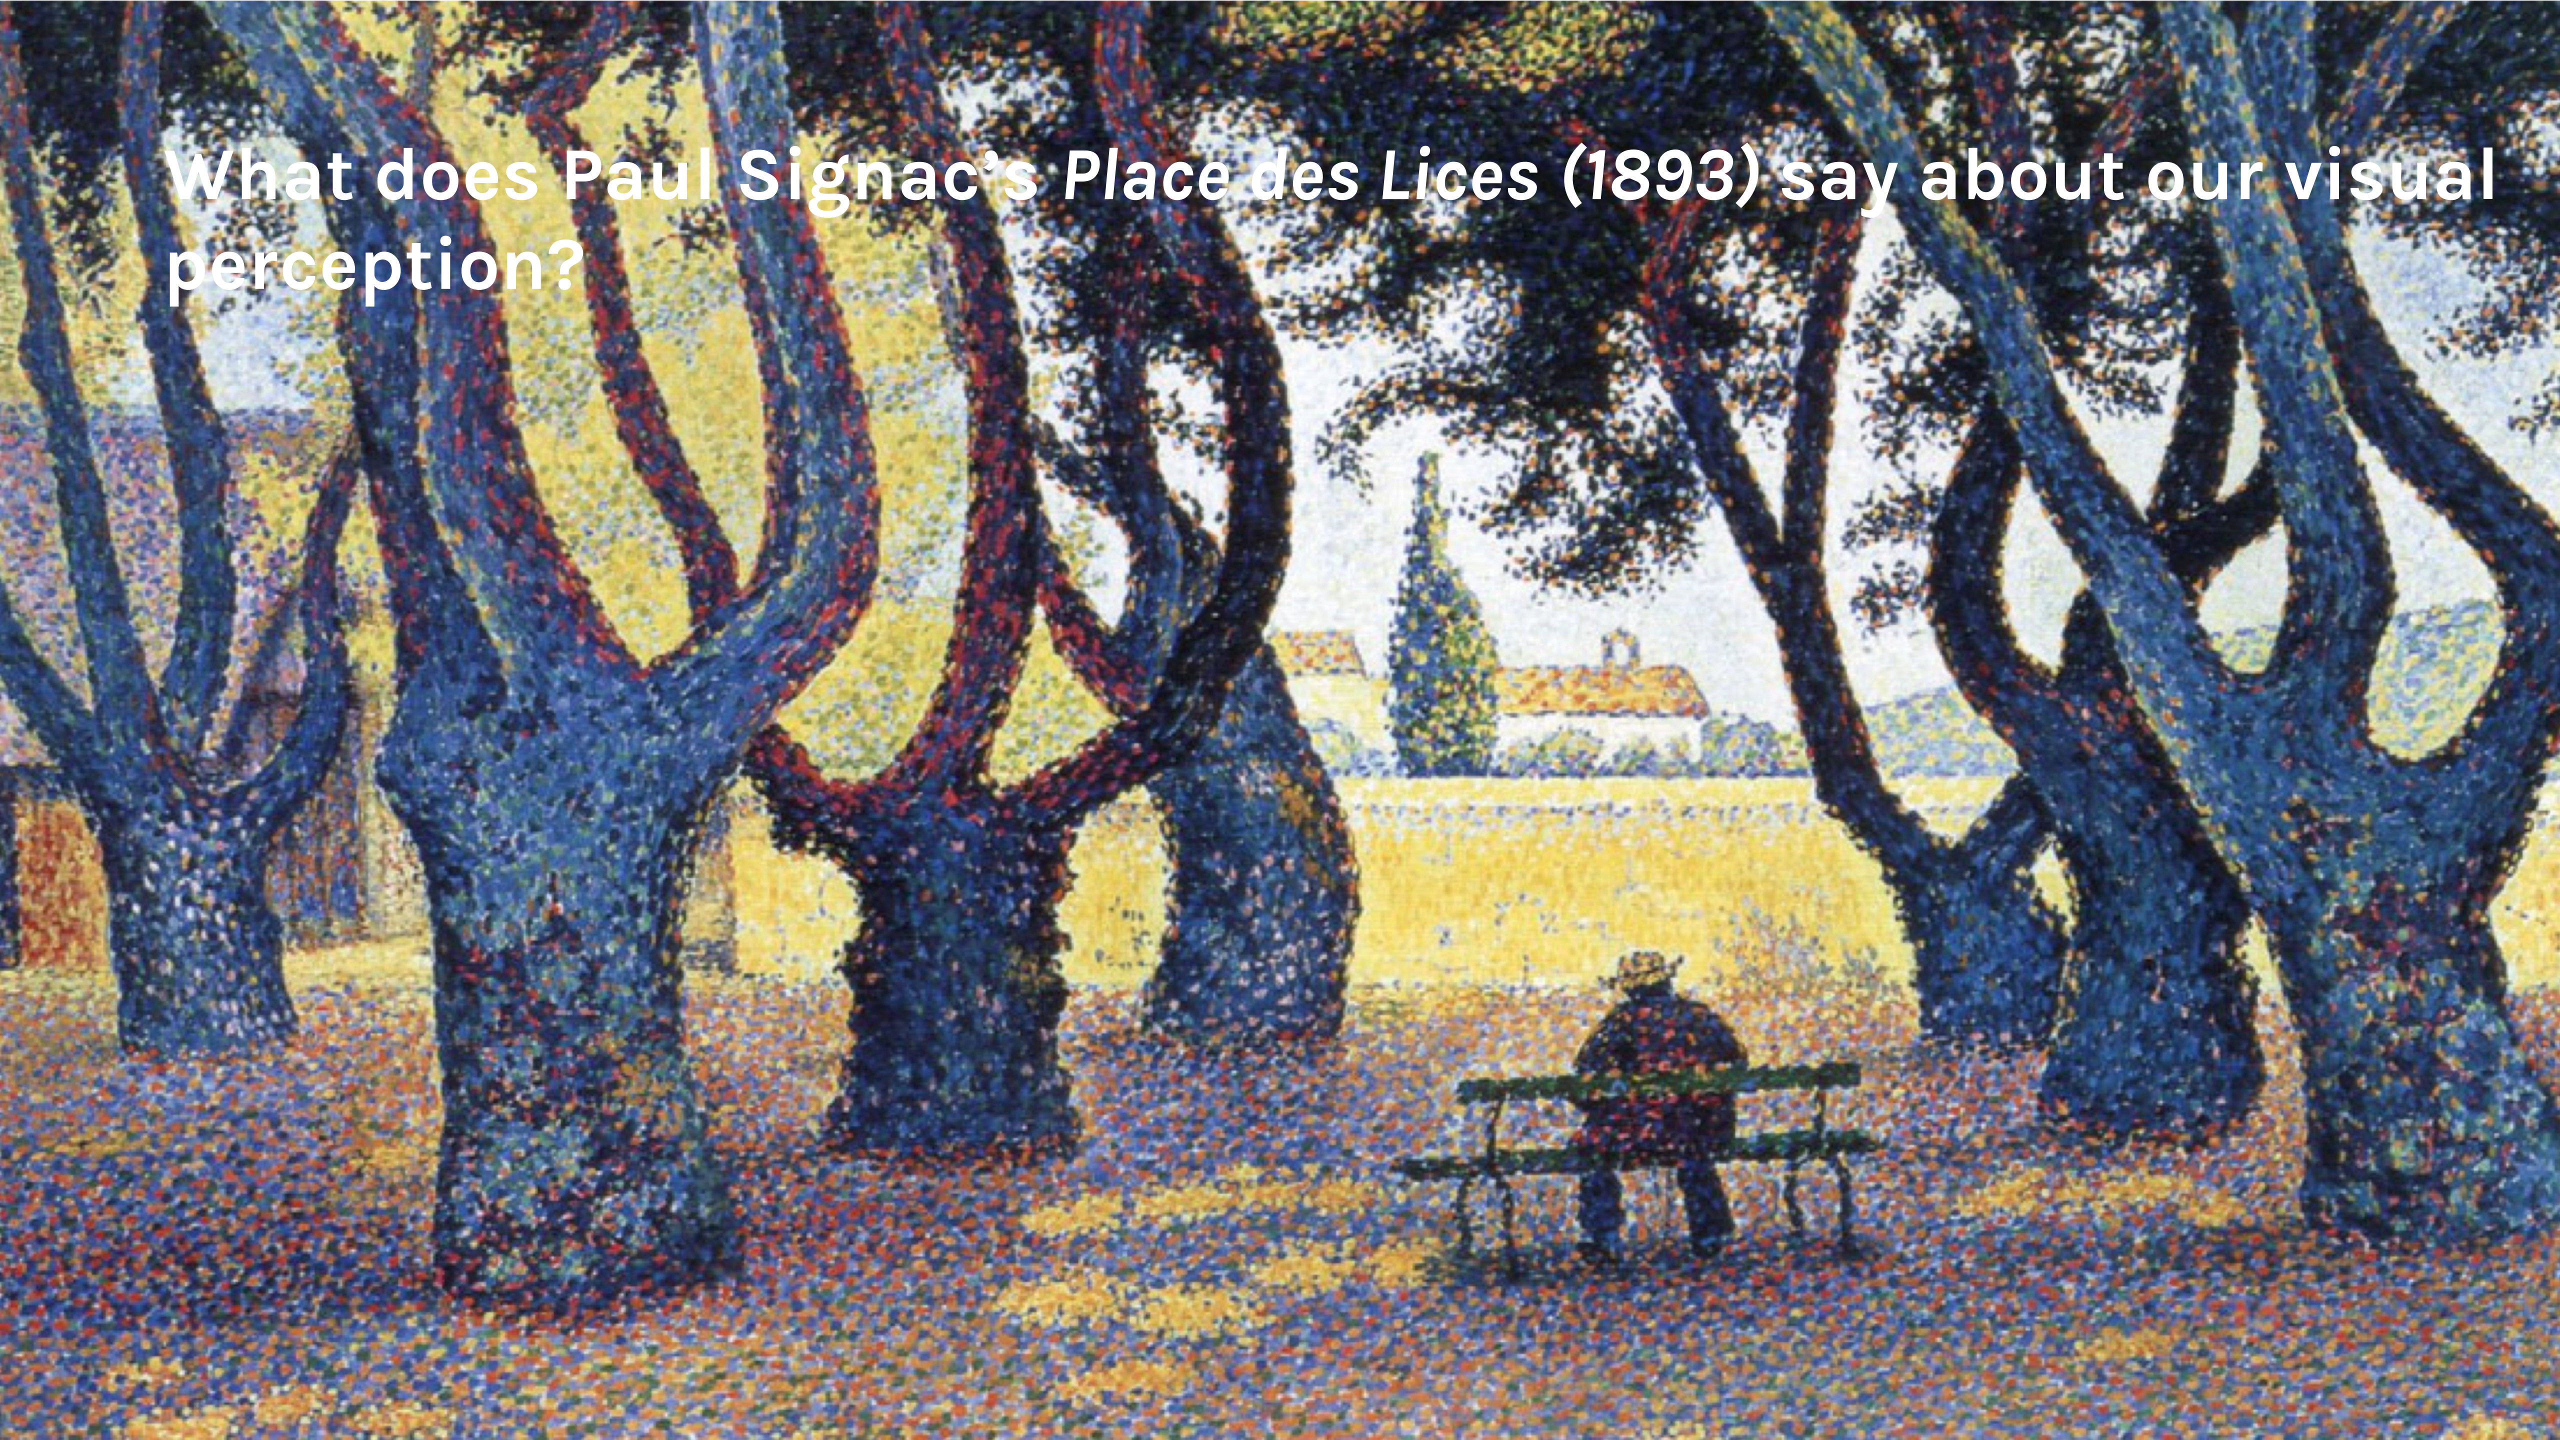

Why does vision appear easy to humans?

- Our brains are specialized to do vision.
- Nearly half of the cortex in a human brain is devoted to doing vision
  (cf. motor control ~20-30%, language ~10-20%)
- *“Vision has evolved to convert the ill-posed problems into solvable ones by adding premises: assumptions about how the world we evolved in is, on average, put together”*
  Steven Pinker, How the Mind Works, 1997
- Gestalt Theory (Laws of Visual Perception),
  Max Wertheimer, 1912

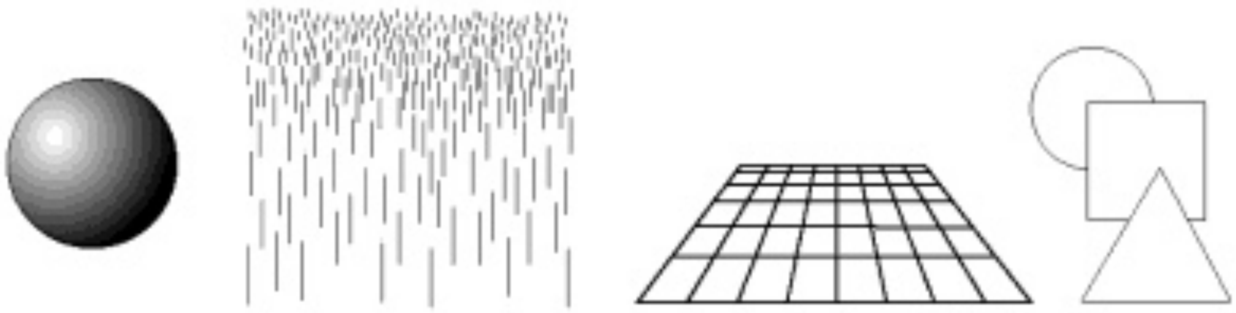

Figures: Steven Pinker, How the Mind Works, 1997

Computer Vision

- "Vision is a process that produces from images of the external world a description that is useful to the viewer and not cluttered with irrelevant information" ~David Marr

- <u>The goal of Computer Vision:</u>
  To develop artificial machine vision systems that make inferences related to the scene being viewed through the images acquired with digital cameras.

Marr’s observation: Studying vision at three different levels

- Vision as an information processing task
  [David Marr, 1982]

- Three levels of understanding:

1. Computational theory
   – What is computed? Why it is computed?
2. Representation and Algorithm
   – How it is computed?
   – Input, Output, Transformation
3. Physical Realization
   – Hardware

Visual Modules and the Information Flow

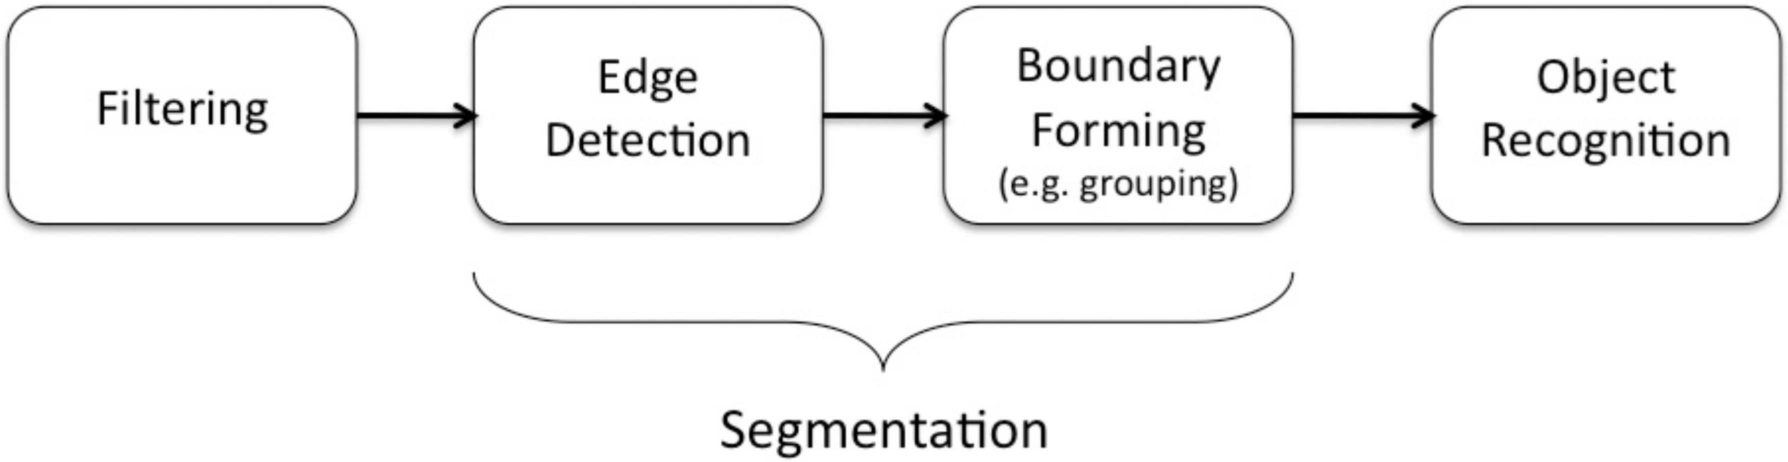

- Visual perception as a data-driven, bottom-up process (traditional view since D. Marr)
- Unidirectional information flow
- Simple low-level cues >> Complex abstract perceptual units

Visual Modules and the Information Flow

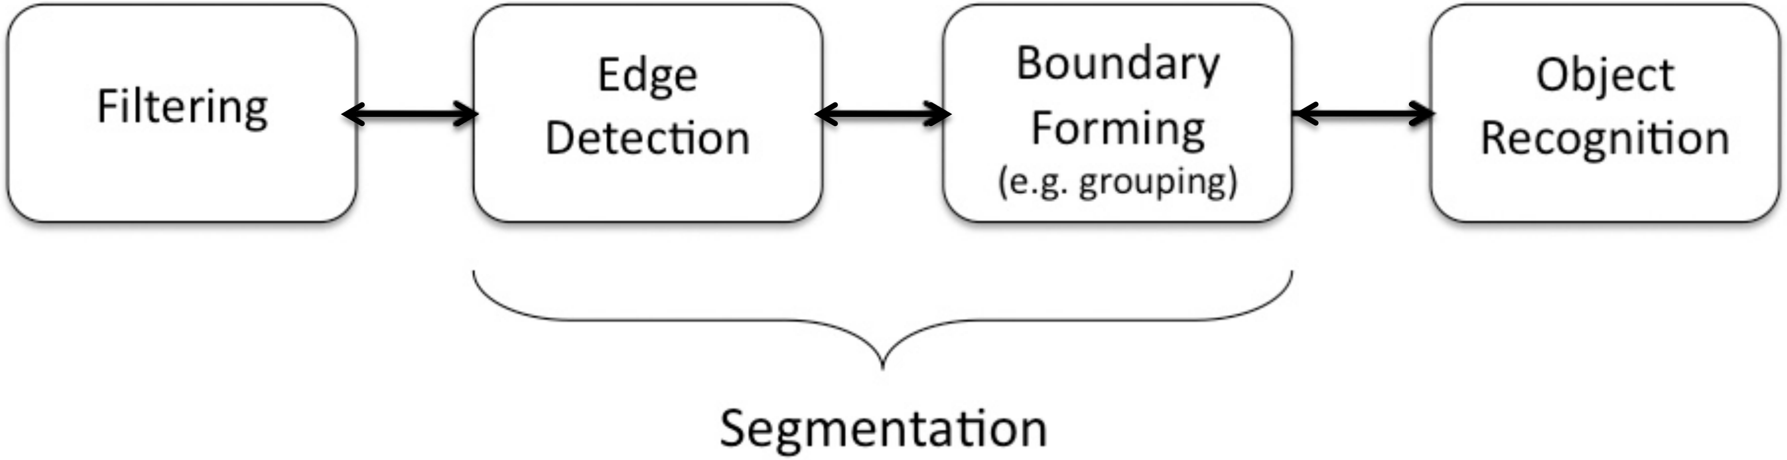

- Vision modules can be categorized into three groups according to their functionality:
  - Low-level vision: filtering out irrelevant image data
  - Mid-level vision: grouping pixels or boundary fragments together
  - High-level vision: complex cognitive processes

Fundamentals of Image Processing

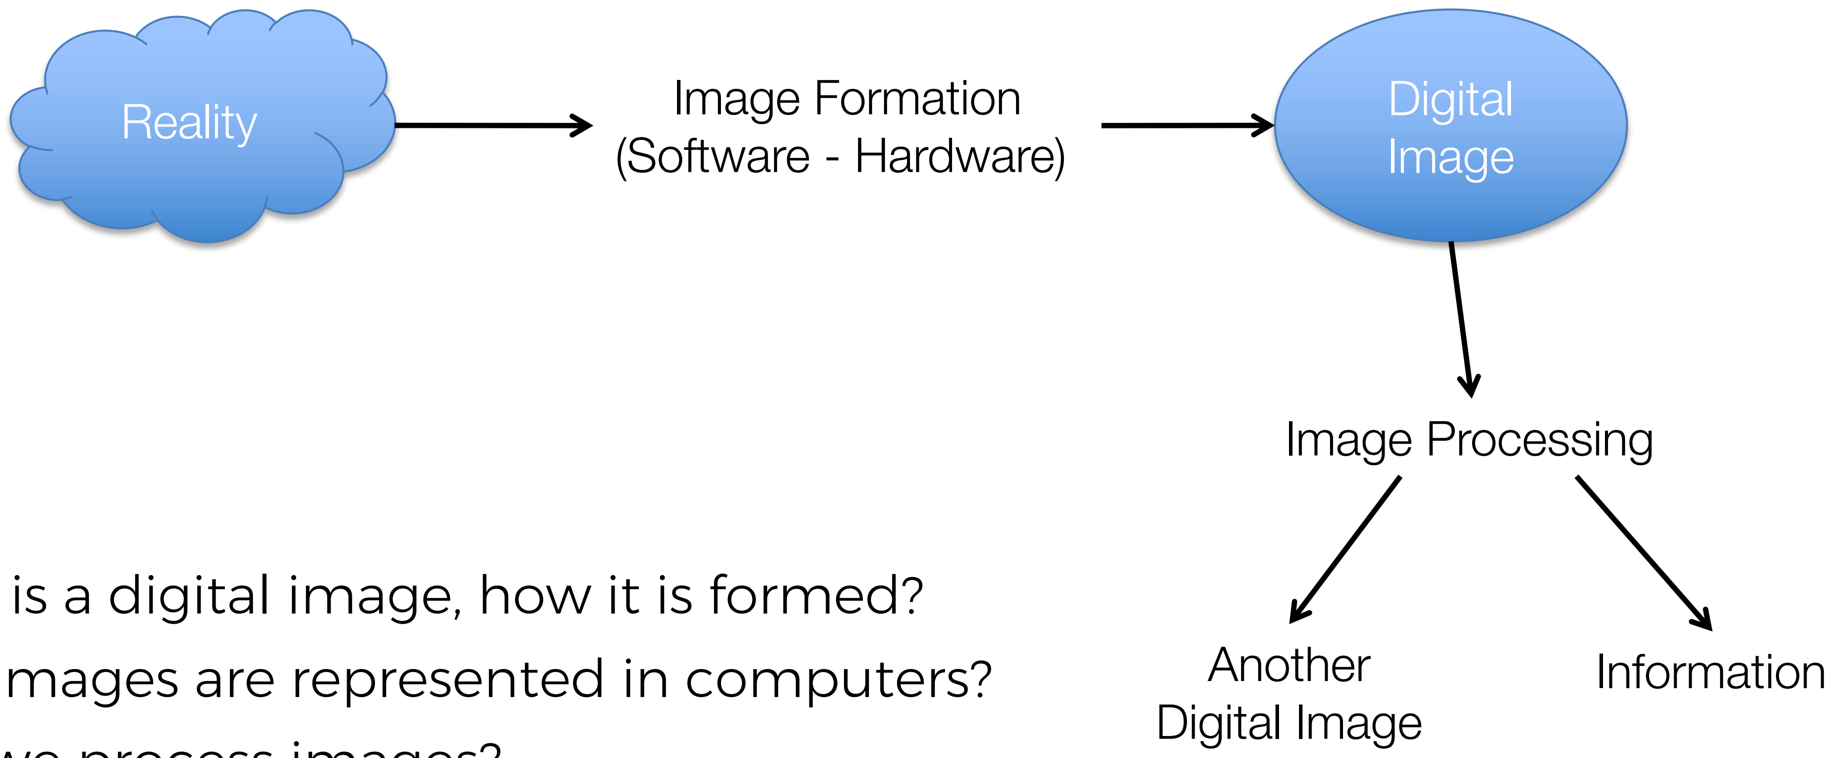

- What is a digital image, how it is formed?
- How images are represented in computers?
- Why we process images?
- How we process images?

Image Formation

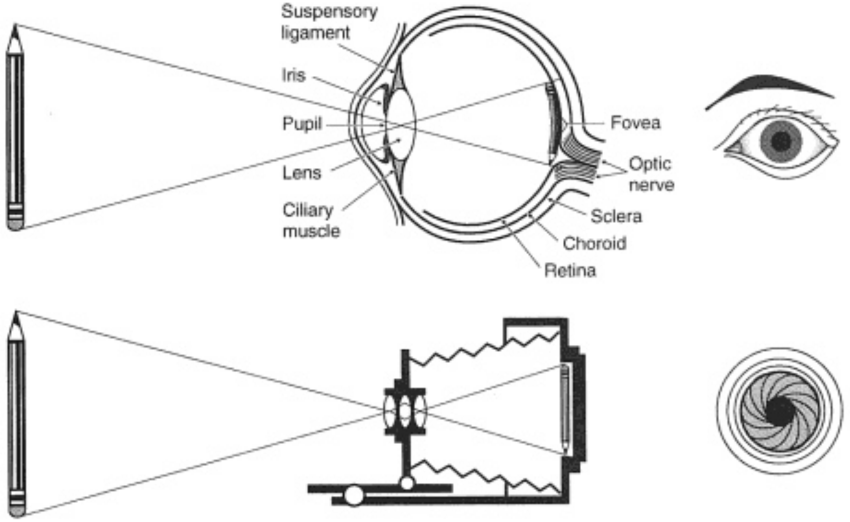

Three Dimensional World → Two Dimensional Image Space

- What is measured in an image location?
  - brightness
  - color

<< viewpoint
illumination conditions
local geometry
local material properties

Figures: Francis Crick, The Astonishing Hypothesis, 1995

Image Formation

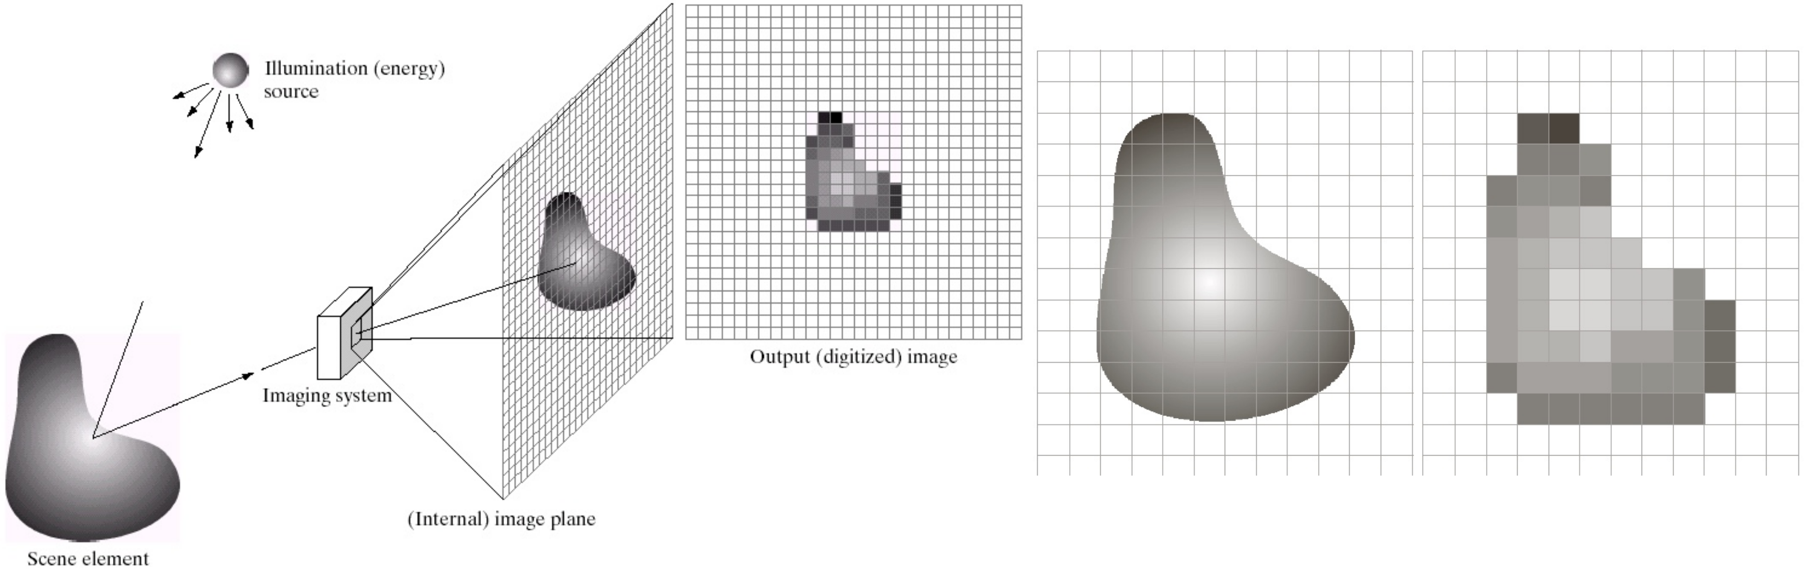

- Discretization
  - in image space - sampling
  - In image brightness - quantization

Figures: Gonzalez and Woods, Digital Image Processing, 3rd Edition, 2008

Image Representation

- **Digital image:** 2D discrete function f
- **Pixel**: Smallest element of an image f(x,y)

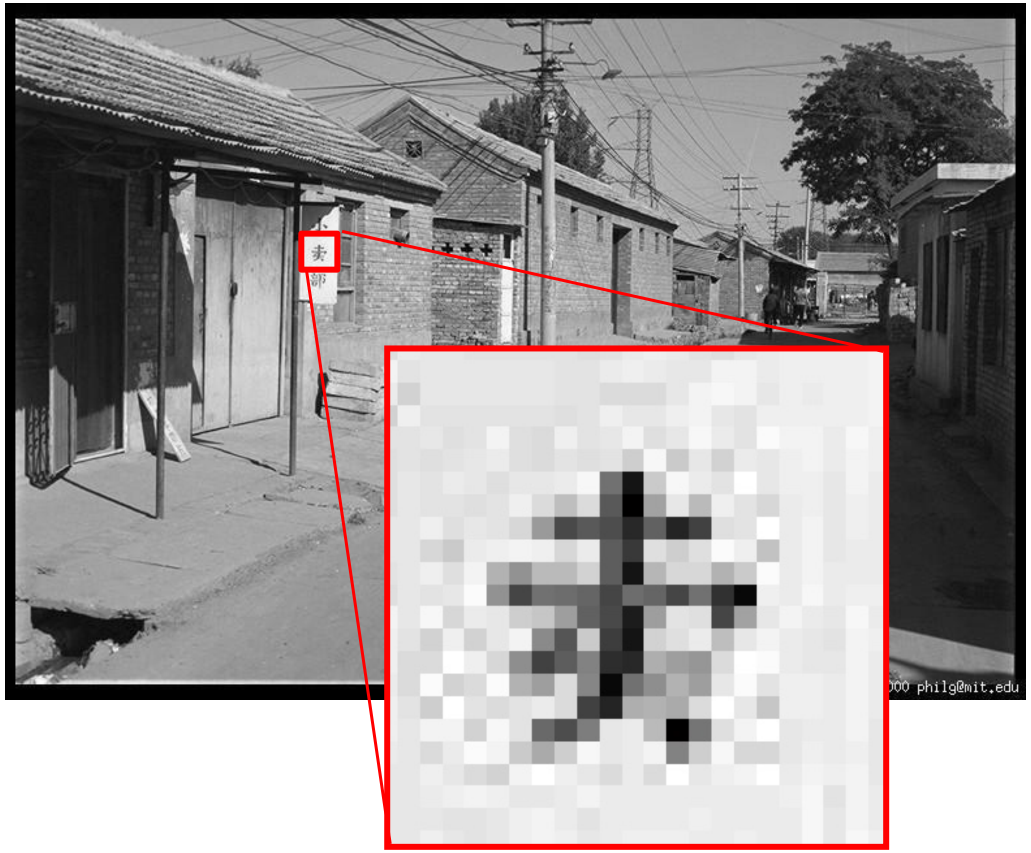

Figure: M. J. Black

Image Representation

- **Digital image**: 2D discrete function f
- **Pixel**: Smallest element of an image f(x,y)

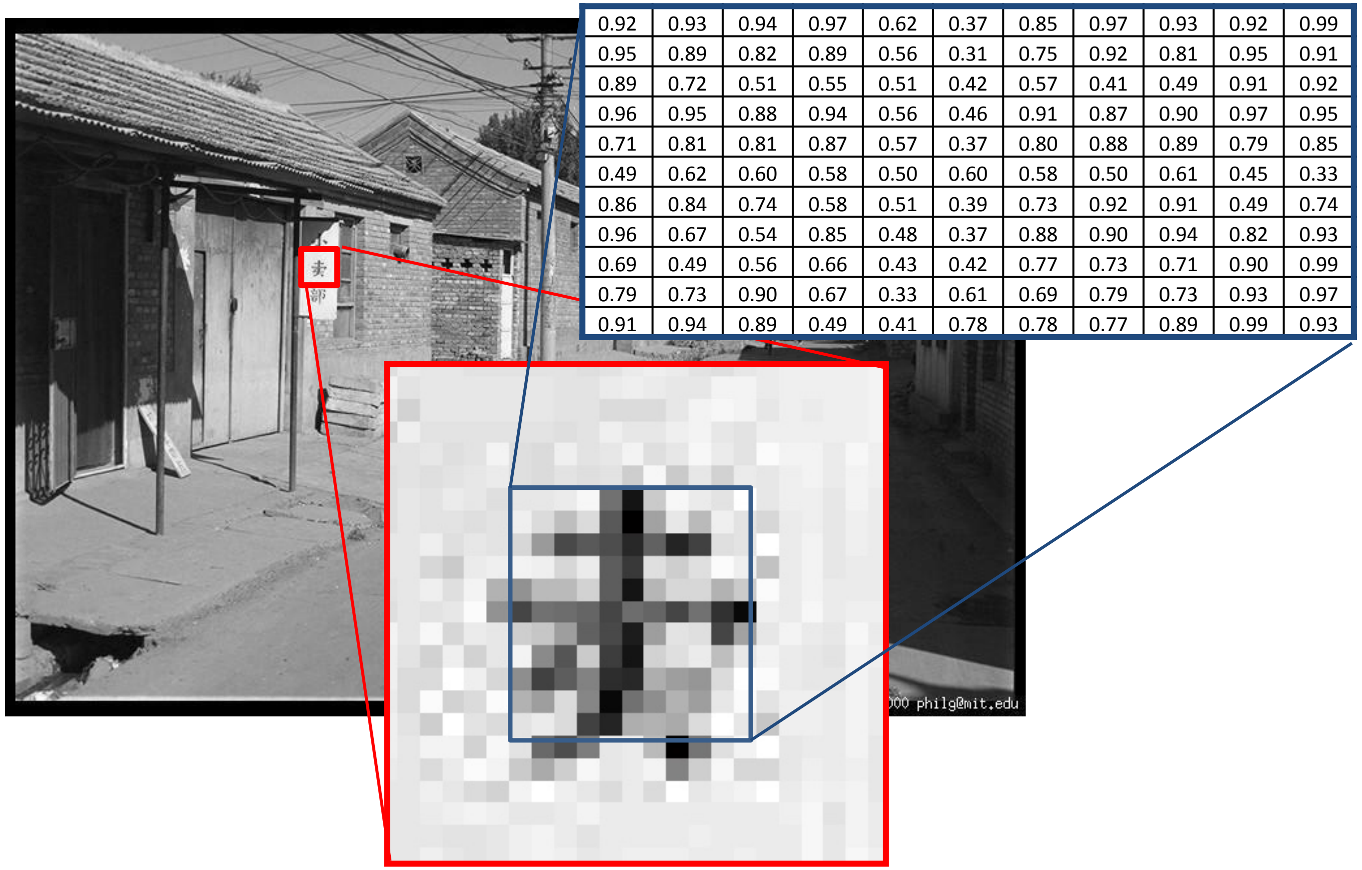

0.92	0.93	0.94	0.97	0.62	0.37	0.85	0.97	0.93	0.92	0.99
0.95	0.89	0.82	0.89	0.56	0.31	0.75	0.92	0.81	0.95	0.91
0.89	0.72	0.51	0.55	0.51	0.42	0.57	0.41	0.49	0.91	0.92
0.96	0.95	0.88	0.94	0.56	0.46	0.91	0.87	0.90	0.97	0.95
0.71	0.81	0.81	0.87	0.57	0.37	0.80	0.88	0.89	0.79	0.85
0.49	0.62	0.60	0.58	0.50	0.60	0.58	0.50	0.61	0.45	0.33
0.86	0.84	0.74	0.58	0.51	0.39	0.73	0.92	0.91	0.49	0.74
0.96	0.67	0.54	0.85	0.48	0.37	0.88	0.90	0.94	0.82	0.93
0.69	0.49	0.56	0.66	0.43	0.42	0.77	0.73	0.71	0.90	0.99
0.79	0.73	0.90	0.67	0.33	0.61	0.69	0.79	0.73	0.93	0.97
0.91	0.94	0.89	0.49	0.41	0.78	0.78	0.77	0.89	0.99	0.93

Figure: M. J. Black

Sample Problems and Techniques

- Edge Detection
- Image Denoising
- Image Smoothing
- Image Deblurring
- Image Segmentation
- Intrinsic Images
- Visual Saliency
- Semantic Segmentation

- PDEs and Variational models
- MRFs
- Graph Theory
- Sparse Coding
- Deep Learning

Image Filtering

- Filtering out the irrelevant information

$$\underbrace{f(x)}_{\text{observed image}} = \underbrace{u(x)}_{\text{desired image}} + \underbrace{n(x)}_{\text{irrelevant data}}$$

- Image denoising, image sharpening, image smoothing, image deblurring, etc.
- Edge detection

Edge Detection

Canny edge detector

- Edges: abrupt changes in the intensity
  - Uniformity of intensity or color
- Edges to object boundaries

Image Filtering

- <u>Difficulty:</u> Some of the irrelevant image information have characteristics similar to those of important image features

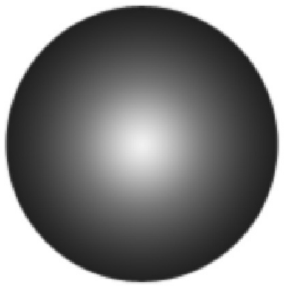

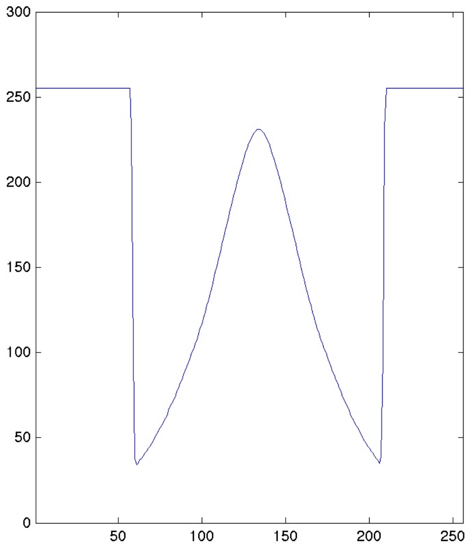

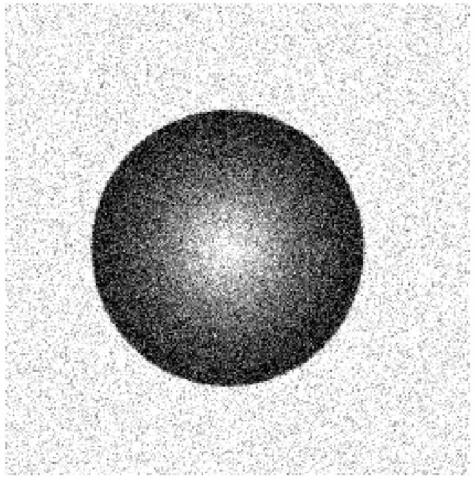

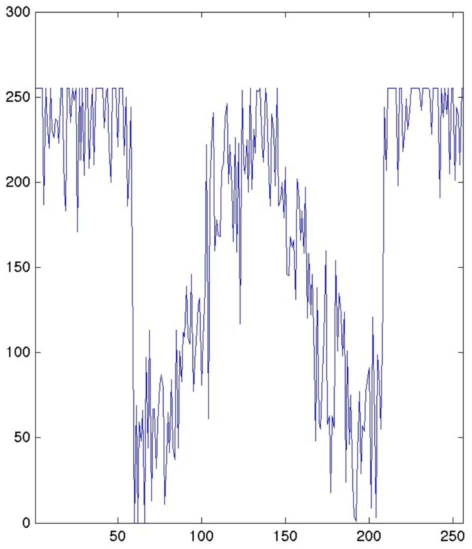

Image Smoothing - A Little Bit of History

- Gaussian Filtering / linear diffusic $\frac{\partial u}{\partial t} = \nabla \cdot (\nabla u) = \nabla^2 u$
  - the most widely used method

- mid 80's – unified formulations
  - methods that combine smoothing and edge detection
  - Geman & Geman'84, Blake & Zisserman'87, Mumford & Shah'89, Perona & Malik'90

Image Denoising

- Images are corrupted with 70% salt-and-pepper noise

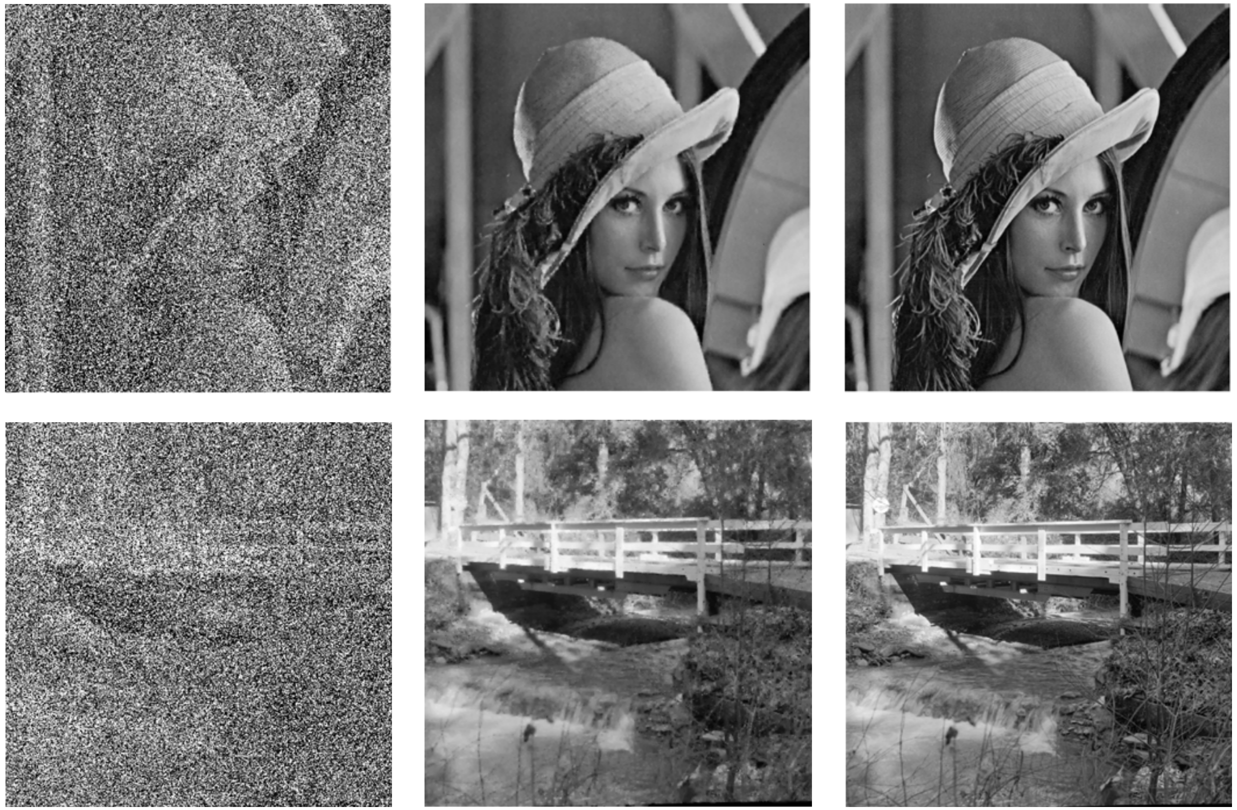

Noisy input | Recovered image | Original image

What do these examples demonstrate?

R. H. Chan, C.-W. Ho, and M. Nikolova, Salt-and-Pepper Noise Removal by Median-Type Noise Detectors and Detail-Preserving Regularization. IEEE TIP 2005

Non-local Means Denoising

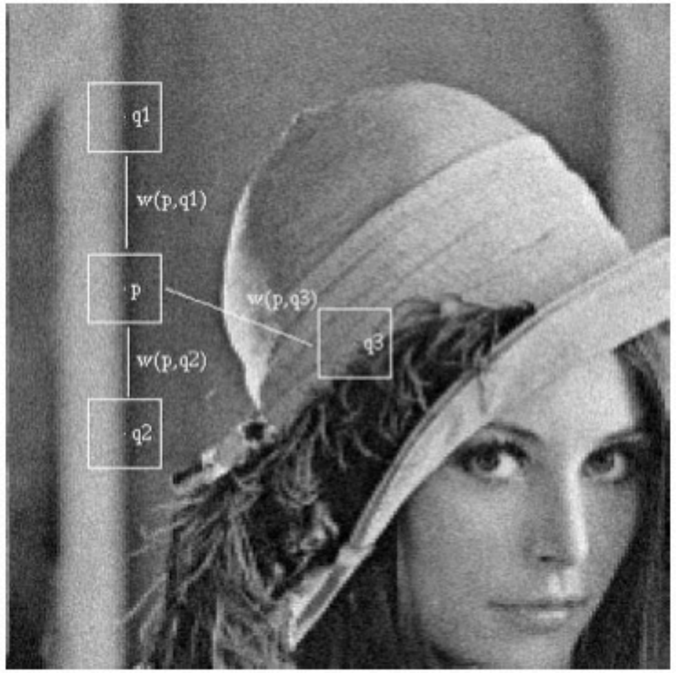

Figure 1. Scheme of NL-means strategy. Similar pixel neighborhoods give a large weight, w(p,q1) and w(p,q2), while much different neighborhoods give a small weight w(p,q3).

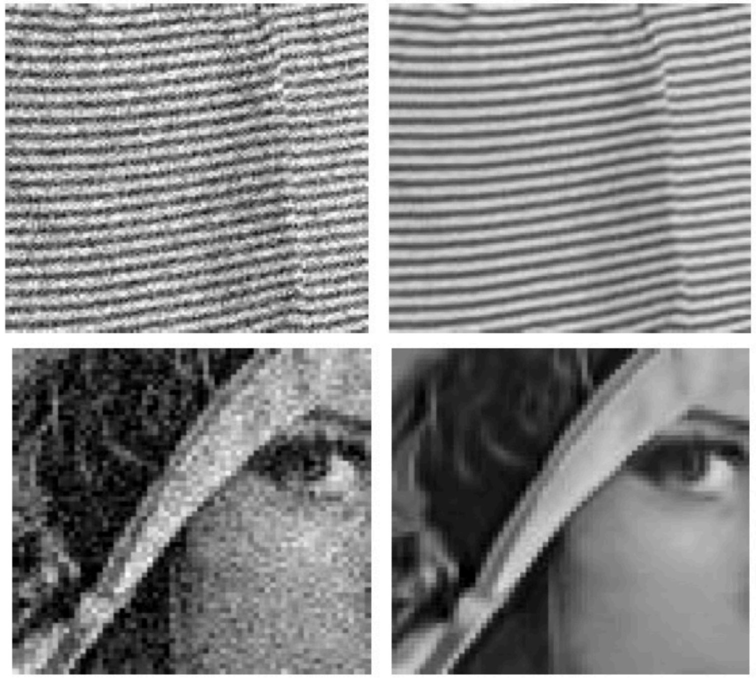

Preserve fine image details and texture during denoising

A. Buades, B. Coll, J. M. Morel, A non-local algorithm for image denoising, CVPR, 2005

Context-Guided Filtering

- Use local image context to steer filtering

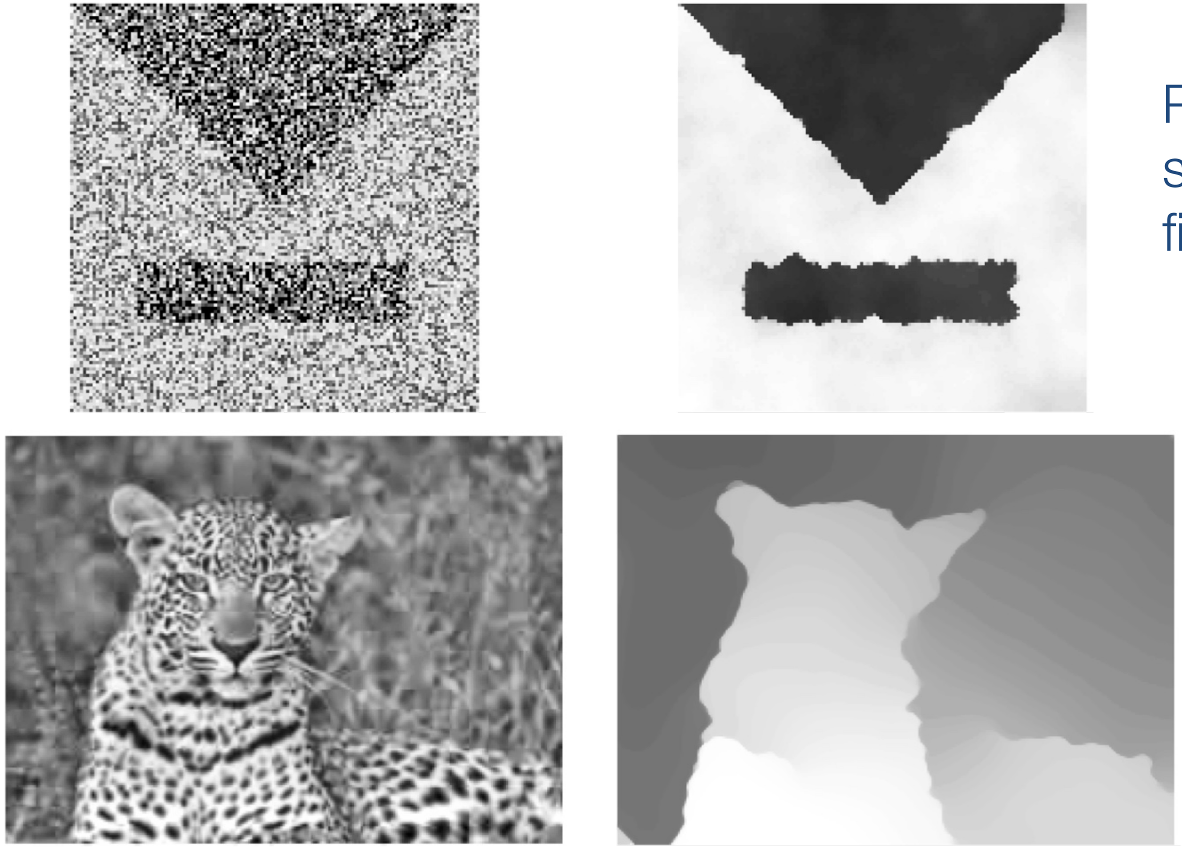

Preserve main image structures during filtering

E. Erdem and S. Tari, Mumford-Shah Regularizer with Contextual Feedback, JMIV, 2009

Image Smoothing

L. Xu, C. Lu, Y. Xu, J. Jia, Image Smoothing via L0 Gradient Minimization, ACM Trans. Graphics 2011
(SIGGRAPH Asia 2011)

Image Smoothing

L. Karacan, E. Erdem, A. Erdem, Structure Preserving Image Smoothing via Region Covariances, ACM Trans. Graphics 2013 (SIGGRAPH Asia 2013)

Image Deblurring

- Remove blur and restore a sharp image

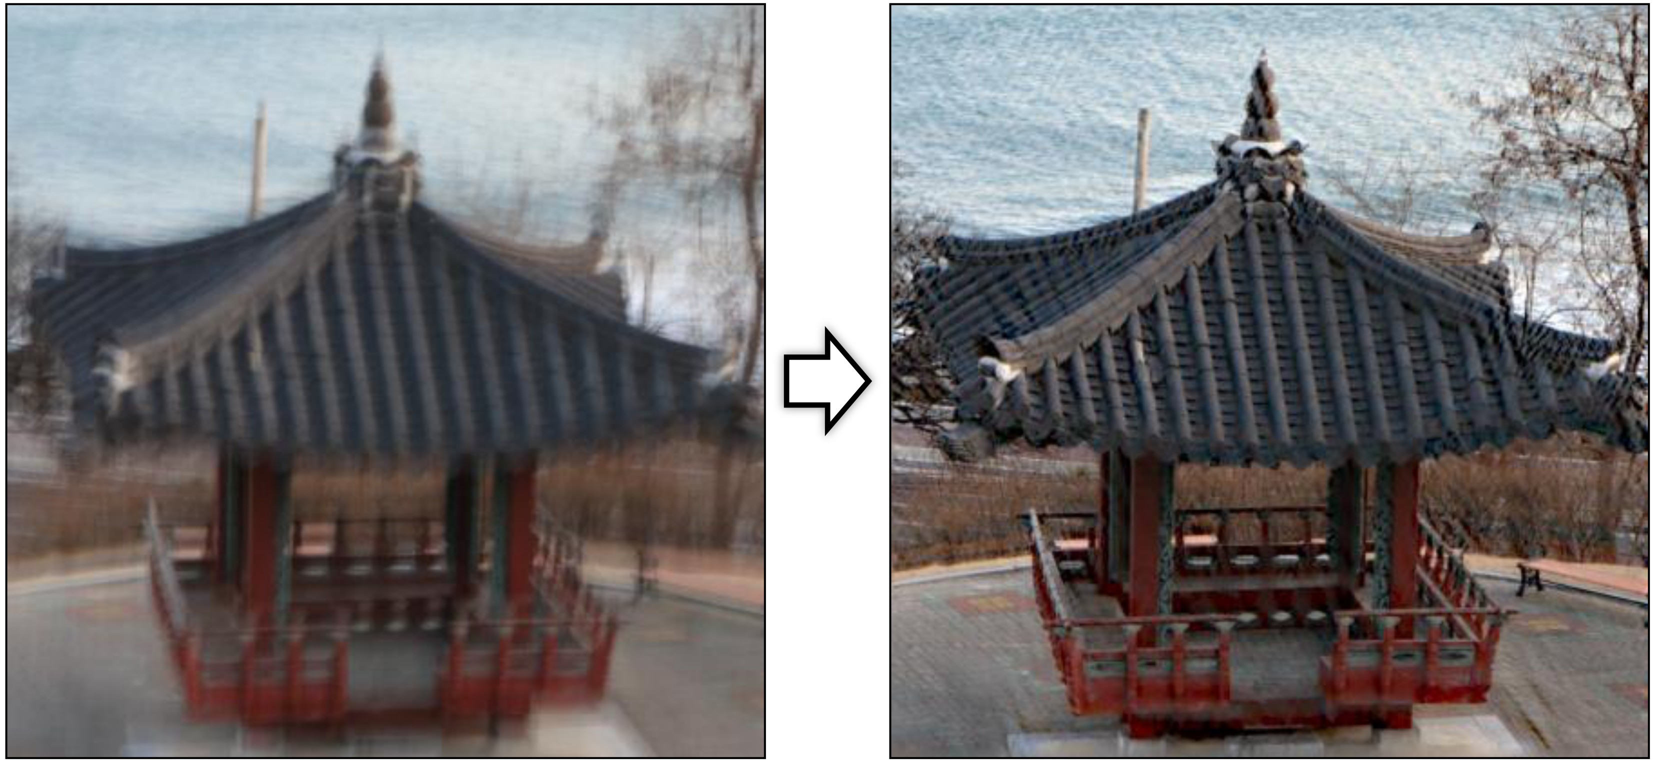

from a given blurred image

find its latent sharp image

Slide credit: Lee and Cho

Image Deblurring

- Remove blur and restore a sharp image

Input blurred image

Levin et al. CVPR 2010

Slide credit: Lee and Cho

Intrinsic Image Decomposition

- Decompose an image into reflectance and shading layers.

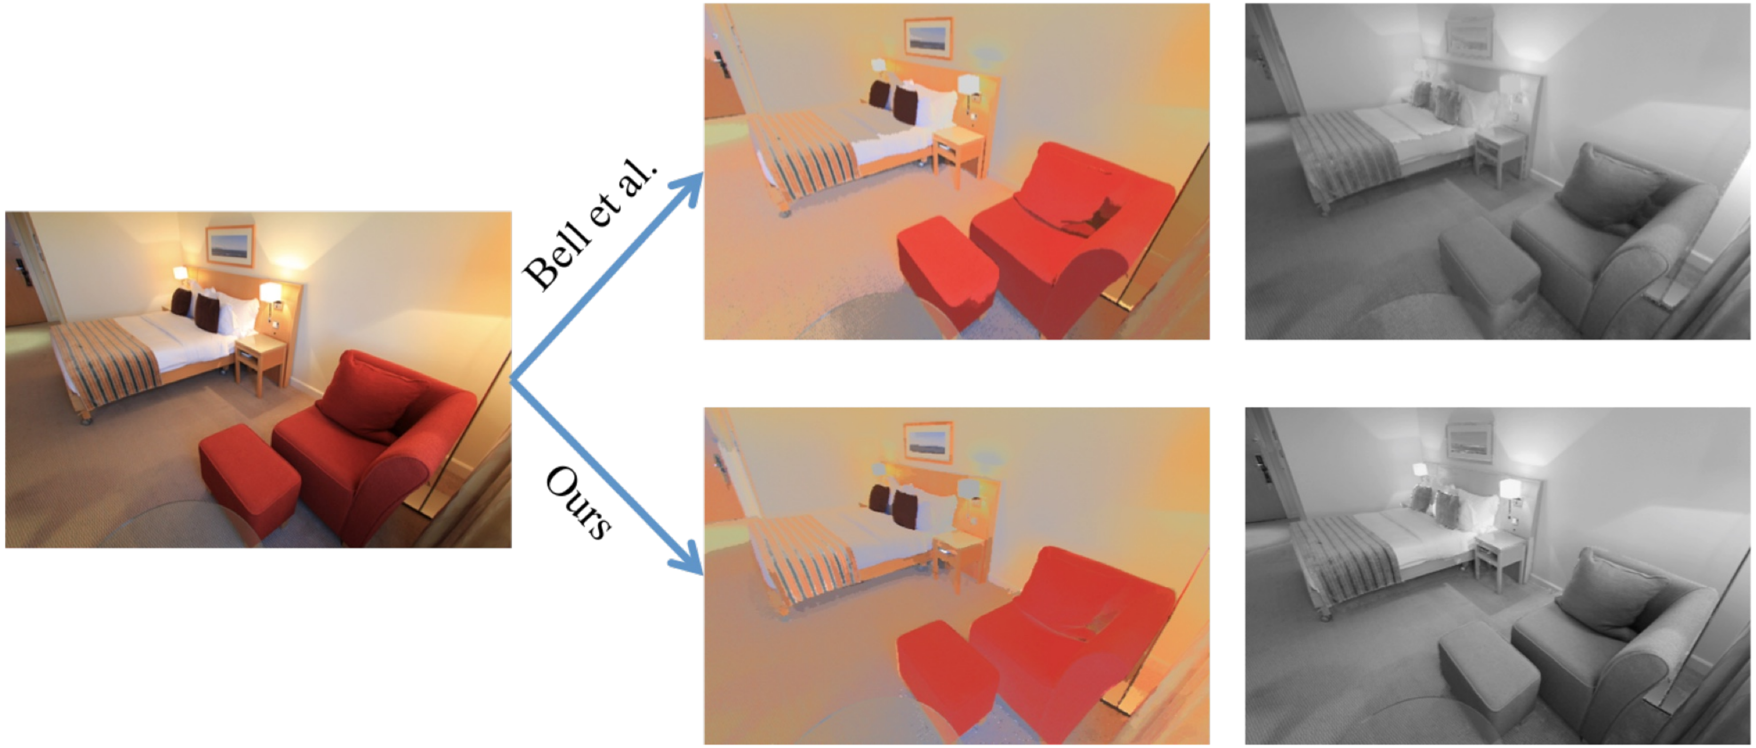

T. Zhou, P. Krähenbühl, A. A. Efros,
Learning Data-driven Reflectance Priors for Intrinsic Image Decomposition, ICCV 2015

Visual Saliency

- The problem of predicting where people look at images

The squares shows where
the observers looked
in eye tracking experiments

E. Erdem and A. Erdem, Visual saliency estimation by nonlinearly integrating features using region covariances, Journal of Vision 2013

Image Retargetting

- automatically resize an image to arbitrary aspect ratios while preserving <u>important image features</u>

How we define the importance?

S. Avidan and A. Shamir, Seam Carving for Content-Aware Image Resizing, SIGGRAPH, 2007

Sparse Coding

- The problem of finding a small number of representative atoms from a dictionary which when combined with right weights represent a given signal.

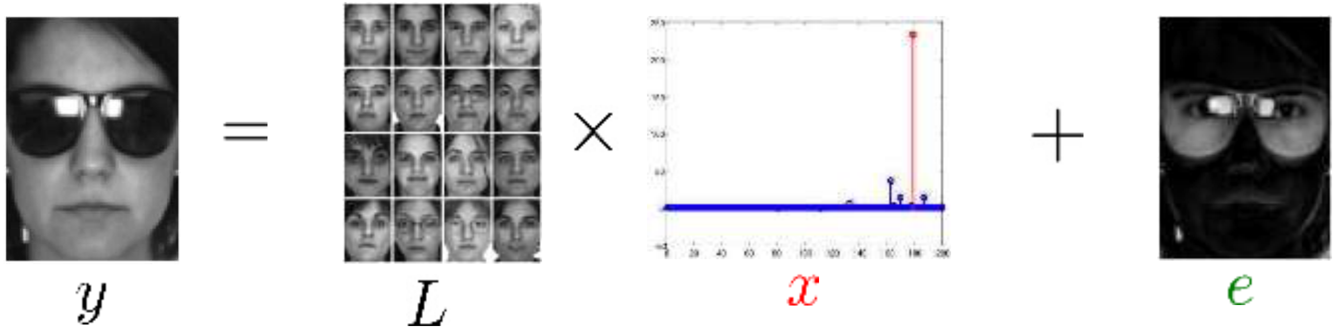

$$y = Lx + e$$

Credit: Yi Ma

Image Inpainting

- Reconstructing lost or deteriorated parts of images

What do
these examples
demonstrate?

M. Bertalmio, G. Sapiro, V. Caselles and C. Ballester, Image Inpainting, SIGGRAPH, 2000

Image Segmentation

- Partition an image into meaningful regions that are likely to correspond to objects exist in the image

Figures: A. Erdem

Grouping of pixels

according to what criteria?

high-level object specific knowledge matters!

Snakes

- Curve Evolution - parametric curve formulation

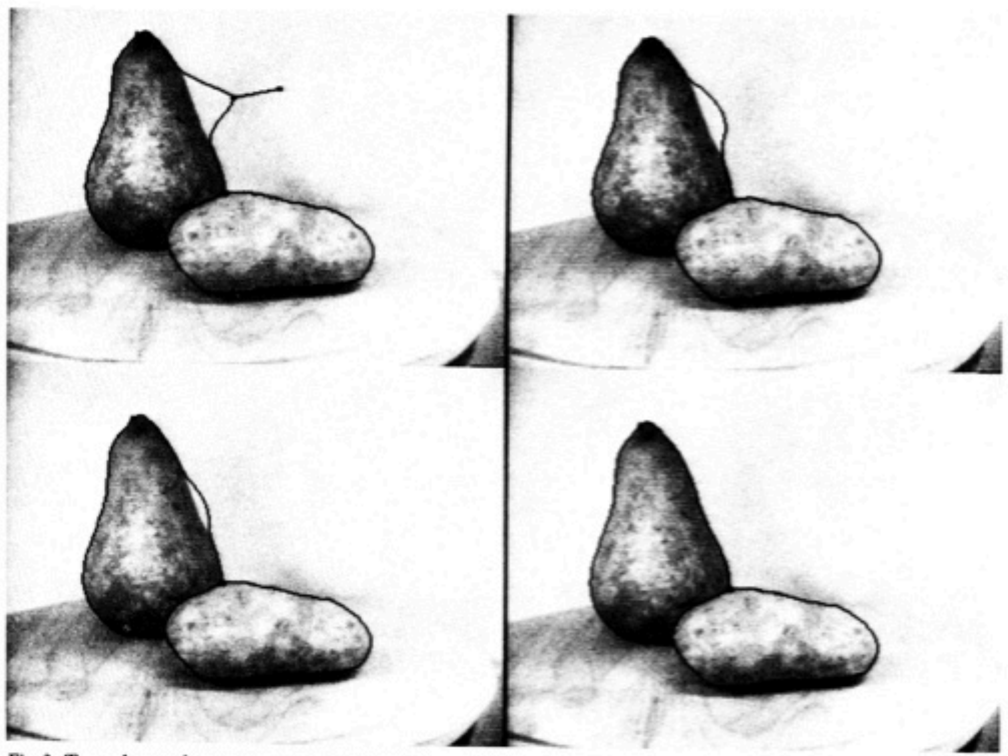

Fig. 3. Two edge snakes on a pear and potato. Upper-left: The user has pulled one of the snakes away from the edge of the pear. Others: After the user lets go, the snake snaps back to the edge of the pear.

M. Kass, A. Witkin, and D. Terzopoulos, Snakes: Active Contour Models, IJCV, 1988

Active Contours Without Edges

- Curve Evolution – a level-set based curve formulation

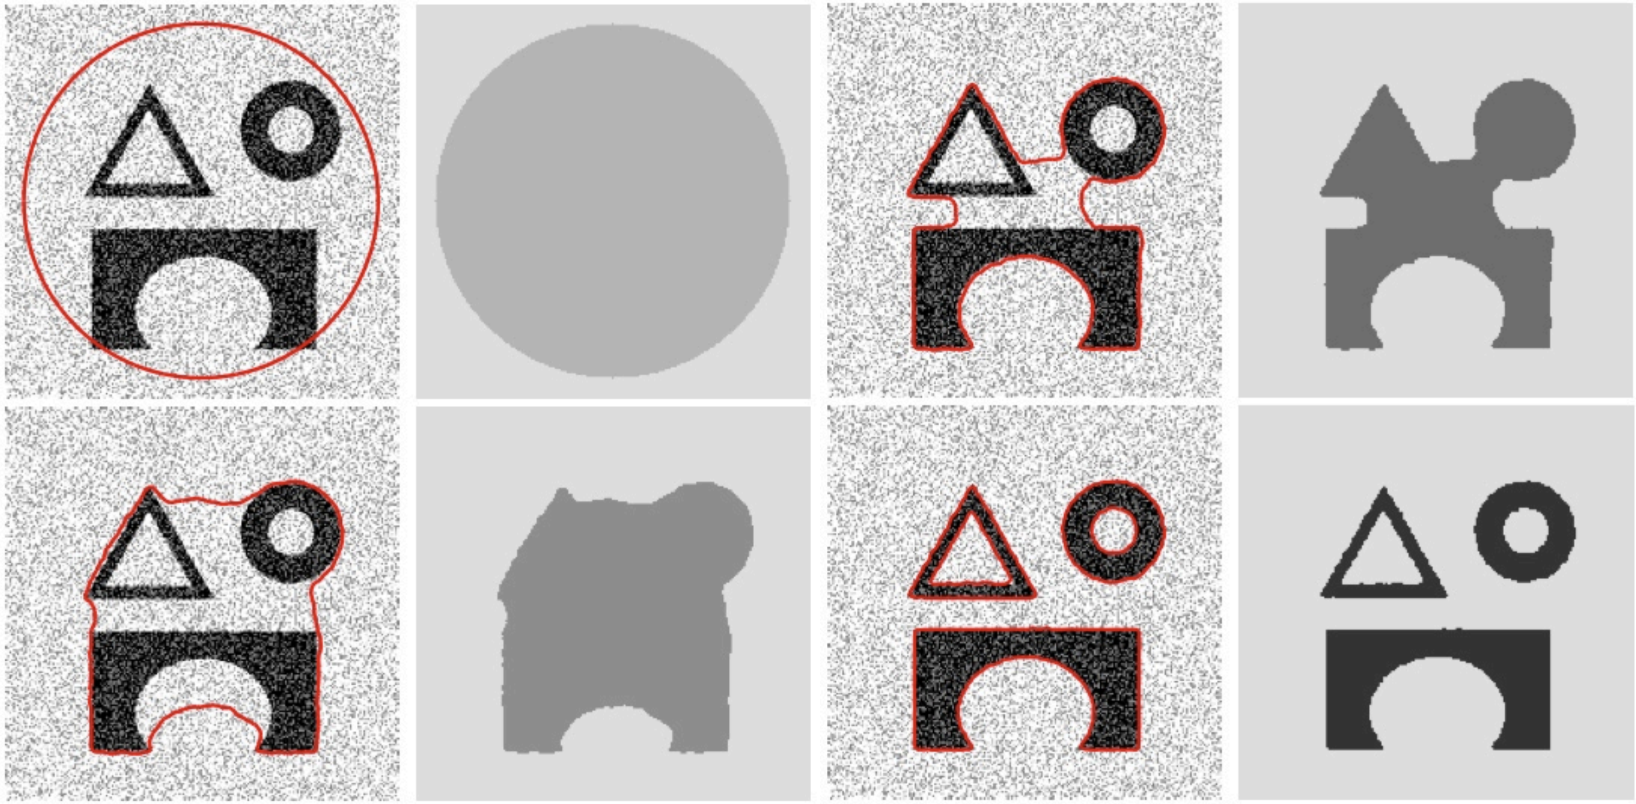

T. Chan and L. Vese. Active Contours Without Edges, IEEE Trans. Image Processing, 2001

Normalized Cuts

- A graph-theoretic formulation for segmentation

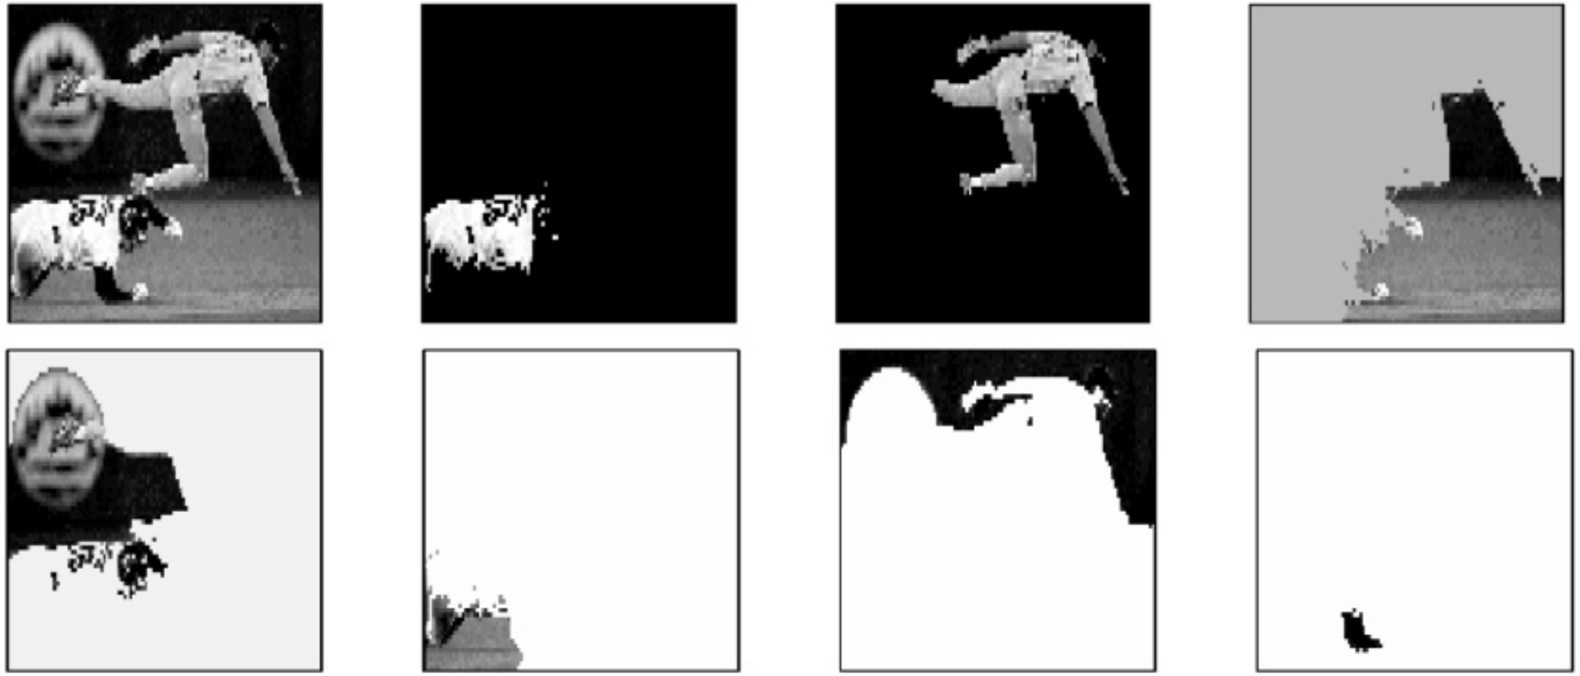

J. Shi and J. Malik, Normalized Cuts and Image Segmentation, IEEE Trans. Pattern Anal. Mach. Intel.

Normalized Cuts

Slide credit: S. Lazebnik

From contours to regions

- gPb-owt-ucm segmentation algorithm

P. Arbelaez, M. Maire, C. Fowlkes and J. Malik, Contour Detection and Hierarchical Image Segmentation, IEEE Trans Pattern Anal. Mach. Intell. 33(5):898-916, 2011

Graphical Models in Vision

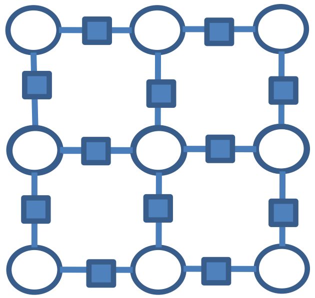

4-connected;
pairwise MRF

$$E(x) = \sum_{i,j \in N_4} \theta_{ij}(x_i,x_j)$$

$\boxed{\text{Order 2}}$

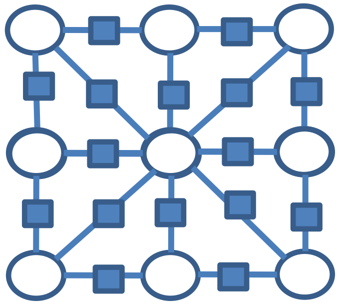

higher(8)-connected;
pairwise MRF

$$E(x) = \sum_{i,j \in N_8} \theta_{ij}(x_i,x_j)$$

$\boxed{\text{Order 2}}$

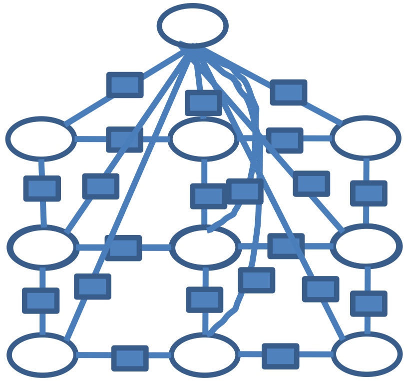

MRF with
global variables

$$E(x) = \sum_{i,j \in N_8} \theta_{ij}(x_i,x_j)$$

$\boxed{\text{Order 2}}$

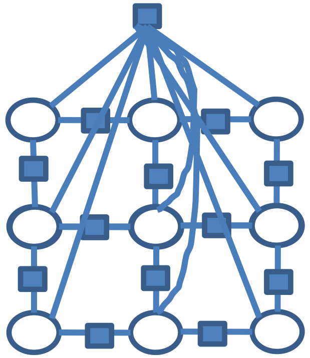

Higher-order MRF

$$E(x) = \sum_{i,j \in N_4} \theta_{ij}(x_i,x_j) + \theta(x_1,\ldots,x_n)$$

$\boxed{\text{Order n}}$

Semantic Segmentation

- The problem of joint recognition and segmentation

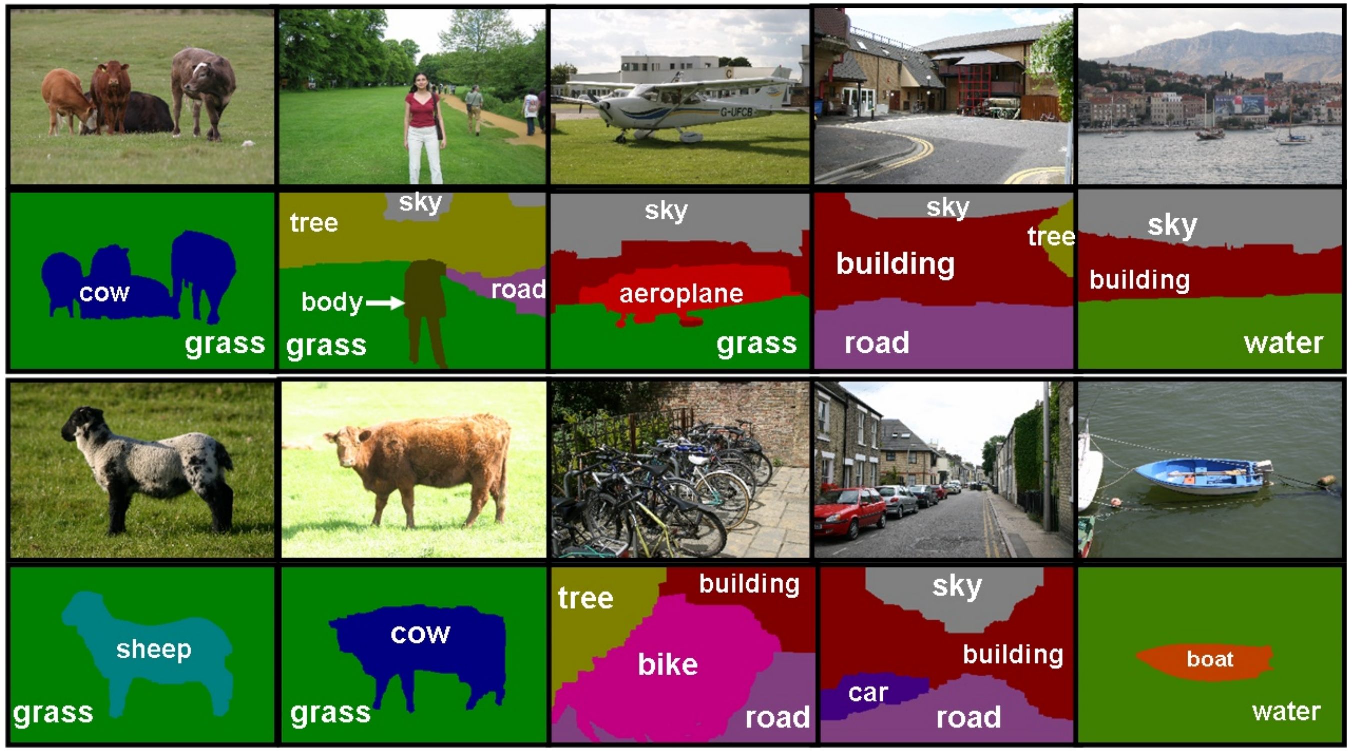

Object classes	Building	Grass	Tree	Cow	Sheep	Sky	Aeroplane	Water	Face	Car
Bike	Flower	Sign	Bird	Book	Chair	Road	Cat	Dog	Body	Boat

[TextonBoost; Shotton et al, '06]

C. Rother

Semantic Segmentation

- The problem of joint recognition and segmentation

Carreira et al., Semantic Segmentation with Second-Order Pooling , ECCV, 2012

Top-down Saliency

- Task-oriented models (e.g. searching for a target object from a specific category)

A. Kocak et al., Top down saliency estimation via superpixel-based discriminative dictionaries, BMVC 2014

Top-down Saliency

- Task-oriented models (e.g. searching for a target object from a specific category)

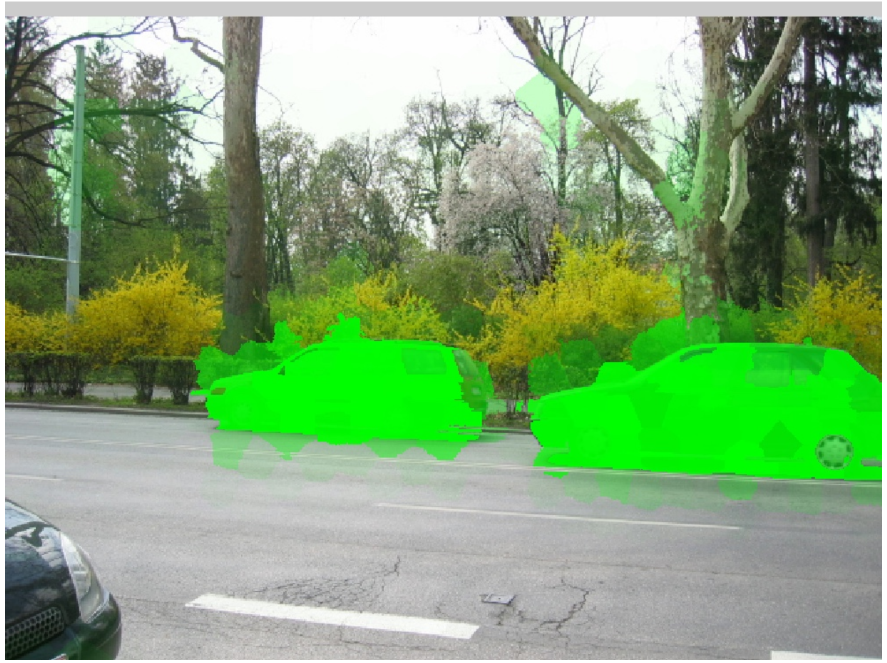

A. Kocak et al., Top down saliency estimation via superpixel-based discriminative dictionaries, BMVC 2014

Top-down Saliency

- Task-oriented models (e.g. searching for a target object from a specific category)

A. Kocak et al., Top down saliency estimation via superpixel-based discriminative dictionaries, BMVC 2014

Deep Learning

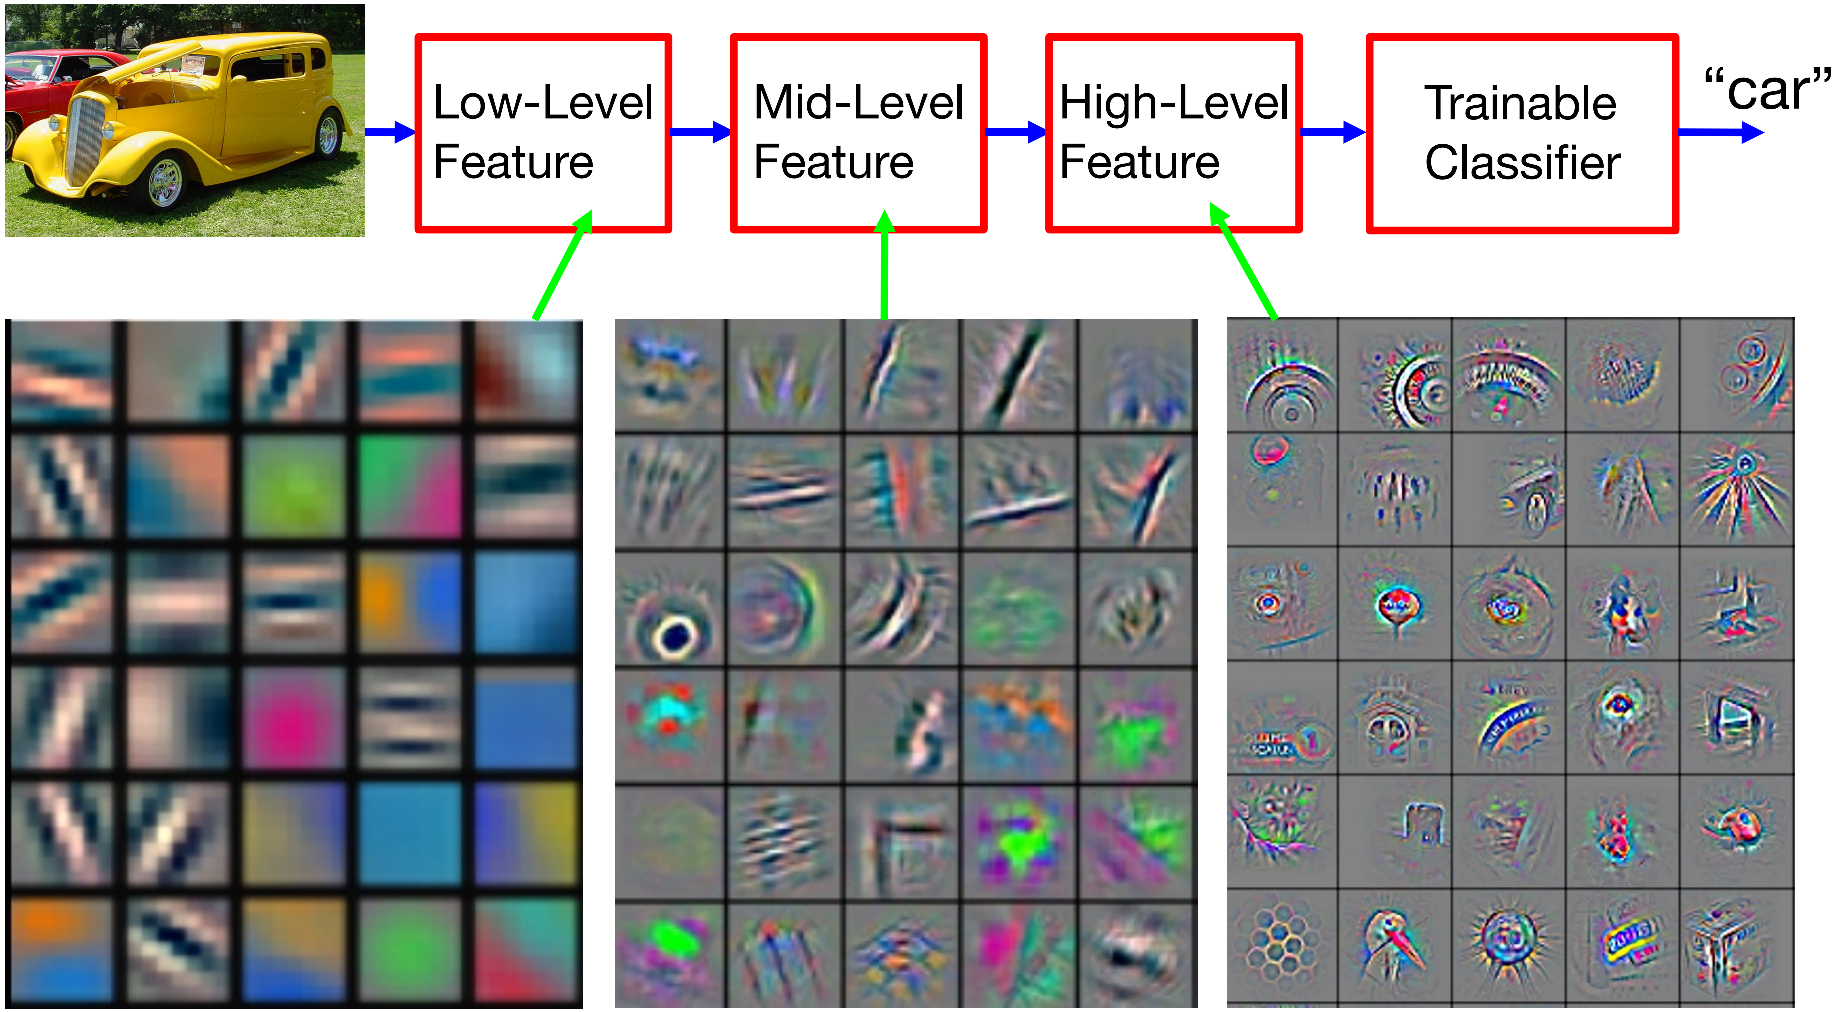

Feature visualization of convolutional net trained on ImageNet from [Zeiler & Fergus 2013]

Slide Credit: Marc'Aurelio Ranzato, Yann LeCun

Deep Learning

- [Krizhevsky et al. NIPS12]
  - 54 million parameters; 8 layers (5 conv, 3 fully-connected)
  - Trained on 1.4M images in ImageNet

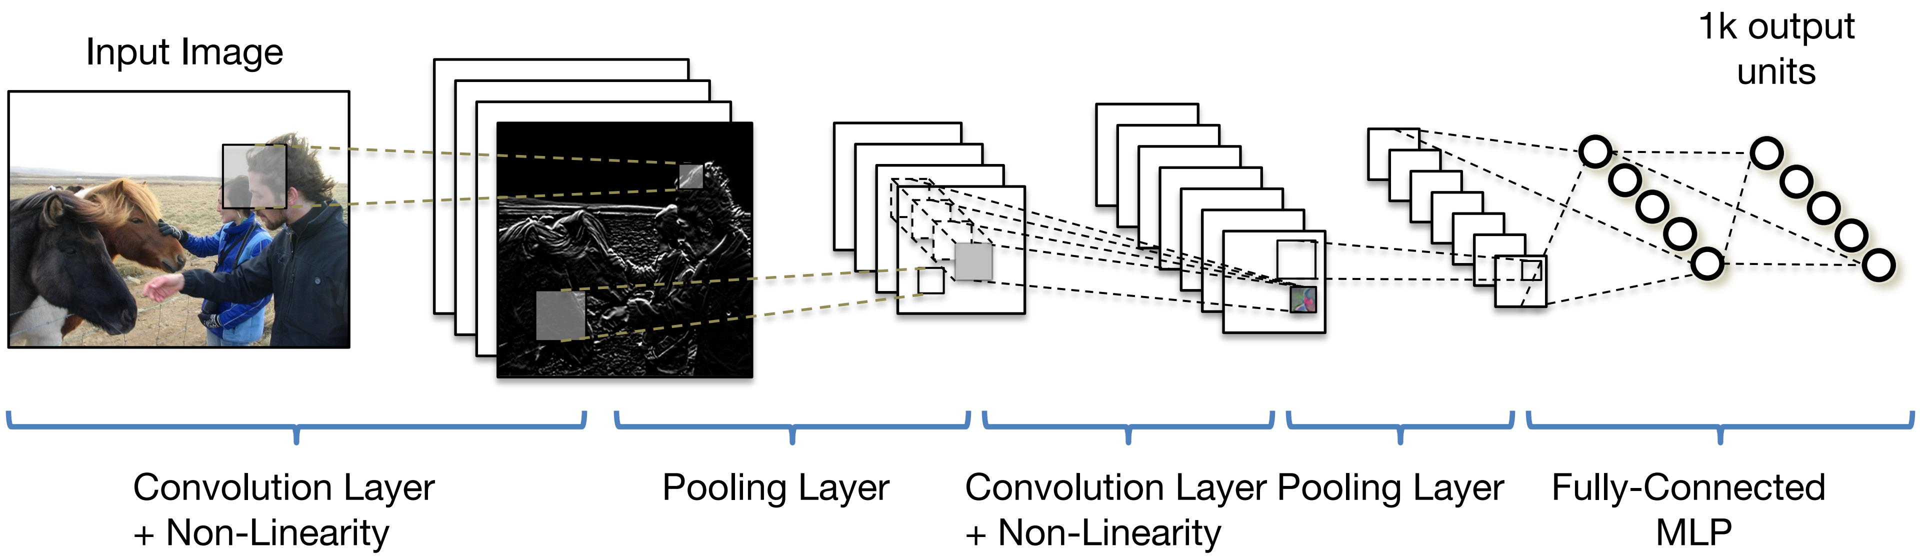

Slide Credit: Dhruv Batra

Semantic Segmentation

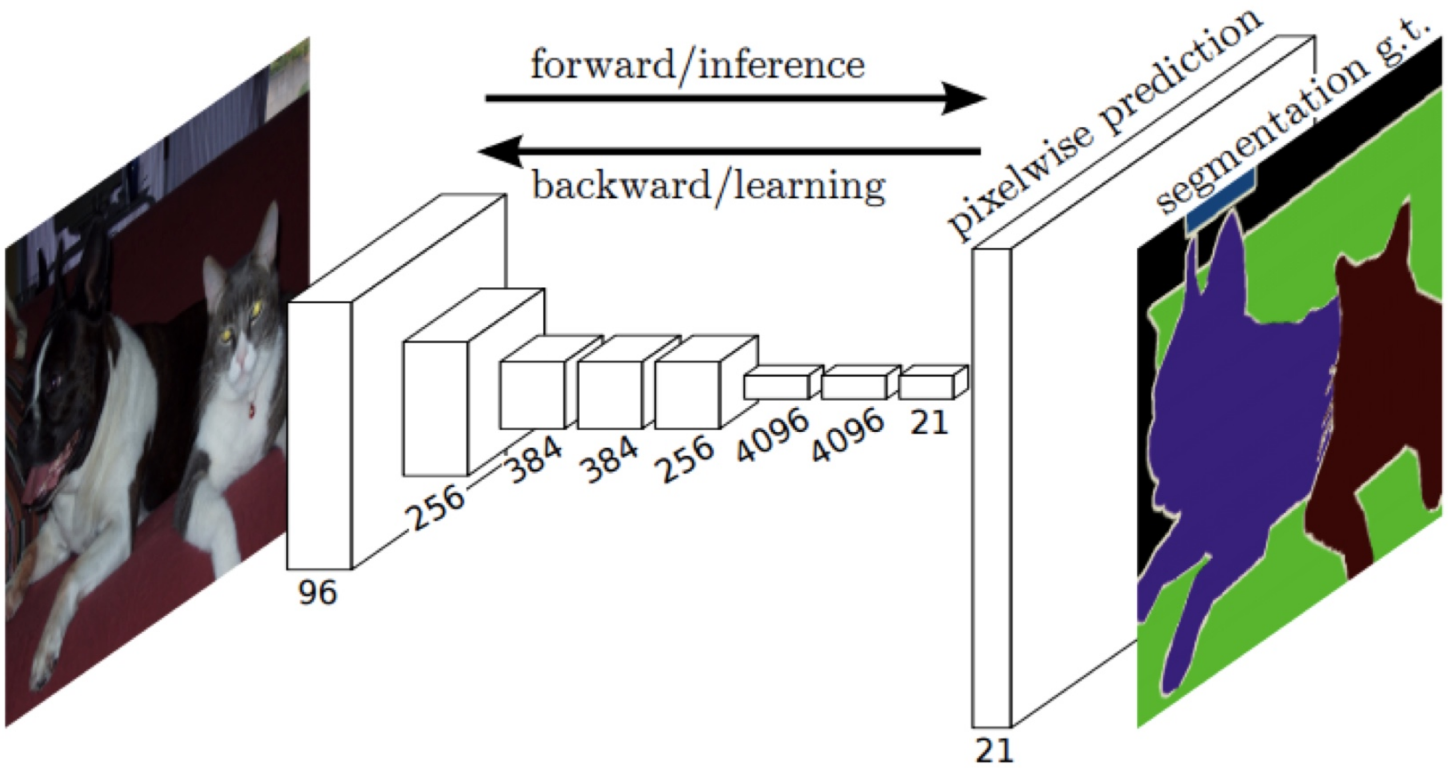

Fully Convolutional Networks for Semantic Segmentation [Long, Shelmer & Darrell 2015]

Instance Segmentation

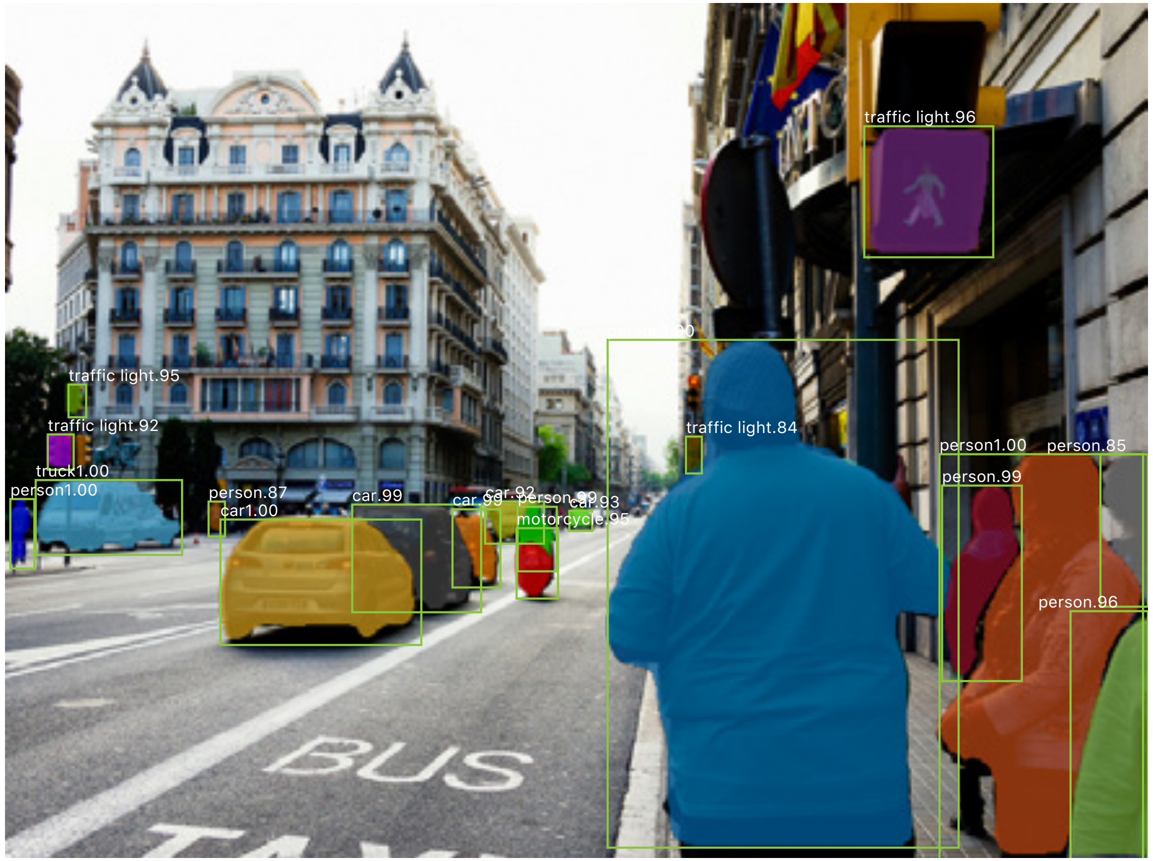

Mask R-CNN [He et al., 2017]

Deep Generative Networks

$4\times$ SRGAN (proposed)

original

Photo-Realistic Single Image Super-Resolution Using a Generative Adversarial Network [Ledig et al., 2017]

Deep Generative Networks

Unpaired Image-to-Image Translation using Cycle-Consistent Adversarial Networks [Zhu et al., 2017]