ICVGIP 2012, IIT Bombay

The eight Indian Conference on Computer Vision, Graphics and Image Processing

Tutorial

Topics in computational visual recognition

Subhransu Maji and Nikhil Rasiwasia

Tutorial outline

- Overview (this)

- Image representation (60 mins, 9:15 - 10:30)

  - motivation, local features, global features, break

- Learning (90 mins, 10:30 - 12:30)

  - discriminative models, tea-break, generative models, break

- Object detection and recognition (90 mins, 12:30 - 2:00)

  - Dalal & Triggs, lunch-break, PASCAL challenge, *poselets* and their applications, tea-break

- Cross-modal search (60 mins, 2:30 - 3:30)

lunch-break 60 mins, break 15 mins, tea-break 20-30 mins

Subhransu Maji
http://ttic.uchicago.edu/~smaji

IIT Kanpur (B.Tech CSE)
University of California, Berkeley (PhD)
Toyota Technological Institute at Chicago
(Research Asst. Prof.)

large scale learning
image classification
feature representation
object detection
pose estimation
segmentation
...

Nikhil Rasiwasia
http://www.svcl.ucsd.edu/~nikux/

IIT Kanpur (B. Tech EE)
University of California, San Diego (Ms+PhD)
Yahoo! Labs, Bangalore (Scientist)

image annotation
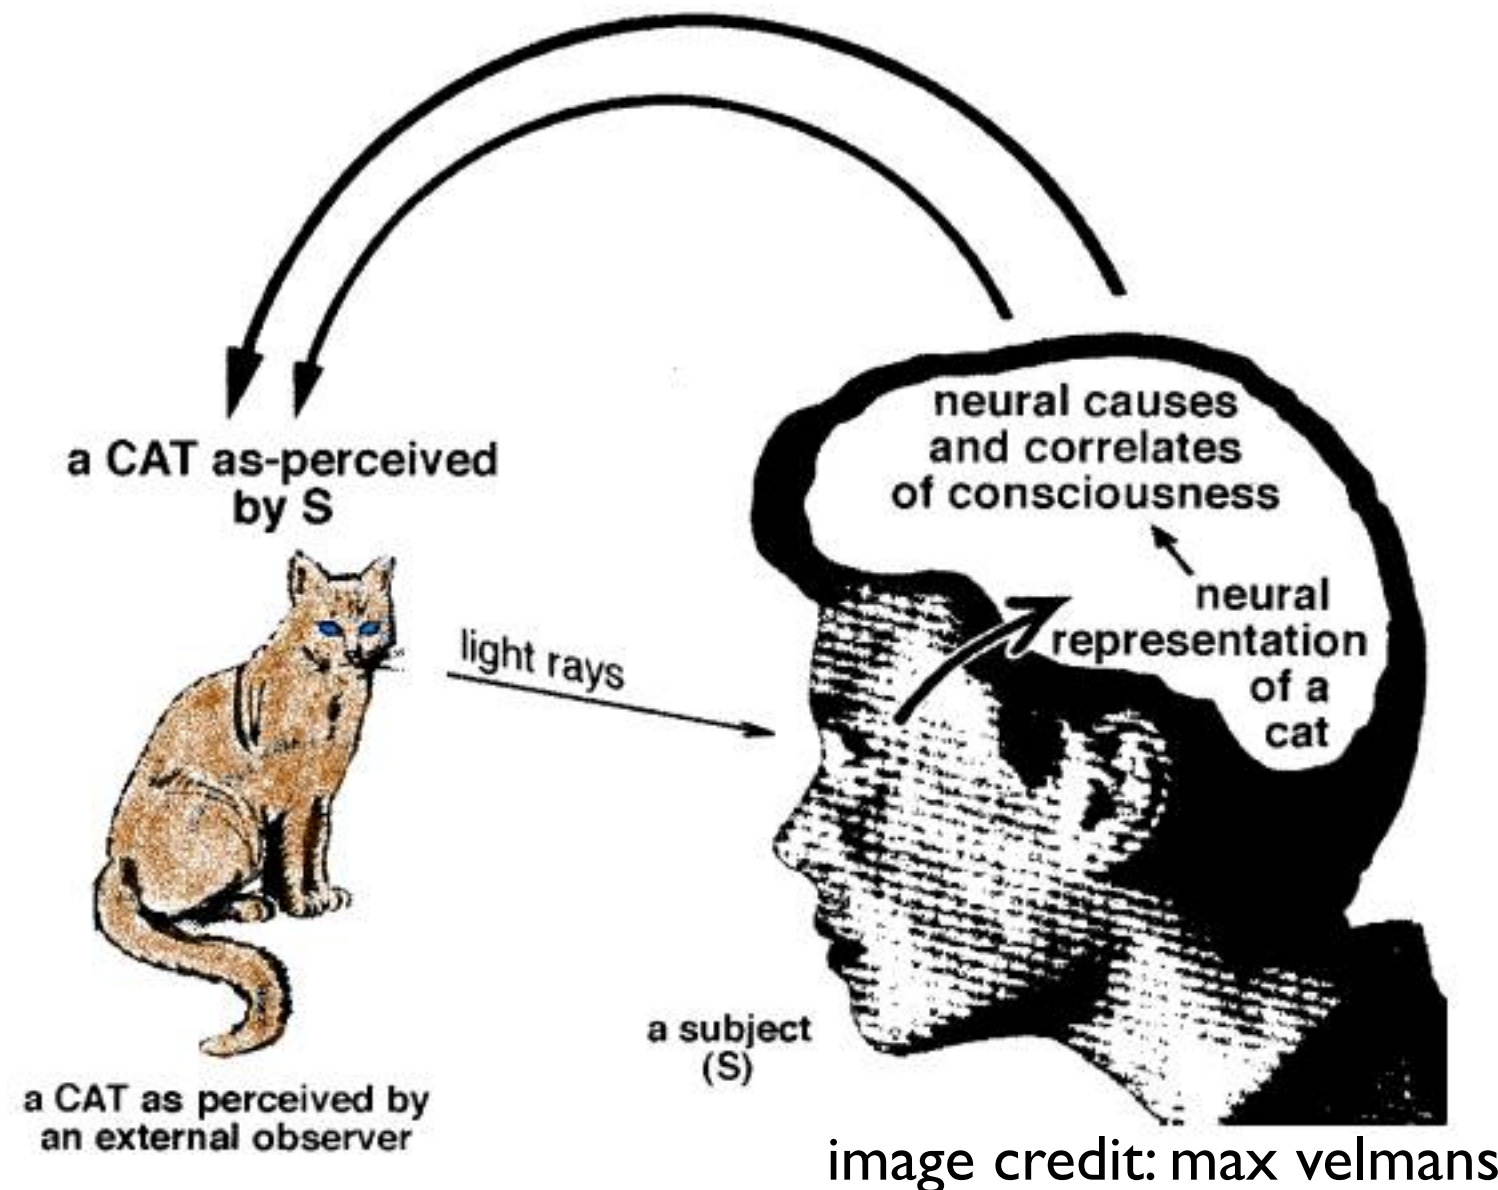image retrieval
scene classification
object detection
segmentation
....

- What background do you have in:

  - Computer vision ?

  - Machine learning ?

  - Related area: speech, nlp, etc.

- What are you currently doing?

  - Student (Bachelors, Masters, PhD)

  - Industry

a CAT as-perceived by S

a CAT as perceived by an external observer

light rays

a subject (S)

neural causes and correlates of consciousness

neural representation of a cat

image credit: max velmans

high dimensional data

hard inverse problem

have to rely on physics, geometry, statistics, learning, etc

Computer vision applications in the real world

image credit : www.visionxinc.com
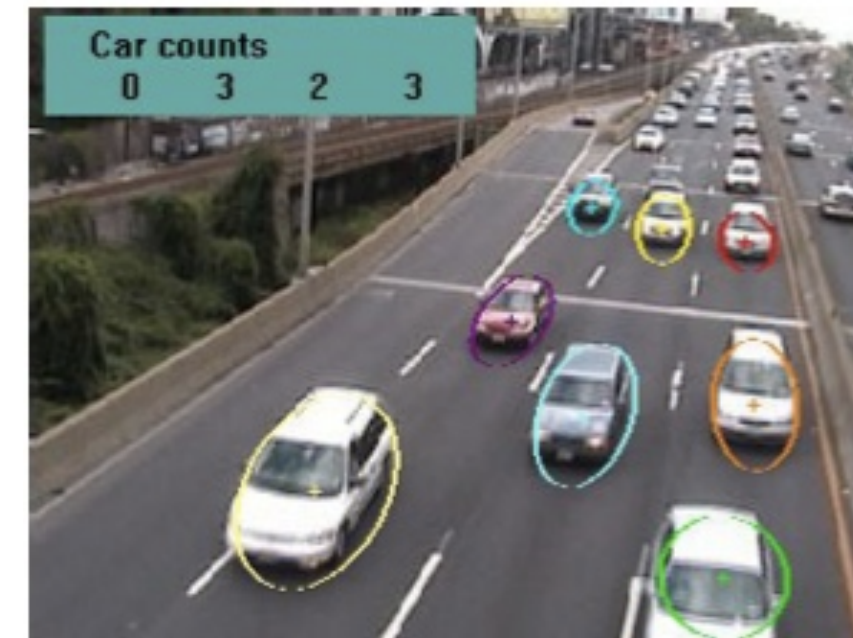
OCR @ ATM machines

machine inspection

automatic checkout

medical imaging

automotive safety

traffic monitoring

Examples from "Computer Vision and Algorithms", Richard Szeliski

What are we going to cover (and what not)?

- We will cover some aspects of high and low-level vision

  - image representation, learning models for visual categories for localizing and estimating their properties from images, cross-modal search and retrieval

- We will not cover:

  - geometry : image formation, structure from motion, etc.

  - low-level image representation : texture representation, boundary detection, optical flow, etc

  - mid-level image analysis: image segmentation, grouping

  - many topics in high-level image analysis