Conclusions

What we didn’t cover

Low-level

- Computational photography
  - Deblurring
  - Special cameras to capture light fields, variable depth of focus
- Physics-based vision
  - Seeing underwater, in a fog

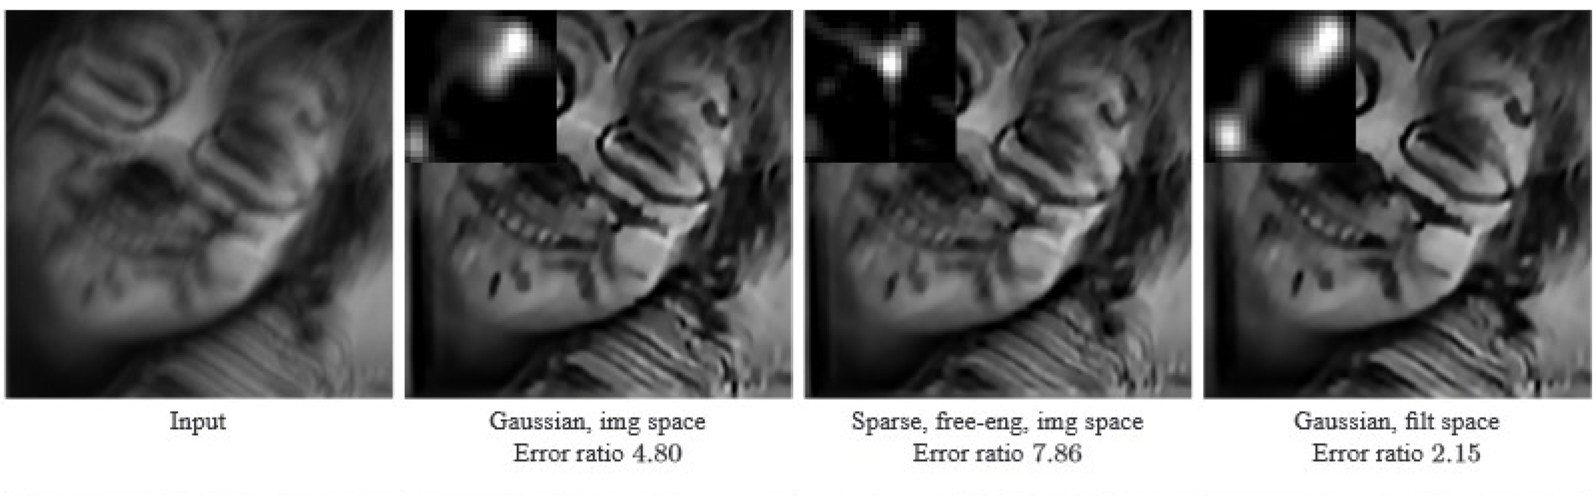

(From Bill Freeman)

Color/Lightness constancy

- What is the color of a material?
- What is its lightness?
- What are its other reflectance properties?

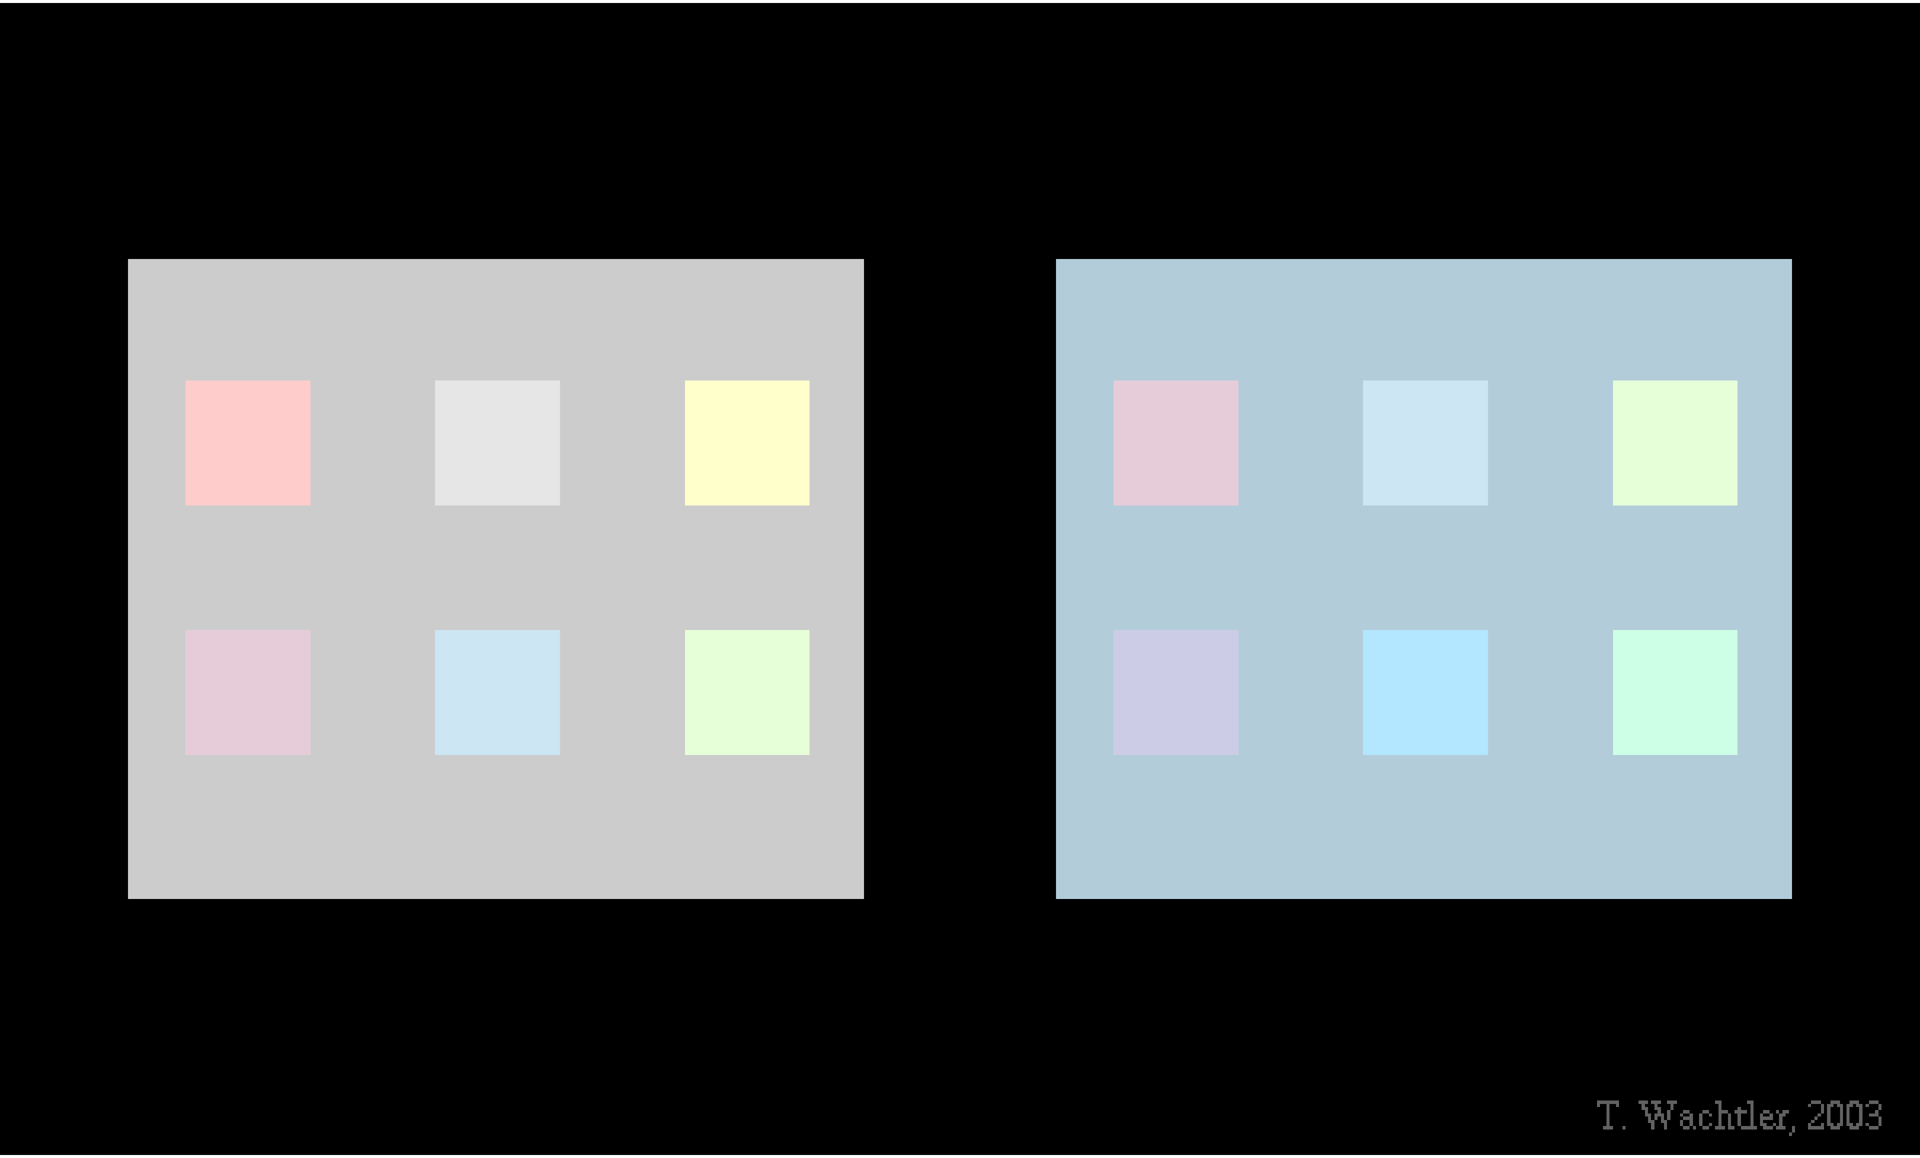

http://neuro.physik.uni-marburg.de/~wachtler/cc.html

Lighting

- Shape from shading
- Photometric stereo
- Lighting insensitive recognition
- Interaction between lighting and motion

Lighting affects appearance

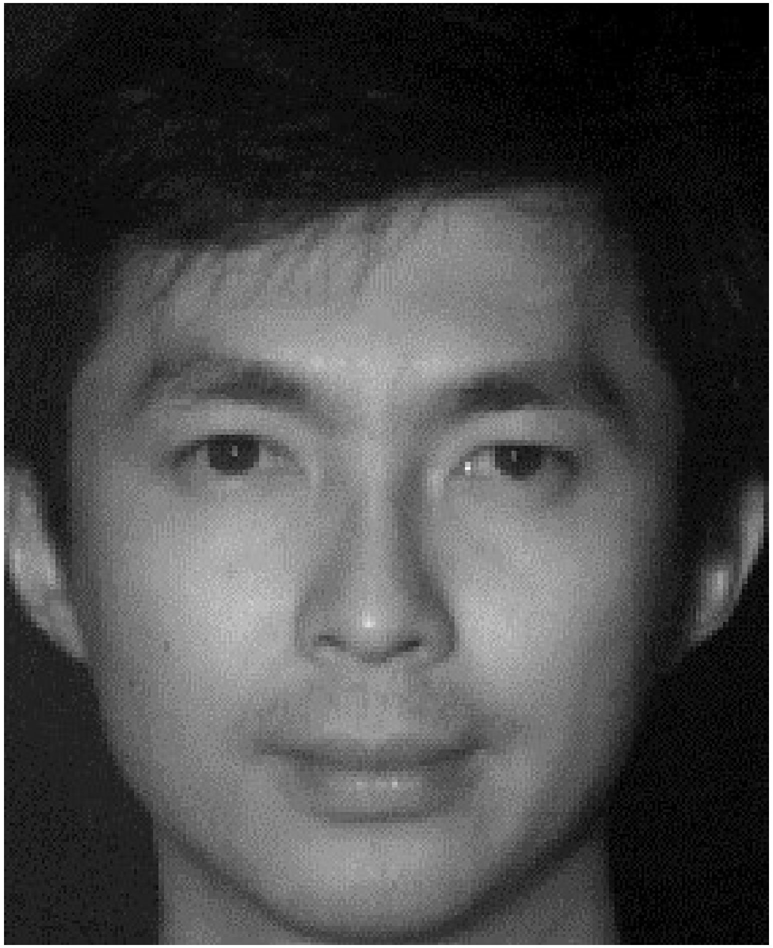 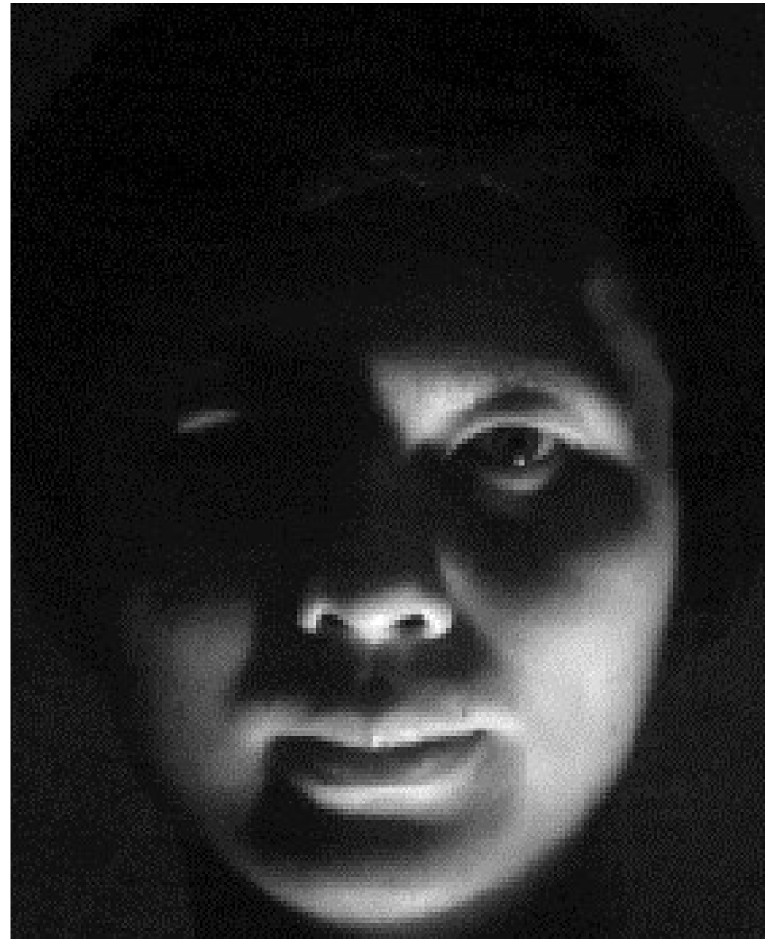

Video Processing

- Tracking – people and objects
- Tracking – hands, faces
- Video summarization
  - http://www.vision.huji.ac.il/video-synopsis/
- Understanding action

Large-scale reconstruction

- Tomasi-Kanade algorithm
- Bundle adjustment and other optimization
- http://grail.cs.washington.edu/rome/

3D Sensing

- Kinect
- How do we use 3D models

(From Jaime Shotton)

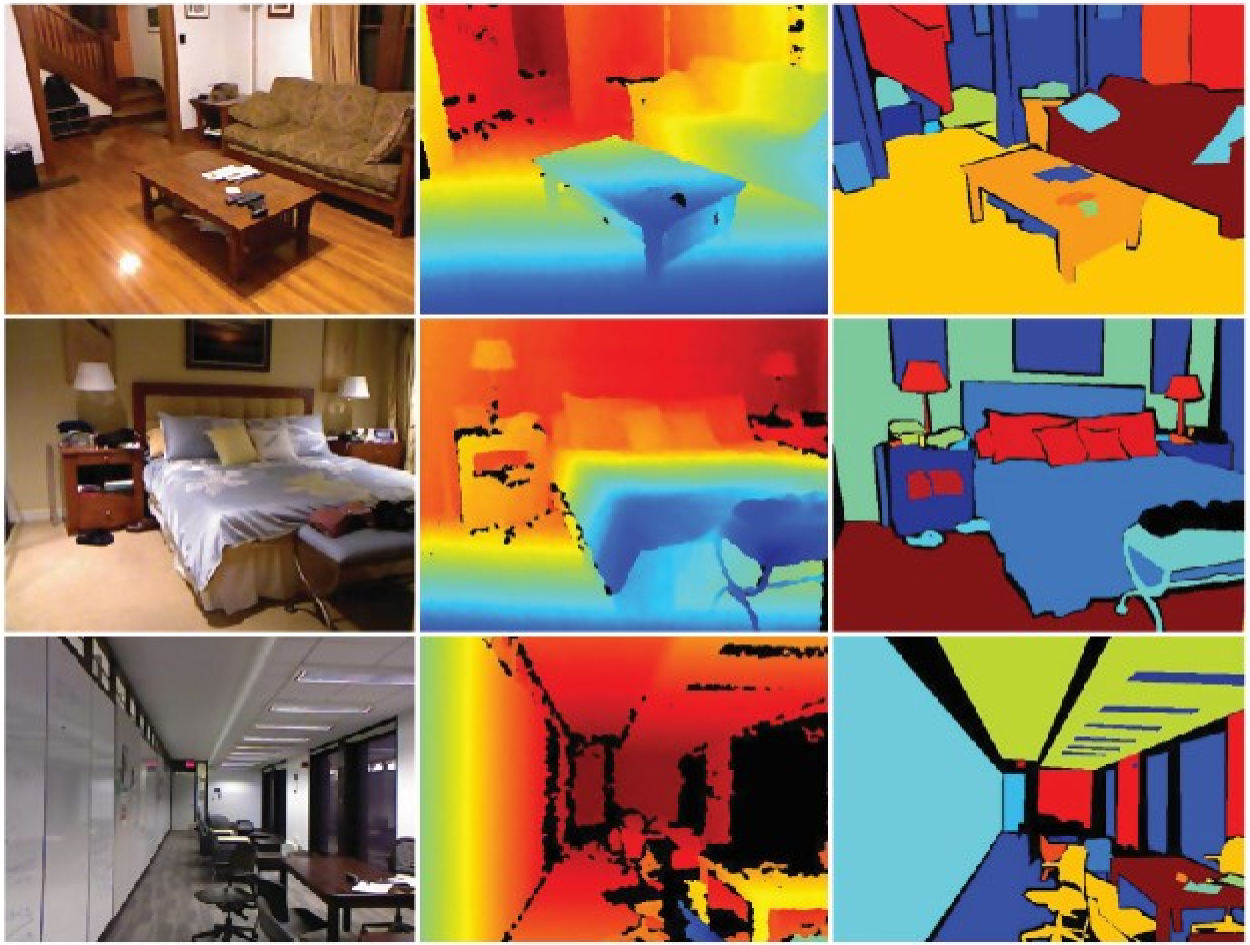

ples from the dataset in [44]. From the left to the right: RGB images, depth images and annotated object regions with l

Understanding an Image

- That is, understanding the properties of an image, rather than reconstructing depth or recognizing objects.
- Pose
- Scene understanding
- Relationships between objects (holding, touching, talking to, …)

Medical Image Analysis

- Segmentation
- Shape understanding
- Identification of key features
  - Count cells
  - Find tumors

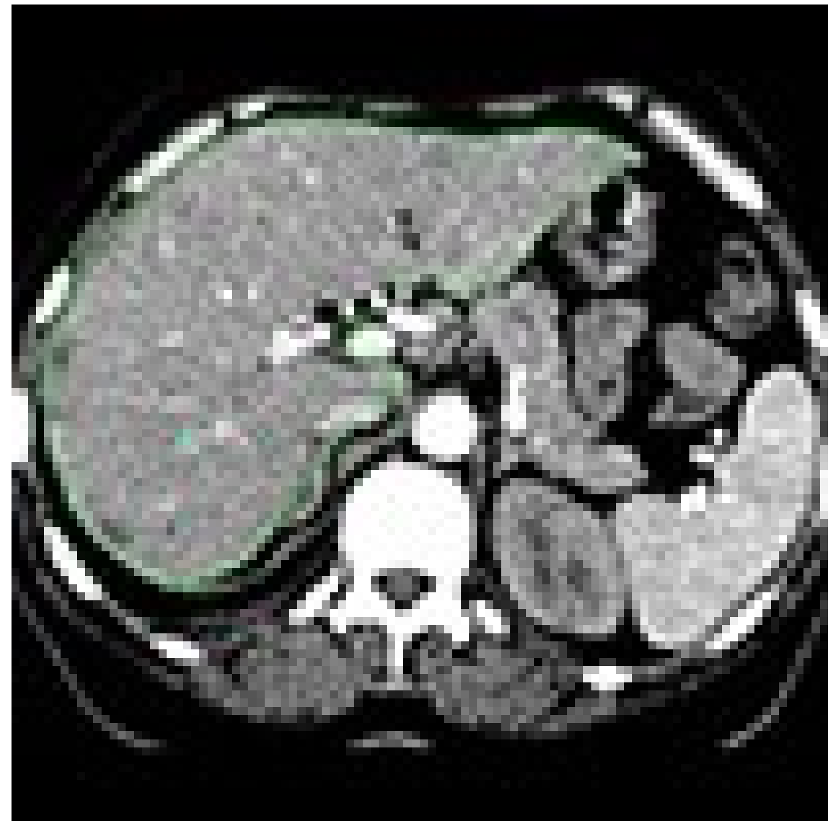

(From Haibin Ling)

Document understanding

- OCR
- Segmentation
  - Find lines
  - Pictures, plots
  - Text spotting in the wild
  - Handwriting

(From Serge Belongie)

Computational modeling of human vision

- Modeling illusory contours

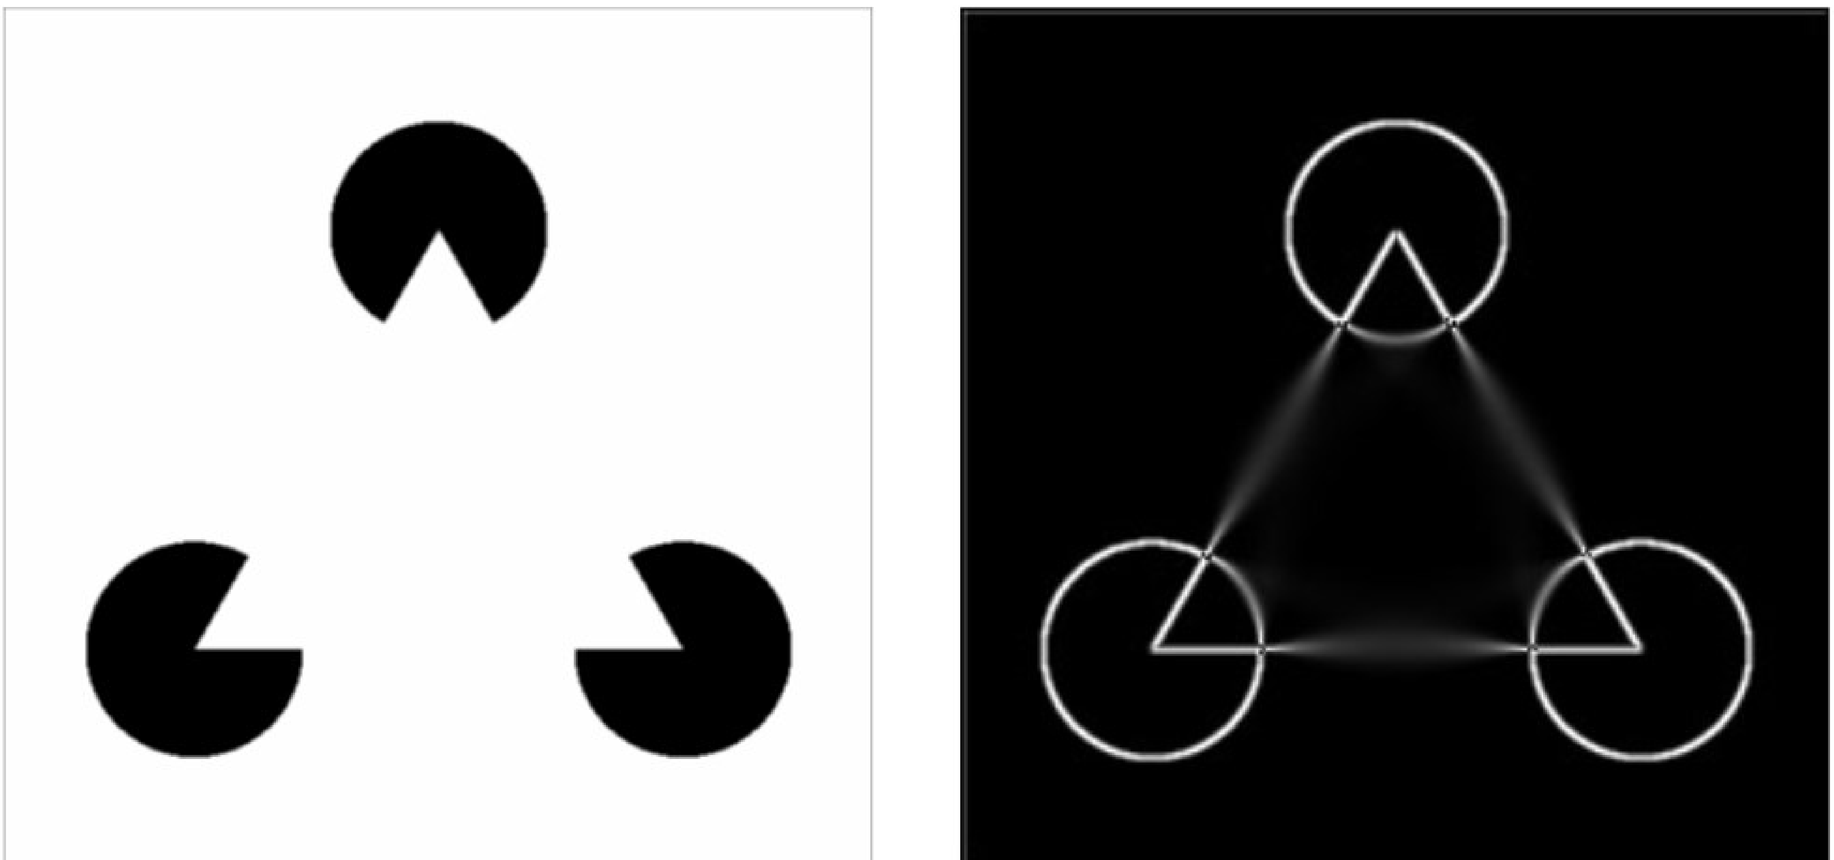

(From Lance Williams)

State of vision

- Research
  - Lots of big data
    - Machine learning.
      - Deep learning for features
    - Large-scale optimization
  - New sensors
    - Using Kinect-type data, building new cameras.
    - Is vision for VR/AR going to be important?
  - Leverage ideas from other fields
    - Sparse coding, manifolds, ideas from machine learning (boosting, CRFs, …)
  - Growing emphasis on engineering
  - Ongoing work on basic problems

Applications

- "Golden age of computer vision"
- Assistive tools for cars
  - Self-driving cars
  - Lane following, smart cruise control, pedestrian detection, warnings about other cars, driver attentiveness.
- Face processing: recognition, detection,
- 3D reconstruction/Motion understanding
  - Kinect, clothes fitting, avatars, models for VR
- Computational photography
- Video search

Many fundamental problems remain

- Representation
  - How should we represent an image or 2D or 3D shape?
    - Representations of images: pixels, superpixels, fourier, wavelets, features from deep learning, edges, segmentations.
    - Representation is about making information explicit and about throwing away information.
    - Imagine you are walking in the forest for an hour. What information do you extract visually? How does this compare to what computer vision can do?
  - Should all representations be task-specific (supervised)? Are there generic representations shared between tasks?
- Interaction between lighting, shape and motion
- Do we value experiments too much, and ideas too little?
  - Value of theoretical elegance
- Do we overvalue aggregate performance on datasets, and undervalue explaining critical examples?
  - Graduate student descent

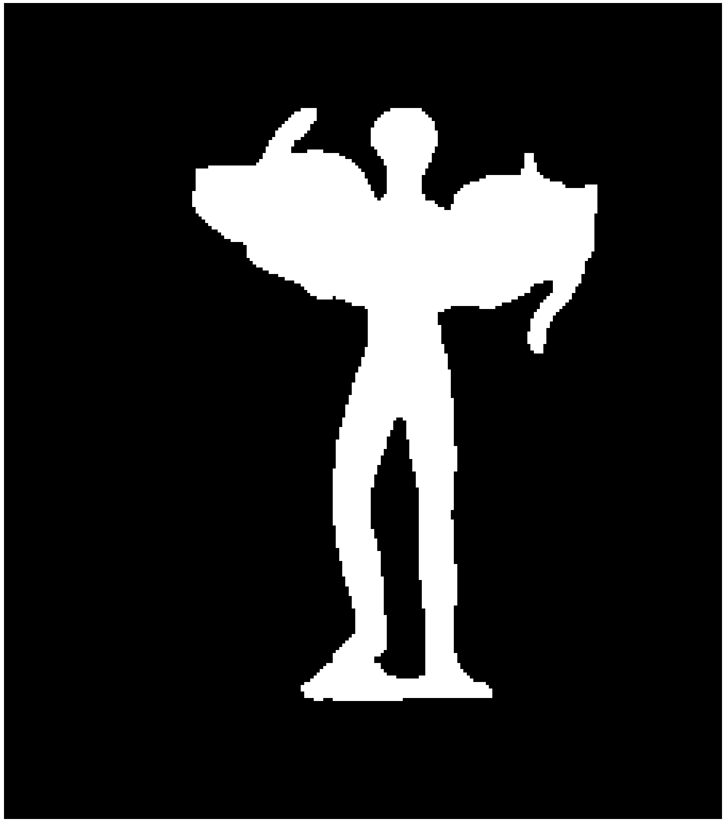

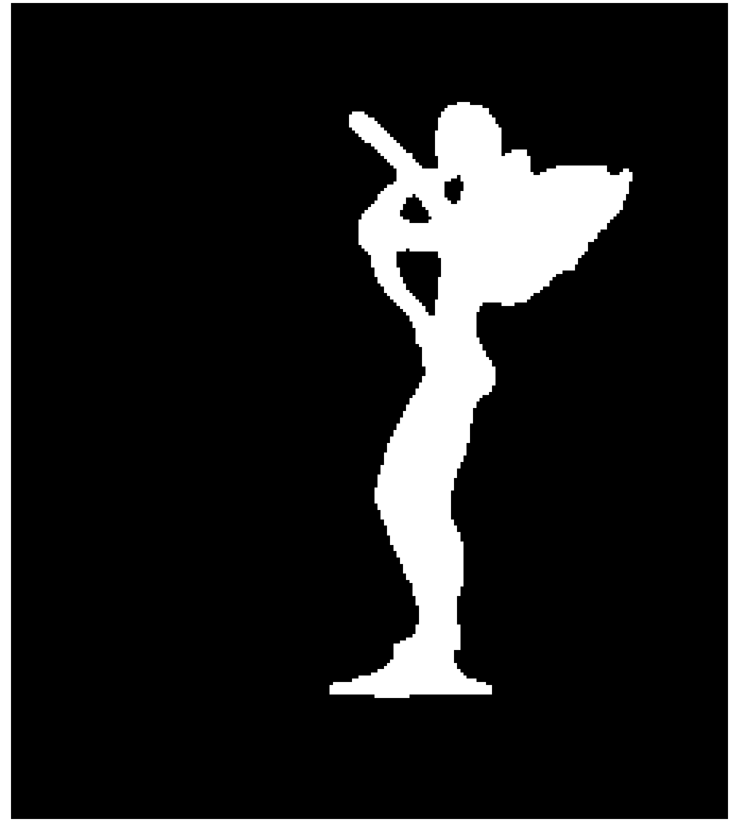

Feedback

- More math? Less?
- Too easy? Too hard?
- Textbook needed? Enough references?
- Problem sets fun? Boring?