Announcements

- Project proposals due Sep 20 (two weeks minus two days)
- Groups OK with instructor approval

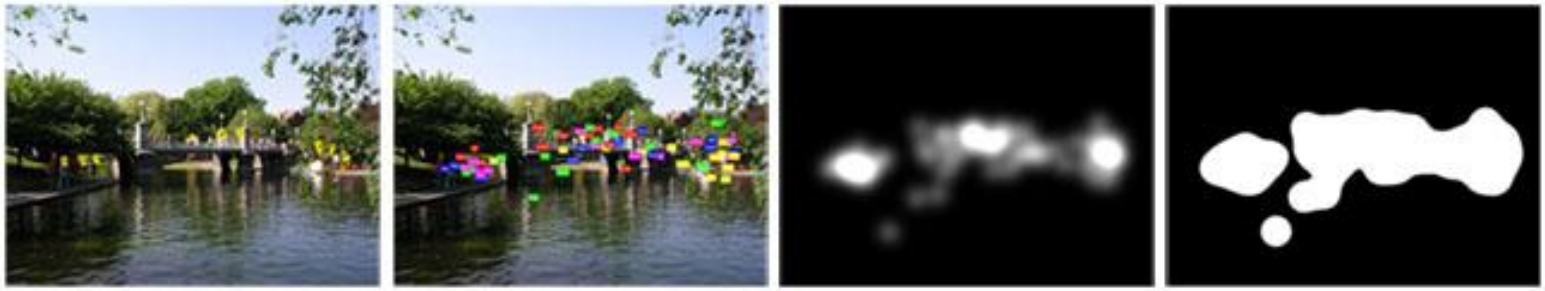

Judd, Ehinger, Durand, Torralba. **Learning to Predict Where Humans Look.** ICCV 2009.

Saliency

CS7670, Noah Snavely, Sep 8, 2011

How does one's gaze move?

- Fixations
- Saccades
  - Eye motion between fixations
  - Fastest movement produced by the human body
- Microsaccades
  - Without them, our vision would fade
  - Improves sharpness of vision?
    - "Eye flickers key for fine detail," BBC News, 13 June 2007, http://news.bbc.co.uk/2/hi/health/6745443.stm

- See also smooth pursuit

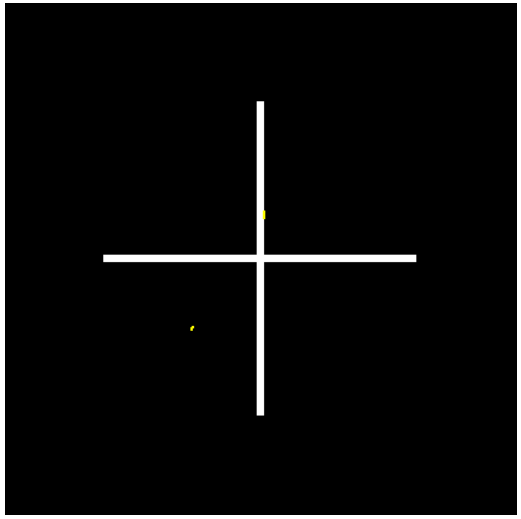

2348 ms @ 250 Hz

1

Free examination.

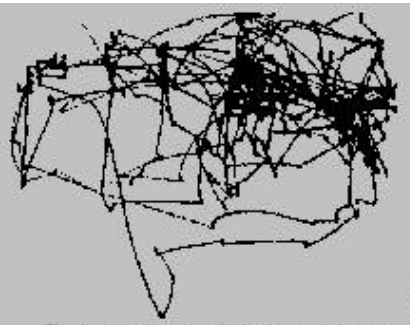

2

Estimate material circumstances of the family

3

Give the ages of the people.

4

Surmise what the family had been doing before the arrival of the unexpected visitor.

5

Remember the clothes worn by the people.

6

Remember positions of people and objects in the room.

7

Estimate how long the visitor had been away from the family.

3 min. recordings of the same subject

Early eye-tracking research, Alfred Yarbus, 1960s

Demo from today's paper

Awareness

- http://www.youtube.com/watch_popup?v=Ahg6qcgoay4

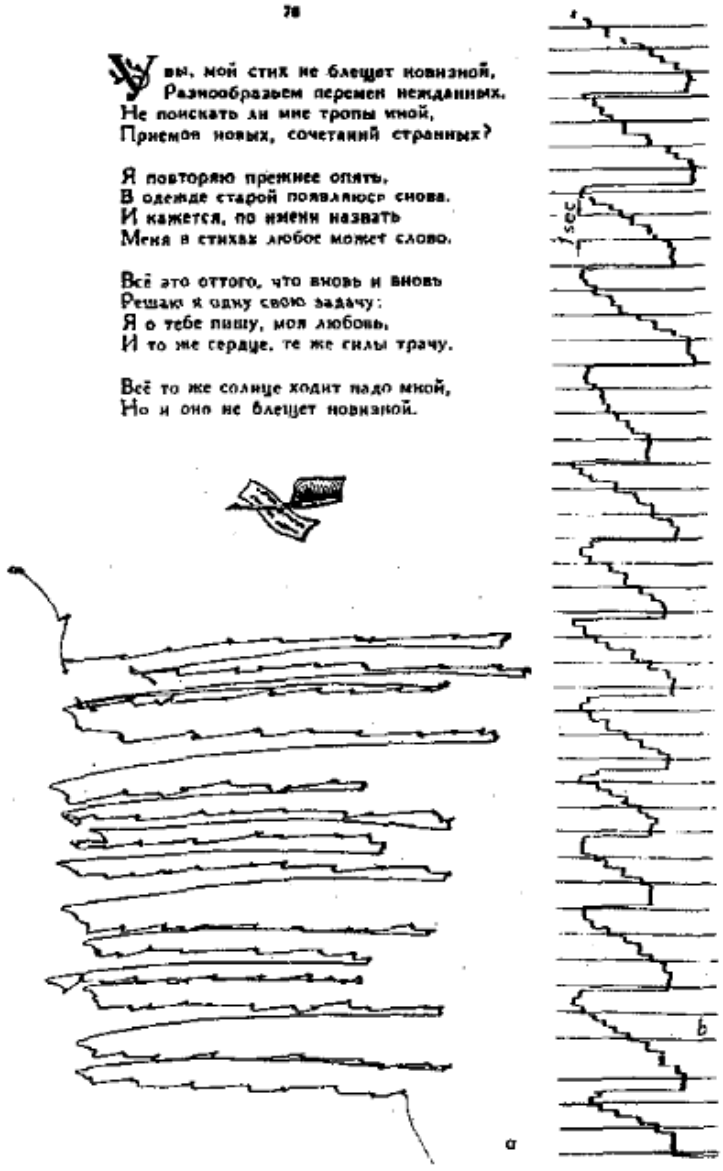

Fig. 125. Record of the eye movements of a subject reading a Shakespeare sonnet. Record on stationary photosensitive paper (a) and on moving phototape of a photokymograph (b).

Eye Movements and Vision, Alfred L. Yarbus

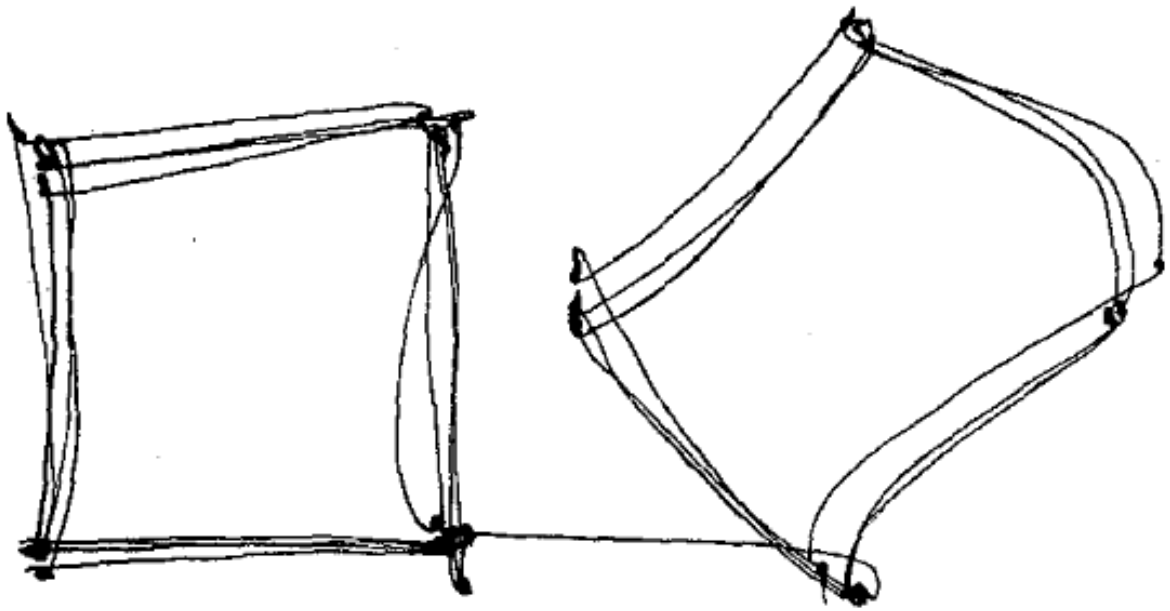

Fig. 83. Record of saccades between the corners of two squares on stationary photosensitive paper.

Eye Movements and Vision, Alfred L. Yarbus

Figure 2.2: Typical eye trackers used for cognitive pupillometry. The eye tracker on the left is an SMI iView X chin-rest style instrument, used primarily for reading and other high-precision applications [108]. The eye tracker on the right is the Polhemus VisionTrak Standard Head Mounted Eye Tracking System [93], used for mobile applications, especially driving and piloting.

Jeff Klingner, PhD thesis

Useability

http://www.useit.com/alertbox/reading_pattern.html

Applications

Abstraction

[Santella et al., SIGGRAPH 2002]

(a) original (b) gaze-based (c) automatic

Automatic cropping

[Santella et al., CHI 2006]

Applications

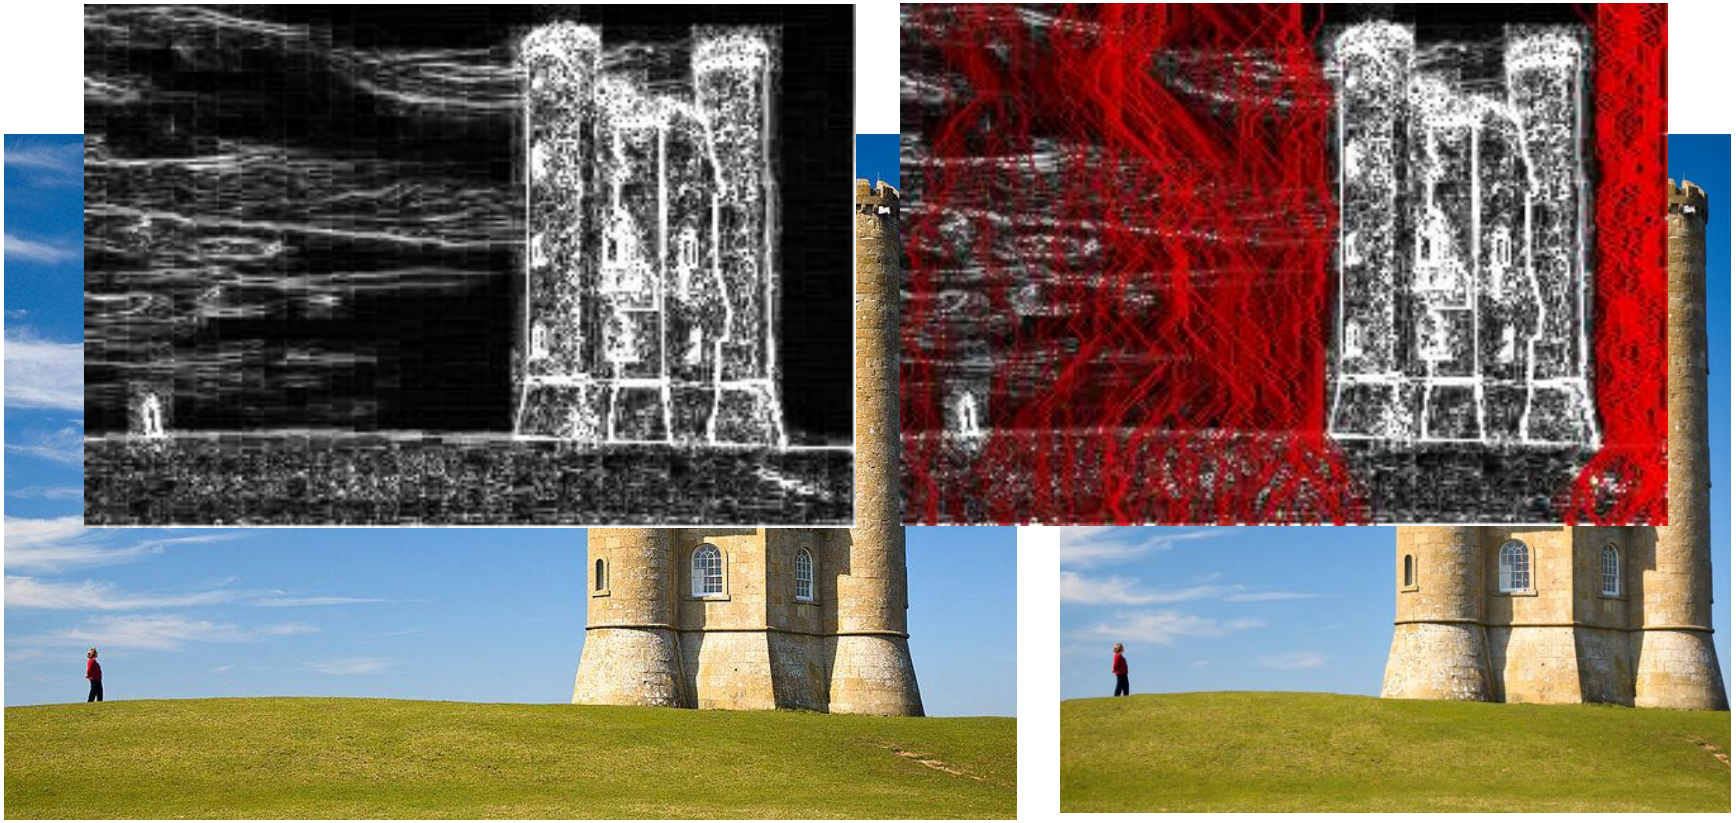

Image Retargeting [Avidan et al., SIGGRAPH 2007]

Applications

- Object recognition?
- Detecting surprising events for surveillance?

What is saliency?

- Measure of conspicuity
- Likelihood of a location to attract a person's attention
- As we've seen, may be task dependent

- Visual system doesn't keep track of every detail

Can you spot the difference?

Airplane

- **Change blindness**
- More examples available at:
  - http://www2.psych.ubc.ca/~rensink/flicker/

Can you spot the difference?

Chopper

- **Change blindness**
- More examples available at:
  - http://www2.psych.ubc.ca/~rensink/flicker/

Can you spot the difference?

http://www.cogsci.uci.edu/~ddhoff/cb.html

Computer graphics applications

For static objects, sensitivities for length estimation are well known. Consider the following pair of lines:

Without the use of a ruler you may have difficulty determining that the right line is 2.7% longer than the line on the left. However, you probably will not have difficulty with the following pair:

But, as we will show, in some conditions observers could not reliably tell these two lengths apart. These limitations can be exploited in tools for processing or synthesizing animations, ones that will work in many of the conditions encountered in normal viewing.

Computer graphics applications

- Ramanarayanan, et al. **Perception of Complex Aggregates**, SIGGRAPH 2008.

How do we predict saliency?

- Probably hopeless
- But people have still tried

- In the past: low-level cues – intensity, contrast, color differences, other gradient information
- Today: mixture of low- and higher-level features + machine learning on a large training set

Judd, Ehinger, Durand, Torralba. **Learning to Predict Where Humans Look.** ICCV 2009.

Prior work

(a) Original image

(b) Hou and Zhang

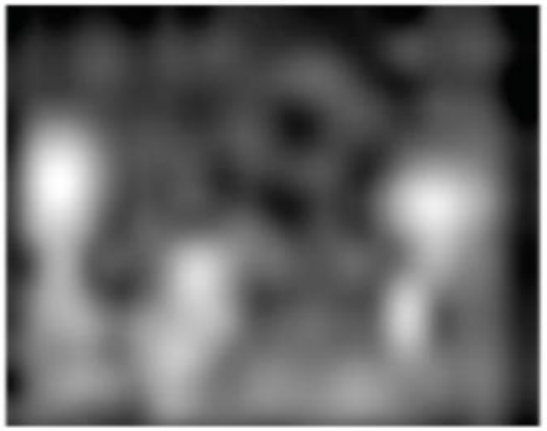

(c) Itti and Koch

- Itti and Koch

L. Itti and C. Koch. A saliency-based search mechanism for overt and covert shifts of visual attention, 2000.

- Ruth Rosenholtz

R. Rosenholtz. A simple saliency model predicts a number of motion popout phenomena. Vision Research 39, 19:3157–3163, 1999.

- Hou and Zhang

X. Hou and L.Q. Zhang. Saliency detection: A spectral residual approach. 2007.

Slide credit: Tilke Judd

Saliency models

- based on biologically plausible linear filters
- measure intensity, illumination, and color contrast
- lots of parameters
- bottom up model

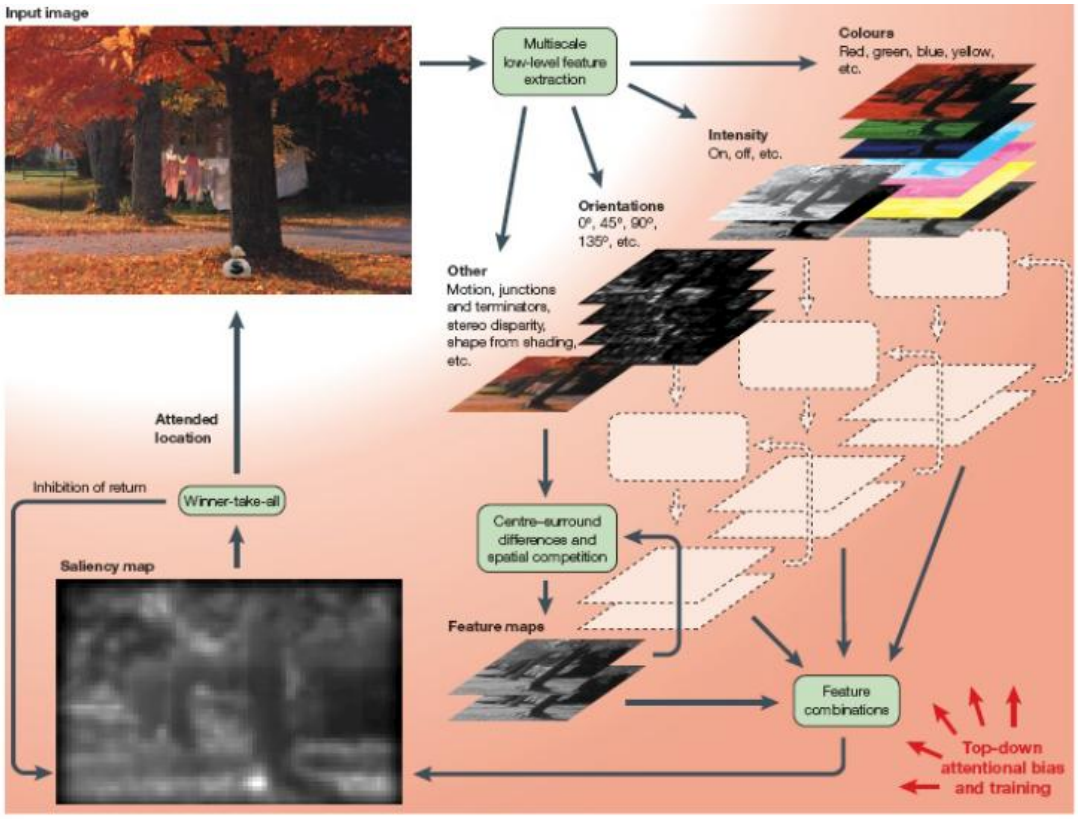

Bottom-up saliency model. From Itti and Koch [2001]

Slide credit: Tilke Judd

Low-level saliency measures don't always predict where people look

(a) Original image (b) Itti and Koch Saliency Map (c) eye tracking locations

Slide credit: Tilke Judd

Where people look

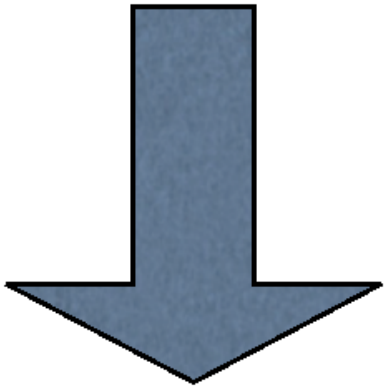

- Top down task and scene dependent cues
- Bottom up saliency cues

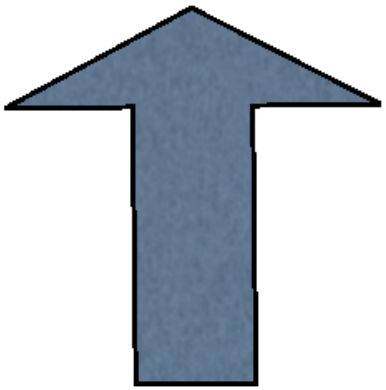

Slide credit: Tilke Judd

This paper

- Goal: Learn a model of where people look directly from eye tracking data
- Steps
  - Collect eye tracking dataset
  - Learn a new model

Slide credit: Tilke Judd

Collecting training data

Natural images of objects and scenes downloaded from Flickr and LabelMe

Slide credit: Tilke Judd

Collect eye tracking data

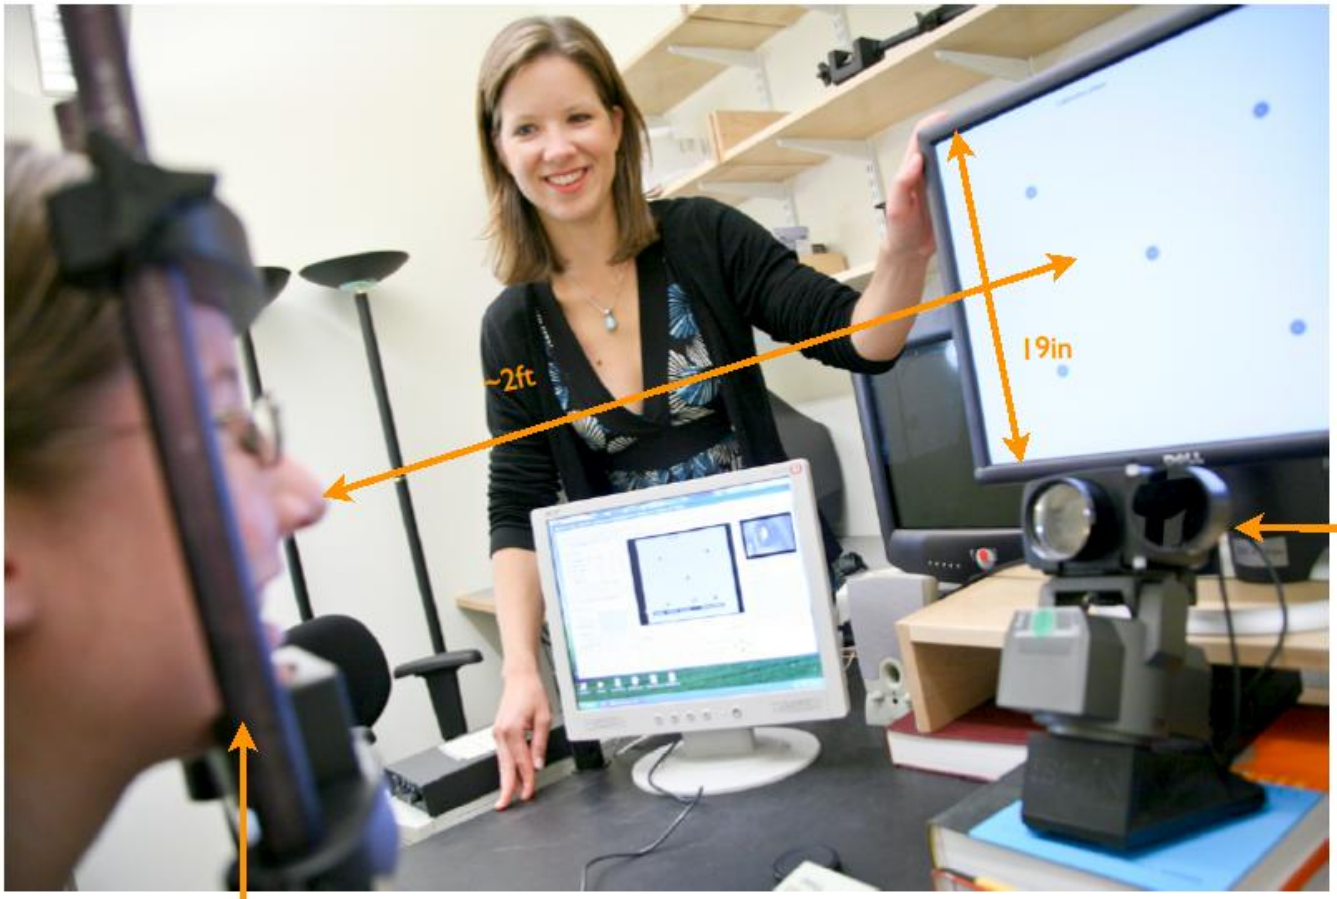

screen resolution
1280x1024

each image shown for
3 seconds

eye tracker measures
location of eye fixation
several times a second.

user rests head in chin rest

[Photo Credit: Jason Dorfman CSAIL website]

15 users on 1003 images

Slide credit: Tilke Judd

Eye tracking data

fixations for one user

Slide credit: Tilke Judd

Eye tracking data

first 5 fixations for 15 users

Slide credit: Tilke Judd

Eye tracking data

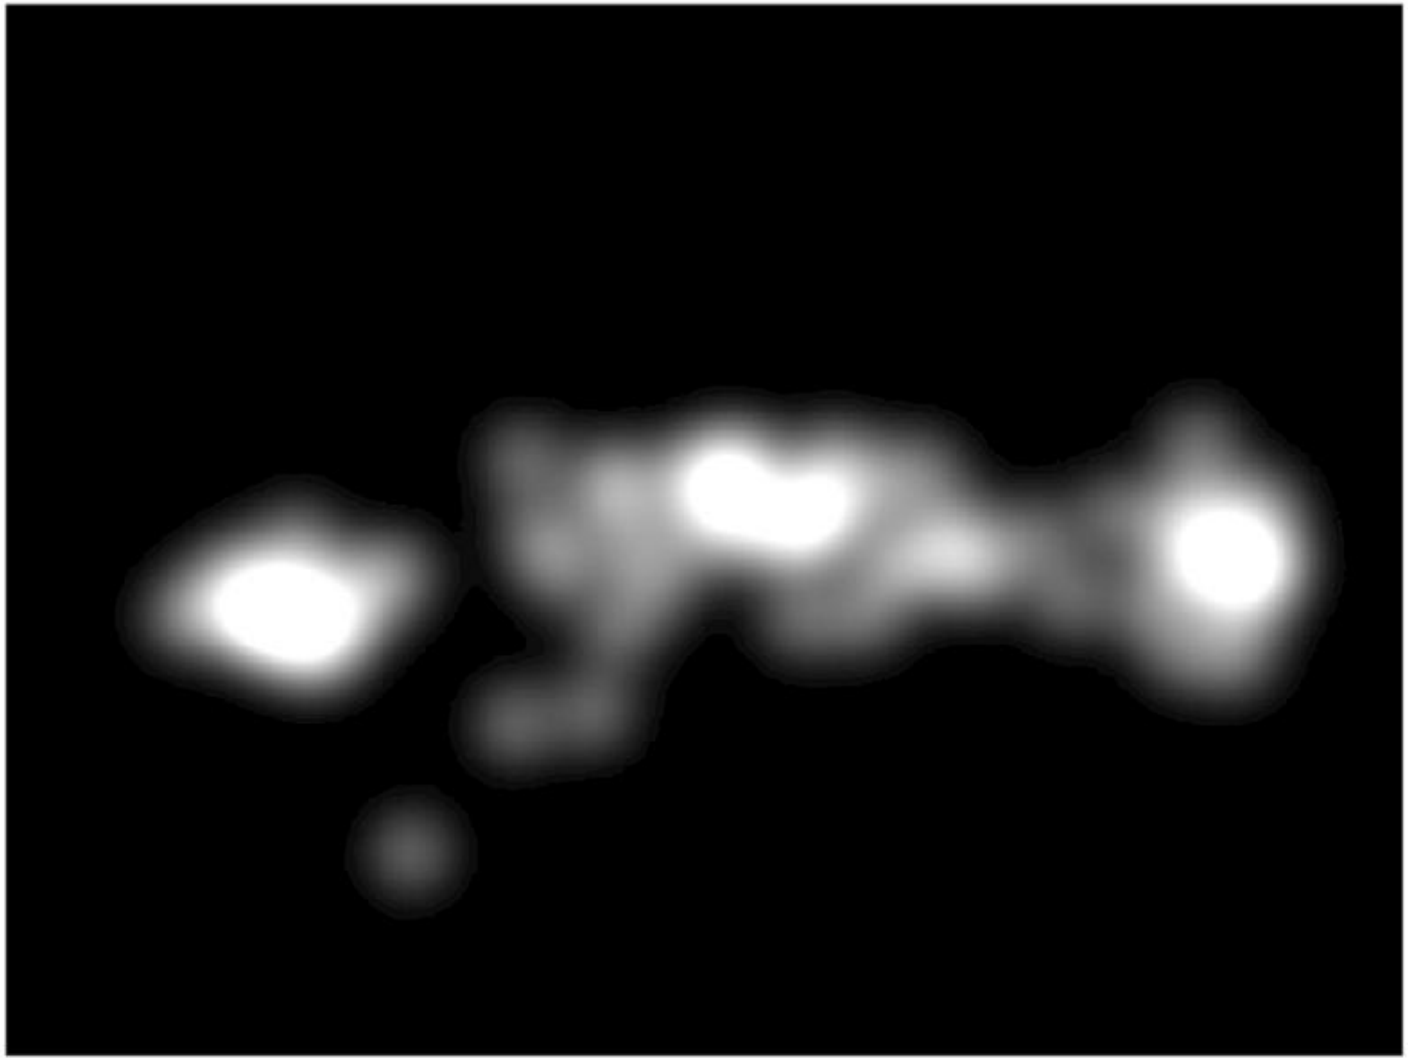

Average fixation locations / continuous saliency map

Slide credit: Tilke Judd

Eye tracking data

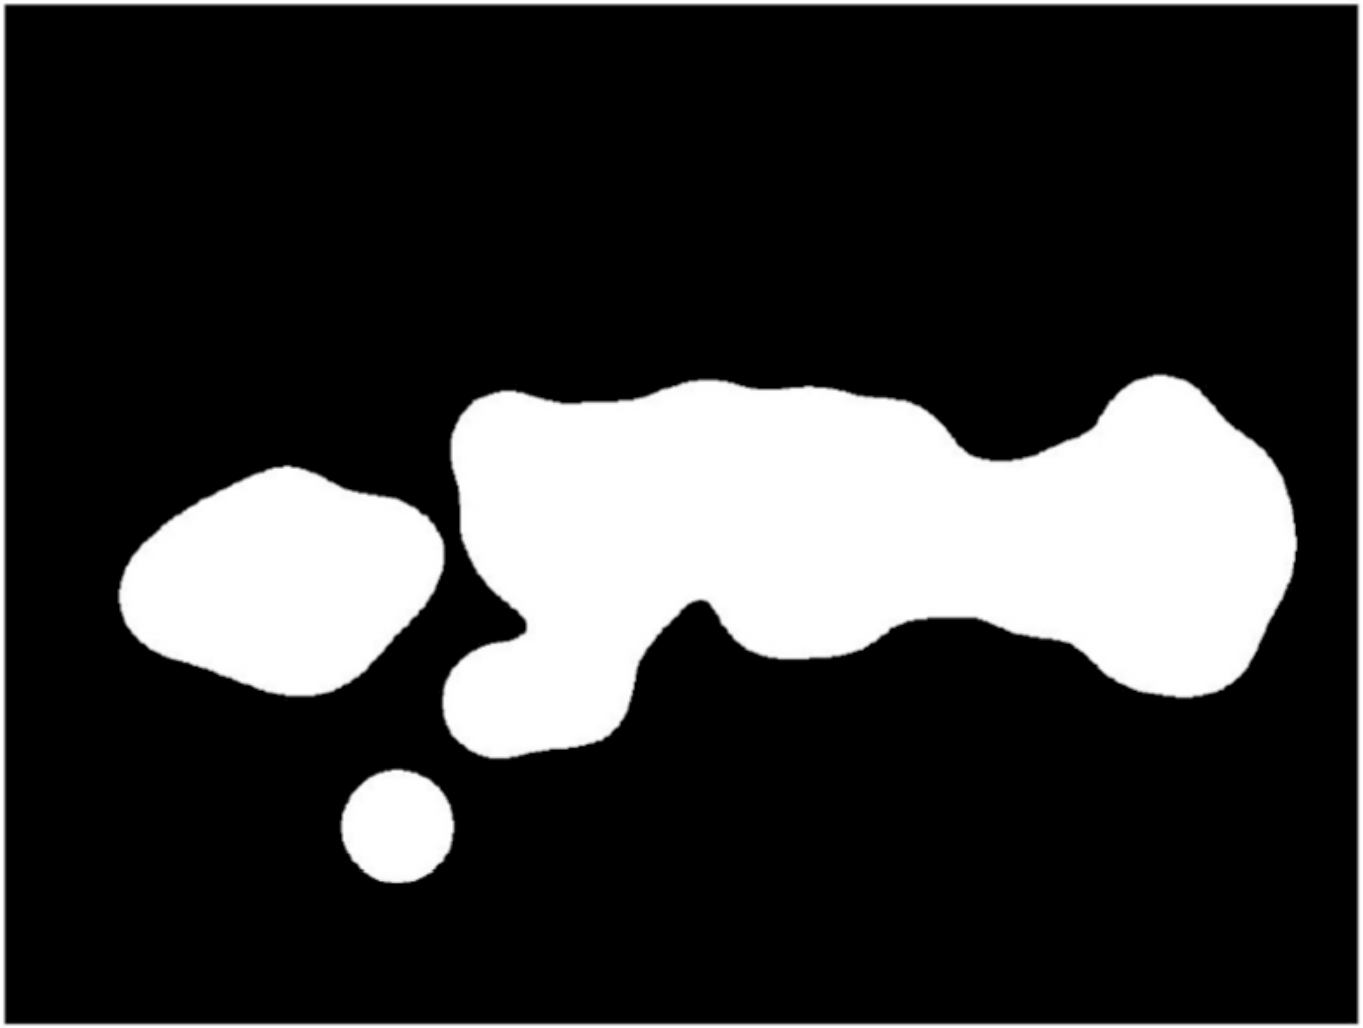

top 20% salient locations

Slide credit: Tilke Judd

Slide credit: Tilke Judd

Slide credit: Tilke Judd

Slide credit: Tilke Judd

Slide credit: Tilke Judd

Slide credit: Tilke Judd

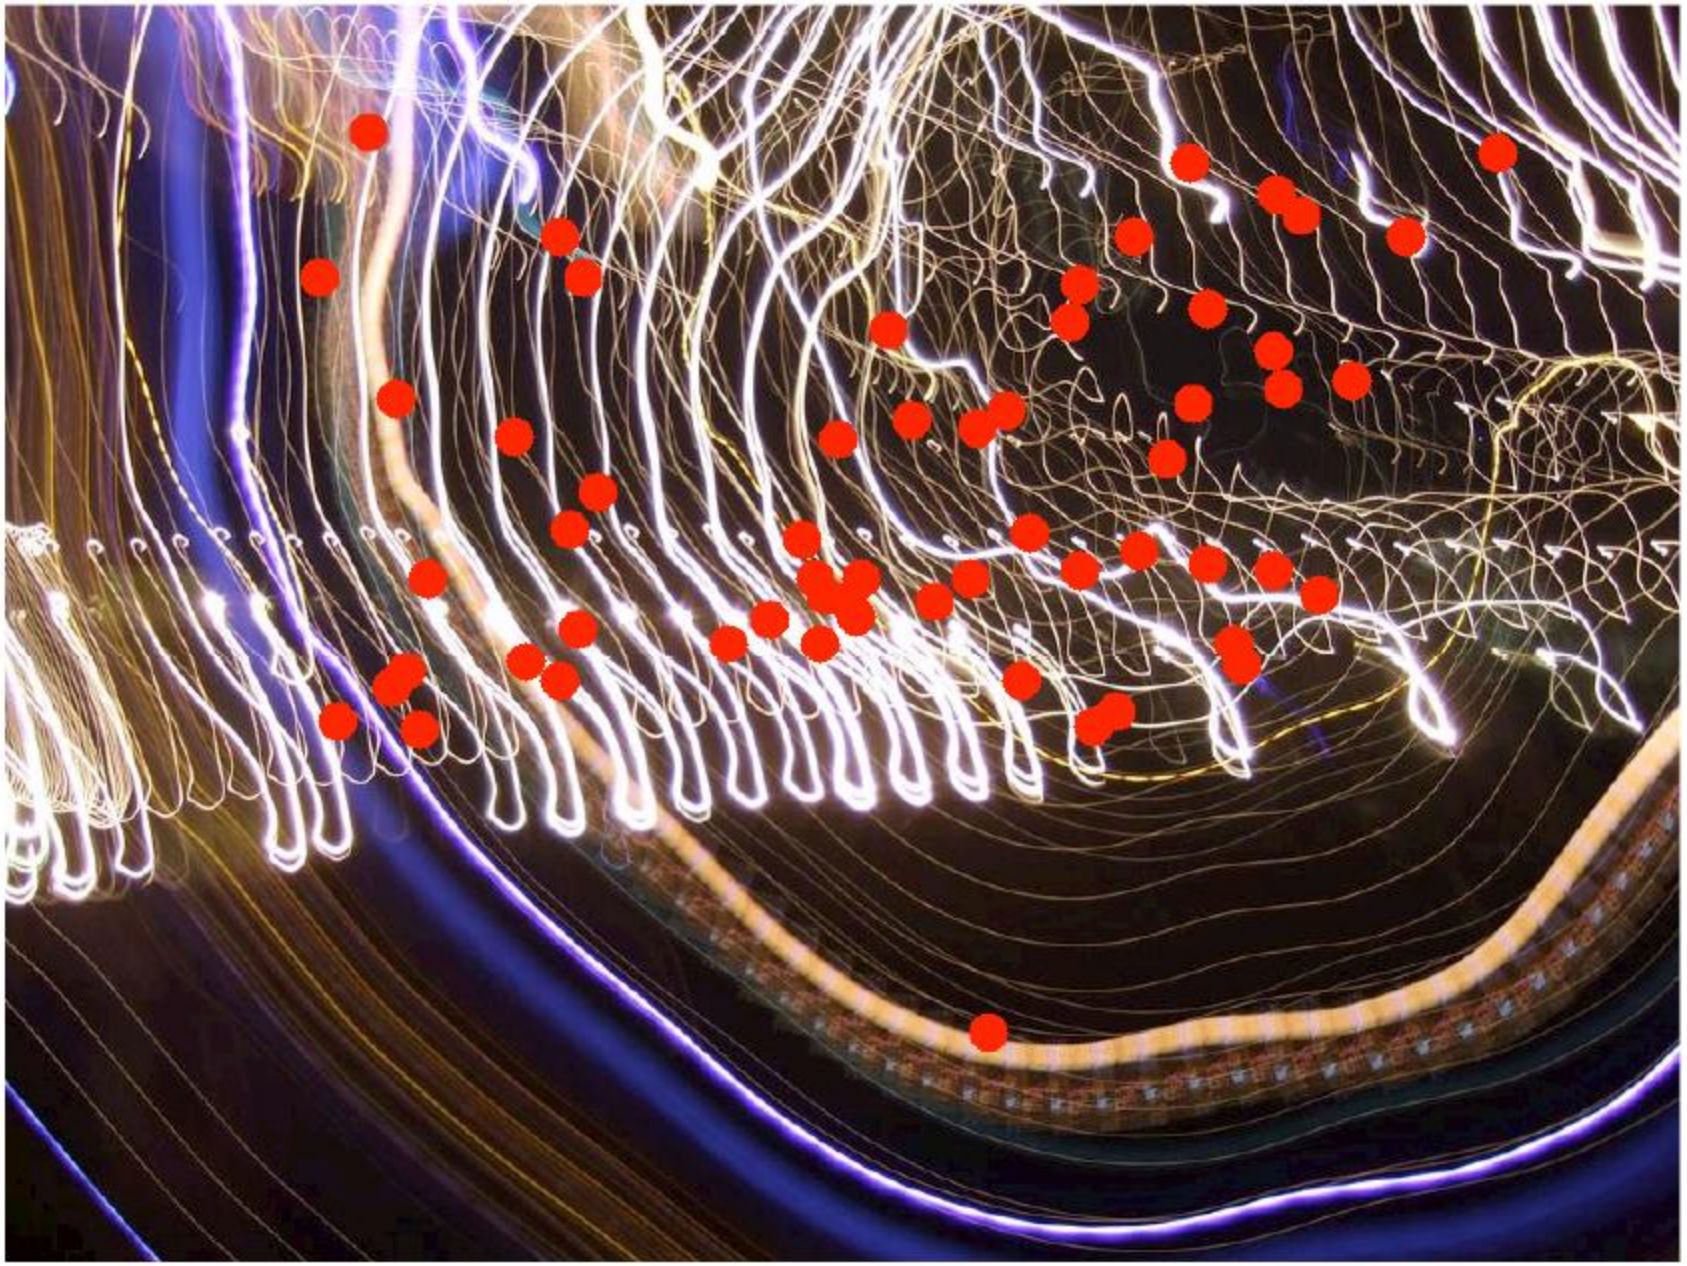

Slide credit: Tilke Judd

How consistent are humans?

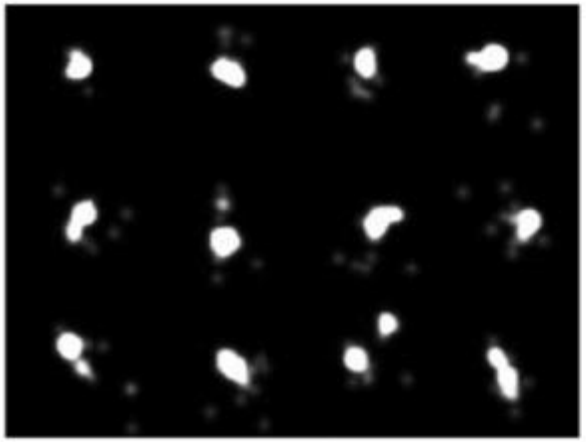

Low entropy saliency maps

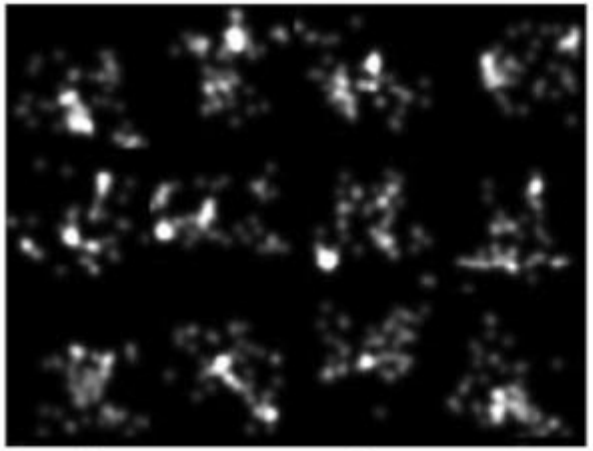

High entropy saliency maps

Slide credit: Tilke Judd

Bias towards the center

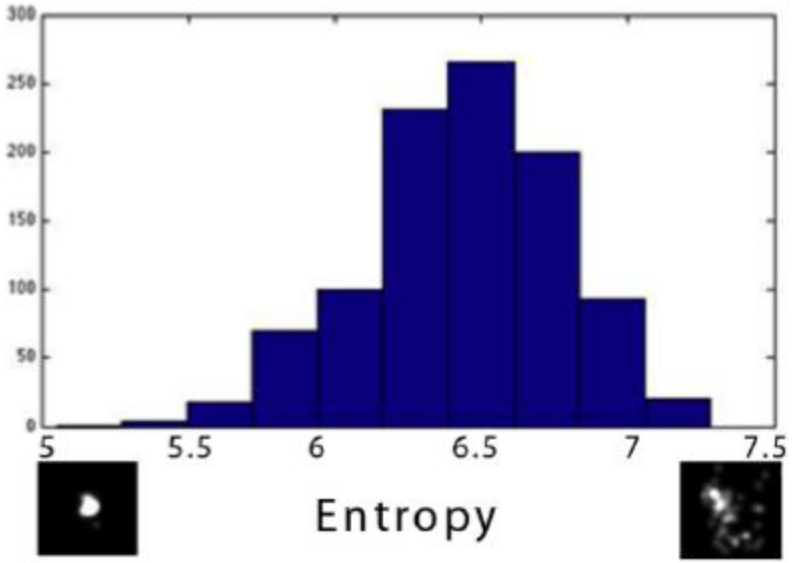

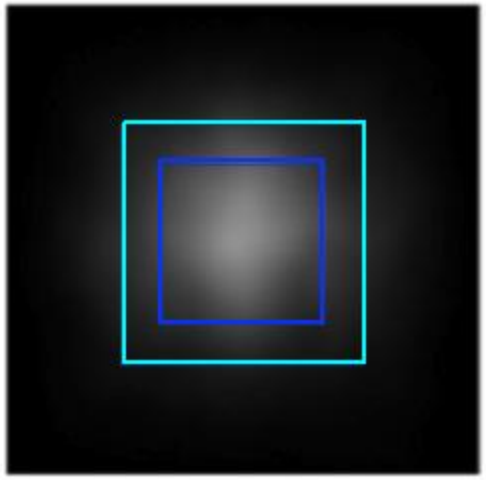

Avg of all saliency maps

40% of fixations within the center 11% of image
70% of fixations within the center 25% of image

Slide credit: Tilke Judd

Human performance

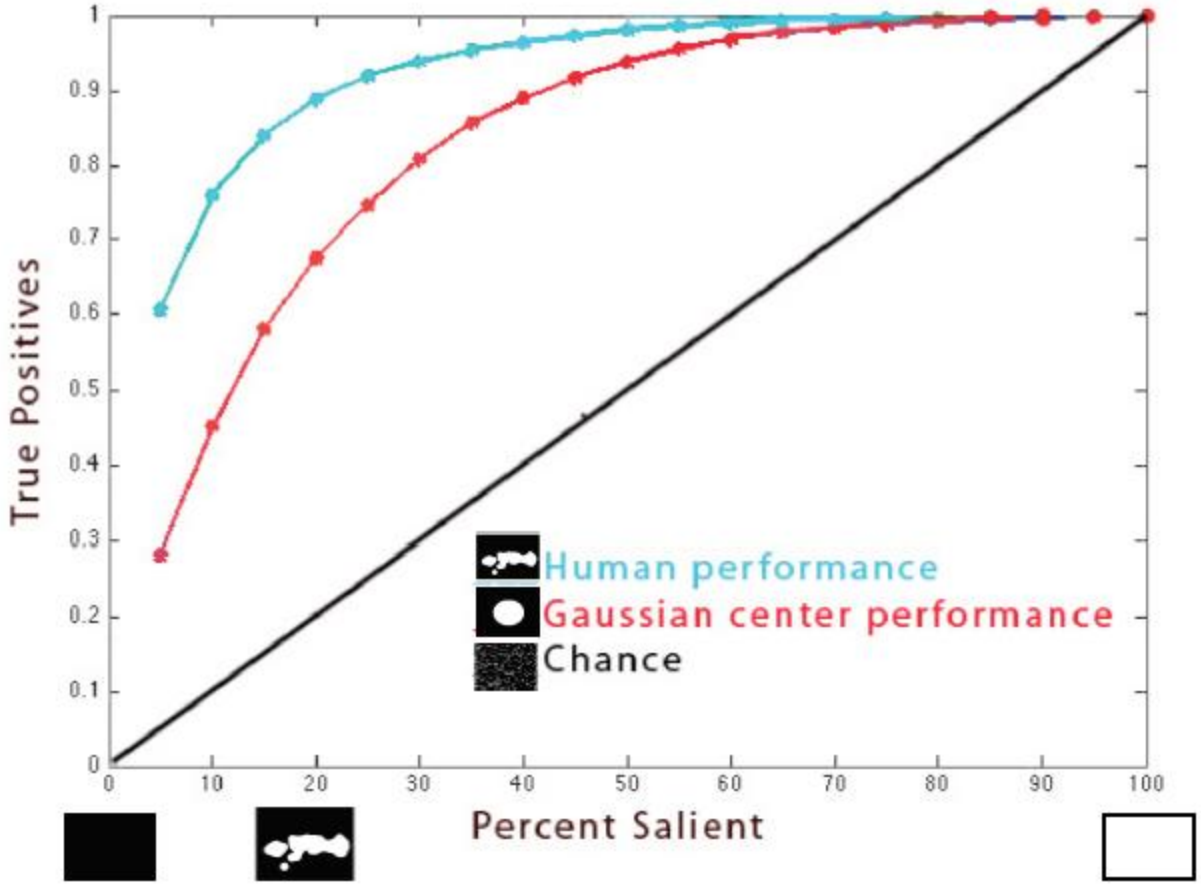

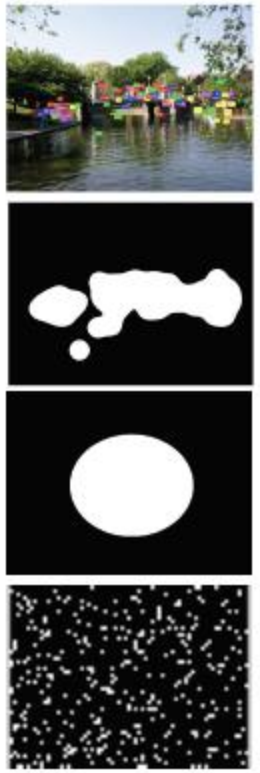

Slide credit: Tilke Judd

Where people look

Faces

People

Text

Body parts

Cars

Animals

Slide credit: Tilke Judd

Size of region of interest

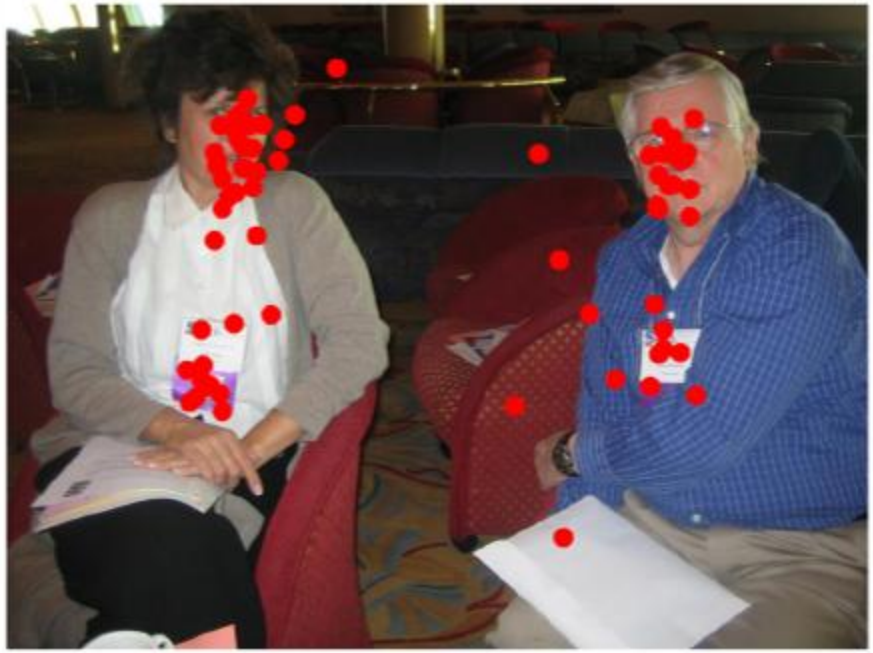

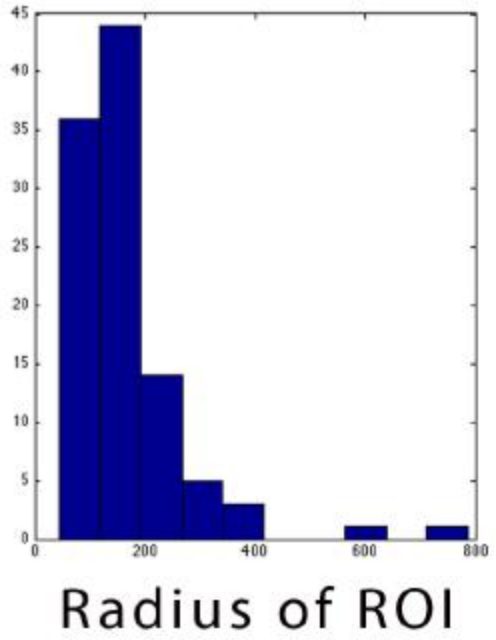

Slide credit: Tilke Judd

Features

- Low level
  illuminance, orientation, color
- Mid level?
  vanishing point, horizon line
- High level
  face detection, object detection

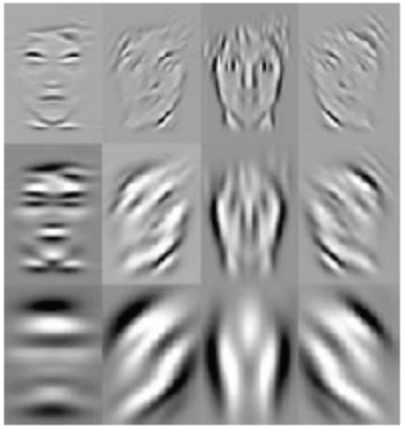

Image filtered with Difference-of-Gaussian(DoG) filters

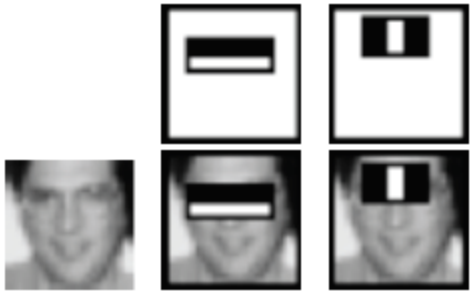

Viola Jones Face detector

Slide credit: Tilke Judd

Features

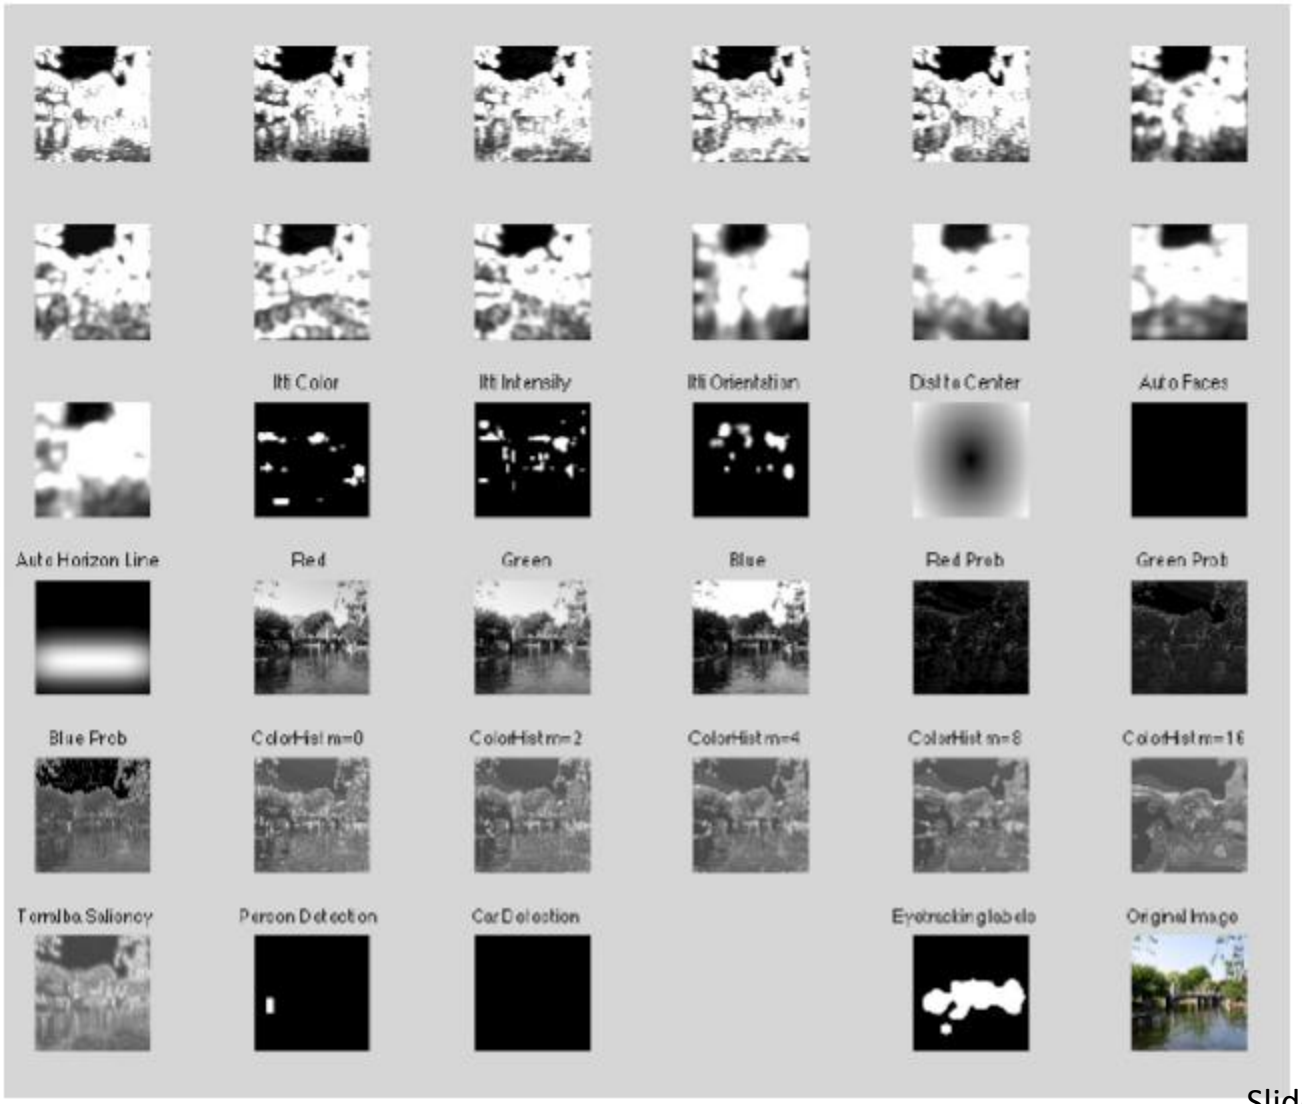

Slide credit: Tilke Judd

Learning

- collect positive salient and negative non salient examples (10 pos : 10 neg)
- Linear support vector machine
- Test on single features and all features

Slide credit: Tilke Judd

Saliency maps from different models

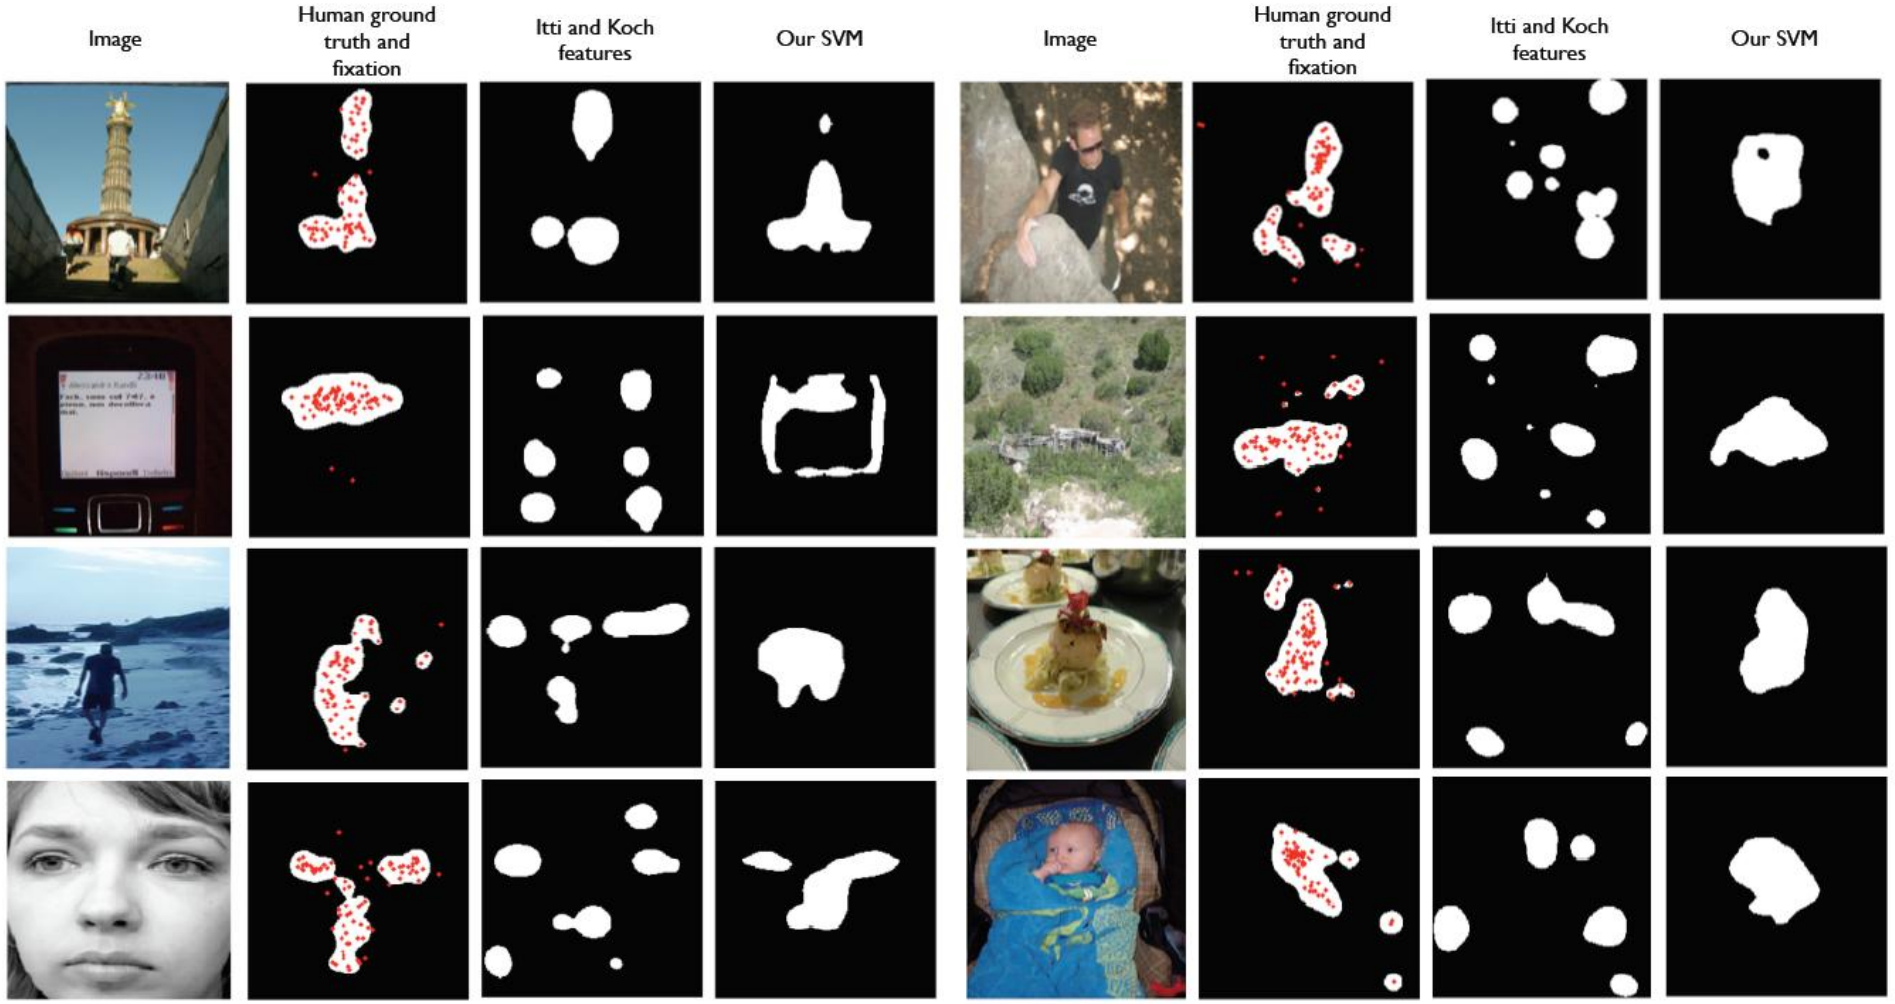

Slide credit: Tilke Judd

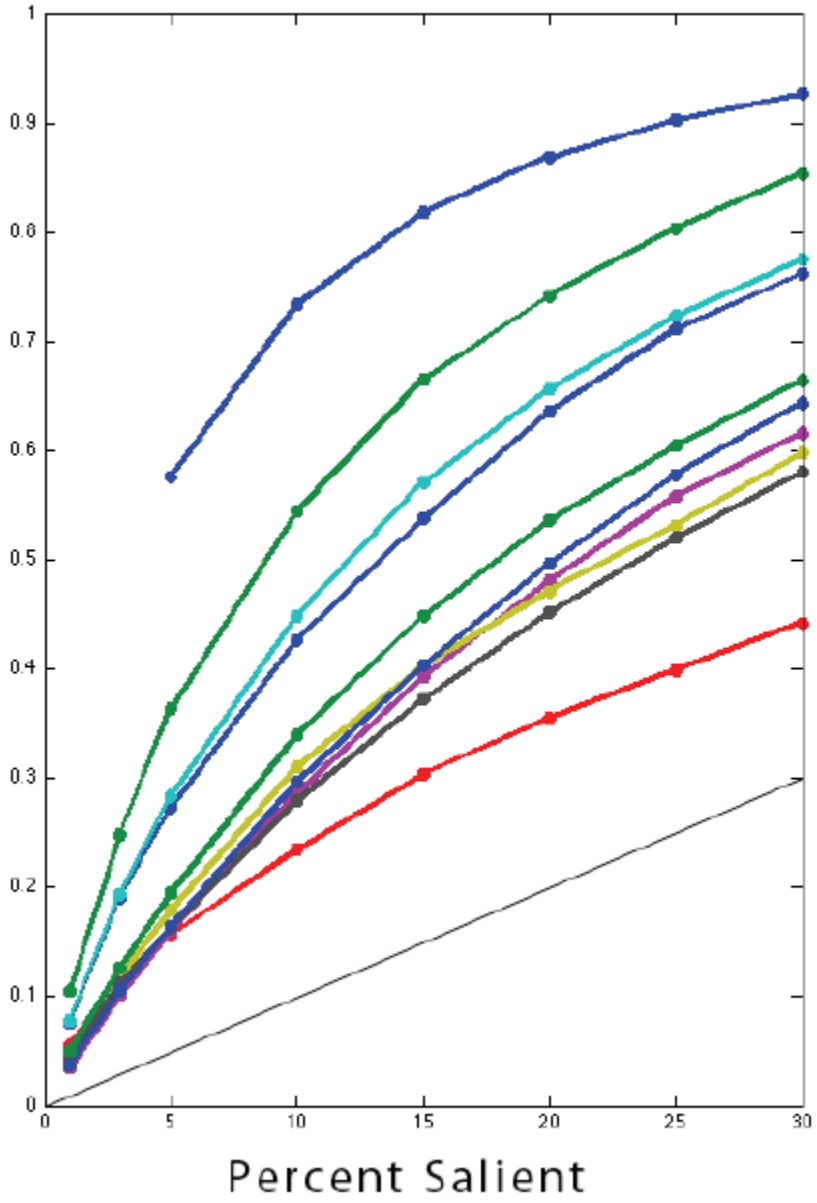

Humans

All Features (our model)

Center
All features without center

Torralba/Rosenholtz
Subbands
Colors
Horizon detector
Itti and Koch

Object detectors
(face, person, car)

Chance

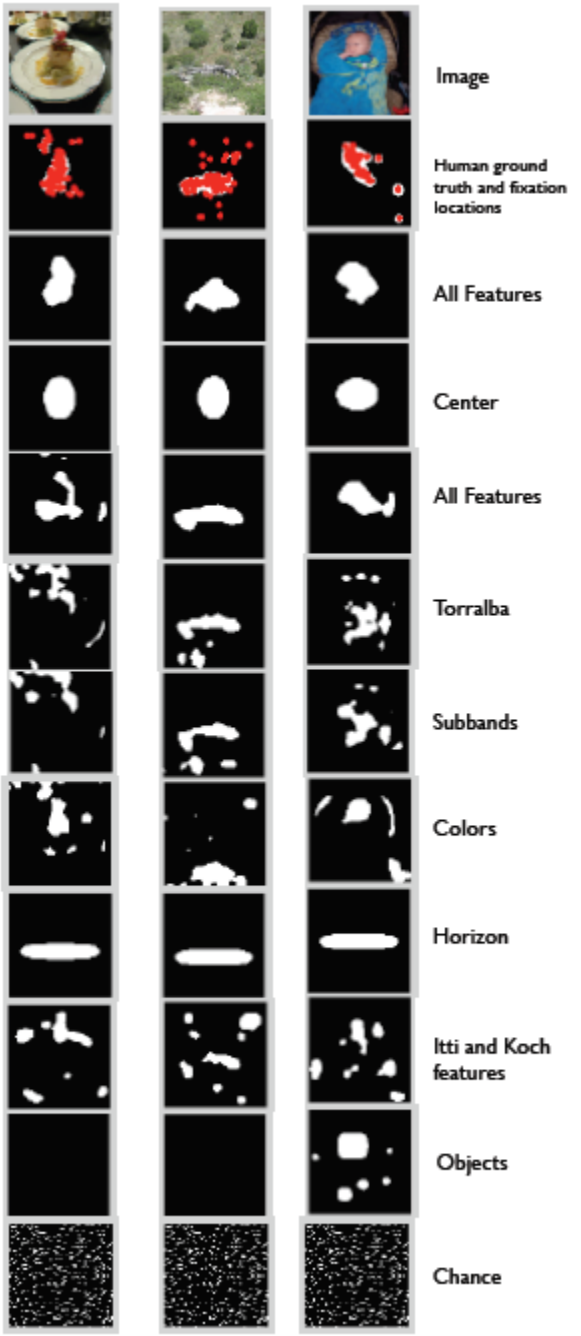

Slide credit: Tilke Judd

Application: Stylization

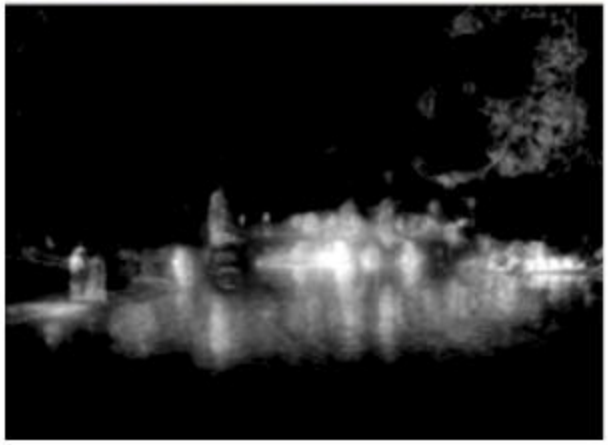

Slide credit: Tilke Judd

Related: what makes something anomalous?

- Boiman and Irani, **Detecting Irregularities in Images and in Video**, ICCV 2005
- http://www.wisdom.weizmann.ac.il/~vision/Irregularities.html

- "Inference by synthesis"