Art, Vision, Probability

Aaron Hertzmann
University of Toronto

What is this?

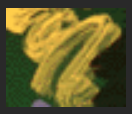

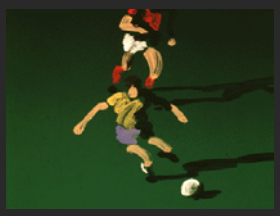

Why does the paint stroke look so much like a leg?
Why does it look good?

How can an photograph that's so abstract be so expressive?

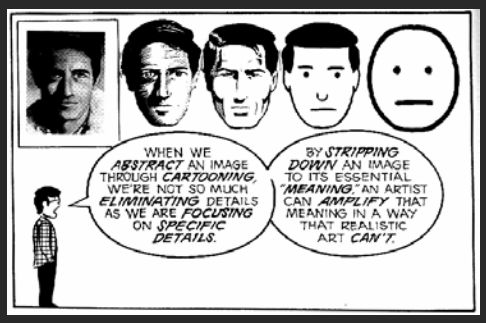

S. McCloud, *Understanding Comics*

The abstraction paradox

The HVS is designed to interpret the world
We give it something totally unlike the world
Yet we perceive it as *more* expressive … why?

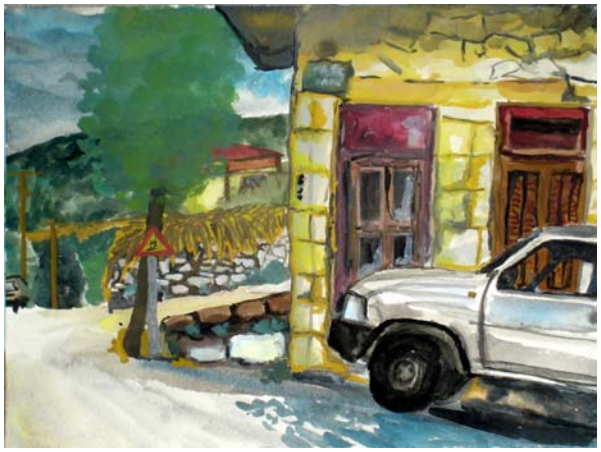

Questions

1. Why does abstraction work?
2. What makes an artwork successful?
3. How can we create a computational model for art and photography?

1. Vision and probability

"The vision problem"

Given light measurements, infer properties of the world

Problems:

- huge set of unknowns
- limited measurements
- indirect link to unknowns

Reconstruction from images

Input video

with: L. Torresani

Non-rigid structure-from-motion

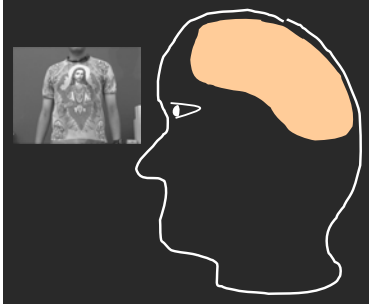

Non-rigid structure-from-motion

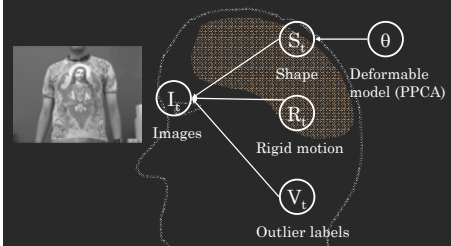

Problem: Given images, solve for everything else … underconstrained: determine both θ and S_t
humans can do it from dots alone (Johansson 1973)

Non-rigid structure-from-motion

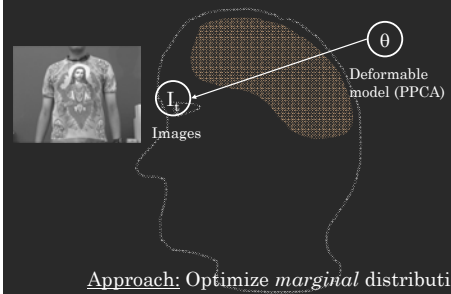

Approach: Optimize *marginal* distribution of θ

Reconstruction from images

Reference frame

Raw video

Lucas-Kanade

Robust algorithm

3D reconstruction

Moral for computer vision

Compute marginals to factor out uncertainty

Recognition and reconstruction must be done jointly!

c.f. W.T. Freeman, A. Yuille, S.C. Zhu, …

2. Human visual processing

Visual encoding

Goal: scene representation to guide actions

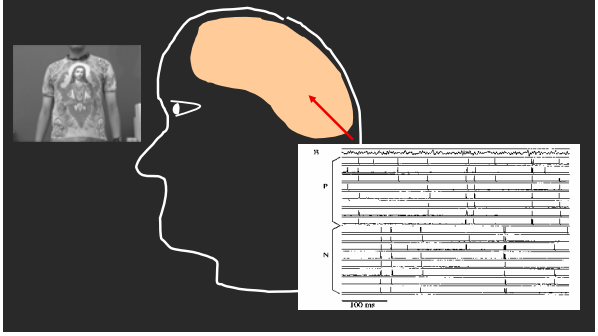

Inference must be *robust*

Robustness to uncertainty

Robustness to limited cues

Neural processing is (probably) probabilistic

See, for example:

- Rao and Olshausen (eds), *Probabilistic Models of the Brain*, 2002
- Hinton, "What kind of graphical model is the brain?" IJCAI 05

Visual processing as Bayesian inference

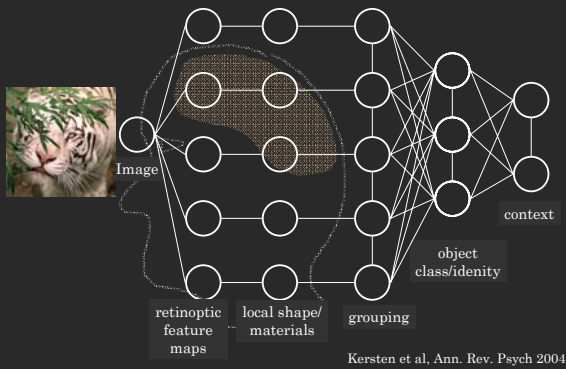

Kersten et al, Ann. Rev. Psych 2004.

Visual processing as Bayesian inference

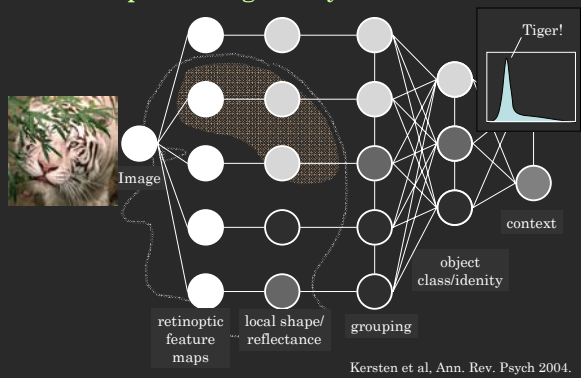

Kersten et al, Ann. Rev. Psych 2004.

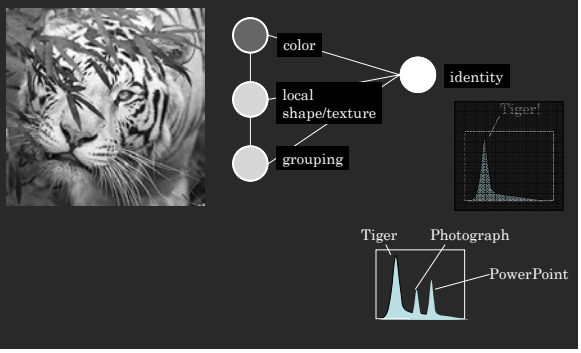

3. Art, vision, probability

Visual ambiguity "illusions"

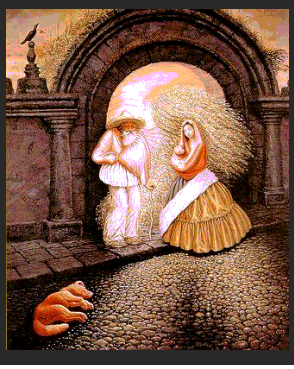

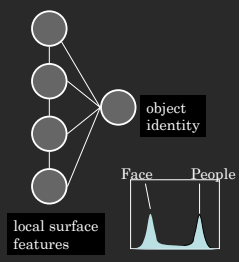

Representation with repeated primitives

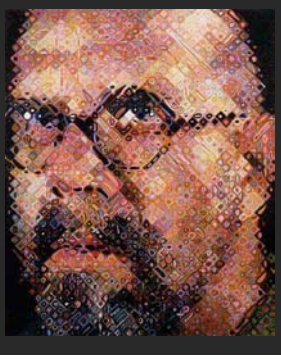

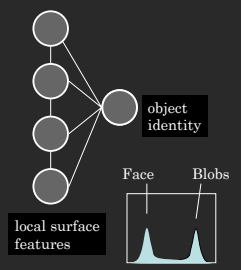

Creating a visual language

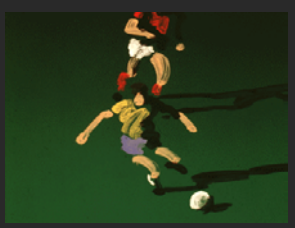

Each interpretation has a compact *code*

Efficient models have high probability in model selection [MacKay]

Composition

Why does line drawing work?

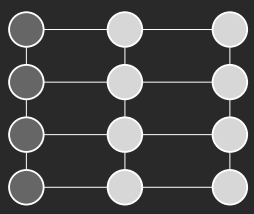

shading/ reflectance

contours/ silhouettes/ edges

grouping/ identity

Segmentation and grouping

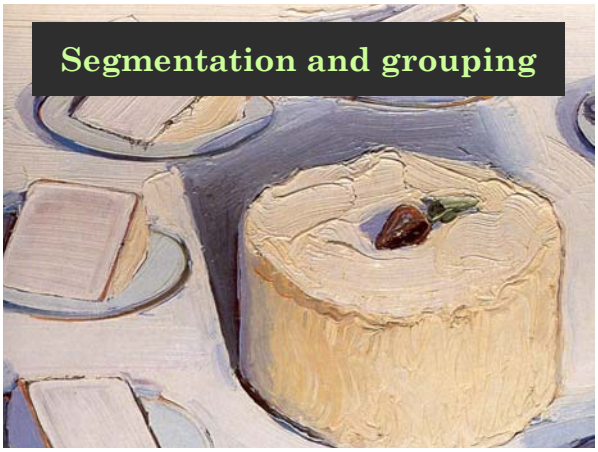

Summary

- The HVS infers scene properties by probabilistic inference
- Art works convey a scene by providing sufficient cues to the visual system
- Ambiguity, economy, and visual languages all correspond to specific probabilistic interpretations
- General-purpose computer vision may require a similar approach

Open questions

1. Can we optimize an image (painting or photo) for a desired effect on the HVS?
2. What kinds of internal representations are useful for modeling art and HVS?
3. What does this have to do with goals of art?
4. What kinds of data do we need from our capture devices and vision algorithms?