



BROWN

Hierarchy and Reusability in Image Analysis

Stuart Geman

Eran Borenstein, Ya Jin, Wei Zhang

- I. Remarks on Computer Vision
- II. Approaches
- III. Bayesian Image Analysis
- IV. Probability Models
- V. Demonstration System: Reading License Plates
- VI. Generalization: Face Detection

I. Remarks on Computer Vision

- Vision is hard
- Why is vision hard?

II. Approaches

III. Bayesian Image Analysis

IV. Probability Models

V. Demonstration System: Reading License Plates


VI. Generalization: Face Detection

License plate images from Logan Airport




Machines *still* can't reliably read license plates

Wafer ID's



Machines can't read fixed-font fixed-scale characters as well as humans

Super Bowl



Machines can't find the bad guys at the Super Bowl

Instantiation



same




Empire style table



twins

Vision is *content sensitive*

“Clutter”



Human Interactive Proofs



Background is *structured*, and made of the *same stuff*!

I. Remarks on Computer Vision

II. Approaches

- Pure learning
- Fodor & Pylyshyn, 1988, and the critique of neural networks
- Observations from the cognitive, neural, and mathematical sciences

III. Bayesian Image Analysis

IV. Probability Models

V. Demonstration System: Reading License Plates

VI. Generalization: Face Detection

Pure learning



Google:

- 2 billion images
- train classifier: pornographic/not pornographic
- good classification
- not nearly as good as human performance

Pure learning

Human learning:

- 1 sample/10 seconds
- 16 hours/day
- 80 years
- < 170 million samples/lifetime

Enough examples?

Did *evolution* have enough examples?

D. Geman: “The interesting limit is N goes to zero, not N goes to infinity”

Fodor & Pylyshyn, 1988, and the critique of neural networks

Properties of human cognition:

* compositionality

roughly: representation through syntactically constrained hierarchy of reusable parts

* productivity


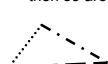

roughly: capable of an infinite number of well-formed actions, thoughts, sentences ...

* systematicity

roughly: invariance

Fodor & Pylyshyn, 1988, and the critique of neural networks

systematicity

If  is a triangle,
then so are  and 

if 'the boy ran home' makes sense, then so does
'the girl ran home'
'john ran home'
'john ran to school'

4 4 4 ~~4~~ 4 4 4
4

Observations from the cognitive, neural, and mathematical sciences

Human brains utilize *strong representations*

- Damassio (simulation=perception)
- Kosslyn ("resolution of the imaging debate")
- Lakeoff, Fauconnier (the role of mental imagery in language understanding)

Observations from the cognitive, neural, and mathematical sciences

Consider the ventral visual pathway:


Structure:	retina ↔ LGN ↔ V1 ↔ V2 ↔ V4 ↔ IT
Topography: topology	highly retinotopic ↔ little
Receptive Field:	small ↔ large
Specificity: high	low ↔ high
Invariance:	low ↔ high

SUMMARY: A hierarchy of less-to-more invariant representations

Observations from the cognitive, neural, and mathematical sciences


Thought experiment:

what if the world is a hierarchy of reusable parts?




Problem: test for 'L'


Observations from the cognitive, neural, and mathematical sciences



God's ROC curve (Neyman-Pearson Lemma):

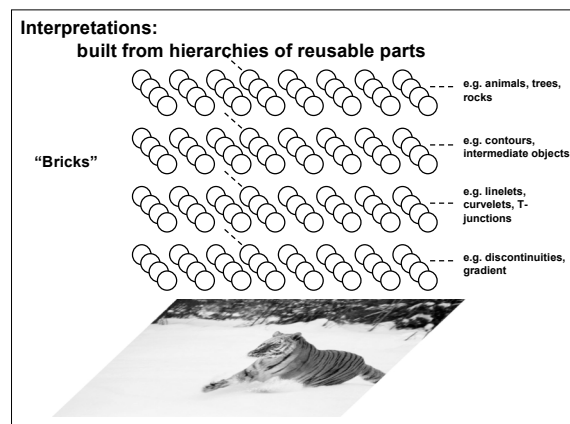
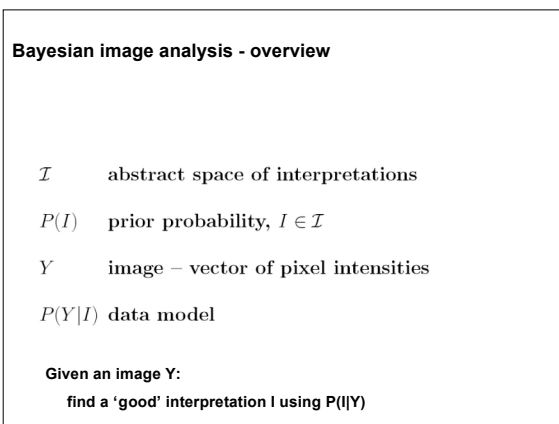
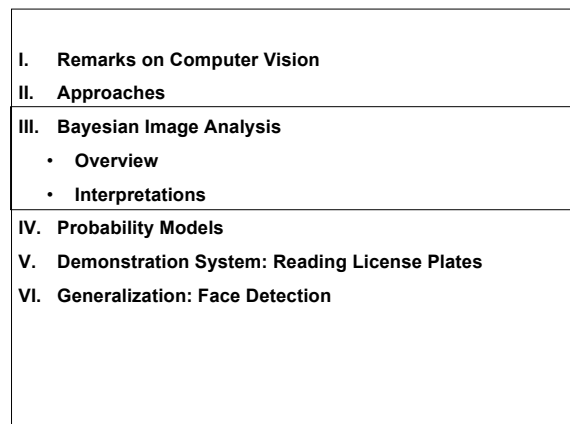
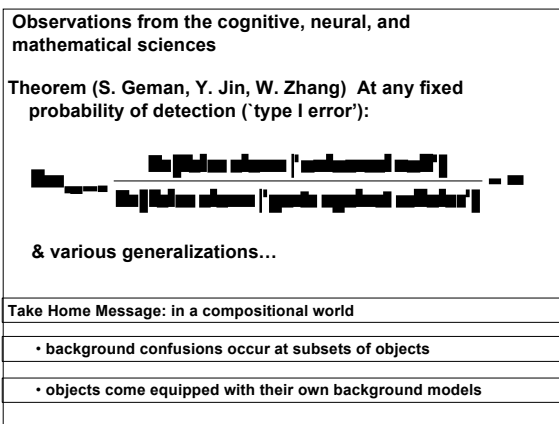

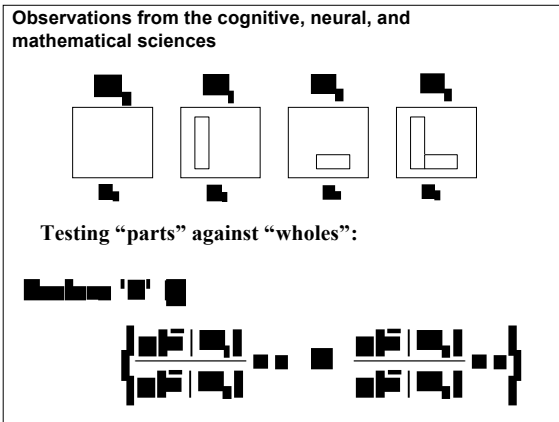



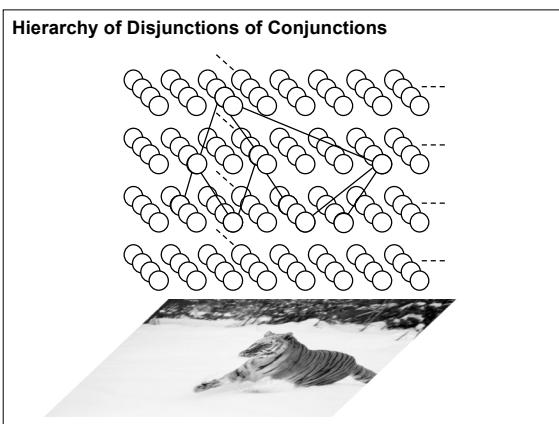
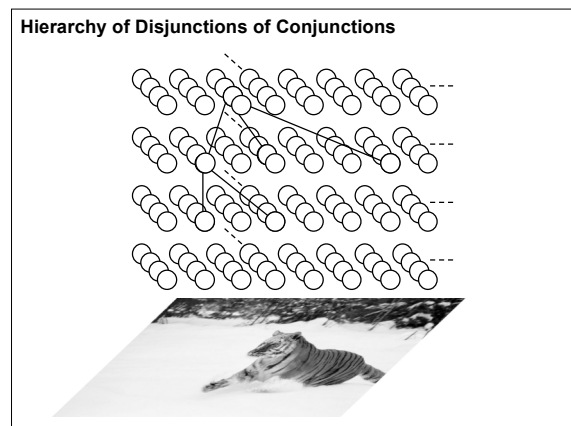
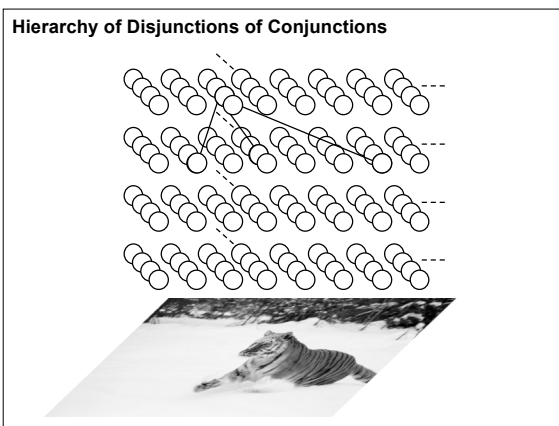
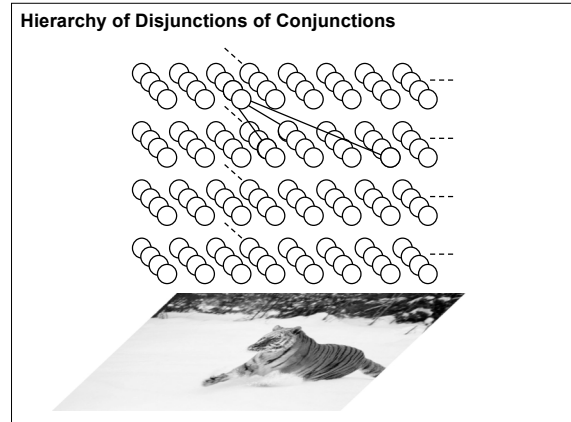
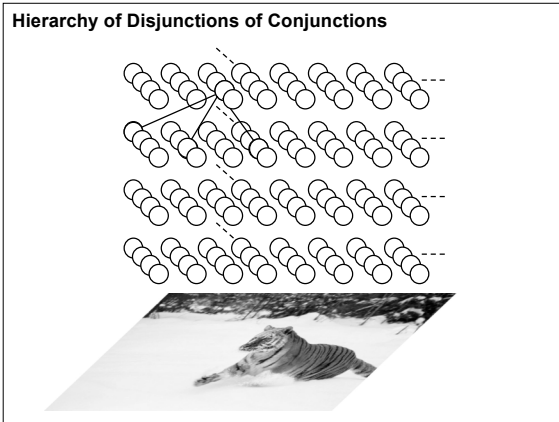
Observations from the cognitive, neural, and mathematical sciences


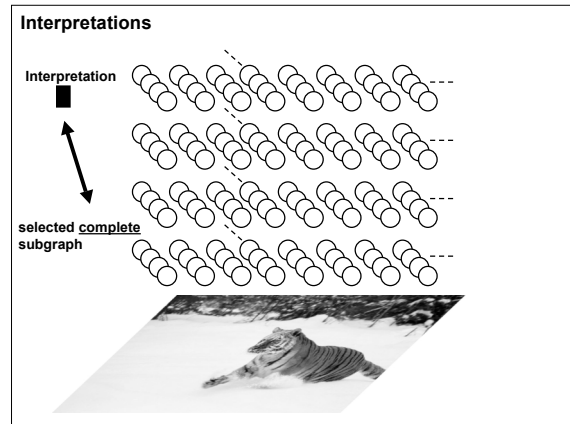



Testing against a universal null (pragmatic):









I. Remarks on Computer Vision

II. Approaches


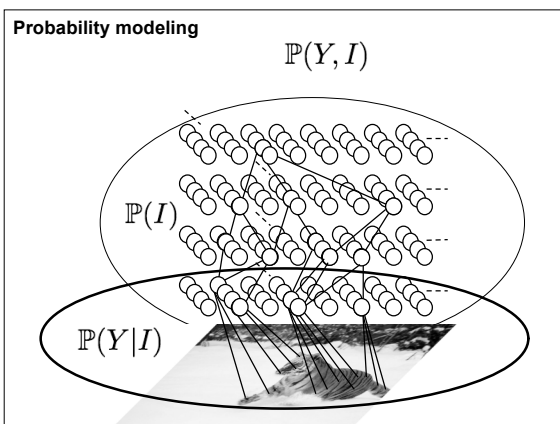
III. Bayesian Image Analysis


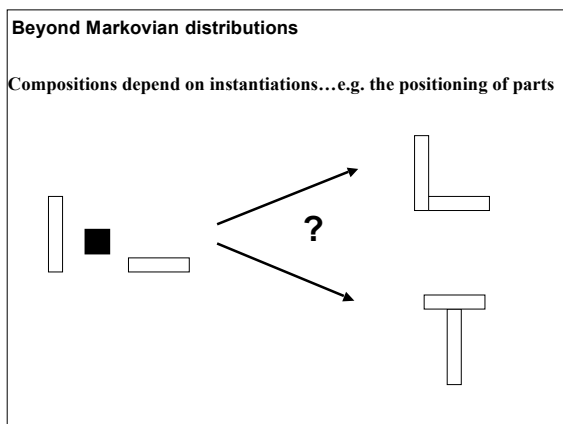
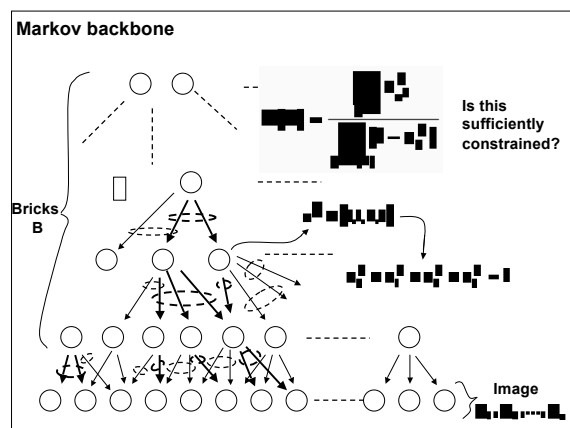
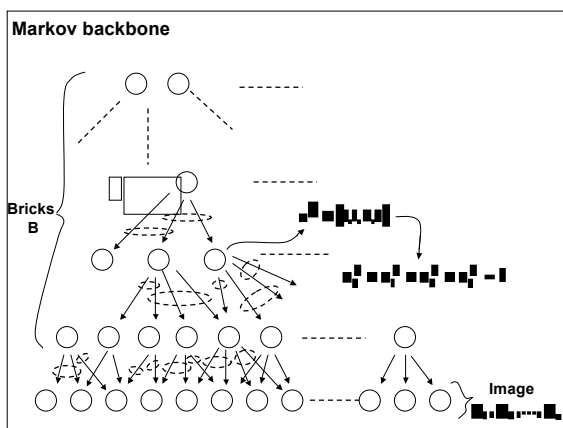
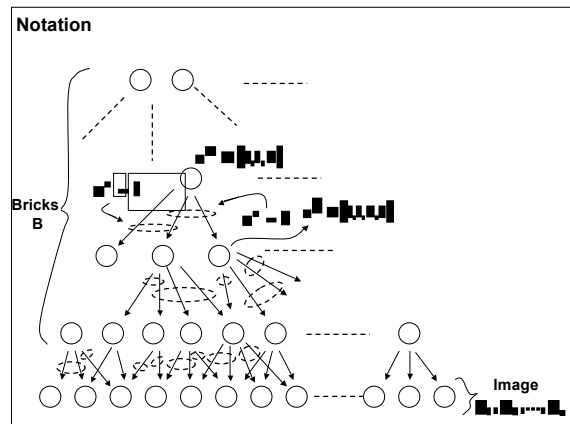
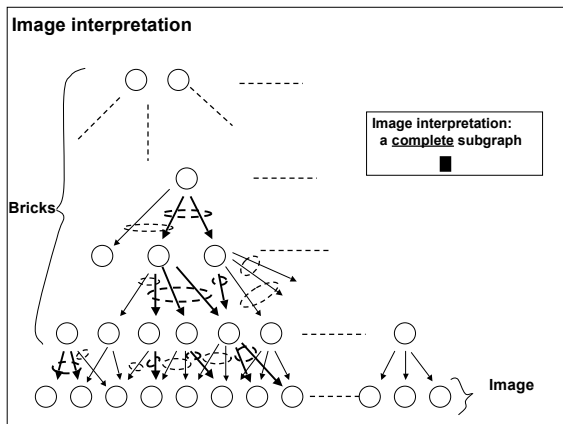
IV. Probability Models

- $P(I)$
- $P(Y|I)$

V. Demonstration System: Reading License Plates

VI. Generalization: Face Detection





Content-Sensitive Perturbation

I $\{x^\beta\}_{\beta \in B}$, state of each brick

β arbitrary, fixed, brick

$a^\beta(I)$ arbitrary function of the progeny of β

$p^c(a^\beta)$ target cond. ($x^\beta > 0$) prob. on a^β

$p^o(a^\beta)$ unperturbed cond. ($x^\beta > 0$) prob. on a^β

Content-Sensitive Perturbation

$$\begin{aligned} p(I) &= p(I|x^\beta=0)p(x^\beta=0) + p(I|x^\beta>0)p(x^\beta>0) \\ &= p(I|x^\beta=0)p(x^\beta=0) + p(I, a^\beta|x^\beta>0)p(x^\beta>0) \\ &= p(I|x^\beta=0)p(x^\beta=0) + p(I|a^\beta, x^\beta>0)p^o(a^\beta)p(x^\beta>0) \end{aligned}$$

$$\begin{aligned} &\rightarrow p(I|x^\beta=0)p(x^\beta=0) + p(I|a^\beta, x^\beta>0)p^c(a^\beta)p(x^\beta>0) \\ &= p(I|x^\beta=0)p(x^\beta=0) + p(I|a^\beta, x^\beta>0)p^o(a^\beta)p(x^\beta>0) \frac{p^c(a^\beta)}{p^o(a^\beta)} \\ &= p(I) \left(\frac{p^c(a^\beta)}{p^o(a^\beta)} \right)^{1_{x^\beta>0}} \end{aligned}$$

Content-Sensitive Perturbation

.... but perturbations interact!

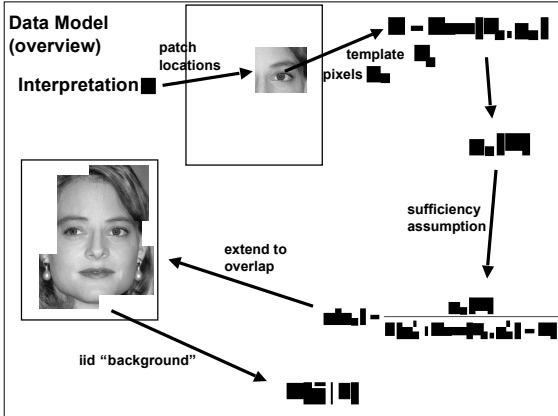

.... nevertheless ...

THEOREM (Wei Zhang):


(1) if $p^c(a^\beta(I)) > 0 \forall \beta \in B \ \& \ I \in \mathcal{I}$,
then $\exists p(I)$ consistent with $\{p^c(a^\beta(I))\}_{\beta \in B}$

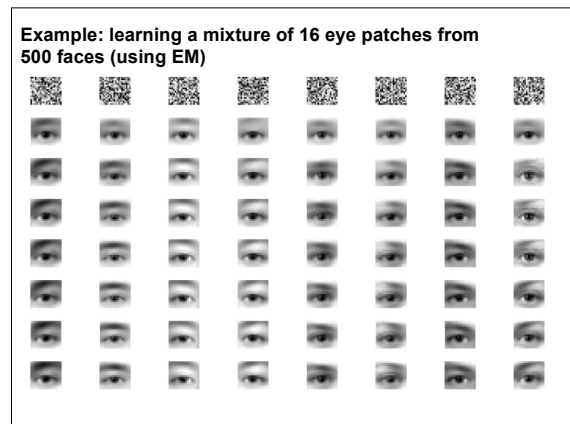



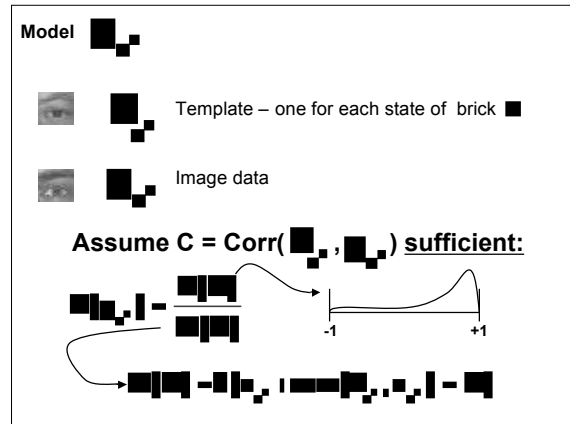

(2) iterative perturbation, over all $\beta \in B$, generates such a distribution.

Probability modeling




Select a terminal brick







Example: learning a mixture of 16 eye patches from 500 faces (using EM)



Example: learning a mixture of 16 nose and mouth patches from 500 faces (using EM)



Example: learning a mixture of 16 nose and mouth patches from 500 faces (using EM)



- I. Remarks on Computer Vision
- II. Approaches
- III. Bayesian Image Analysis
- IV. Probability Models
- V. Demonstration System: Reading License Plates
- VI. Generalization: Face Detection

Test set: 385 images, mostly from Logan Airport



Courtesy of Visics Corporation

Architecture




Image interpretation






Image interpretation




Performance

- 385 images
- Six plates read with mistakes (>98%)
- Approx. 99.5% characters read correctly
- Zero false positives

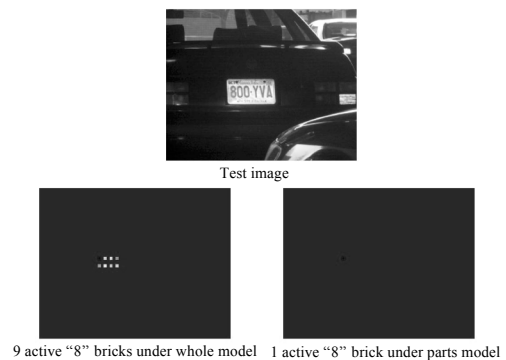
Efficient discrimination: Markov versus Content-Sensitive dist.



Efficient discrimination: Markov versus Content-Sensitive dist.



Efficient discrimination: testing objects against their parts



Summary

Vision is Content Sensitive


Non-Markovian probability models

Background is Structured, and Made of the Same Stuff


Objects come equipped with their own background models

- I. Remarks on Computer Vision
- II. Approaches
- III. Bayesian Image Analysis
- IV. Probability Models
- V. Demonstration System: Reading License Plates
- VI. Generalization: Face Detection

Face Hierarchy



$\mathbb{P}(I)$



Sampling faces from the distribution

$\mathbb{P}(Y, I)$




PATTERN SYNTHESIS


= PATTERN ANALYSIS

Ulf Grenander

Efficient computation: depth-first search




Efficient computation: depth-first search



Number of visits to each pixel. Left: linear scale Right: log scale

Efficient computation: depth-first search



Number of visits to each pixel. Left: linear scale Right: log scale