Making Decisions:
Modeling perceptual decisions

Outline

- Graphical models for belief example completed
- Discrete, continuous and mixed belief distributions
- Vision as an inference problem
- Modeling Perceptual beliefs: Priors and likelihoods
- Modeling Perceptual decisions: Actions, outcomes, and utility spaces for perceptual decisions

Graphical Models: Modeling complex inferences

Nodes store conditional Probability Tables

Arrows represent conditioning

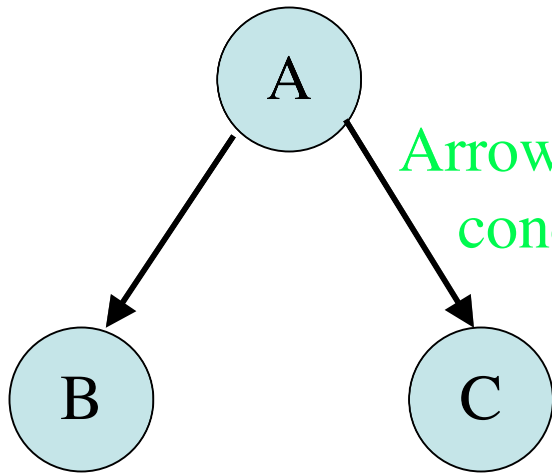

This model represents the decomposition:

$$P(A,B,C) \; = \; P(B|A)\, P(C|A)\, P(A)$$

What would the diagram for $P(B|A,C)\, P(A|C)\, P(\,C)$ look like?

Sprinkler Problem

You see wet grass.
Did it rain?

P(C=F)	P(C=T)
0.5	0.5

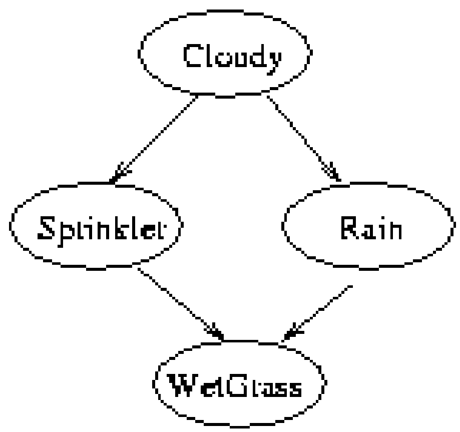

C	P(S=F)	P(S=T)
F	0.5	0.5
T	0.9	0.1

C	P(R=F)	P(R=T)
F	0.8	0.2
T	0.2	0.8

S	R	P(W=F)	P(W=T)
F	F	1.0	0.0
T	F	0.1	0.9
F	T	0.1	0.9
T	T	0.01	0.99

Sprinkler Problem cont'd

4 variables: C = Cloudy, R=rain; S=sprinkler; W=wet grass

What's the probability of rain? Need $P(R|W)$

Given: $P(C) = [0.5\ 0.5]$; $P(S|C)$; $P(R|C)$; $P(W|R,S)$

Do the math:

1) *Get the joint:*

   $P(R,W,S,C) = P(W|R,S)\, P(R|C)\, P(S|C)\, P(C)$

2) *Marginalize:*

   $P(R,W) = \Sigma_S \Sigma_C\, P(R,W,S,C)$

3) *Divide:*

   $P(R|W) = P(R,W)/P(W) = P(R,W)/\ \Sigma_R\, P(R,W)$

Continuous data, discrete beliefs

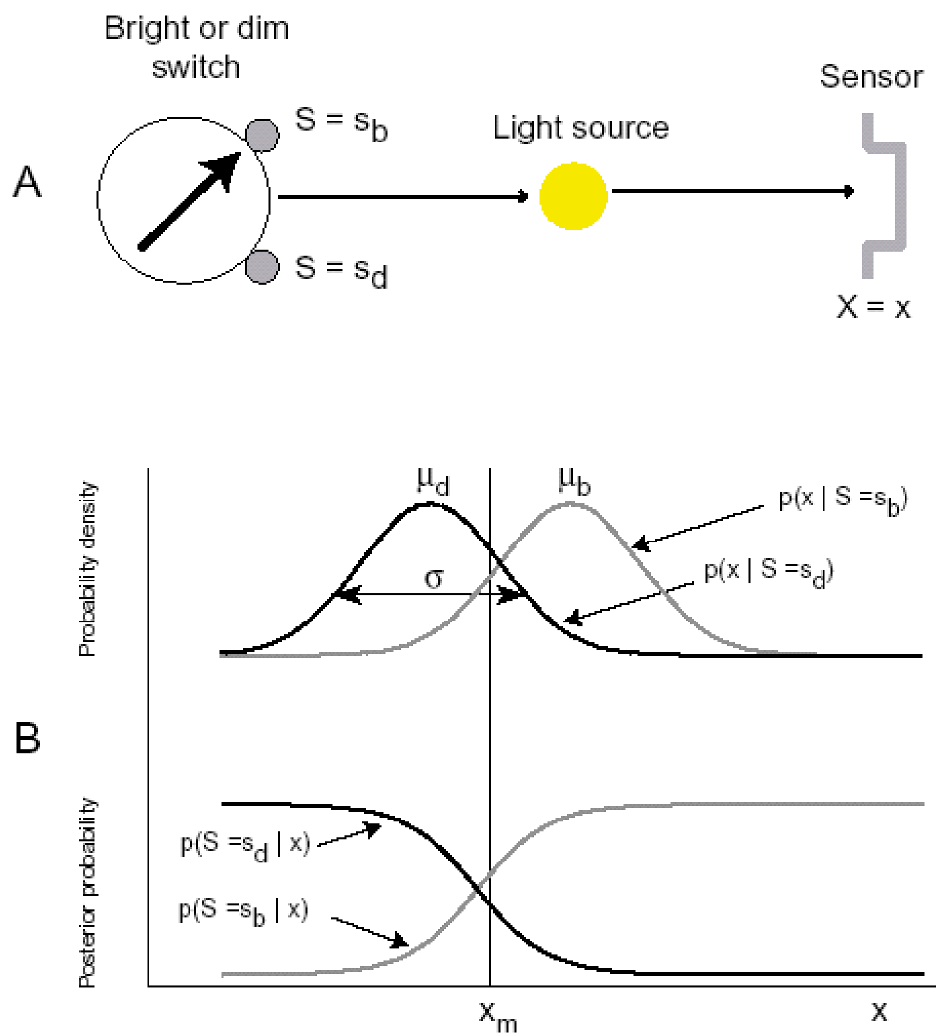

$$P(S = s_b|x) = p(x|S = s_b)(1/3)/p(x), \text{ and } P(S = s_b|x) = p(x|S = s_b)(2/3)/p(x)$$

Continuous data, Multi-class Beliefs

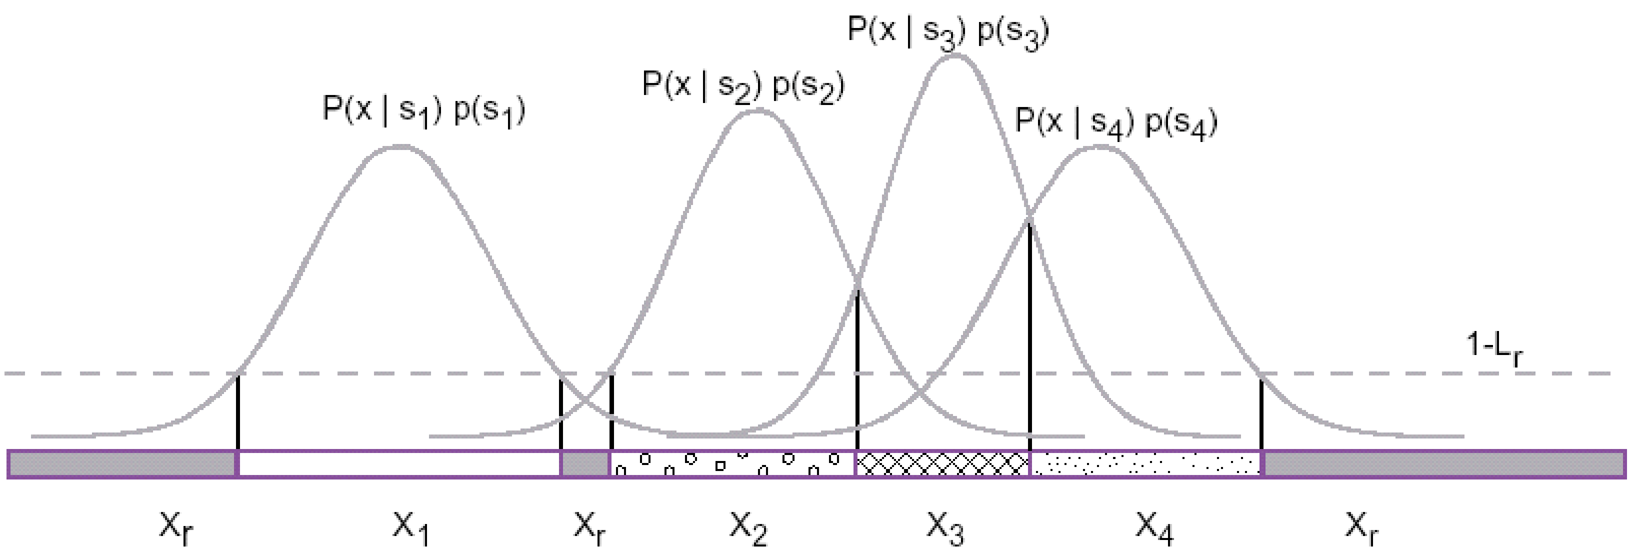

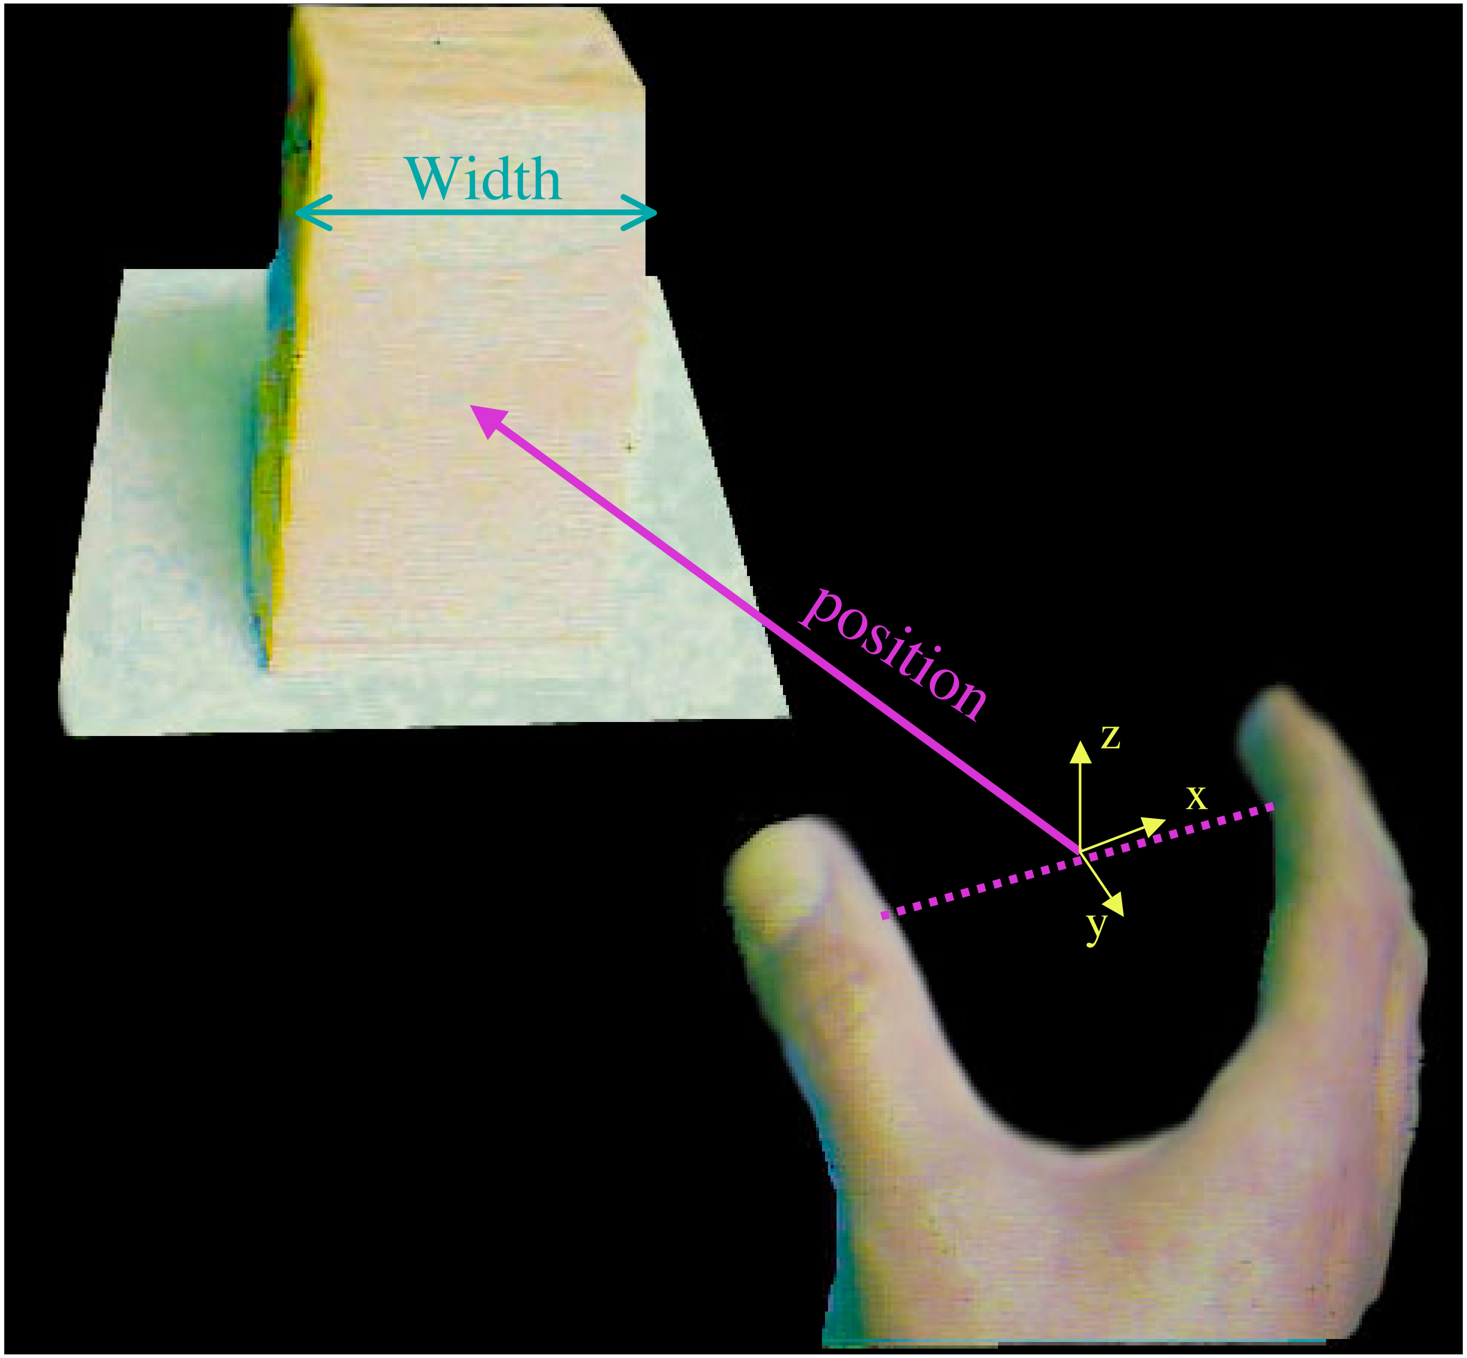
Width
position
z
x
y

The problem of perception

Image Measurement

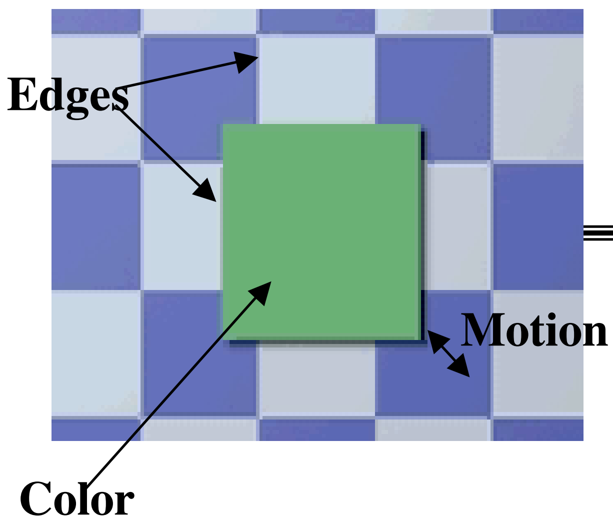

Scene Inference

For Grasping

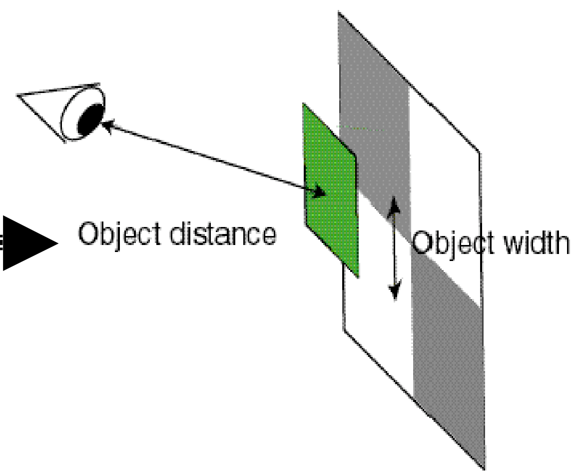

Decisions/Actions

Classify
Identify
Learn/Update
Plan Action
Control Action

Vision as Statistical Inference

Rigid Rotation Or Shape Change?

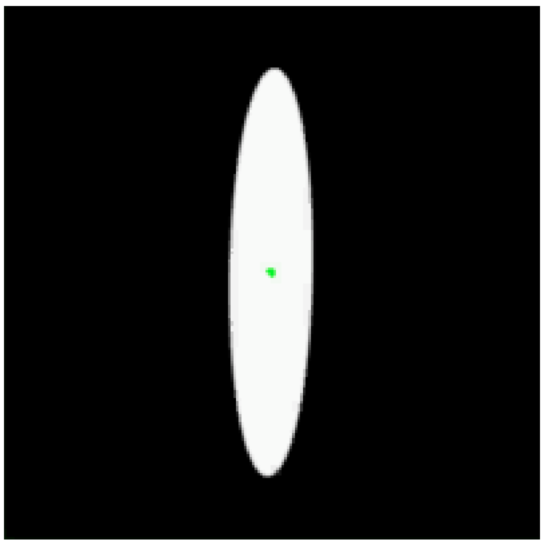

Weiss & Adelson, 2000

What Moves, Square or Light Source?

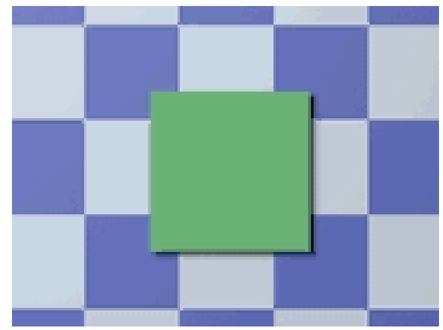

Kersten et al, 1996

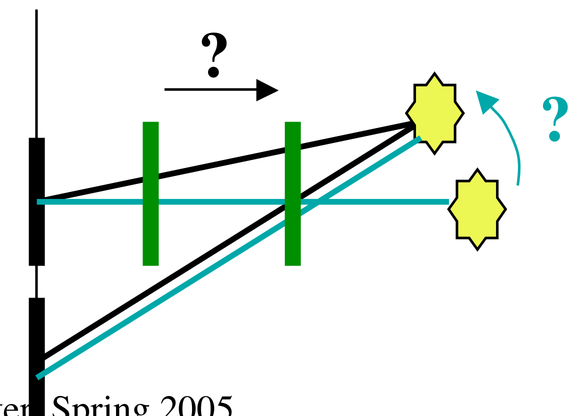

Modeling Perceptual Beliefs

- Observations *O* in the form of image measurements-
  - Cone and rod outputs
  - Luminance edges
  - Texture
  - Shading gradients
  - Etc
- World states *s* include
  - Intrinsic attributes of objects (reflectance, shape, material)
  - Relations between objects (position, orientation, configuration)

Belief equation for perception

$$p(s \mid O) = \overbrace{p(O \mid s)}^{likelihood}\,\overbrace{p(s)}^{prior} / p(O)$$

$$where$$

$$p(O) = \int_{-\infty}^{\infty} p(O \mid s)\, p(s)\, ds$$

Image data

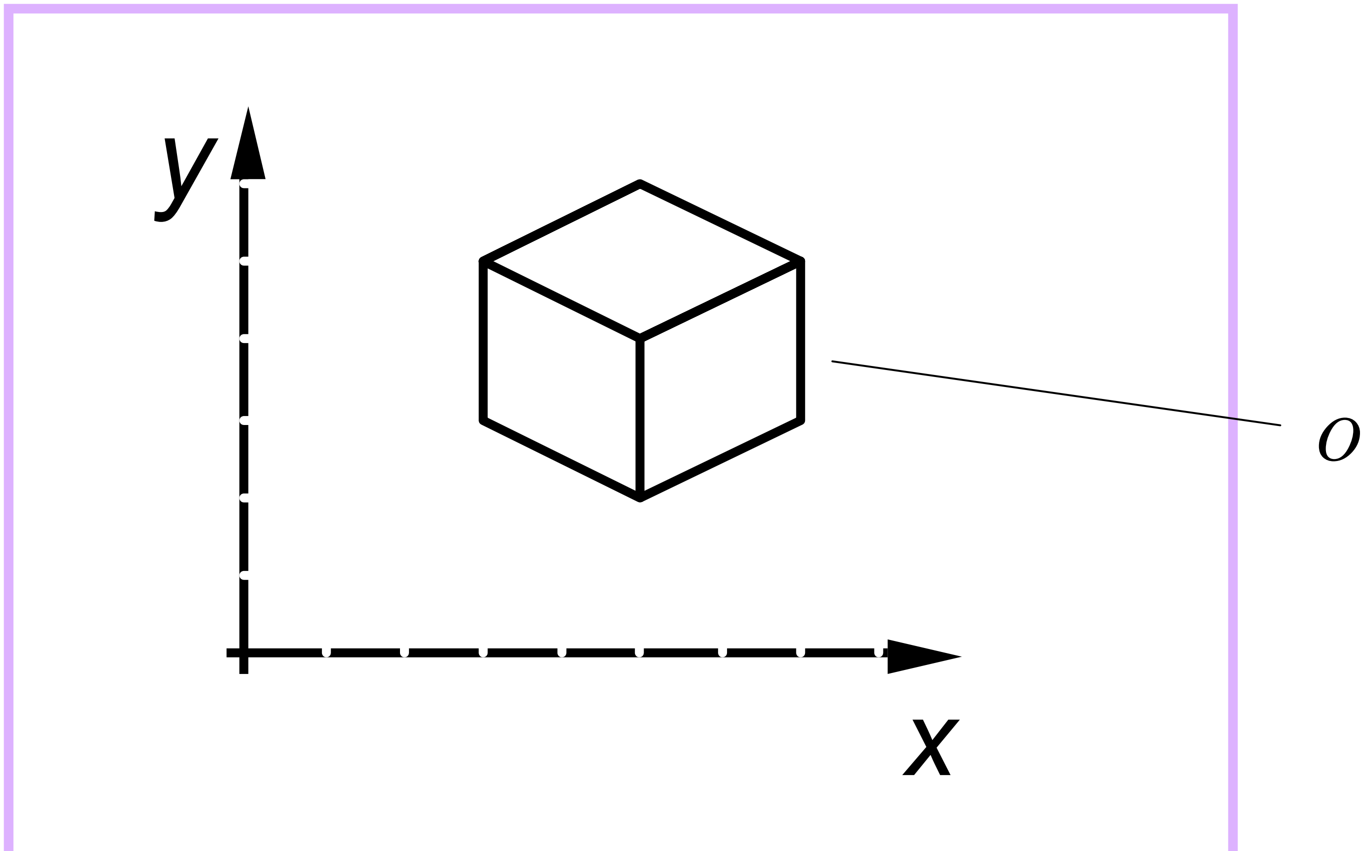

Possible world states consistent with *O*

Many 3D shapes can map to the same 2D image

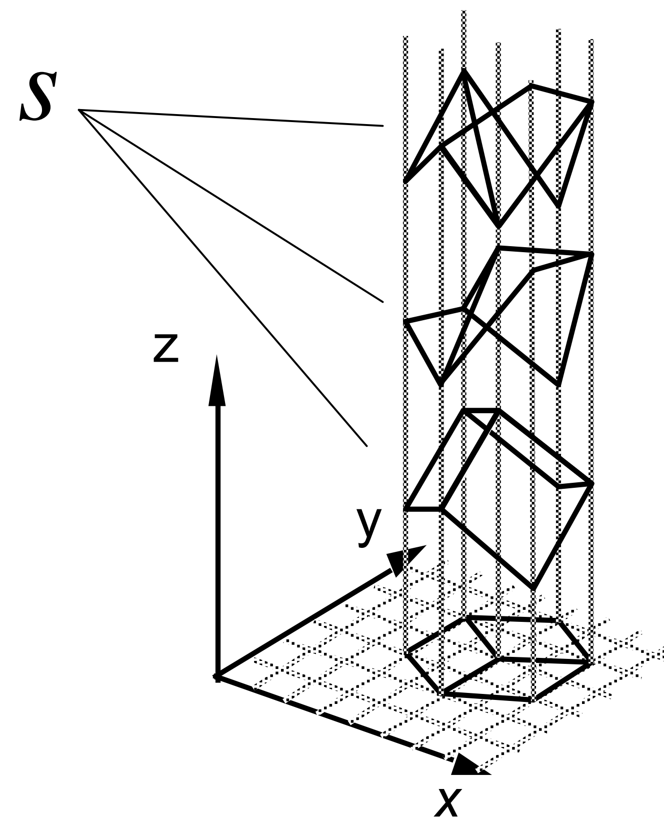

Forward models for perception:

Built in knowledge of image formation

Images are produced by physical processes that can be modeled via a rendering equation:

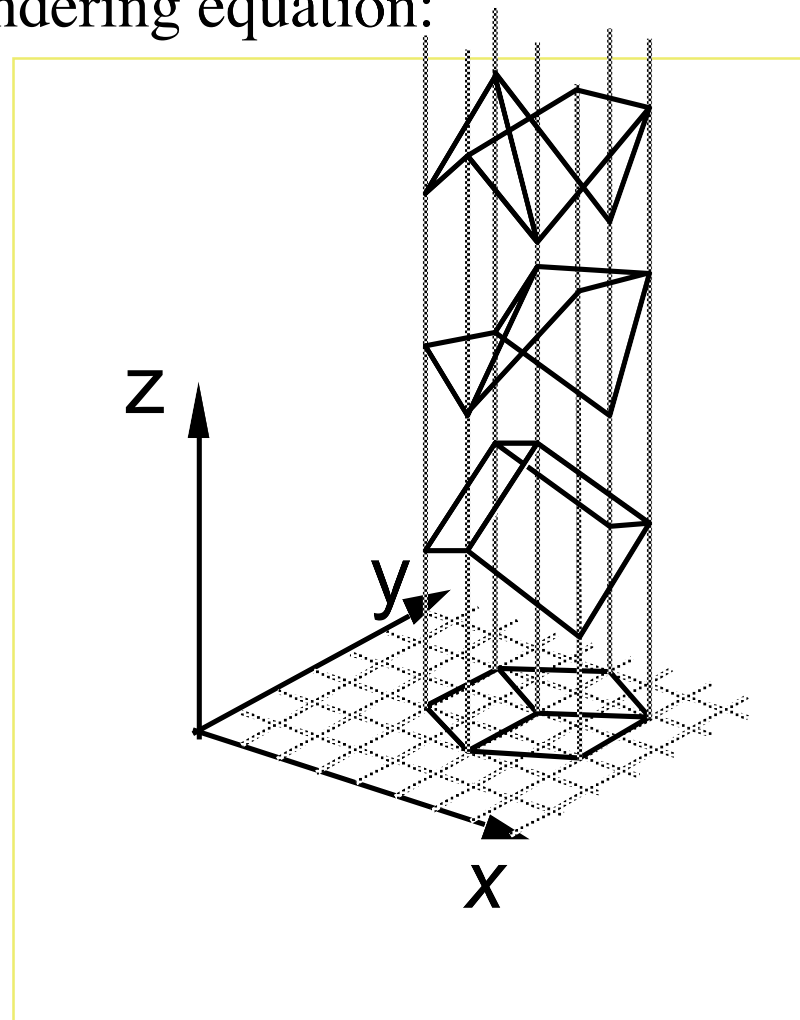

$$I = f(\{Q_i\}, L, V) = f(scene)$$

$\{Q_i\}$ = object descriptions

L = description of the scene lighting

V = viewpoint and imaging parameters (e.g. focus)

Modeling rendering probabilistically:

Likelihood: $p(I \mid scene)$

e.g. for no rendering noise

$$p(I \mid scene) = \delta(I - f(scene))$$

Prior further narrows selection

Priors weight likelihood selections

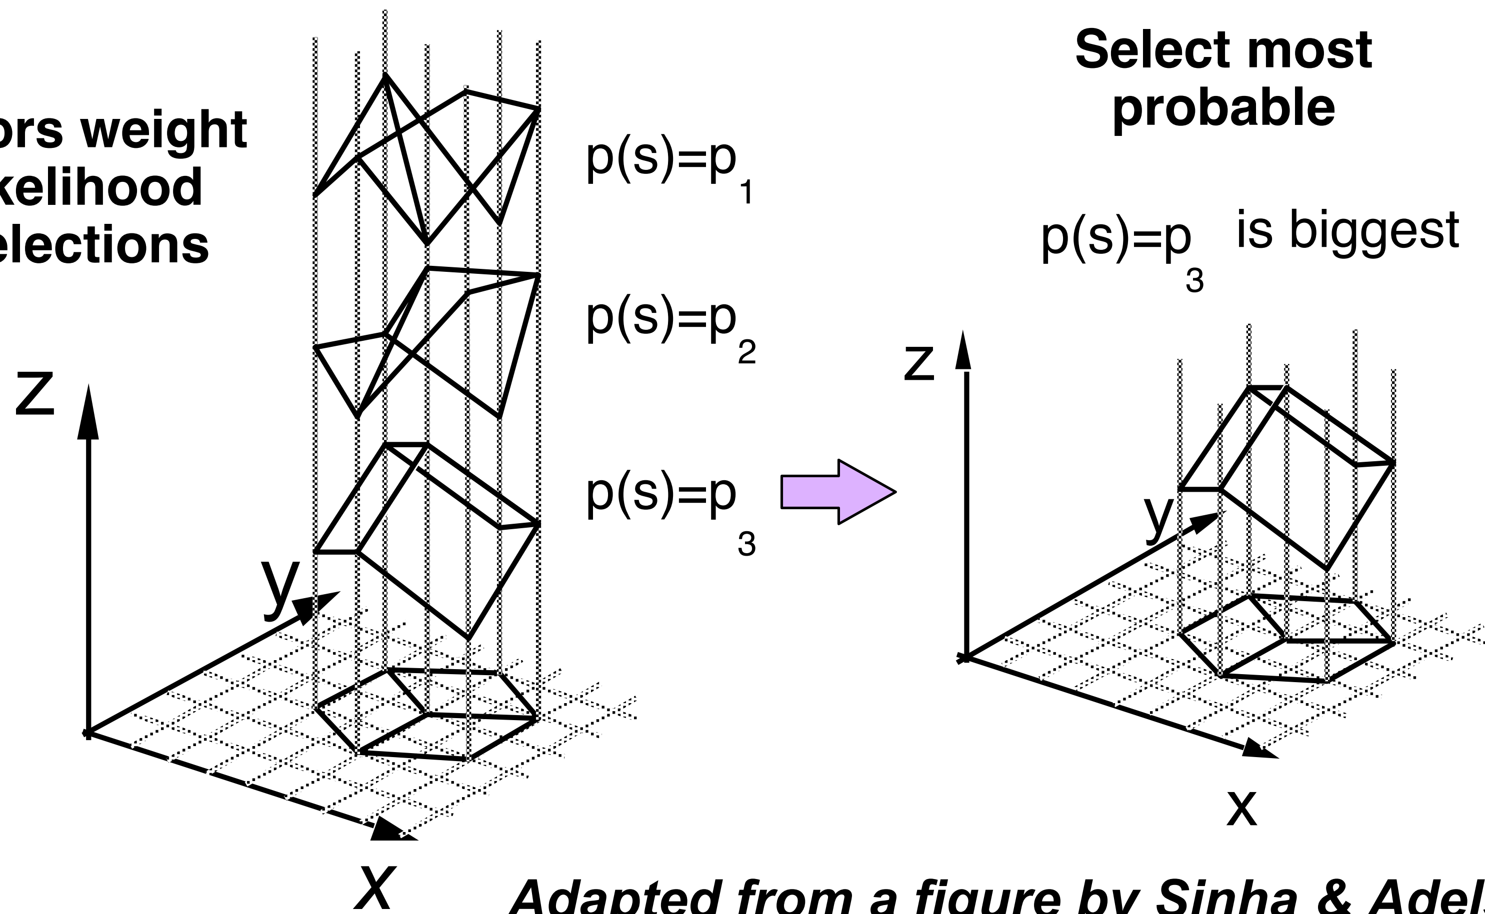

Adapted from a figure by Sinha & Adelson

Graphical Models for Perceptual Inference

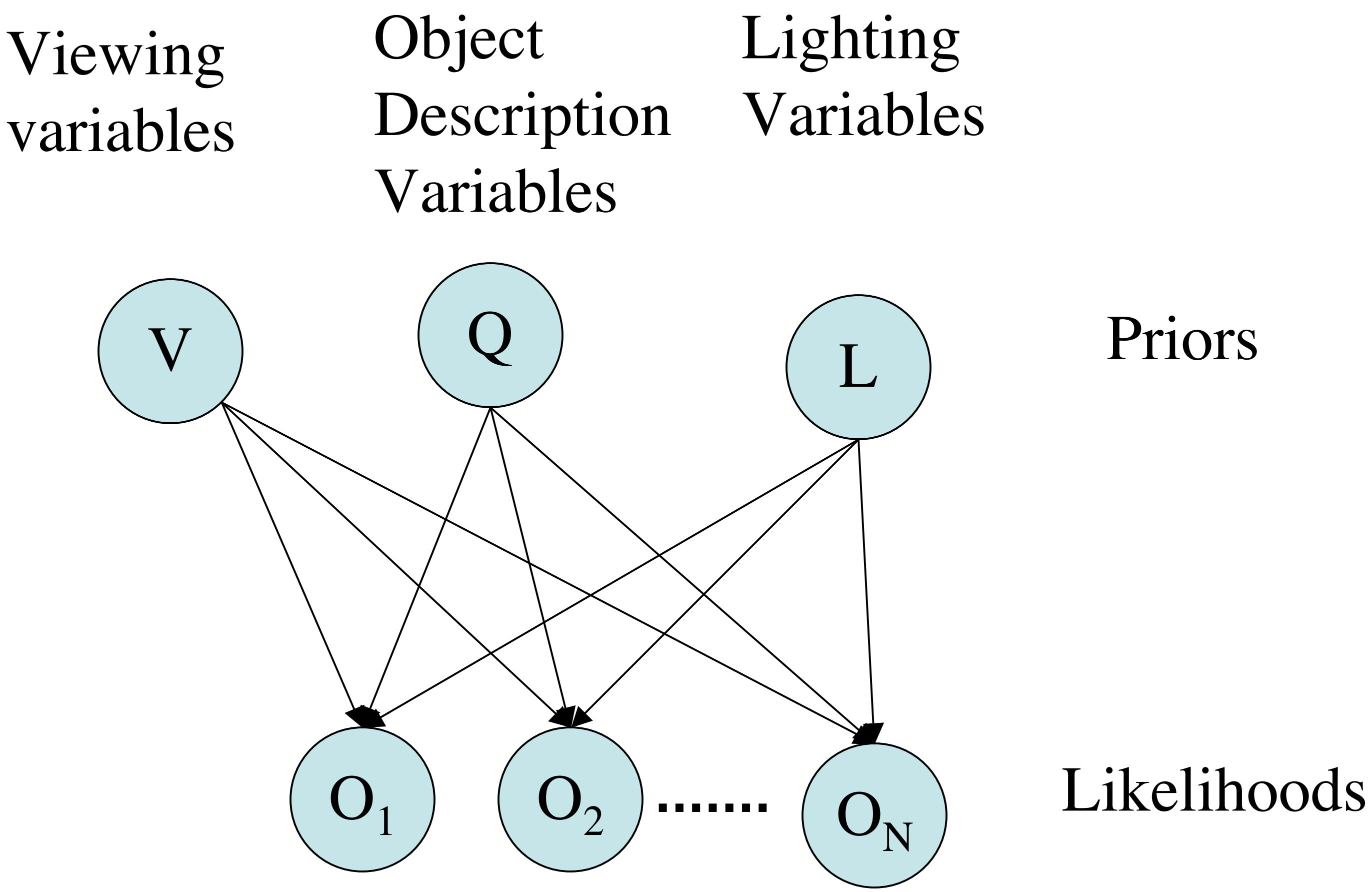

All of these nodes can be subdivided into individual Variables for more compact representation of the probability distribution

There is an ideal world…

- Matching environment to assumptions
  - Prior on light source direction

Prior on light source motion

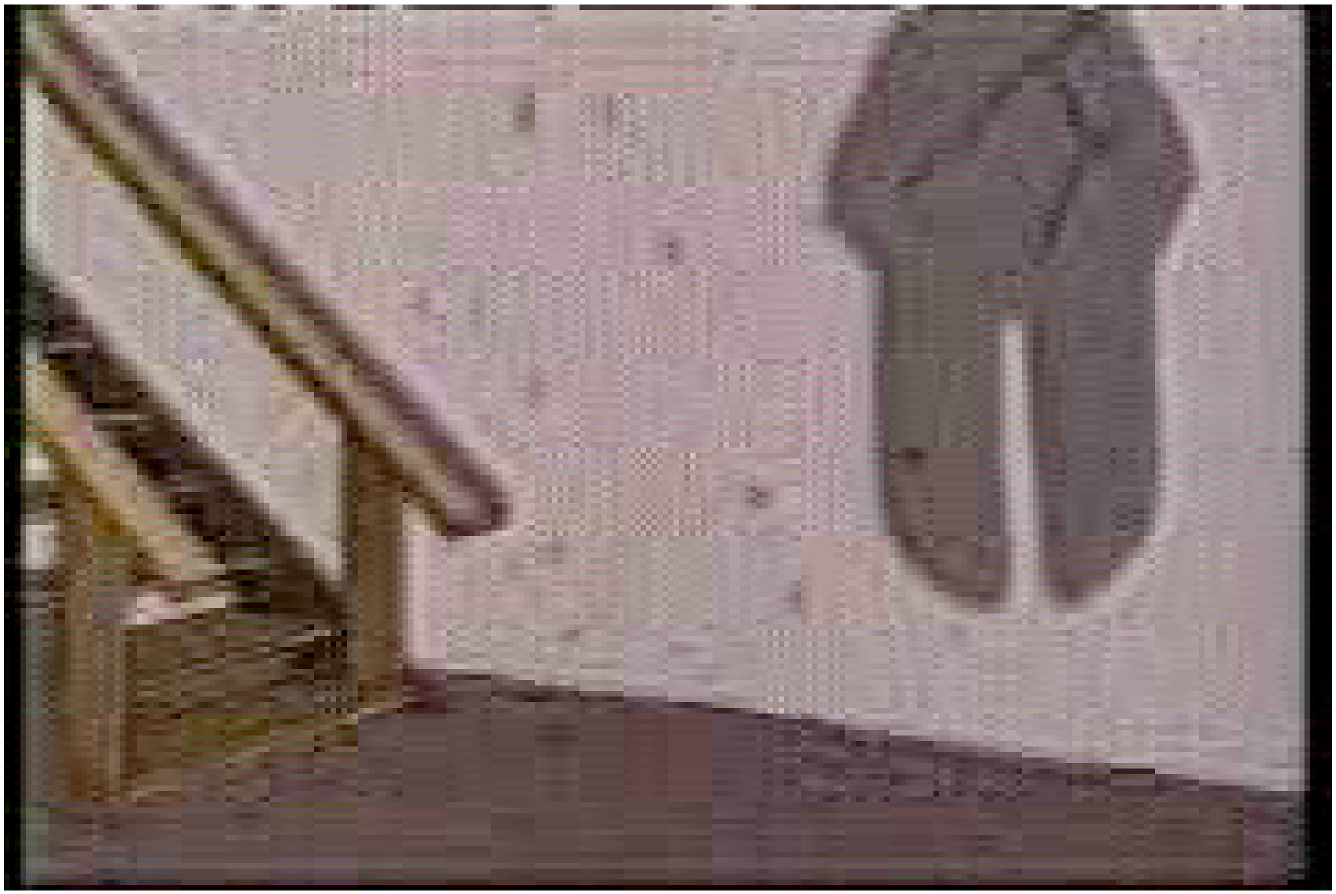

Combining Likelihood over different observations

Multiple data sources, sometimes contradictory.

What data is relevant? How to combine?

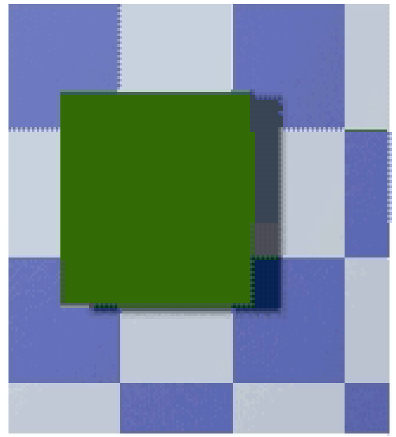

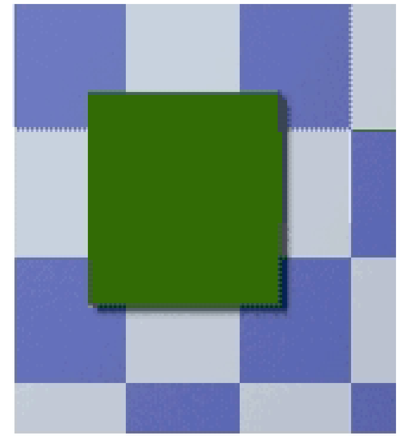

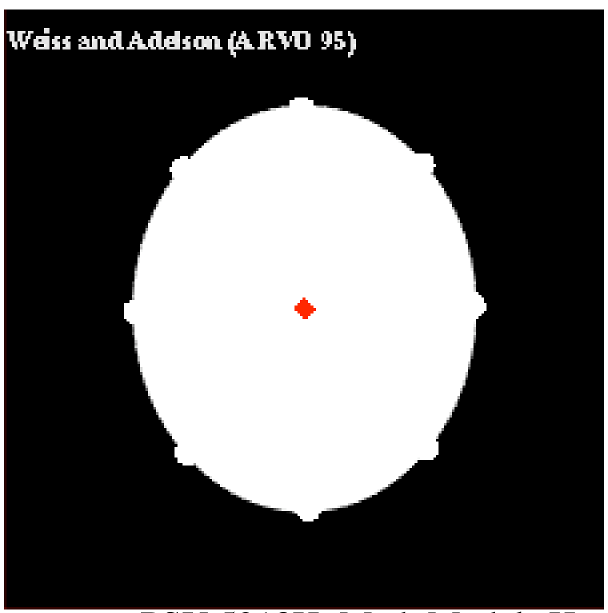

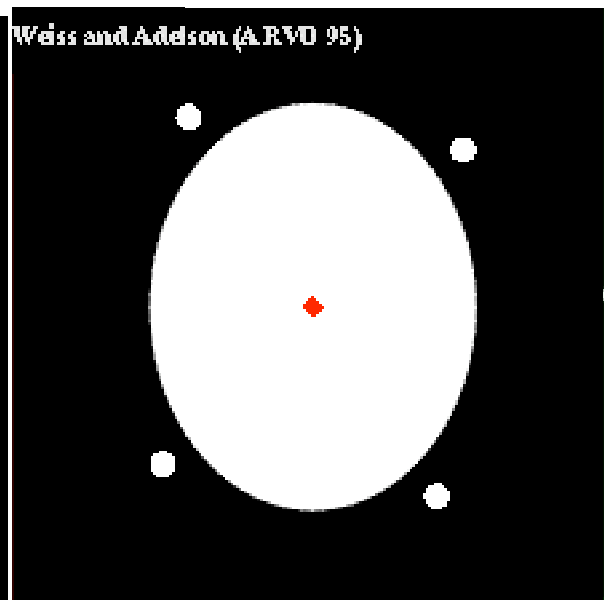

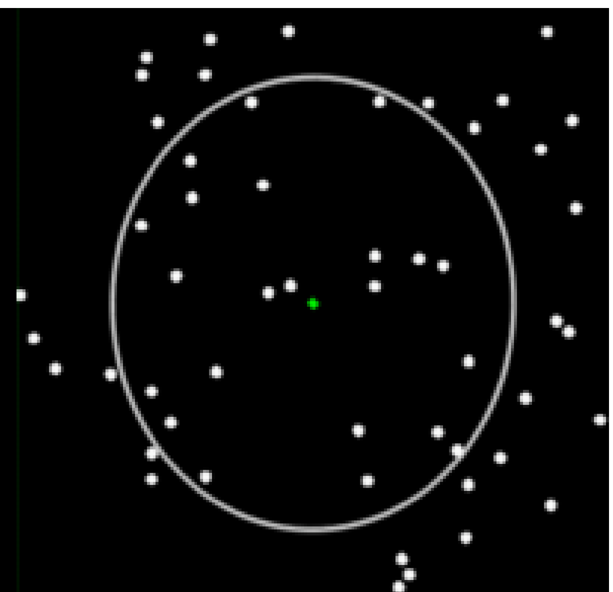

Modeling Perceptual Decisions

- One possible strategy for human perception is to *try to get the right answer.* How do we model this with decision theory?

Penalize responses with an error Utility function:

Given sensory data $\boldsymbol{O} = \{o_1, o_2, \ldots, o_n\}$ and some world property $\boldsymbol{S}$ perception is trying to infer, compute the decision:

$$V(r) = \sum_s P(s|O)U(r,s)$$ Value of response

$$r_{best} = \arg\max_r V(r)$$ Response is best value

Where:

$$U(r,s) = \delta(r-s) \quad \Leftrightarrow \quad \begin{cases} 1 & \text{if} \quad r = s \\ 0 & \text{otherwise} \end{cases}$$

Rigid Rotation
Or Shape Change?

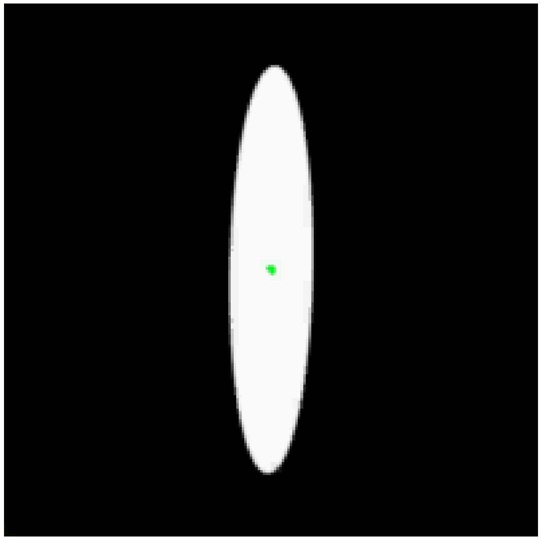

Actions:

$a = 1$ Choose rigid

$a = 0$ Choose not rigid

Utility

U(s,a)	a=0	a=1
$s = 1$	1	0
$s = 0$	0	1

Rigidity judgment

Let s is a random variable representing rigidity

s=1 it is rigid

s=0 not rigid

Let $O_1, O_2, \ldots, O_n$ be descriptions of the ellipse for frames 1,2, … , n

Need likelihood

$$P(O_1, O_2, \ldots, O_n \mid s)$$

Need Prior

$$P(s)$$

"You must choose, but Choose Wisely"

Given previous formulation, can we minimize the number of errors we make?

- *Given:*

  responses a_i, categories s_i, current category s, data o

- *To Minimize error:*
  - Decide a_i **if** $P(a_i \mid o) > P(a_k \mid o)$ **for all** $i \neq k$

    $P(o \mid s_i) P(s_i) > P(o \mid s_k) P(s_k)$

    $P(o \mid s_i) / P(o \mid s_k) > P(s_k) / P(s_i)$

    $P(o \mid s_i) / P(o \mid s_k) > T$

Optimal classifications always involve hard boundaries