Benchmarks for Vision

Exact probabilities of arbitrary features
in a generative model of naturalistic images

xaq pitkow @ columbia university

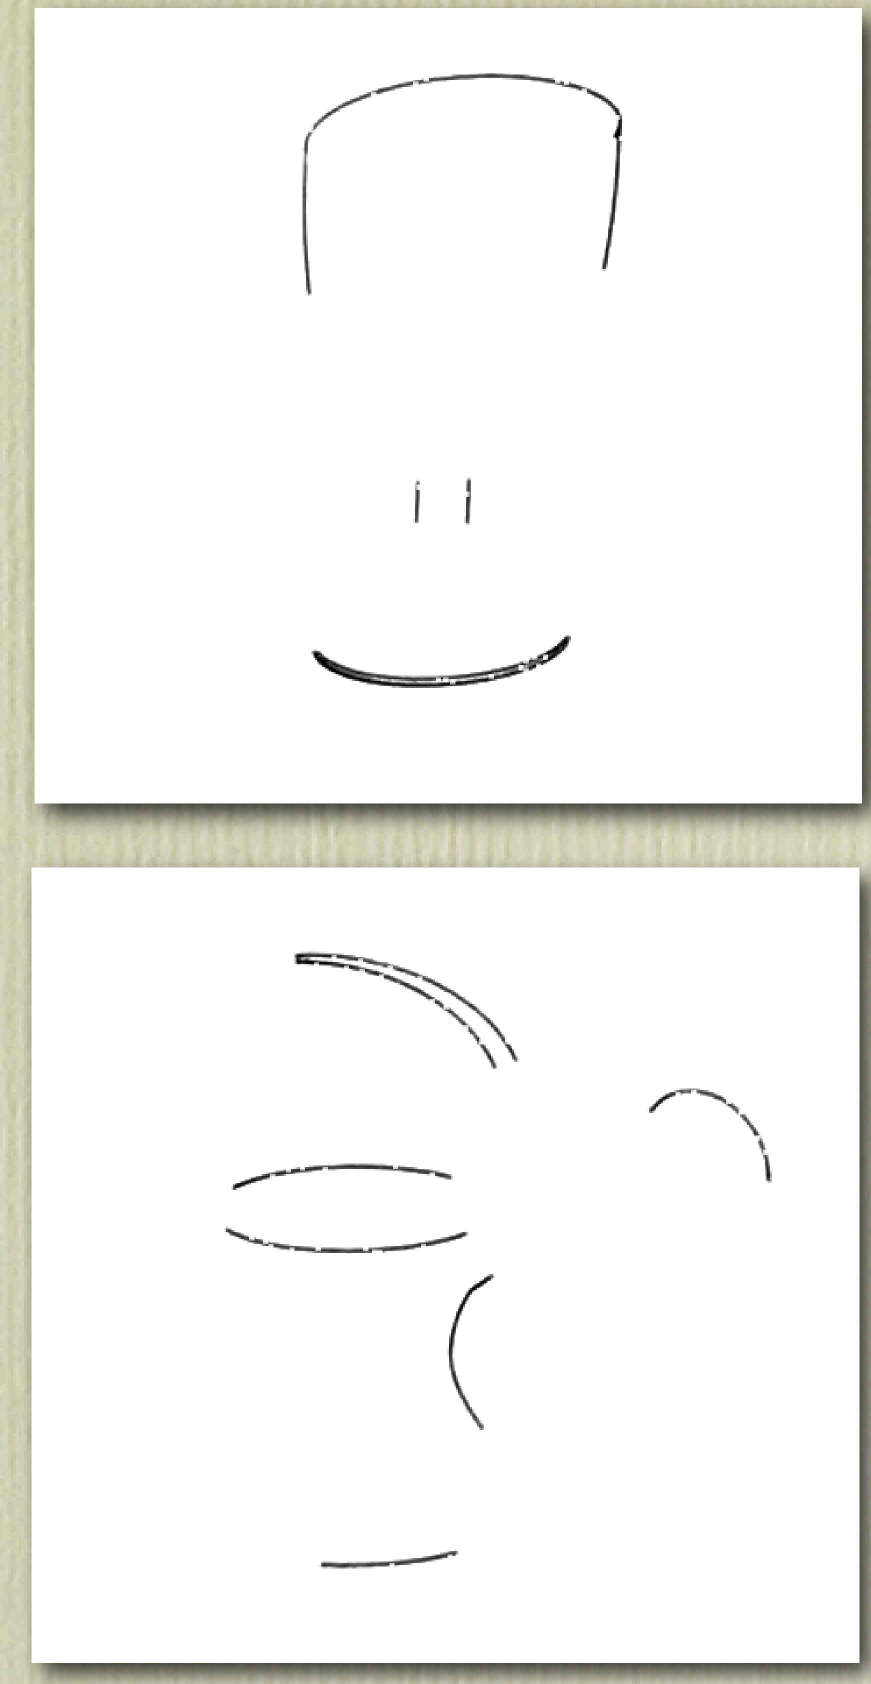

hoffman '91

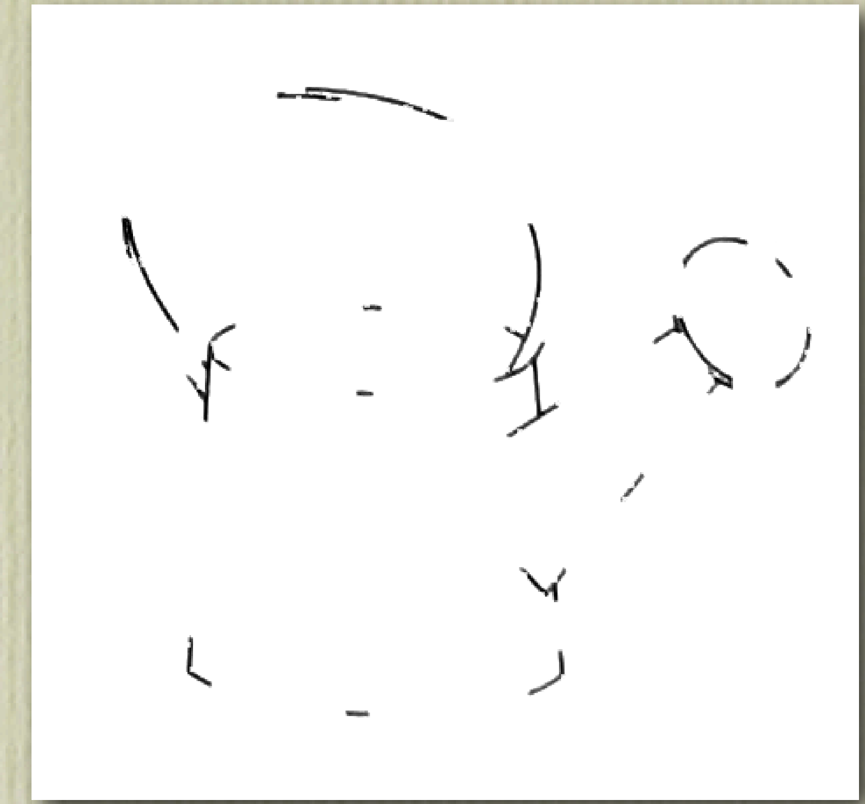

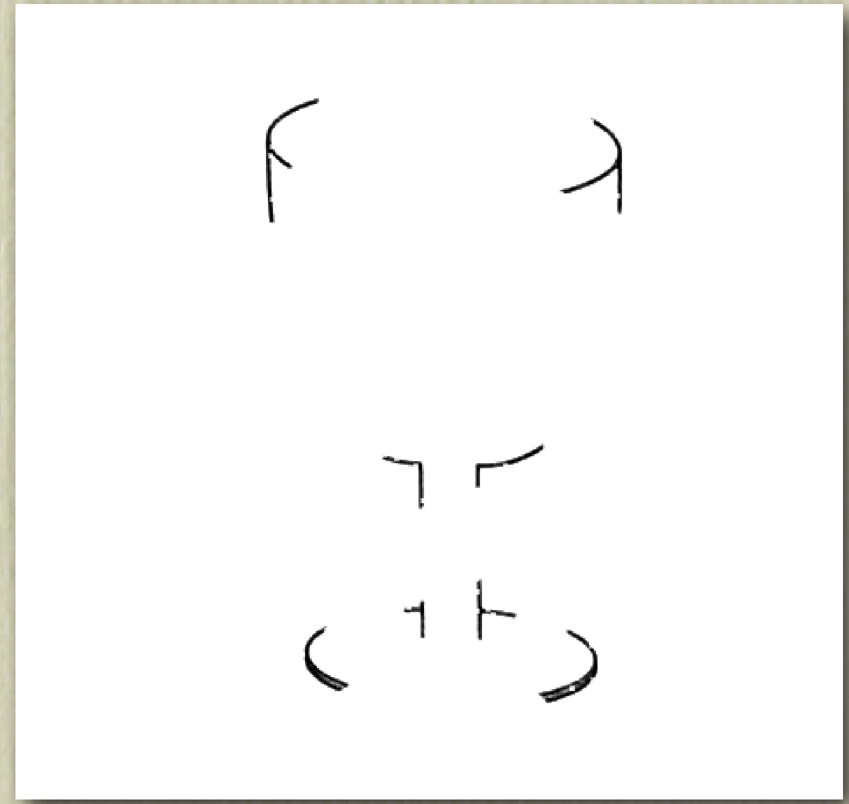

hoffman '91

hoffman '91

- **WANTED: explicit probabilities for interesting features!**
    - impossible for real natural scenes
    - most simple models are too unrealistic
    - need richer yet tractable model

- Minimal description of local features:

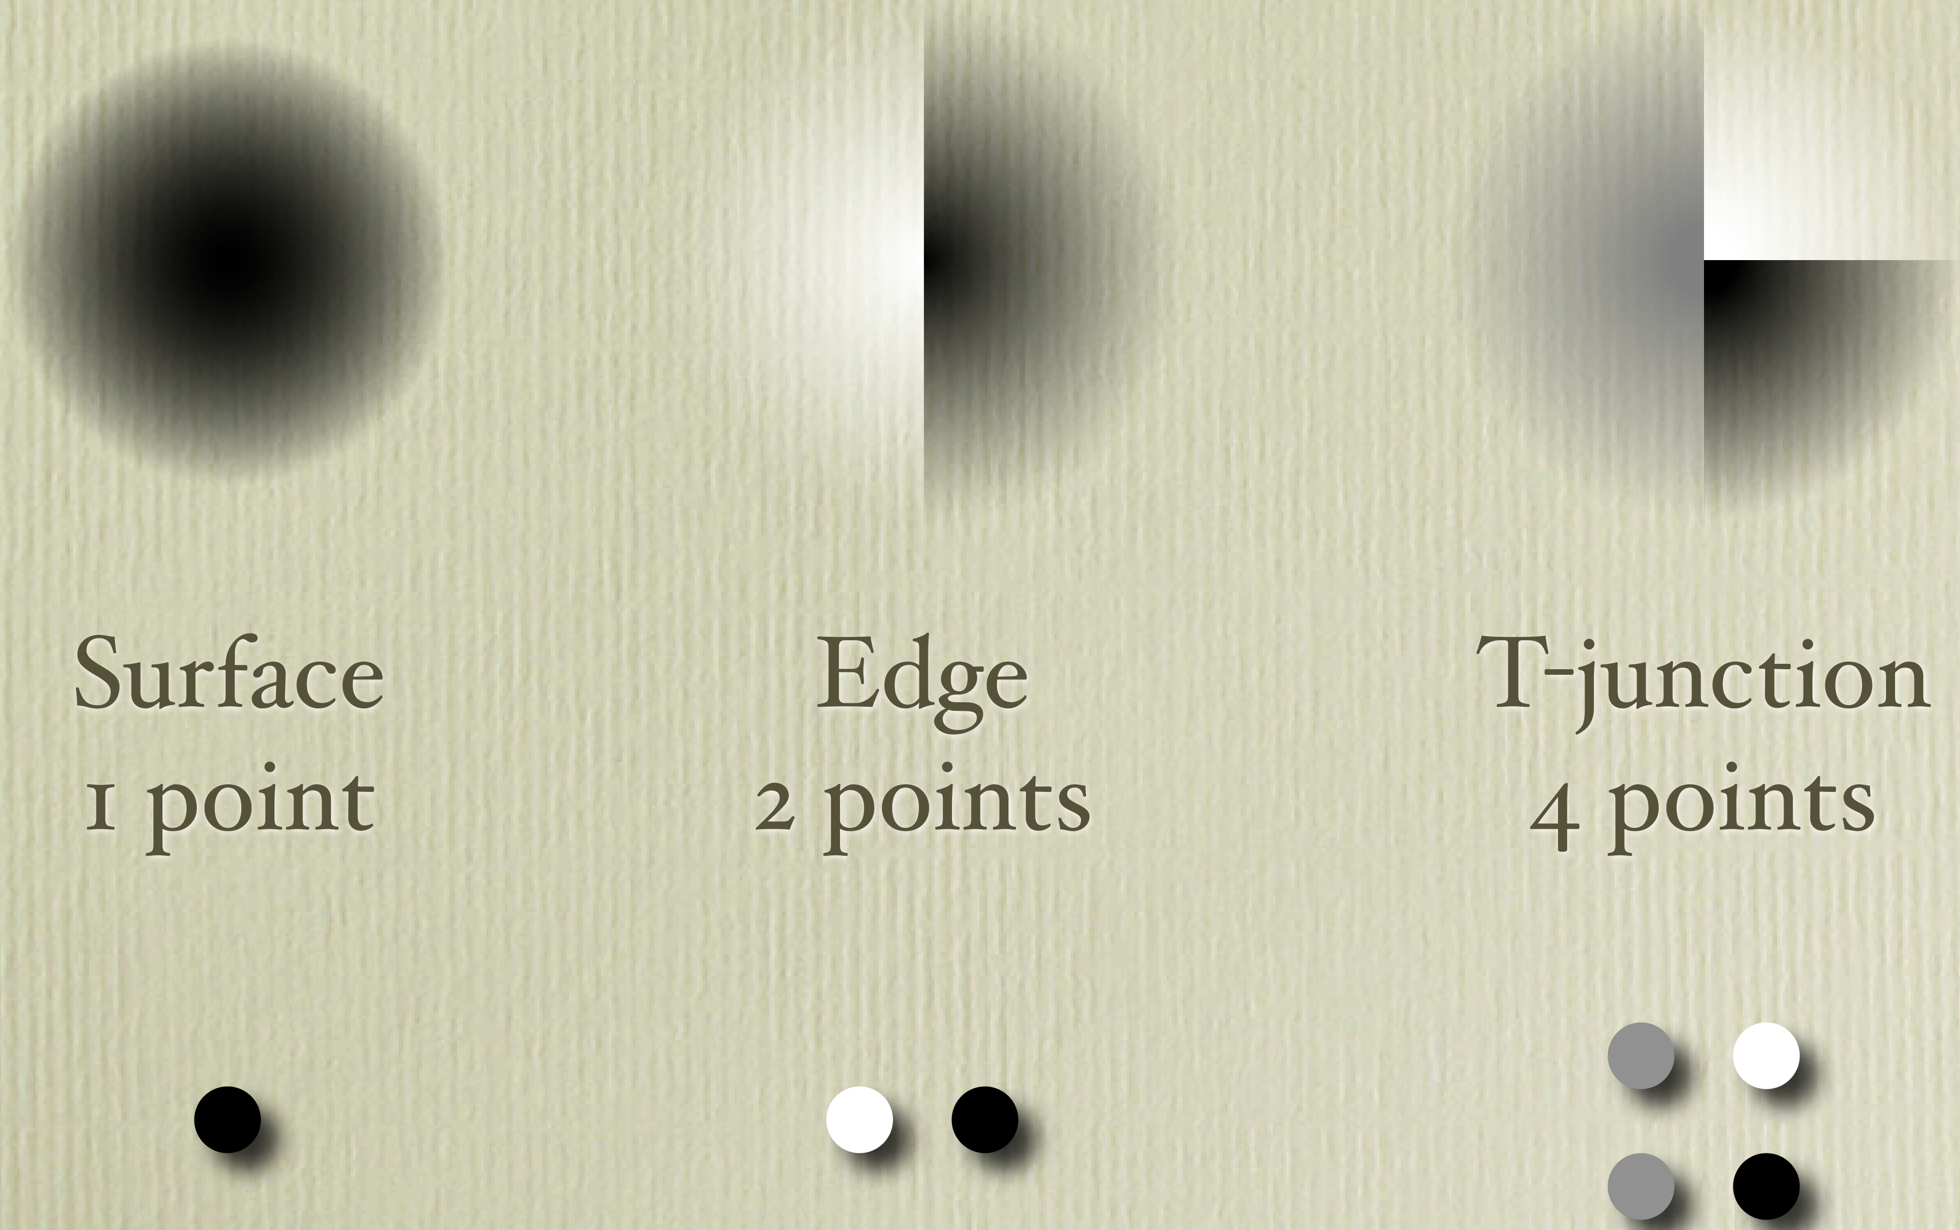

Brightness perception (3 pts)

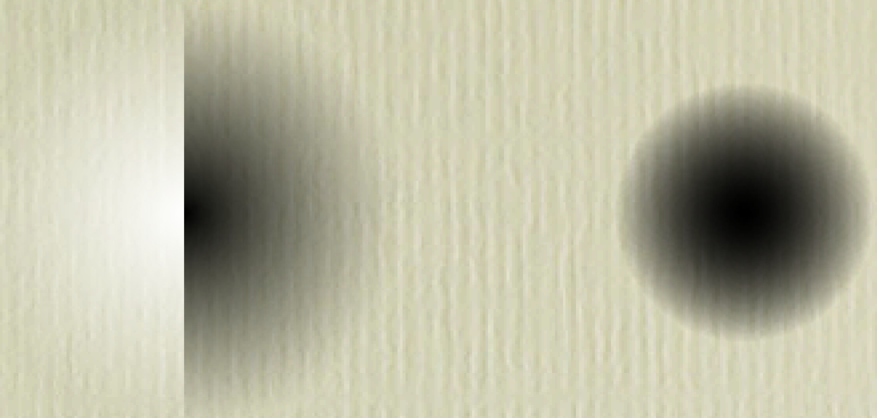

Brightness perception (3 pts)

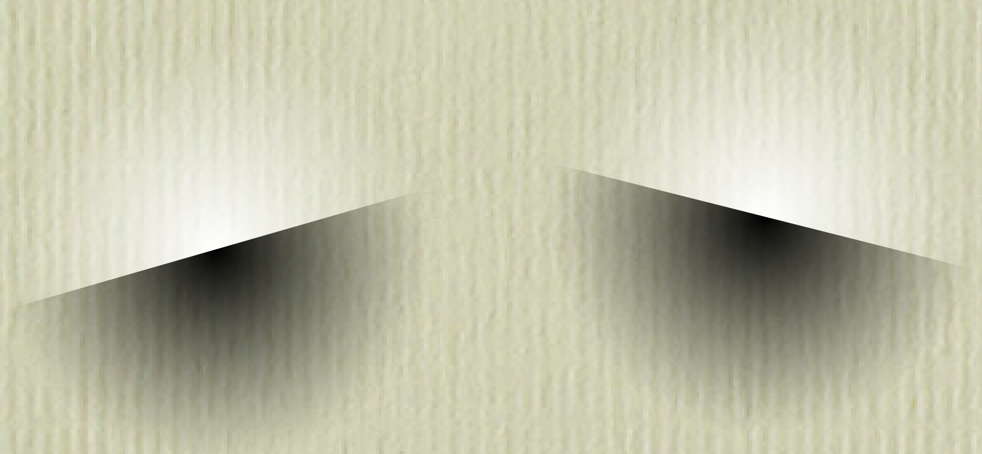

Contour facilitation (4 pts)

Brightness perception (3 pts)

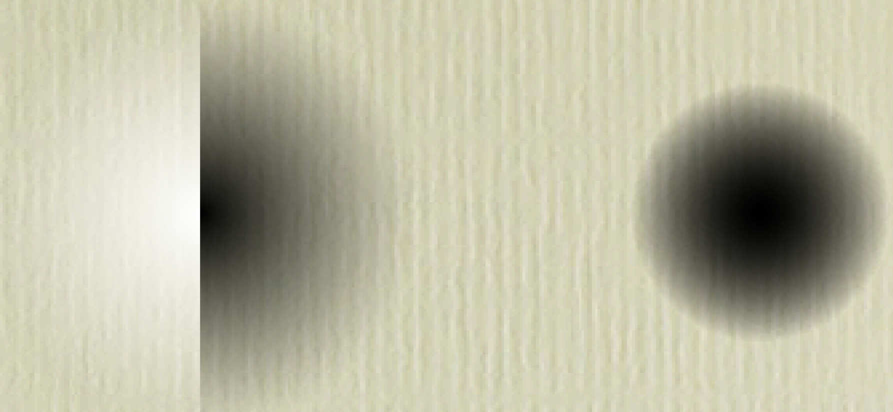

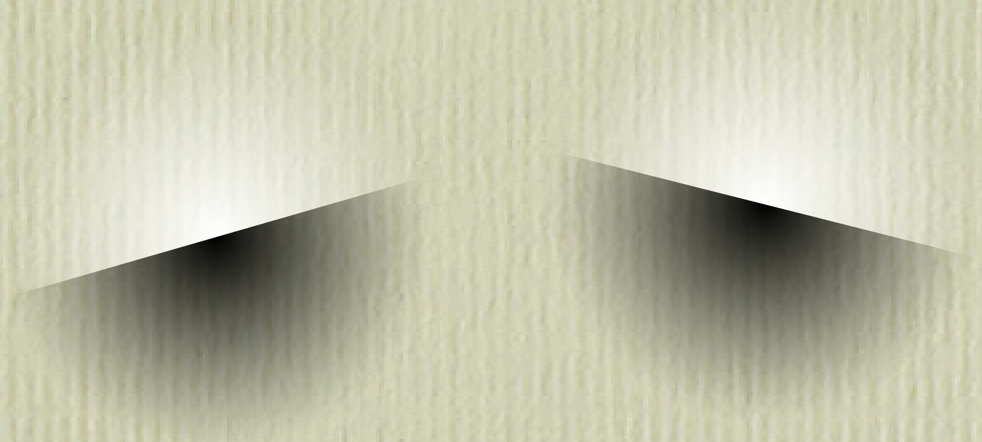

Contour facilitation (4 pts)

Amodal completion (8 pts)

- Natural scene statistics:

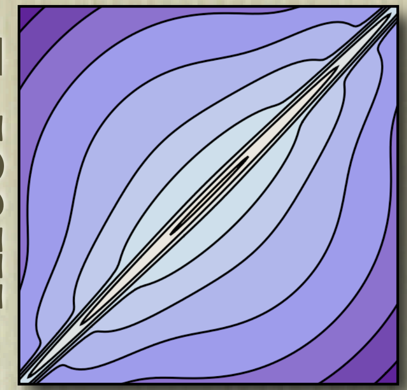

- Natural scene statistics:

- Fill in missing information:

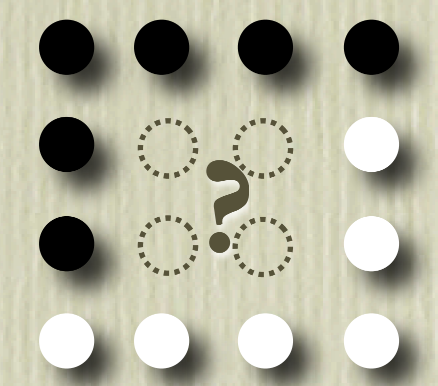

- Natural scene statistics:

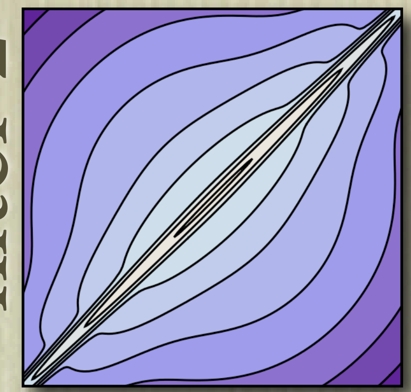

- Fill in missing information:

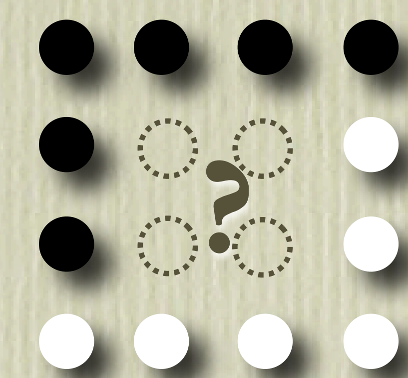

- Relate measurements to object properties:

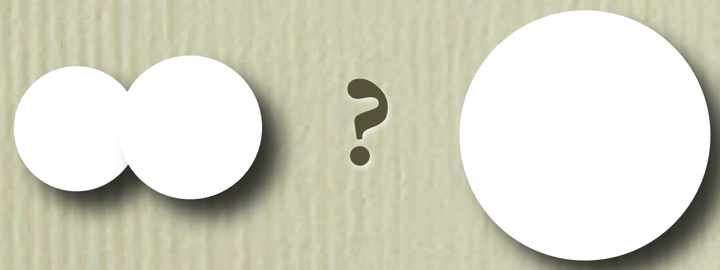

what shapes are there?

which object is in front?

Dead Leaves Model

Drop 'Leaves' with random shape, size, color, location.

Naturalistic because:

- objects
- occlusion
- edges

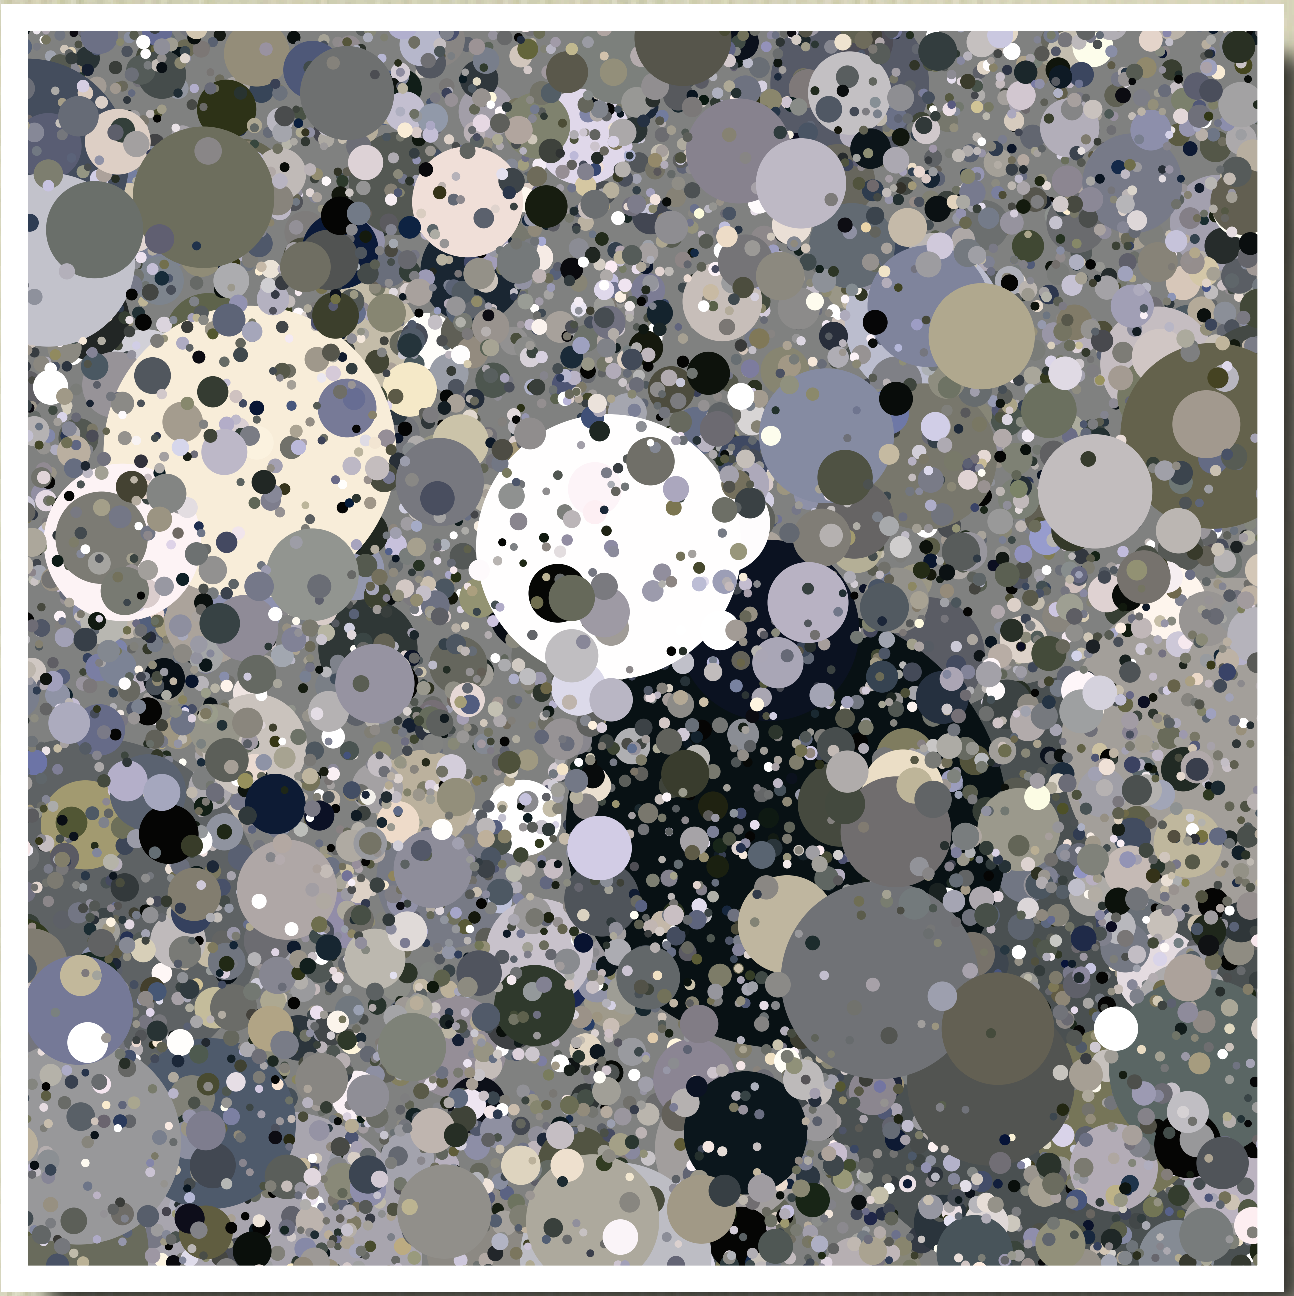

Matheson 1975, Ruderman 1997, Mumford 2001

Goal: Multipoint Probabilities

- Find joint probability of image intensity at arbitrary set of points

$$P(X, Y, Z)$$

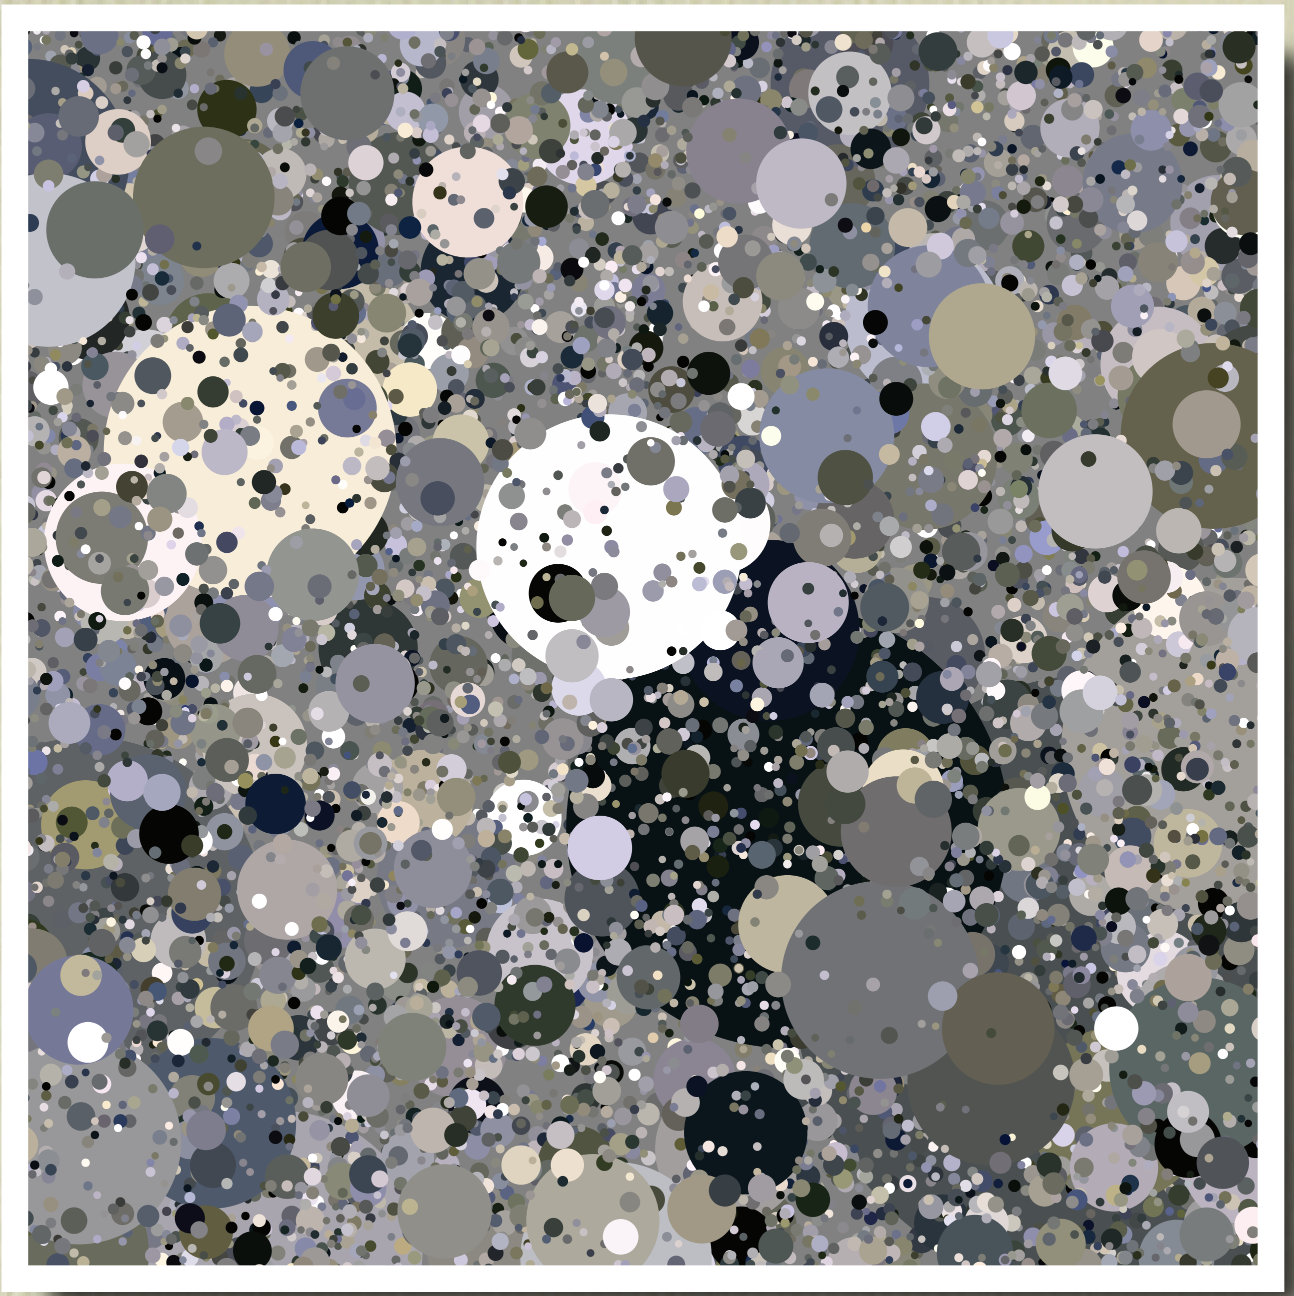

lowercase: position
uppercase: intensity

Goal: Multipoint Probabilities

- Find joint probability of image intensity at arbitrary set of points

$$P(X, Y, Z)$$

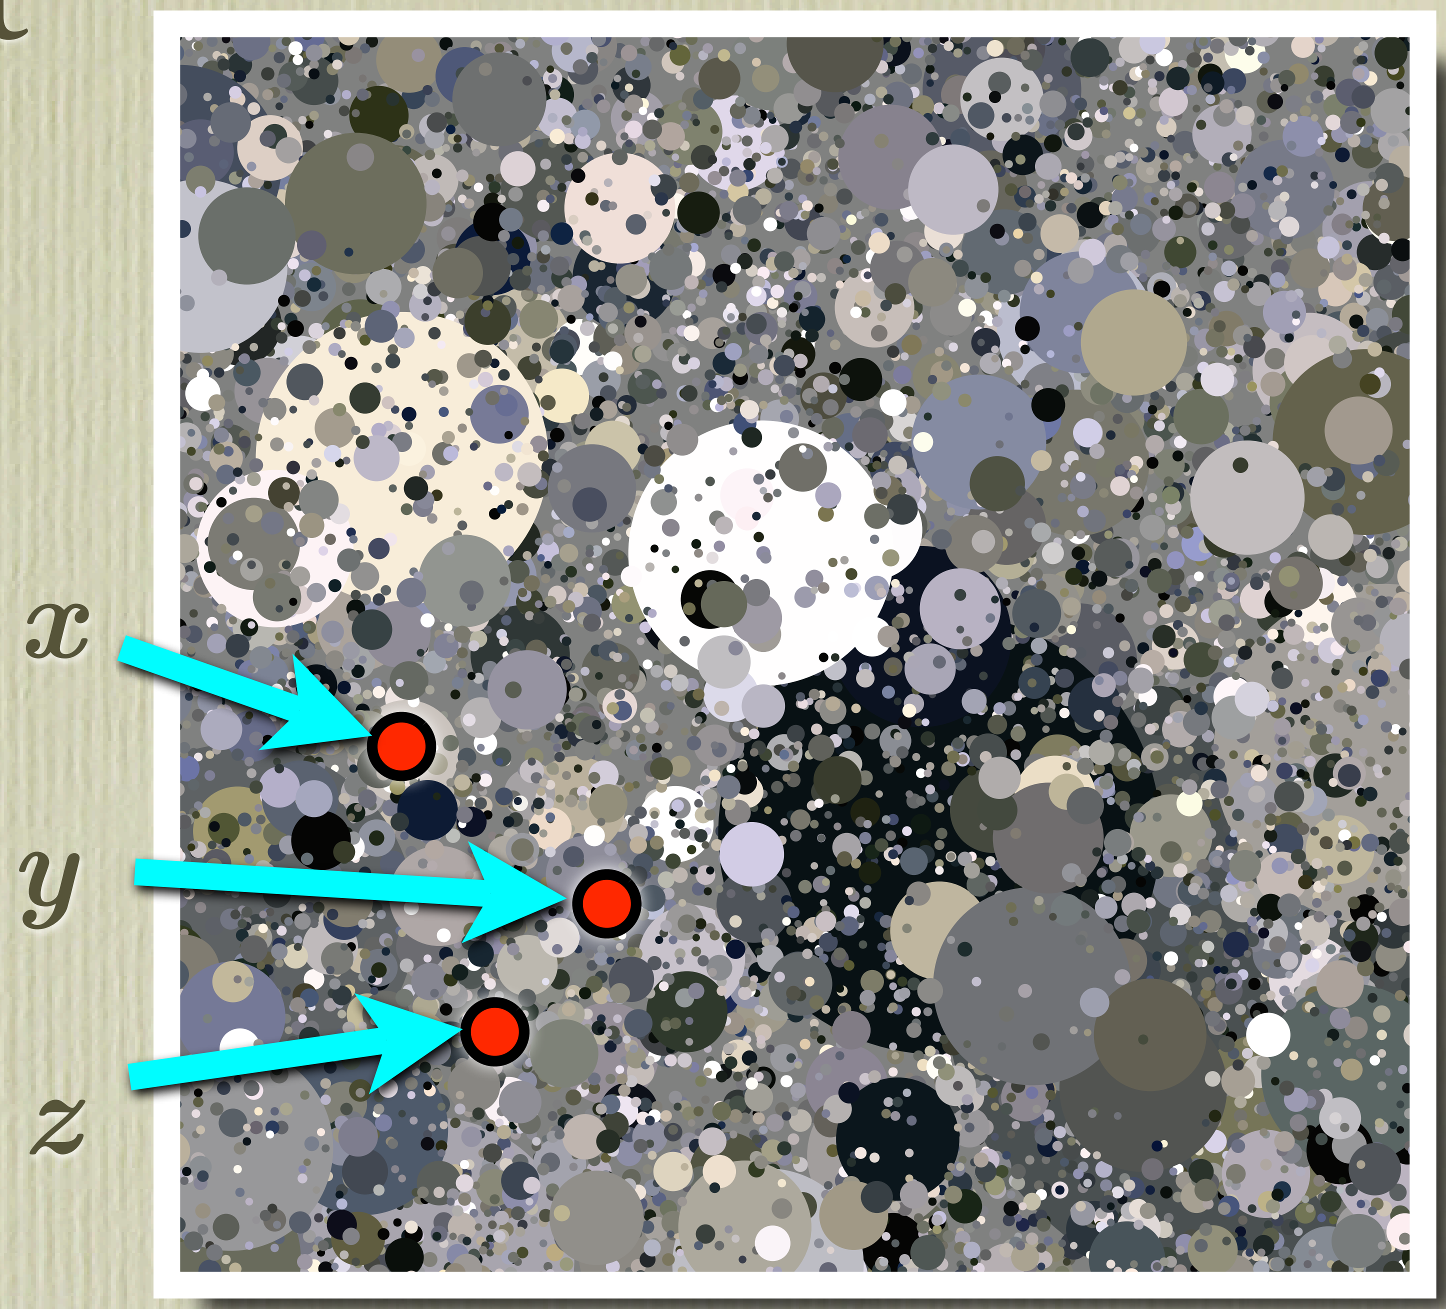

lowercase: position
uppercase: intensity

- The arrangement of the objects determines a "**partition**" of the points.

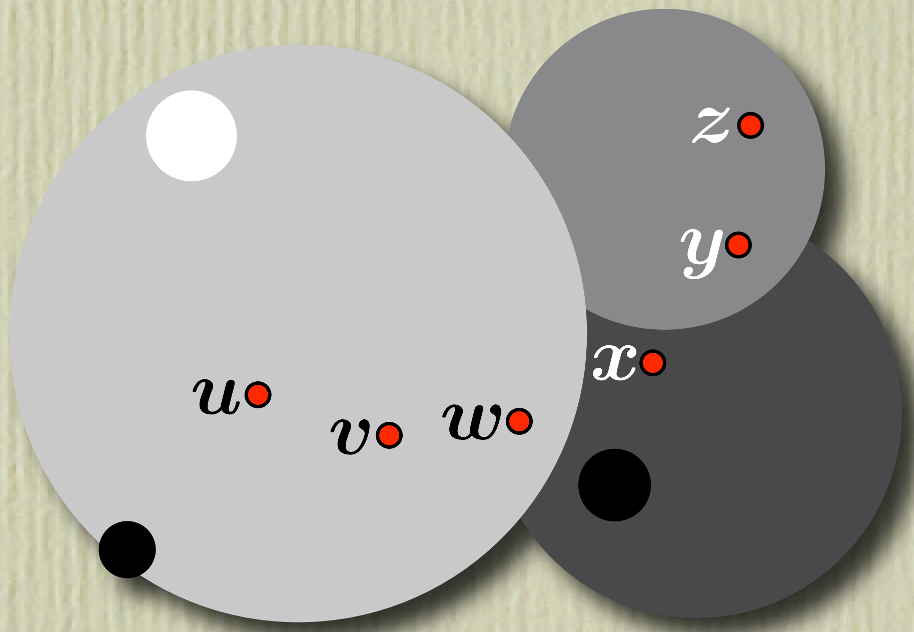

$$\pi = \{\{u, v, w\}, \{x\}, \{y, z\}\}$$

- The arrangement of the objects determines a "**partition**" of the points.

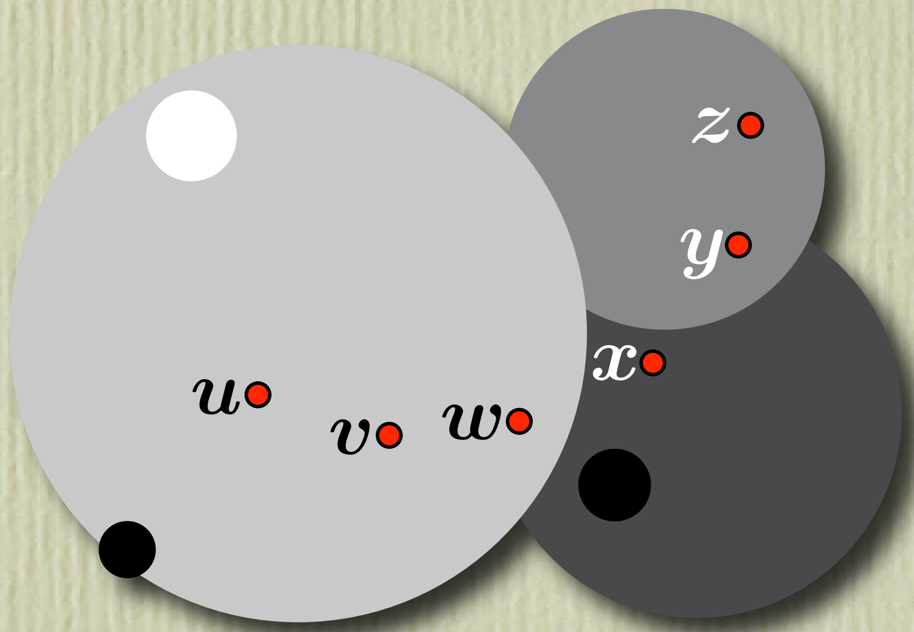

$$\pi = \{\{u, v, w\}, \{x\}, \{y, z\}\}$$

- Points in different objects have **independent** intensities

$$P(U, V, W, X, Y, Z|\pi) = P(U, V, W) \cdot P(X) \cdot P(Y, Z)$$

- The arrangement of the objects determines a "**partition**" of the points.

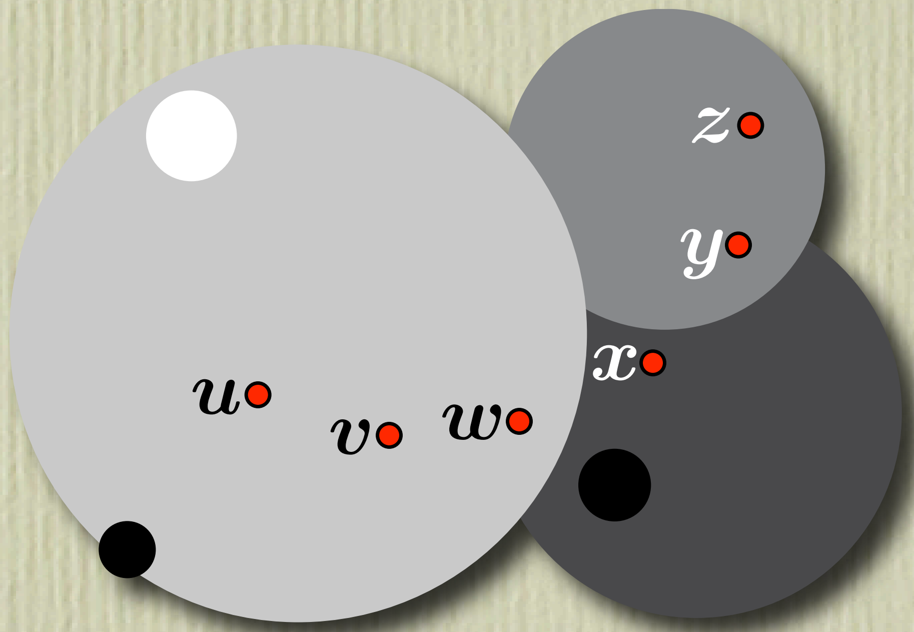

$$\pi = \{\{u, v, w\}, \{x\}, \{y, z\}\}$$

- Points in different objects have **independent** intensities

$$P(U, V, W, X, Y, Z|\pi) = P(U, V, W) \cdot P(X) \cdot P(Y, Z)$$

- Full joint probability is average over all partitions.

$$P(I) = \sum_{\pi} P_\pi P(I|\pi)$$

- What is probability of being in same object? Calculate recursively.

$$P_{\{x,y\}}$$

P = *prob for **any** obj*
Q = *prob for **this** obj*

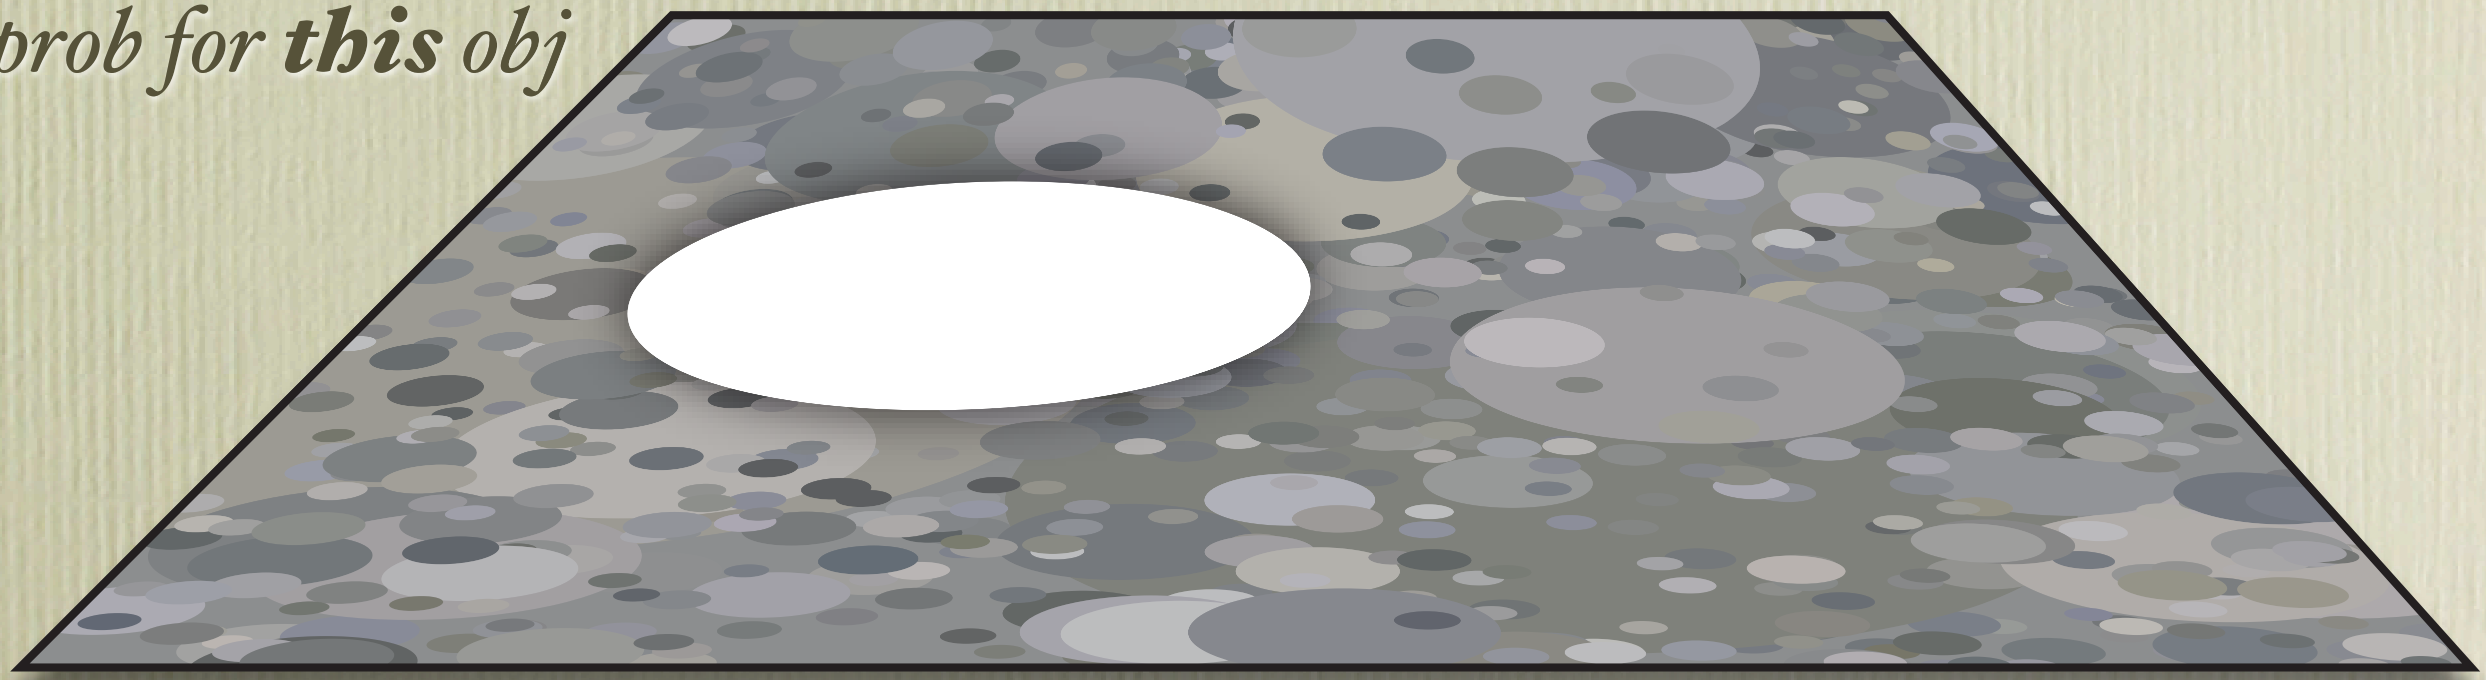

Ruderman 1997

- What is probability of being in same object? Calculate recursively.

$$P_{\{x,y\}}$$

P = *prob for* ***any*** *obj*

Q = *prob for* ***this*** *obj*

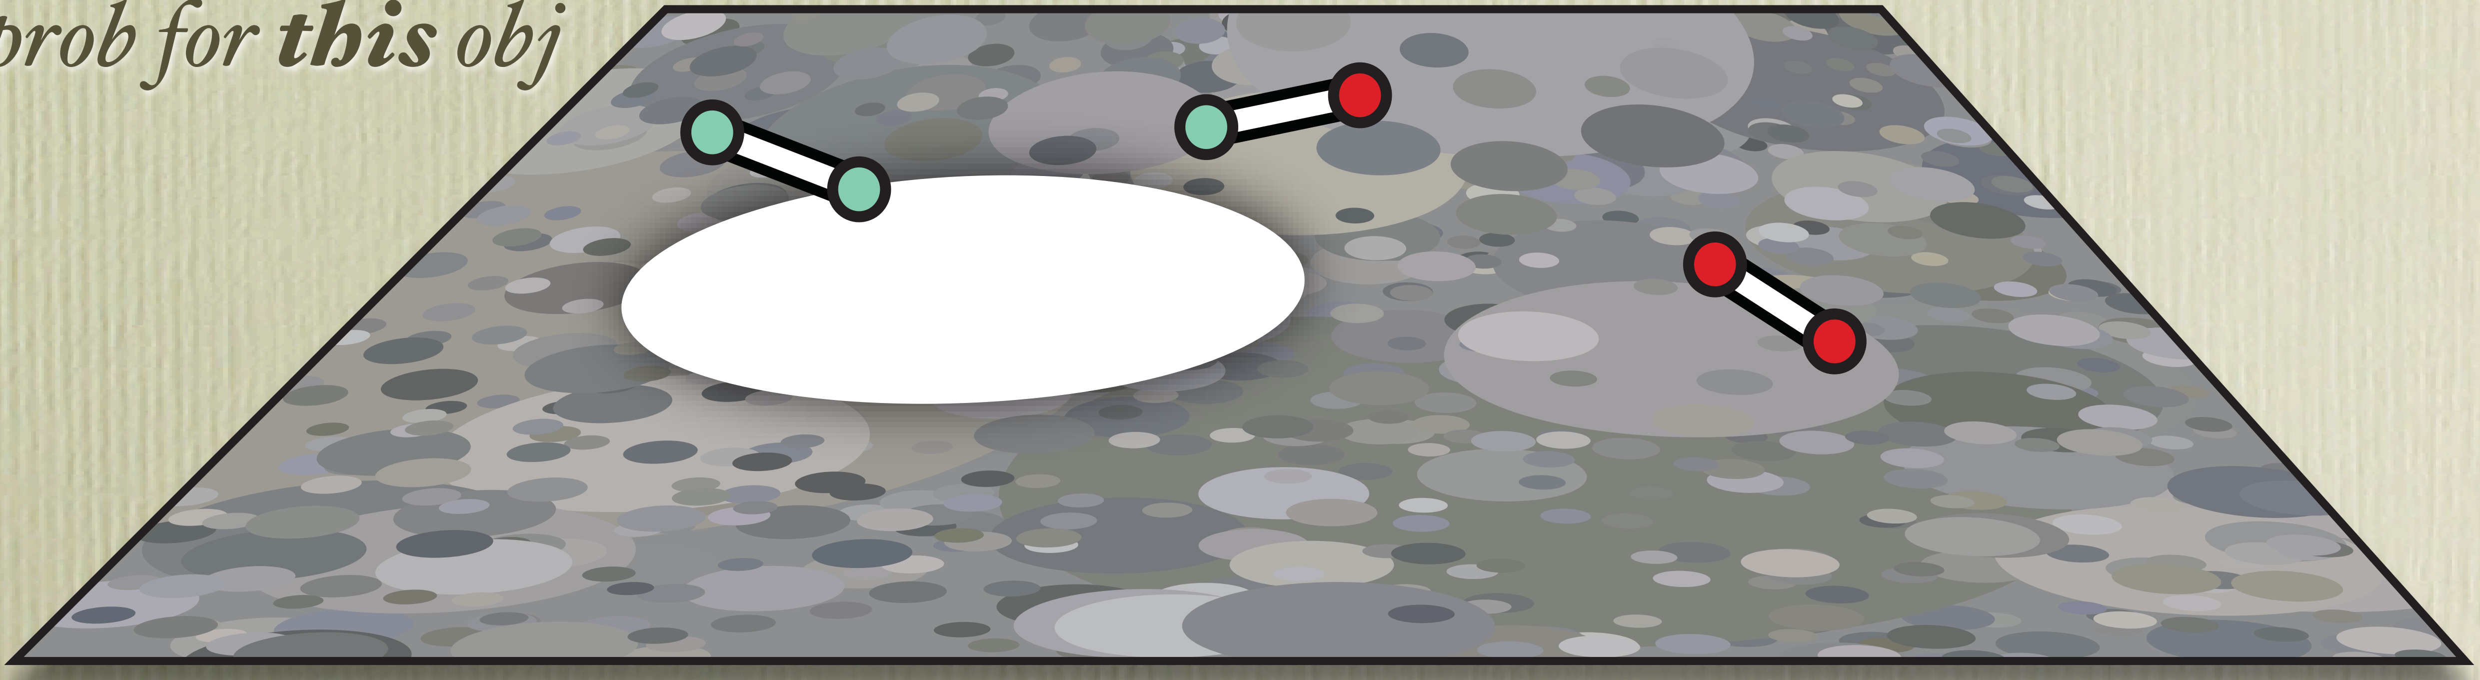

Ruderman 1997

- What is probability of being in same object? Calculate recursively.

$$P_{\{x,y\}}$$

P = *prob for* ***any*** *obj*

Q = *prob for* ***this*** *obj*

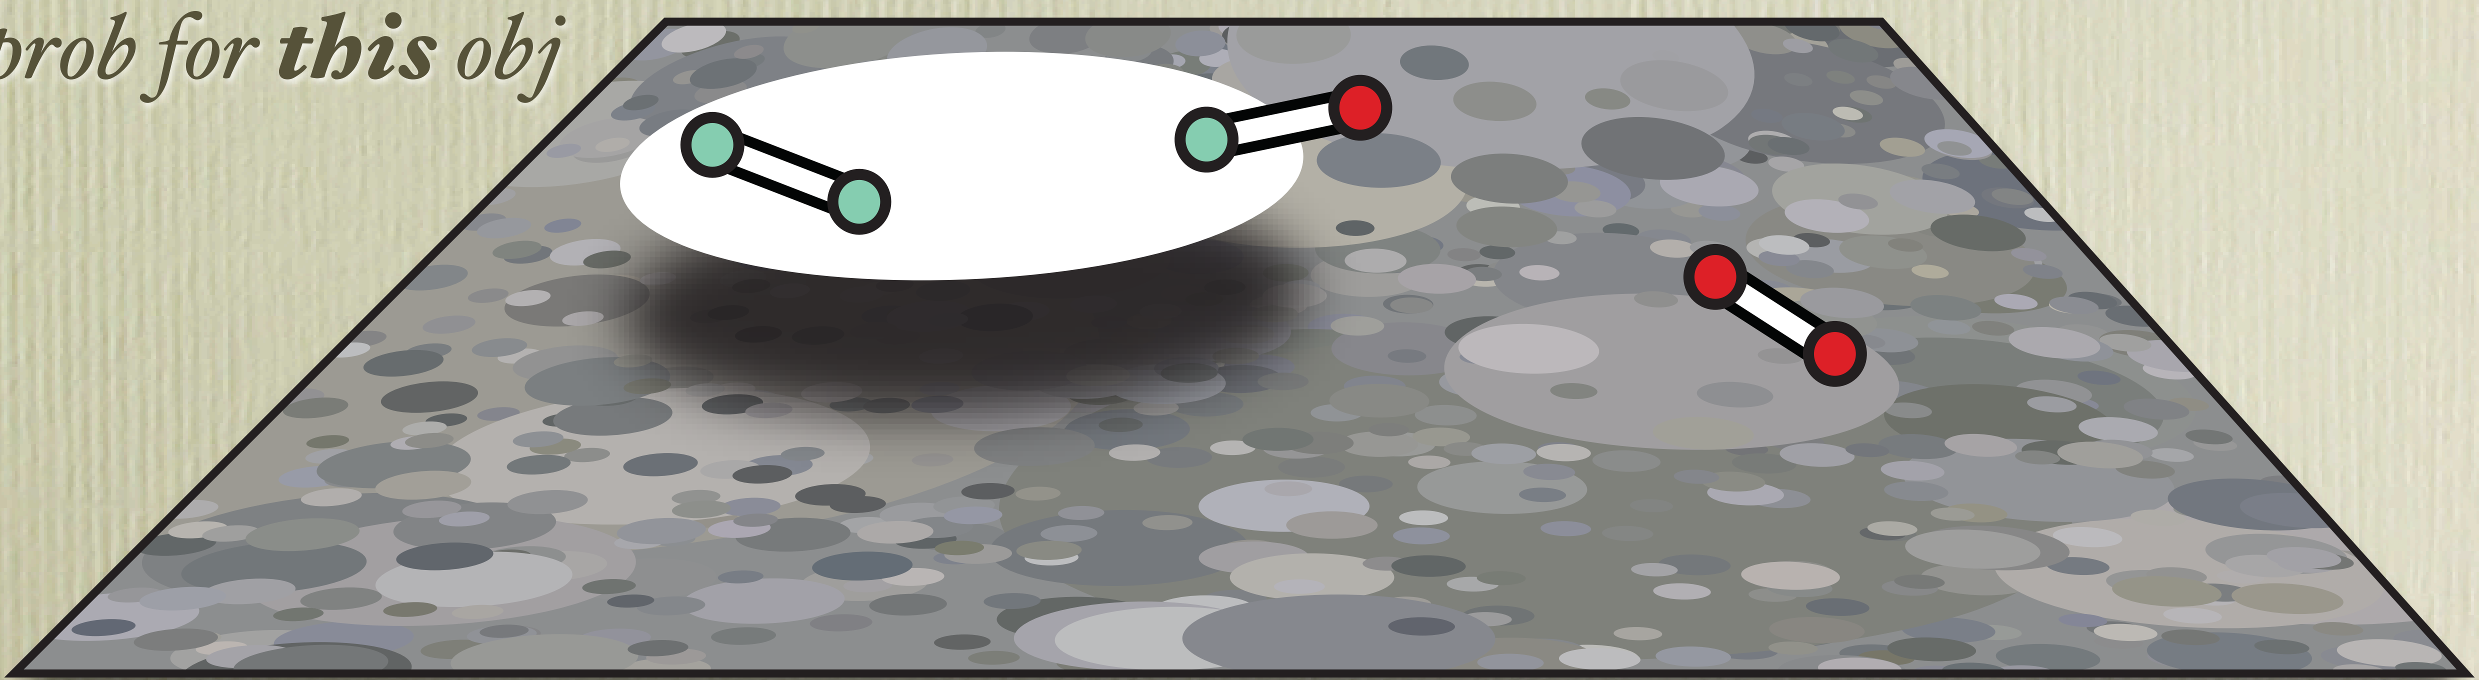

Ruderman 1997

- What is probability of being in same object? Calculate recursively.

$$P_{\{x,y\}} = Q_{11}$$

P = *prob for **any** obj*
Q = *prob for **this** obj*

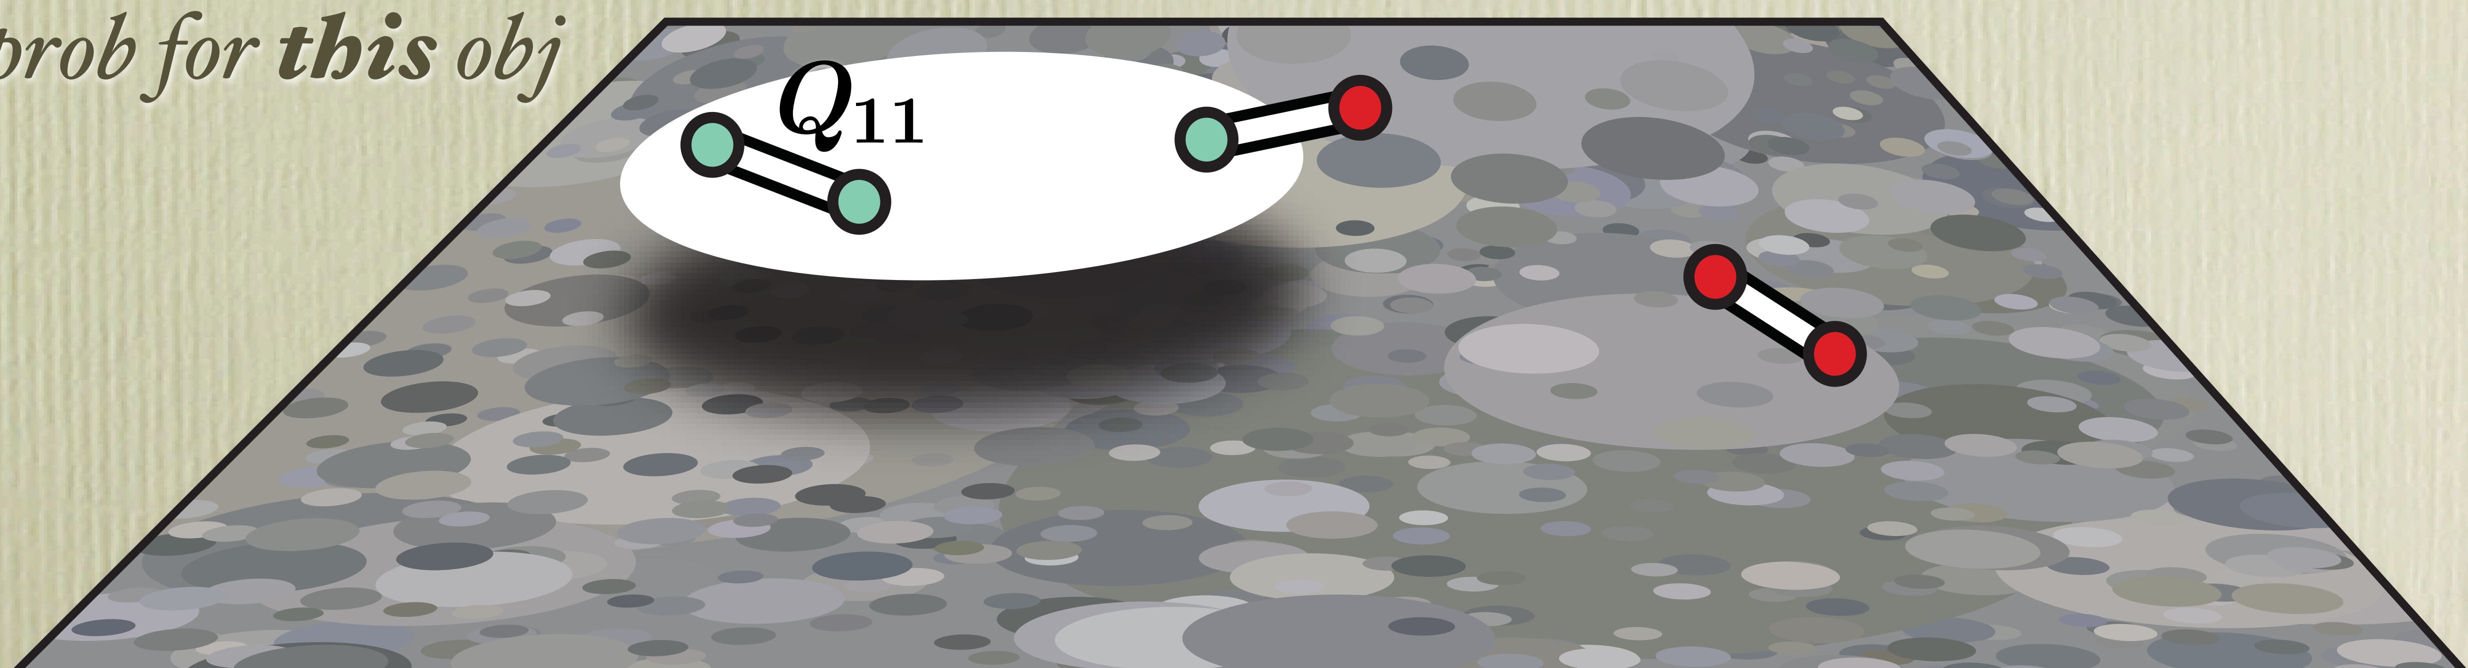

Ruderman 1997

- What is probability of being in same object? Calculate recursively.

$$P_{\{x,y\}} = Q_{11} + Q_{00} P_{\{x,y\}}$$

P = *prob for* ***any*** *obj*

Q = *prob for* ***this*** *obj*

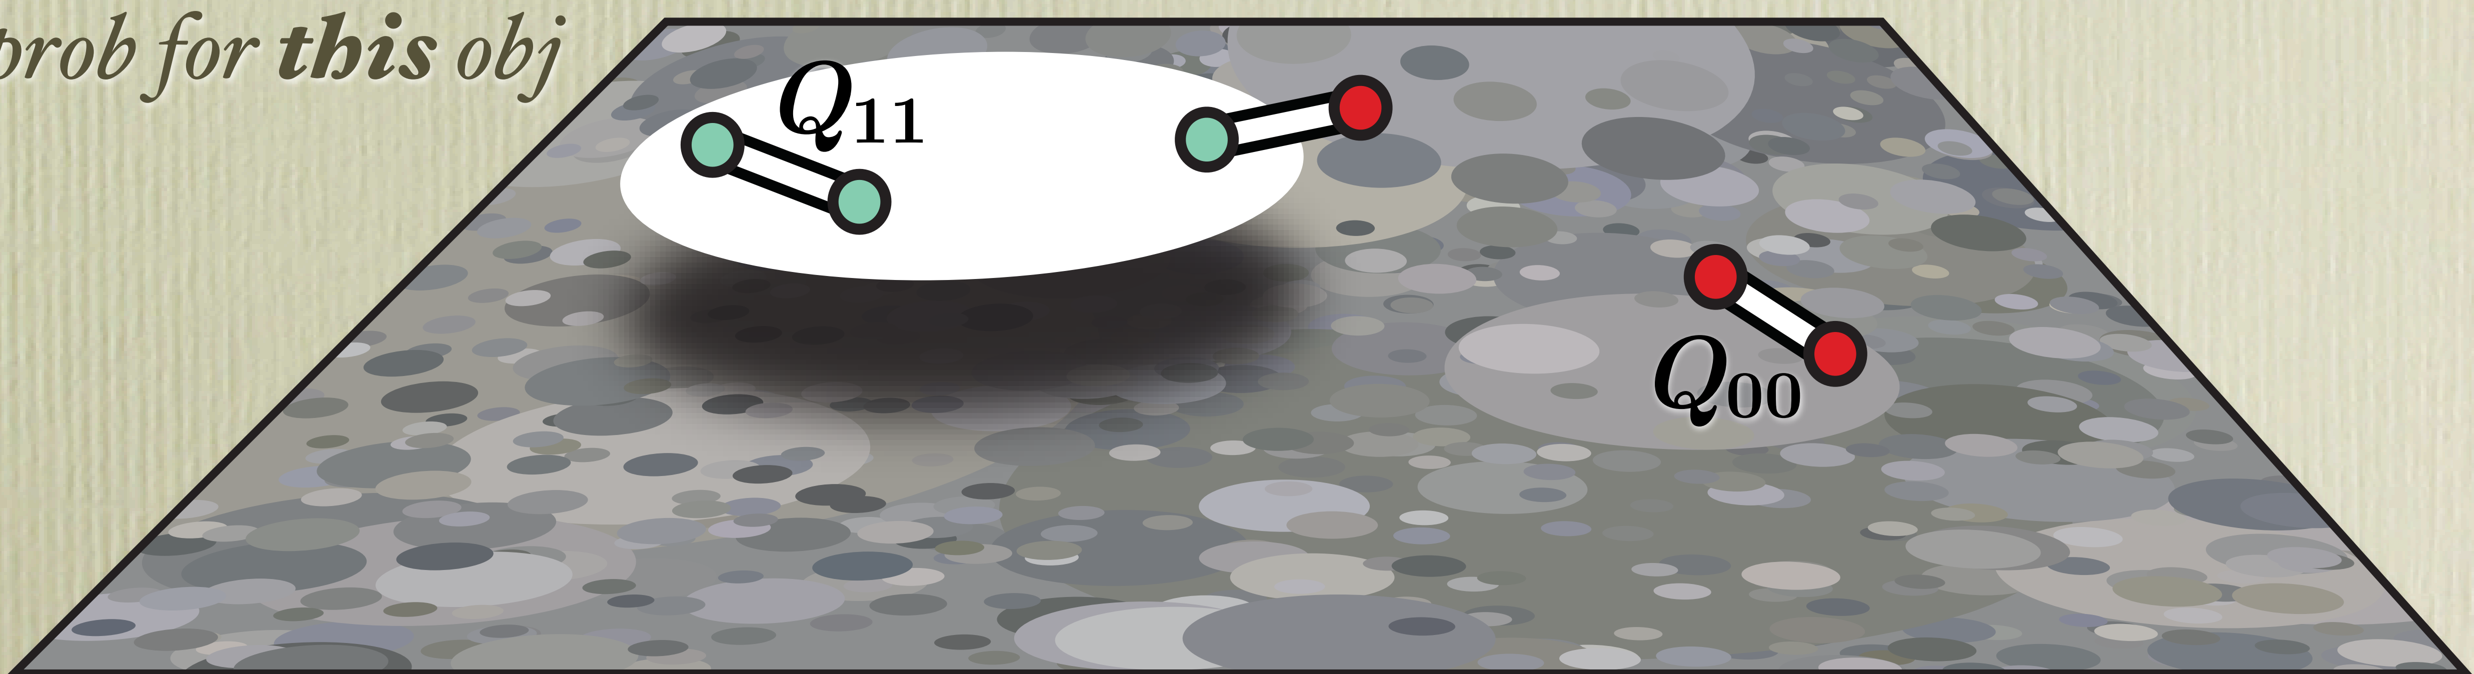

Ruderman 1997

- What is probability of being in same object? Calculate recursively.

$$P_{\{x,y\}} = Q_{11} + Q_{00} P_{\{x,y\}}$$
$$= \frac{Q_{11}}{1 - Q_{00}}$$

P = *prob for **any** obj*

Q = *prob for **this** obj*

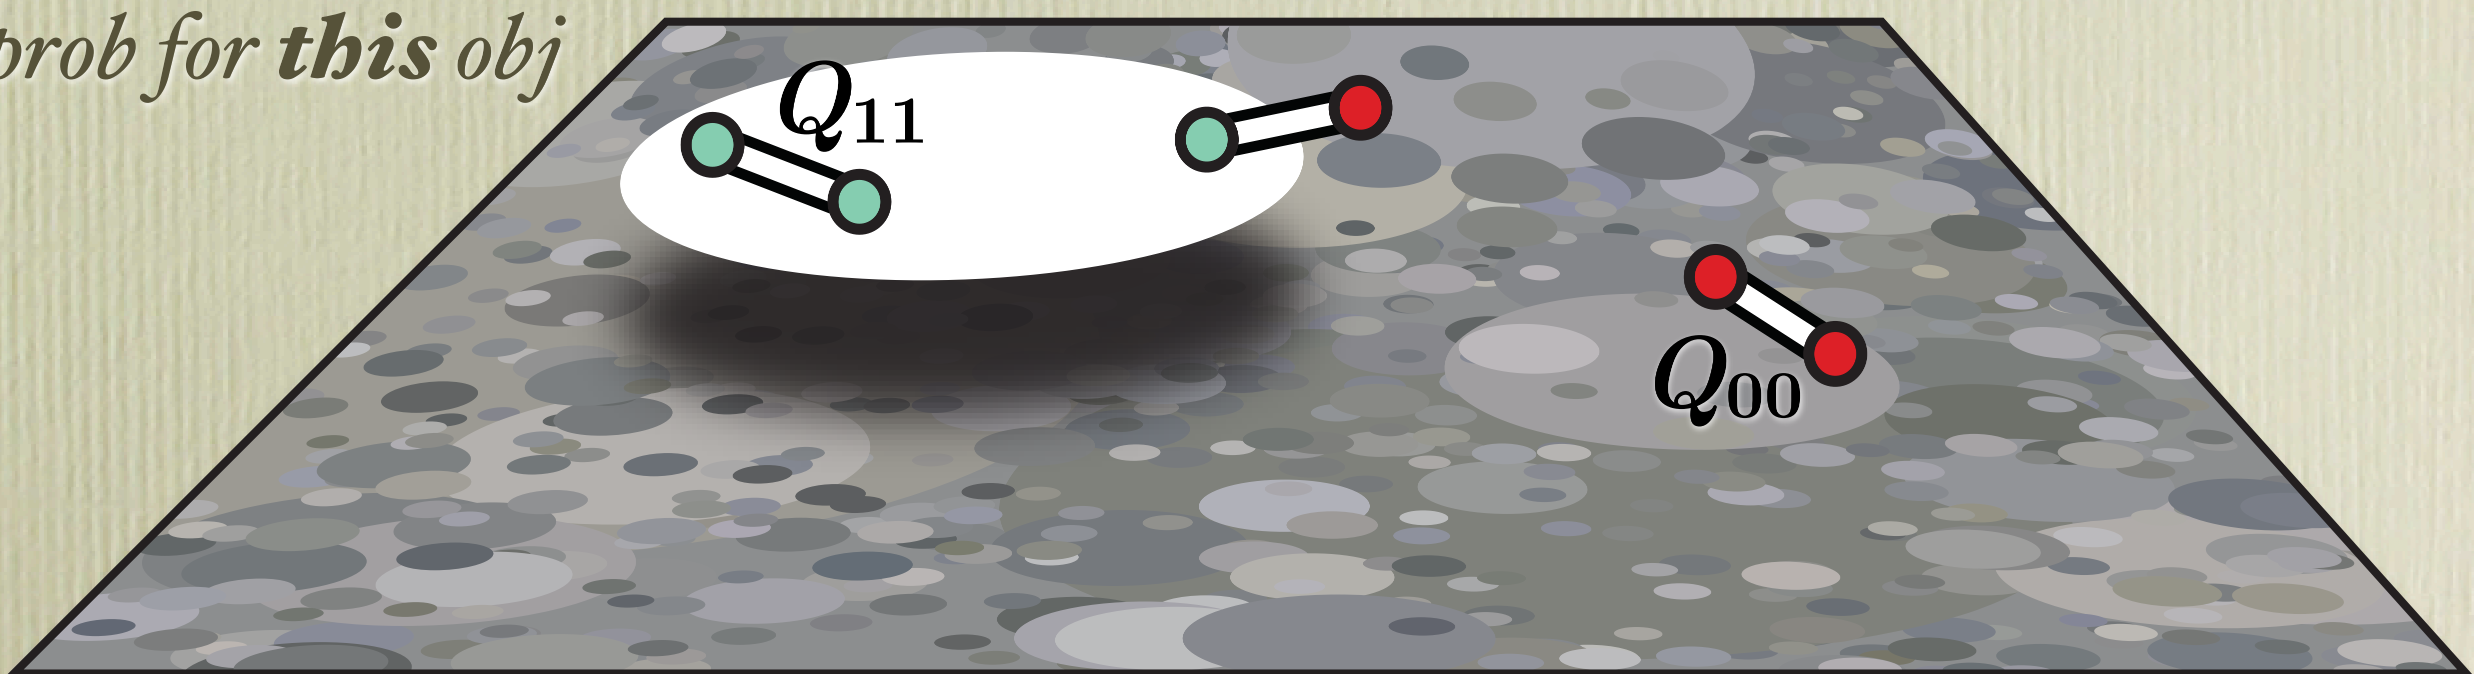

Ruderman 1997

- What is probability of being in same object? Calculate recursively.

$$P_{\{x,y\}} = Q_{11} + Q_{00}P_{\{x,y\}}$$
$$= \frac{Q_{11}}{1 - Q_{00}}$$

P = *prob for* ***any*** *obj*

Q = *prob for* ***this*** *obj*

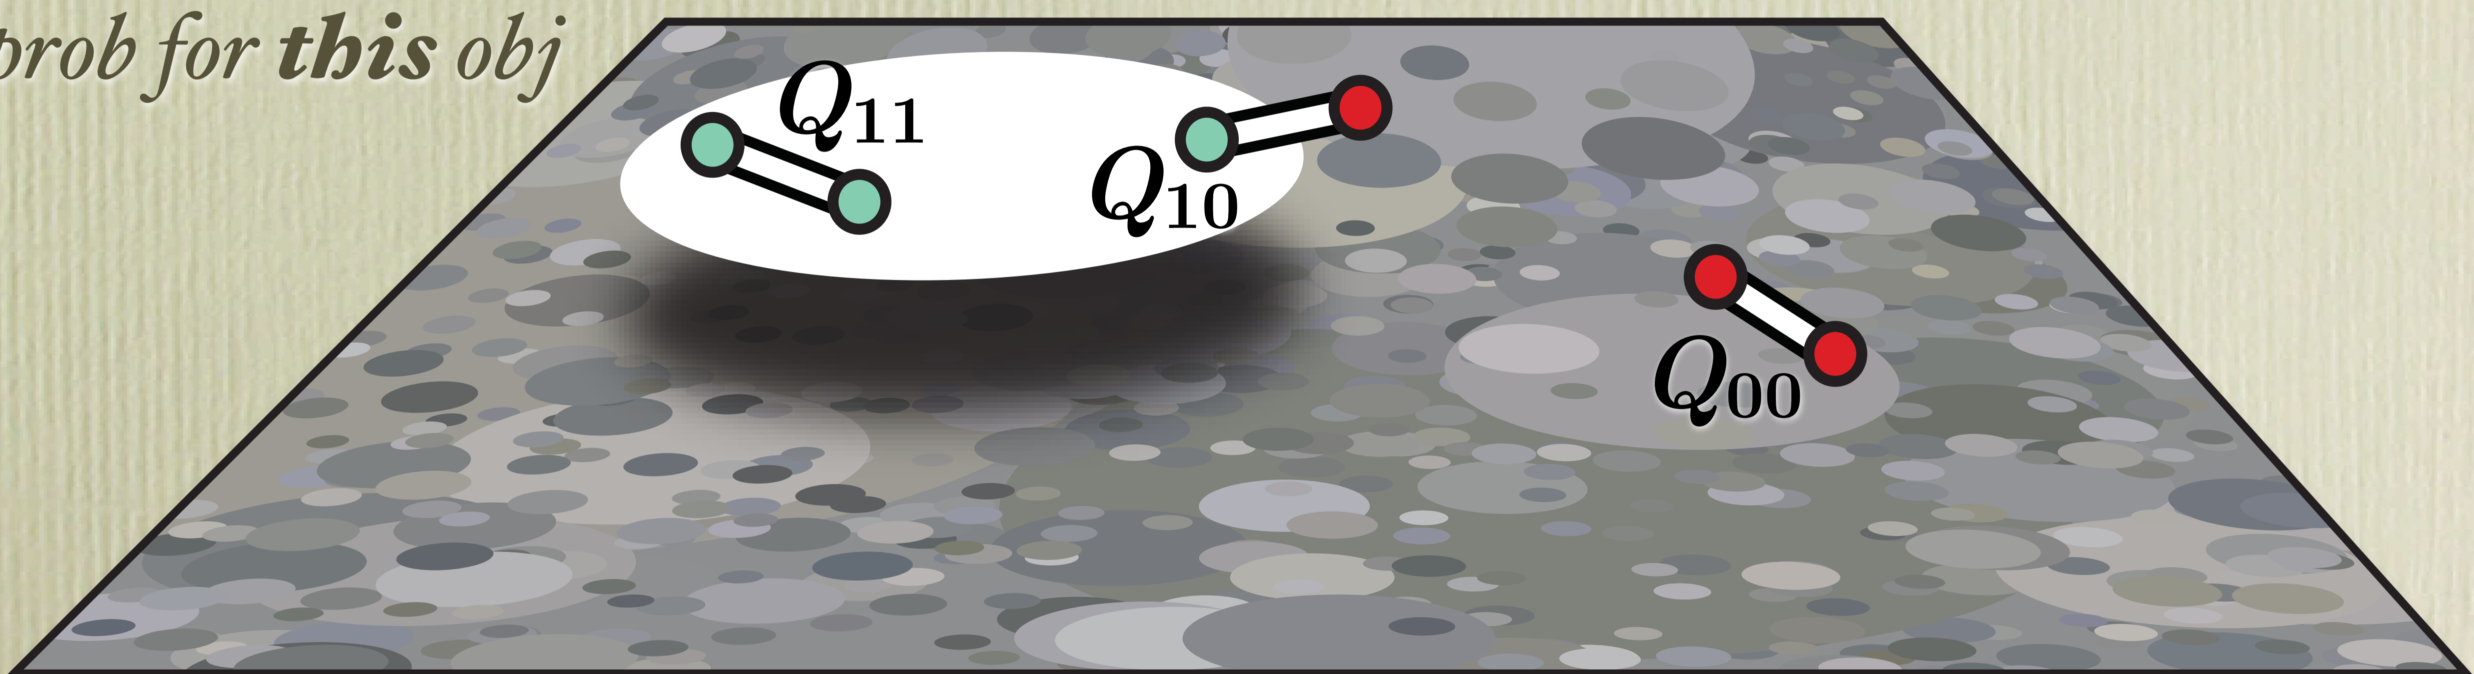

Ruderman 1997

$$\pi = \{\{u, v, w\}, \{x\}, \{y, z\}\}$$

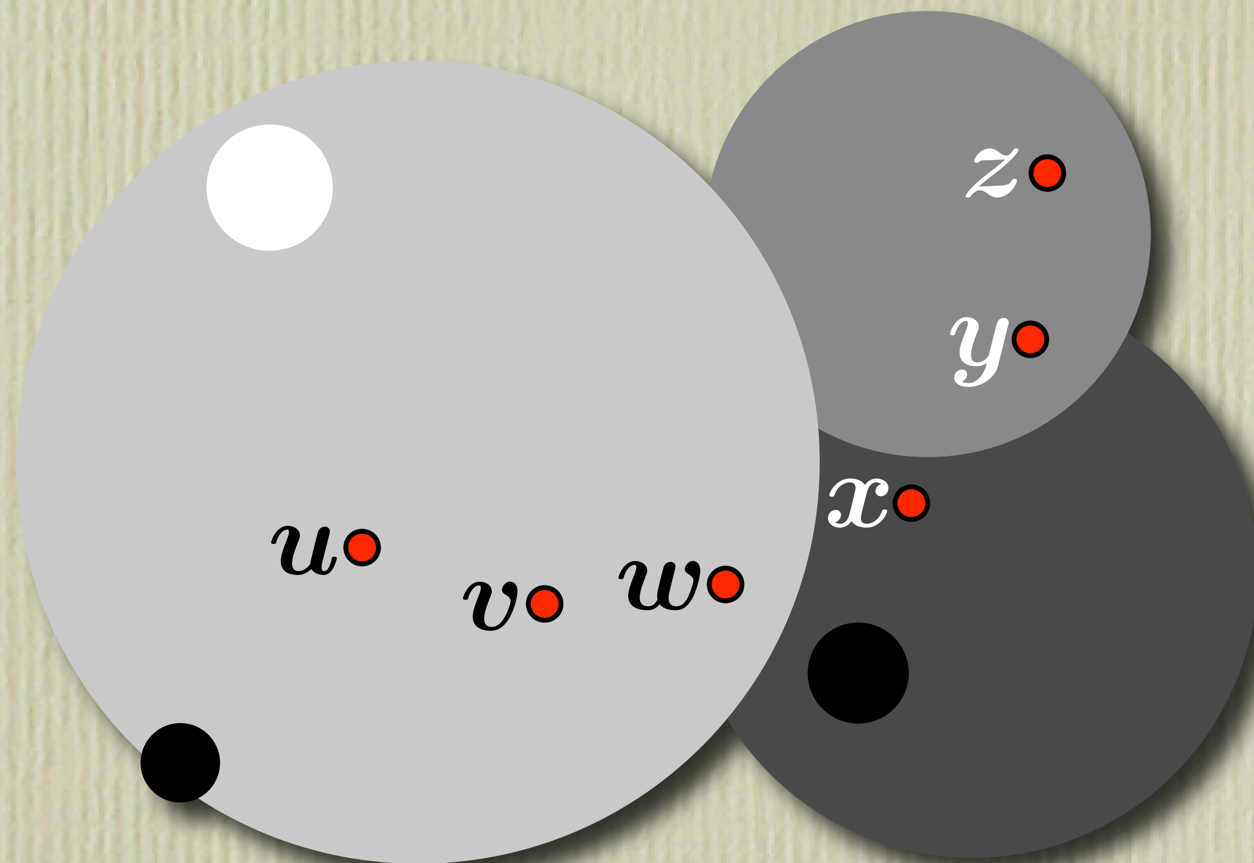

$$P_\pi = Q_{111000} P_{\{x\},\{y,z\}}$$

$$+ Q_{000100} P_{\{u,v,w\},\{y,z\}}$$

$$+ Q_{000011} P_{\{u,v,w\},\{x\}}$$

$$+ Q_{000000} P_\pi$$

$$\pi = \{\{u, v, w\}, \{x\}, \{y, z\}\}$$

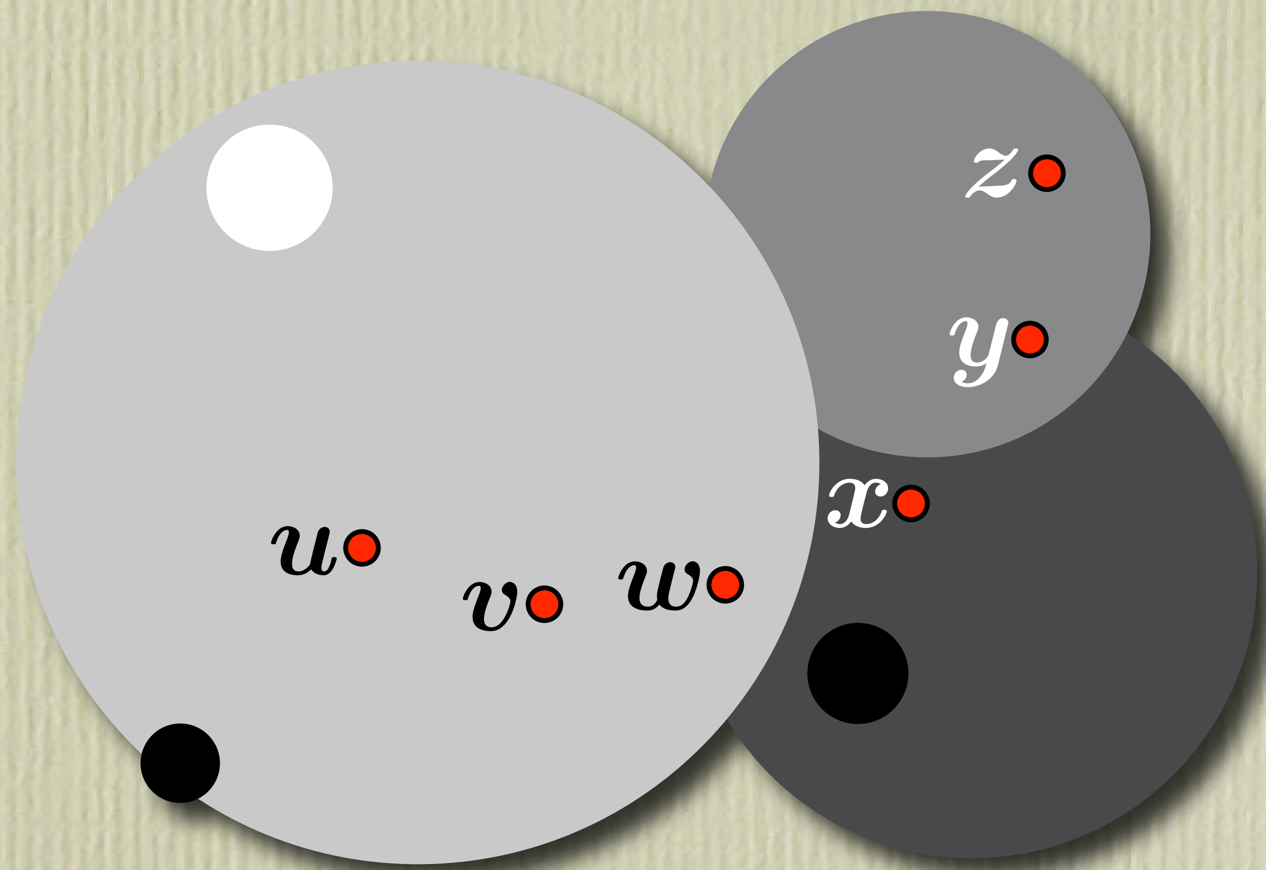

$$
\begin{aligned}
P_\pi = \; & Q_{111000} P_{\{x\},\{y,z\}} \\
& + Q_{000100} P_{\{u,v,w\},\{y,z\}} \\
& + Q_{000011} P_{\{u,v,w\},\{x\}} \\
& + Q_{000000} P_\pi
\end{aligned}
$$

In general:

$$P_\pi = \frac{1}{1 - Q_{\sigma_0}} \sum_{j=1}^{J(\pi)} Q_{\sigma_j} P_{\pi_{\setminus j}}$$

$J(\pi)$	# of sets in π
π_j	*j*th set in π
$\pi_{\setminus j}$	π without *j*th set
σ_j	$000\,\underbrace{111}_{\pi_j}\,0$

Geometry gives Q_σ

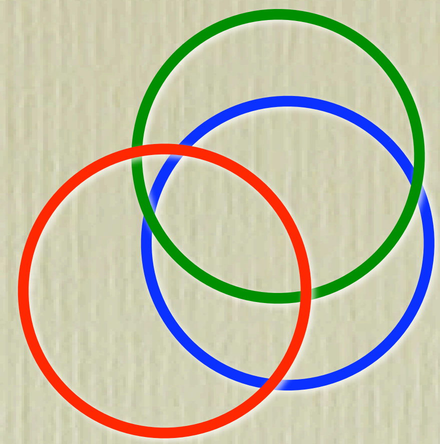

↓

Recursion gives P_π

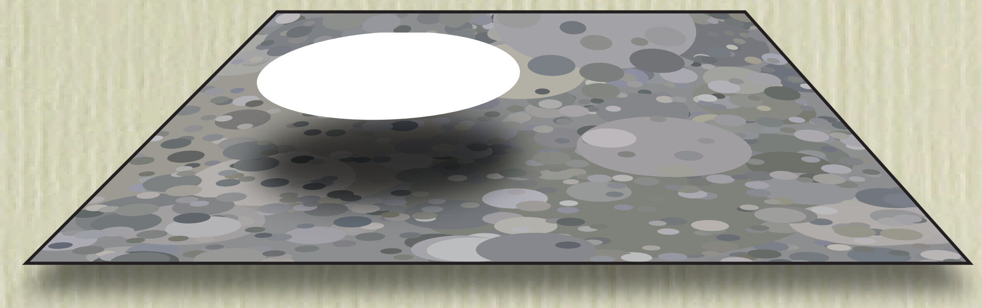

↓

Weighted sum gives $P(I)$

$$P(I) = \sum_\pi P_\pi P(I|\pi)$$

↓

Joint feature probabilities!

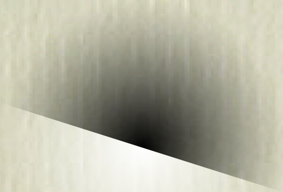

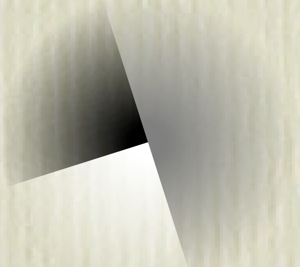

How many points?

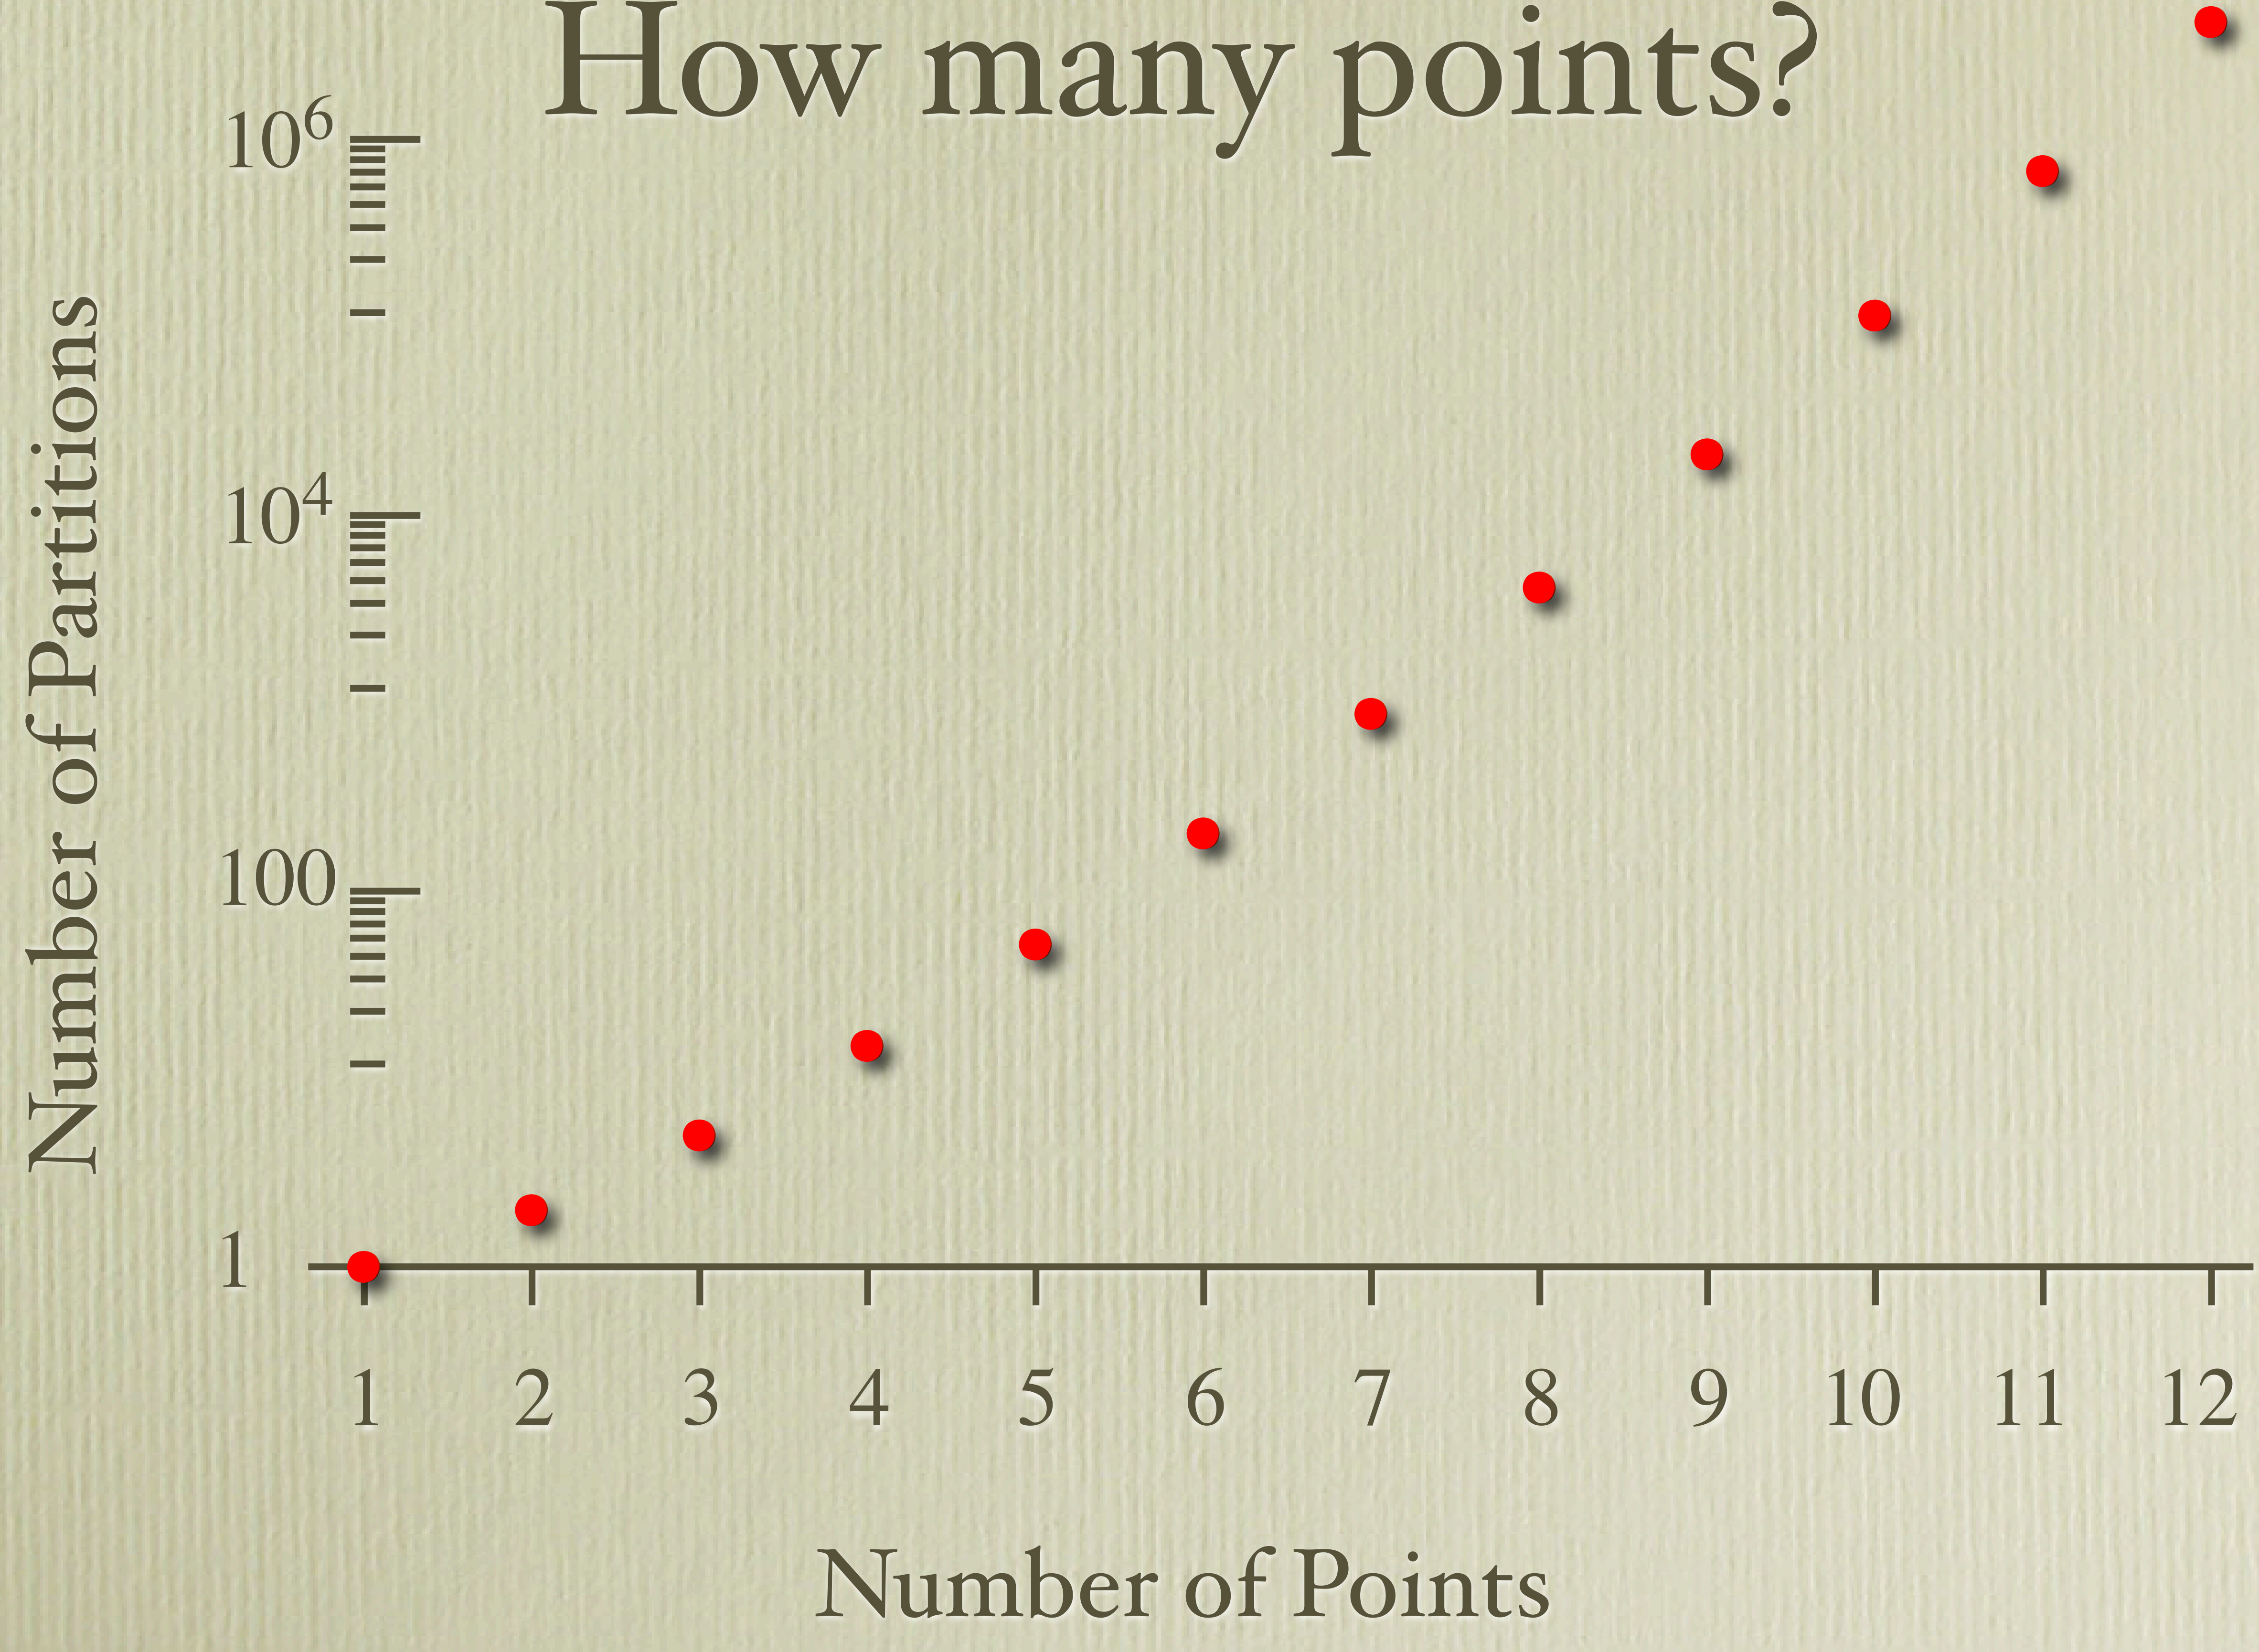

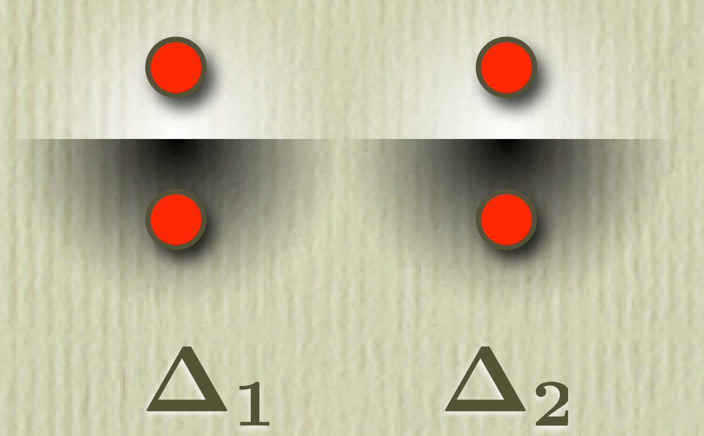
Δ_1 Δ_2

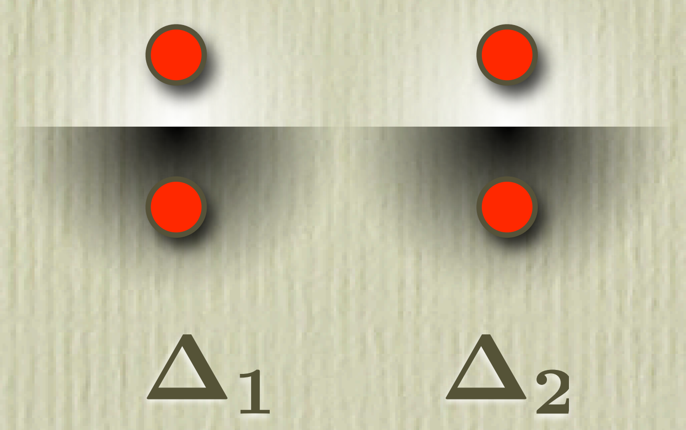

$$\log \frac{P(\Delta_1, \Delta_2)}{P(\Delta_1)P(\Delta_2)}$$

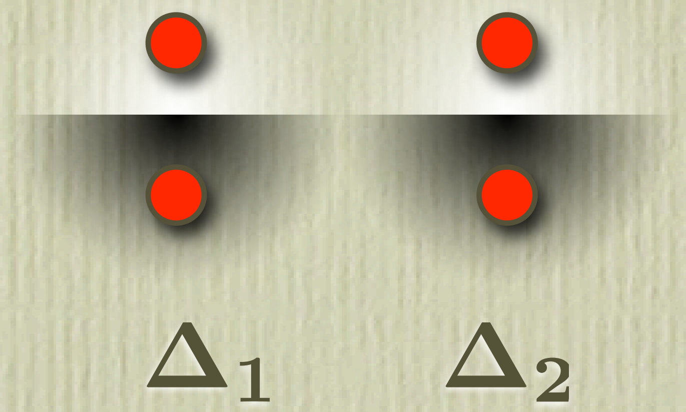
Δ_1
Δ_2

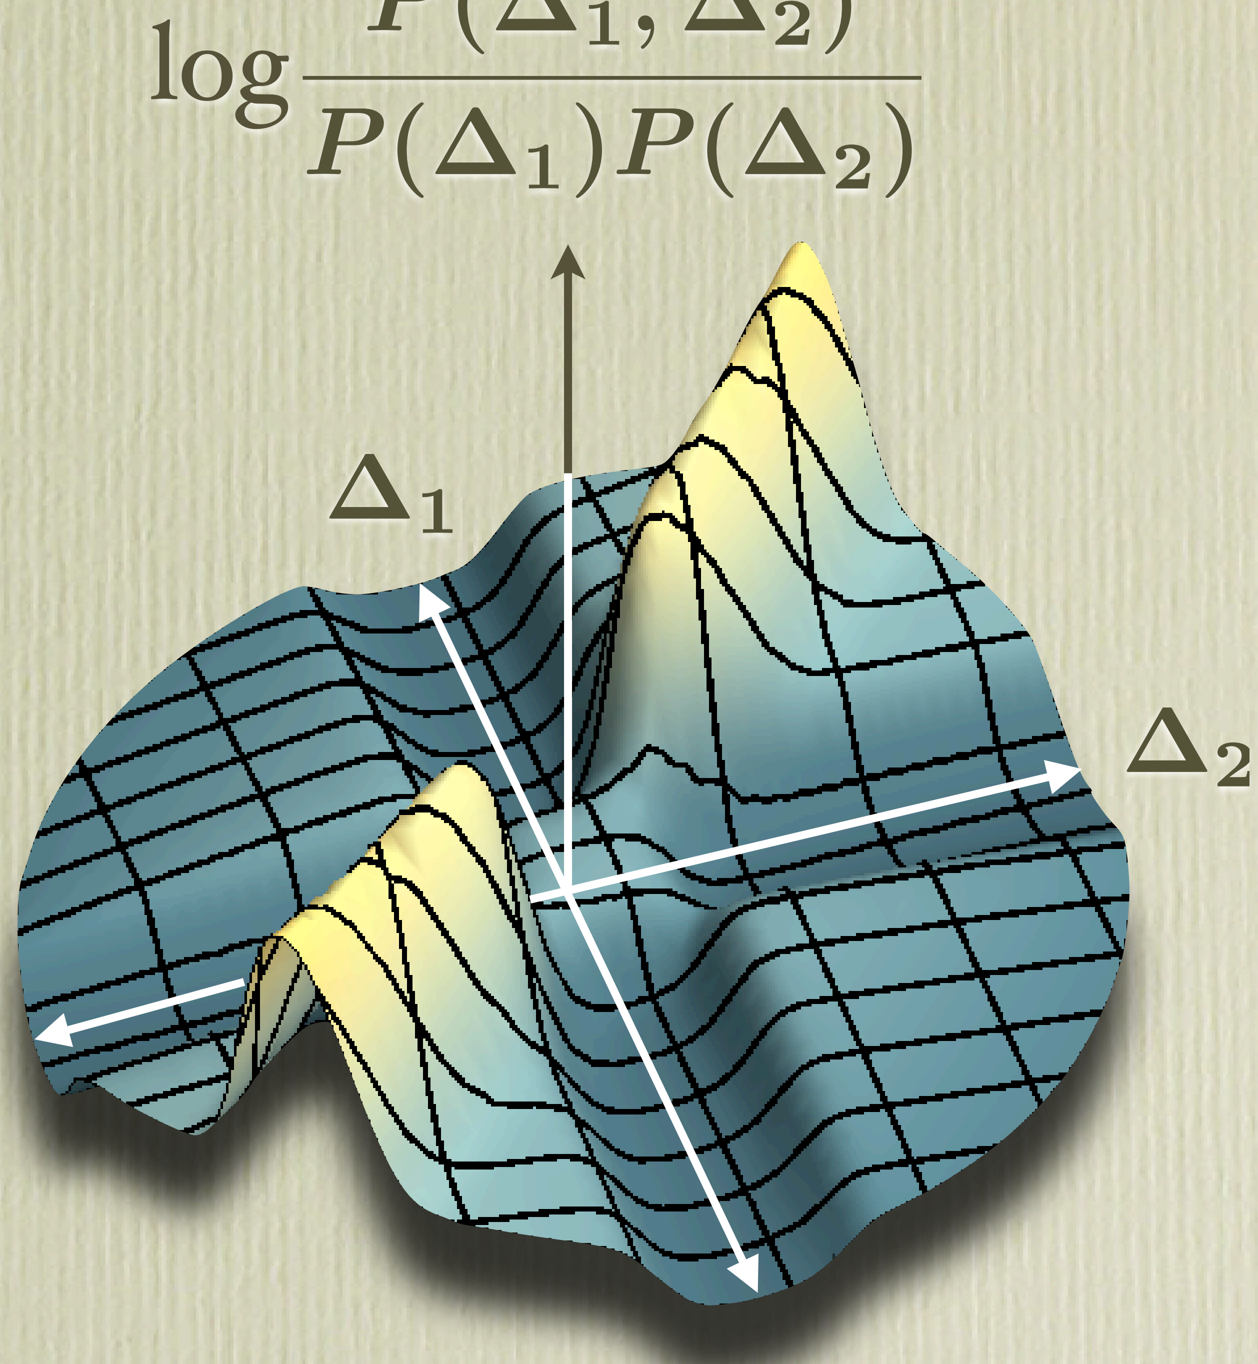
$\log \frac{P(\Delta_1, \Delta_2)}{P(\Delta_1)P(\Delta_2)}$
Δ_1
Δ_2

Thanks to...

- Ken Miller
- Larry Abbott
- Taro Toyoizumi
- Vladimir Itskov

the columbia neurotheory center

- the Swartz foundation

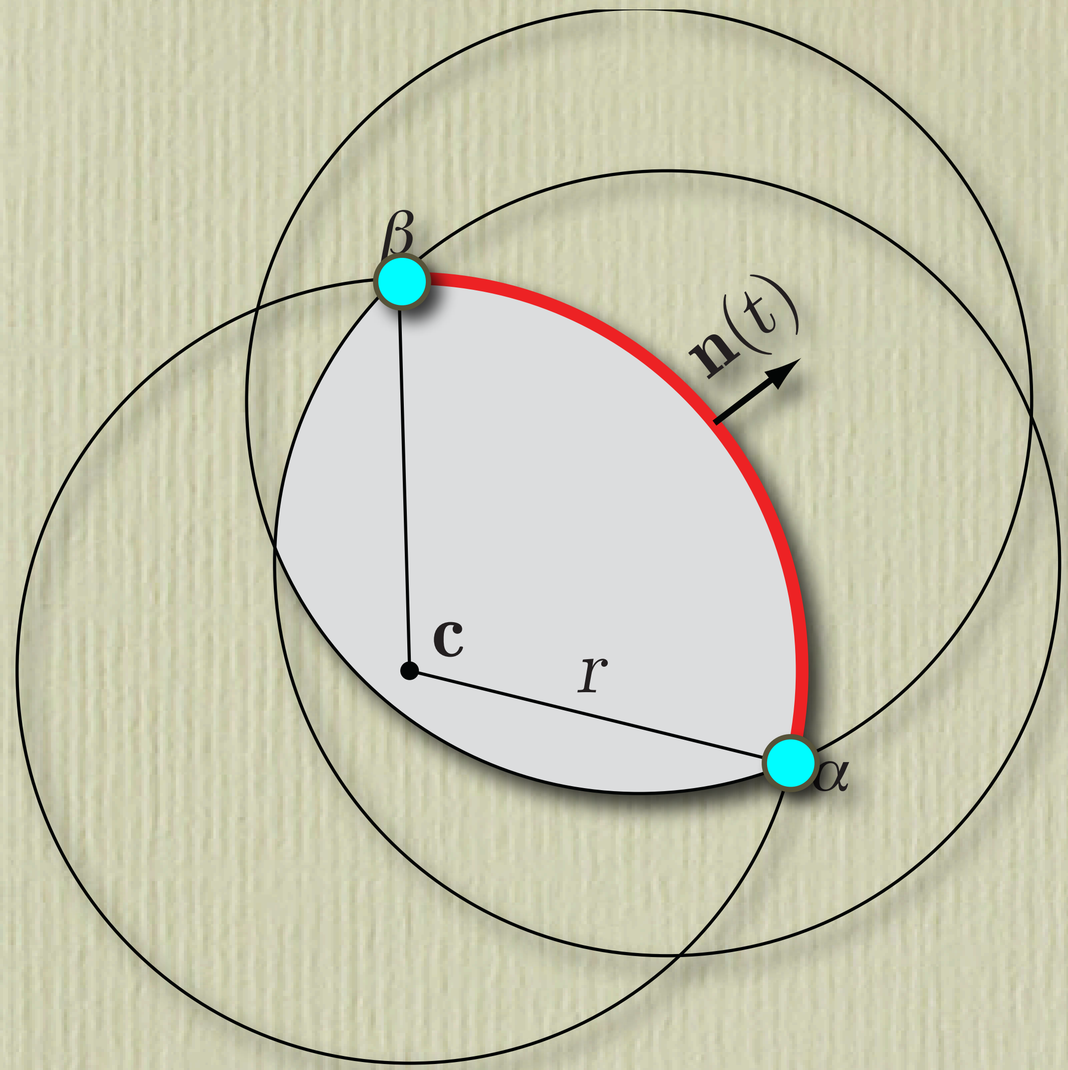

- Calculate areas using the Divergence Theorem

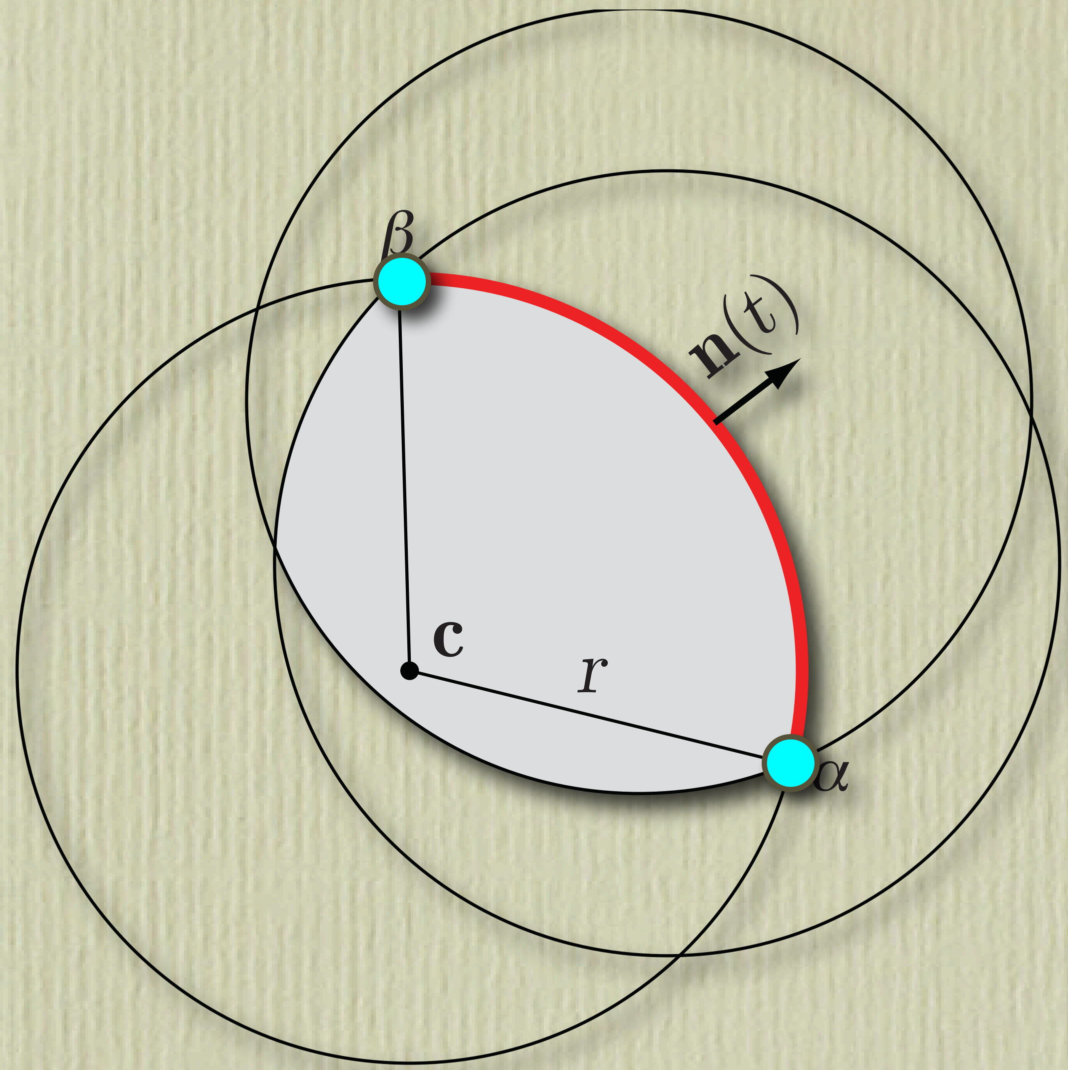

- Calculate areas using the Divergence Theorem

- Average over shape ensemble

$$Q_\sigma = \int dr\, P(r)\, Q_\sigma(r)$$

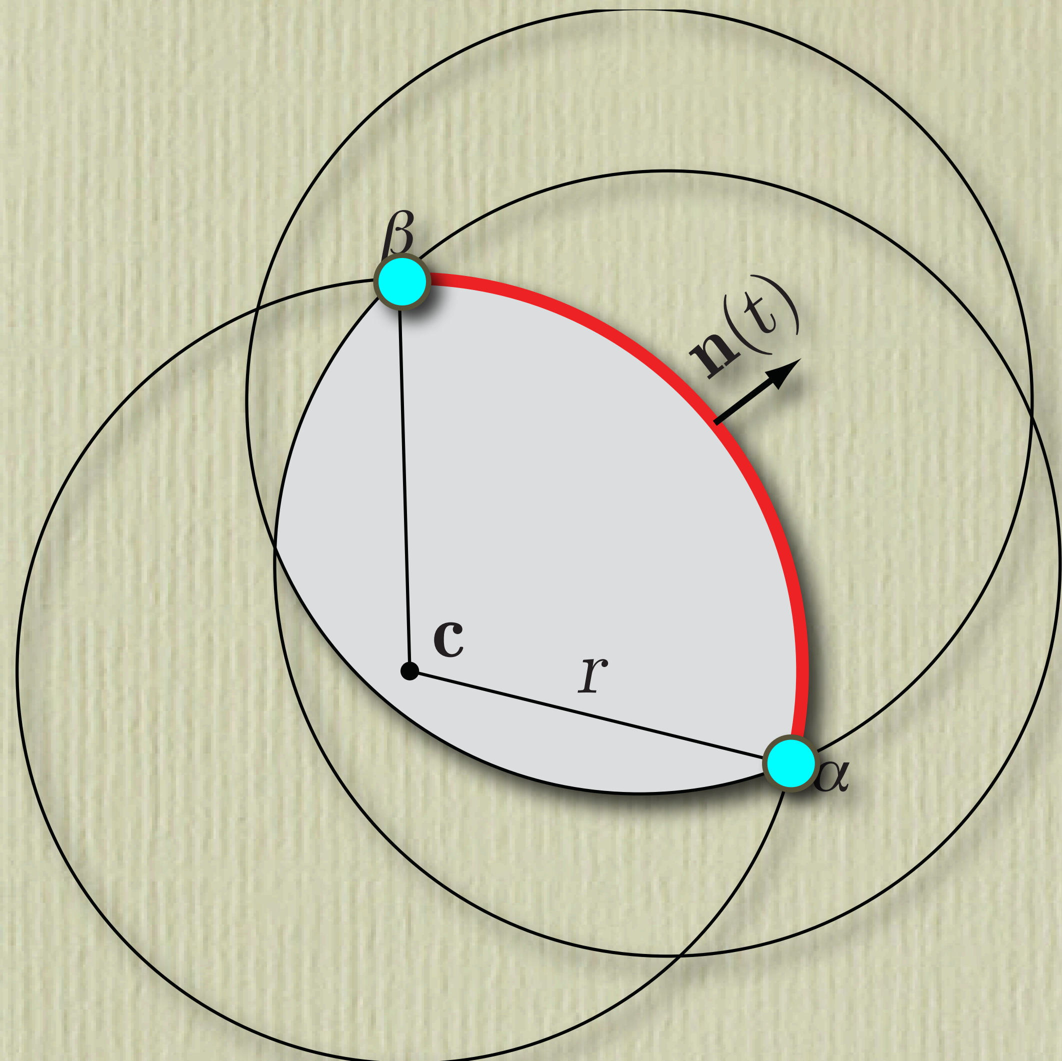

- Calculate areas using the Divergence Theorem

- Average over shape ensemble

$$Q_{\sigma} = \int dr\, P(r)\, Q_{\sigma}(r)$$

- Graph changes with radius