exact feature probabilities in images with occlusion

xaq pitkow

(xaq@neurotheory.columbia.edu)

june 2010, swartz meeting @ yale

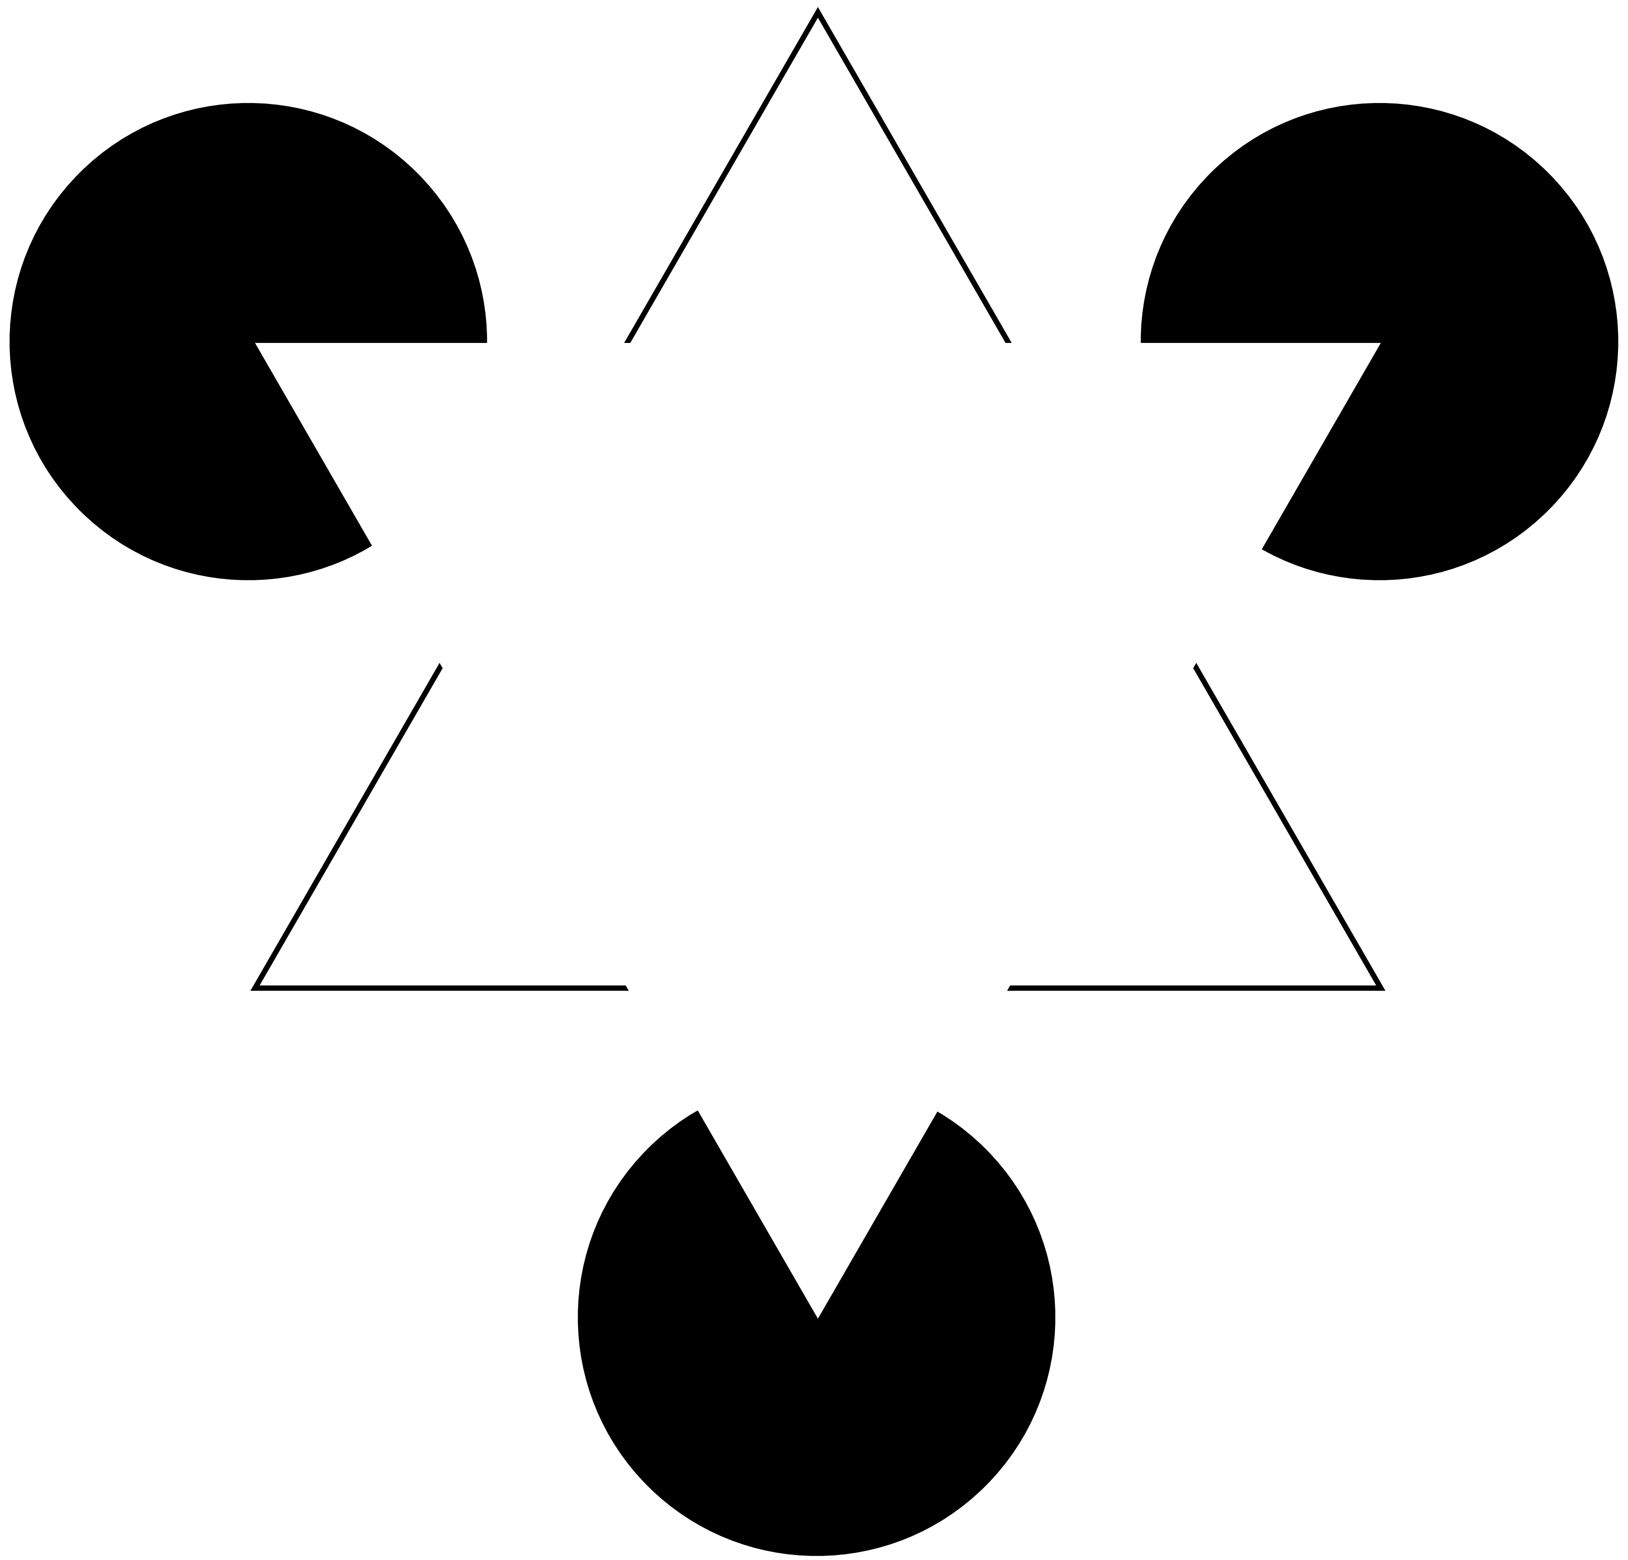

Kanizsa triangle

Linear

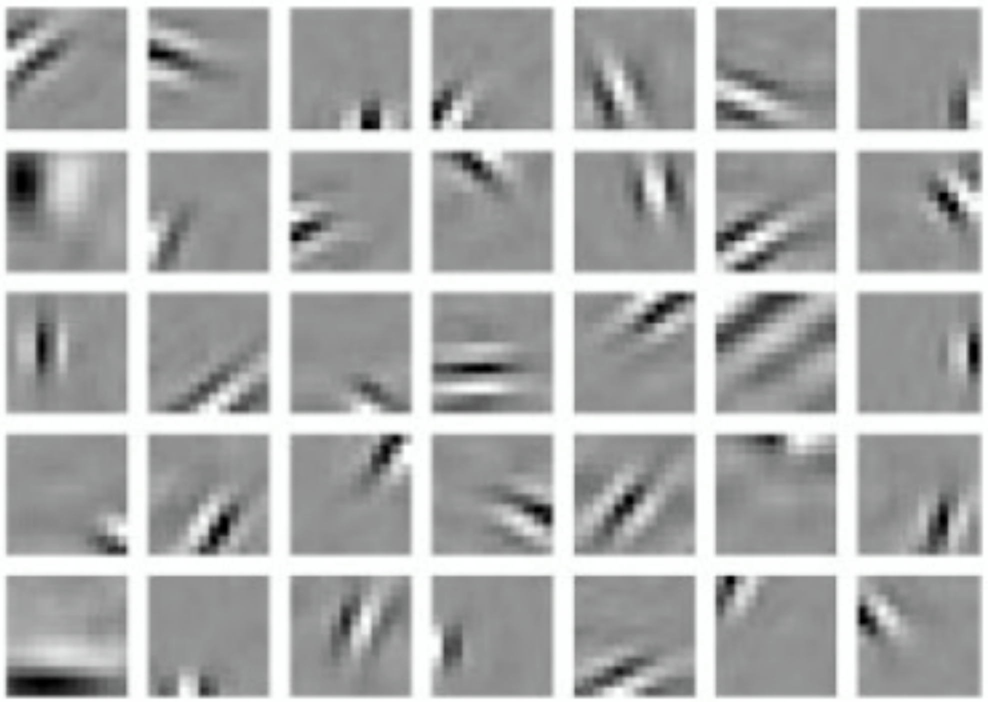

olshausen & field 1996
bell & sejnowski 1997

non-Linear

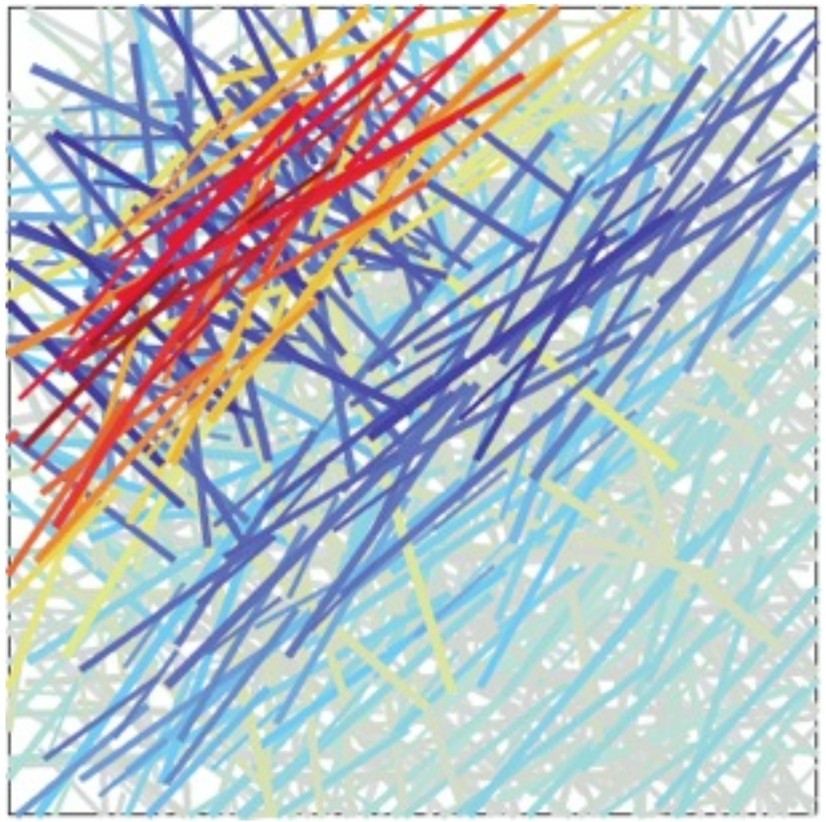

karklin & lewicki 2008
wainwright & simoncelli 2000

Dead Leaves Model

Matheron 1975, Ruderman 1997, Lee, Mumford, Huang 2001

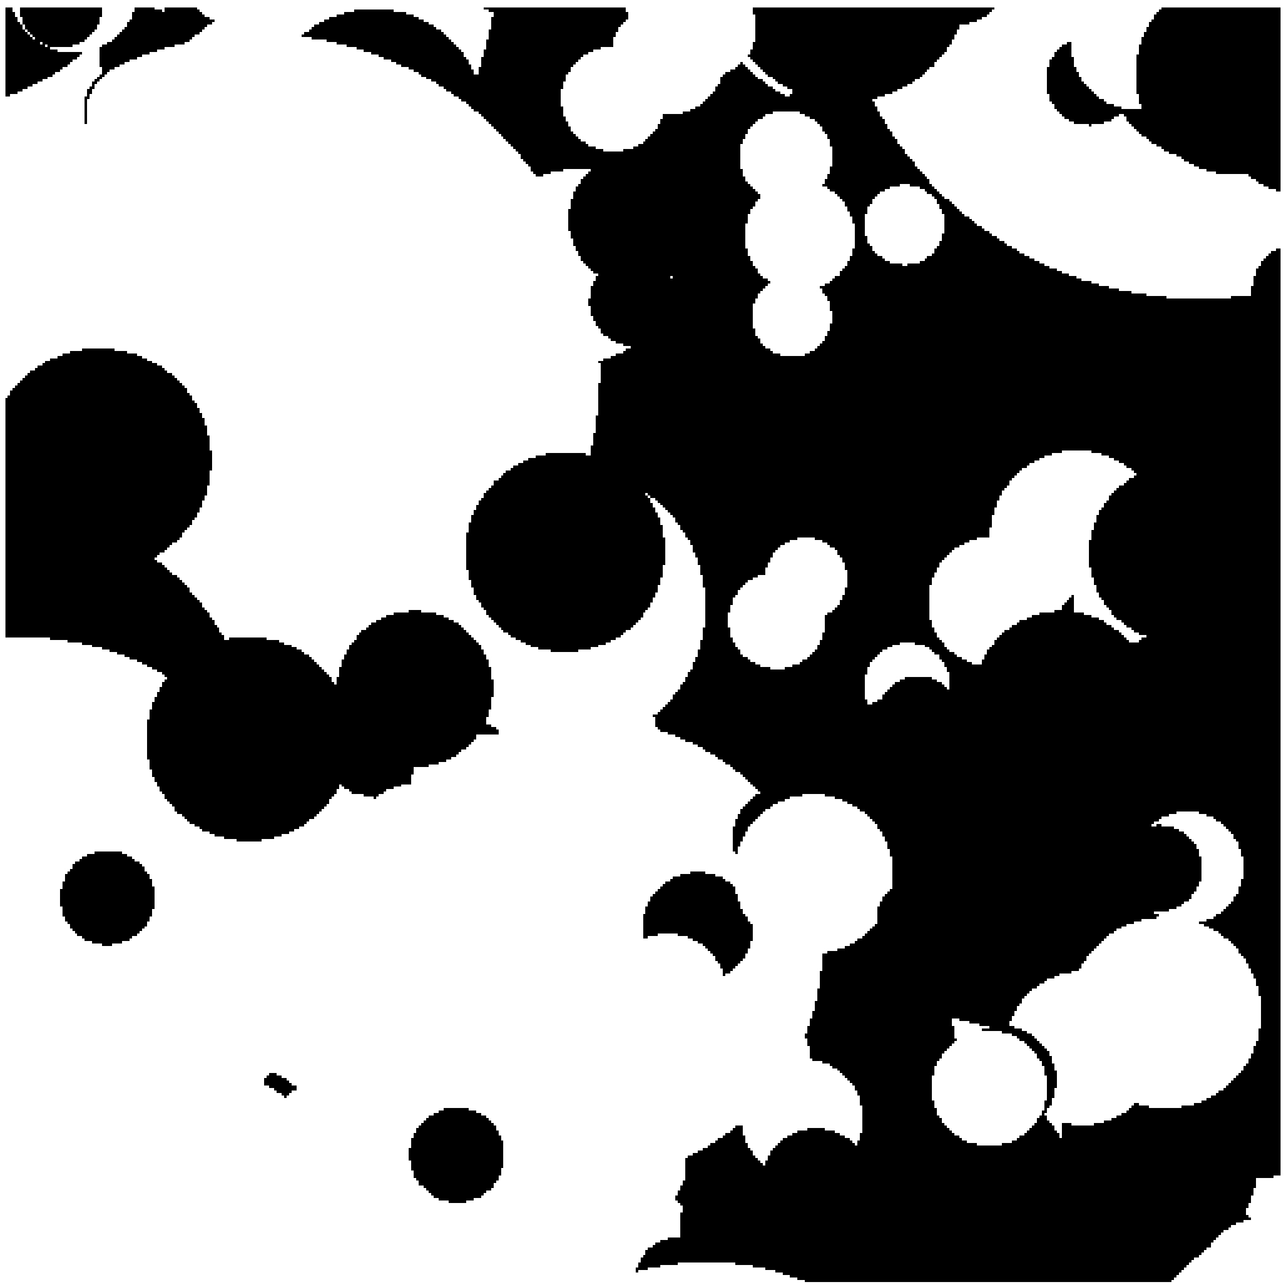

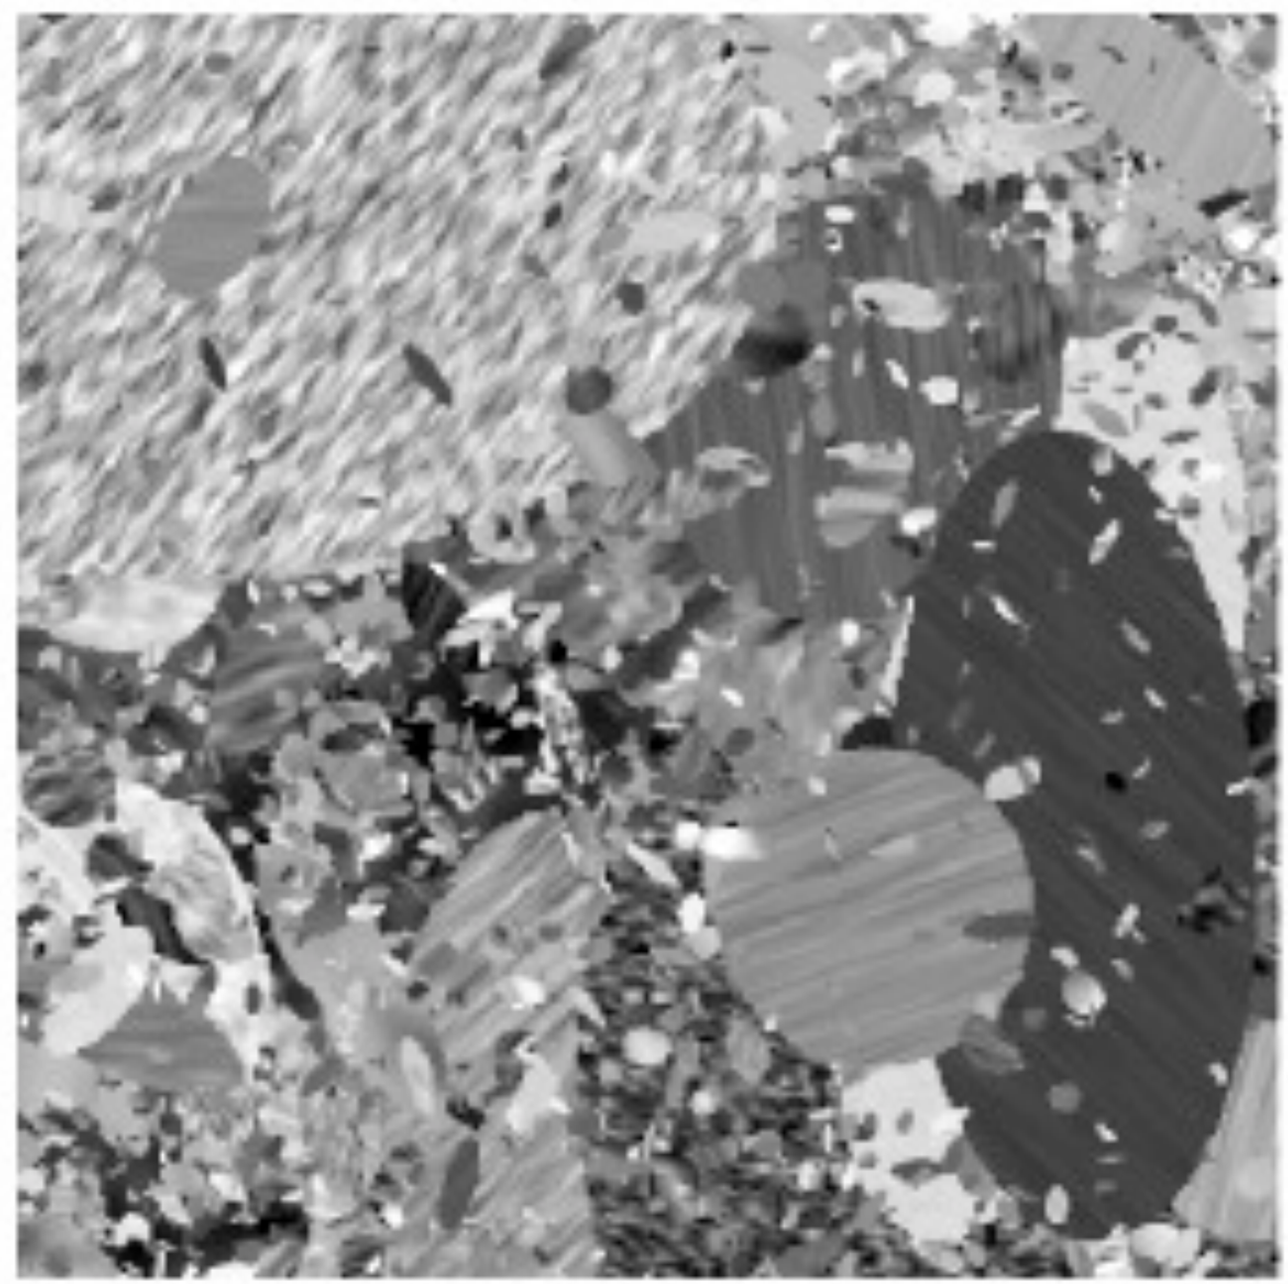

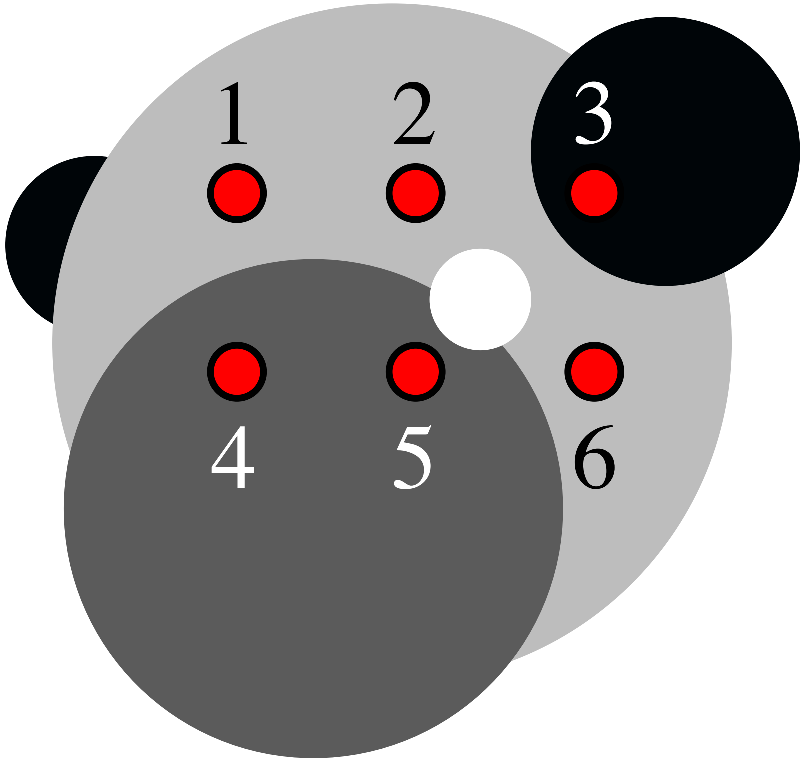

Object membership function:

which pixels are in which object?

$$\mathfrak{מ} = \{\{1,2,6\},\{3\},\{4,5\}\}$$

Pitkow, in review

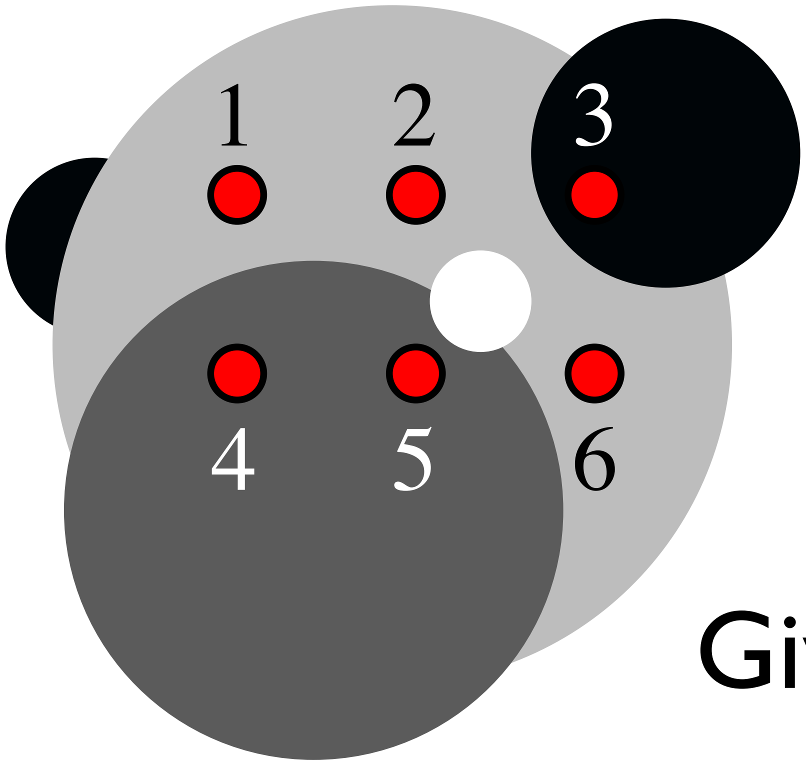

Object membership function:

which pixels are in which object?

$$\text{מ} = \{\{1, 2, 6\}, \{3\}, \{4, 5\}\}$$

Given מ, the joint pixel prob factorizes

$$P(\mathbf{I}|\text{מ}) = \prod_{n=1}^{|\text{מ}|} P(\mathbf{I}_{\text{מ}_n})$$

Pitkow, in review

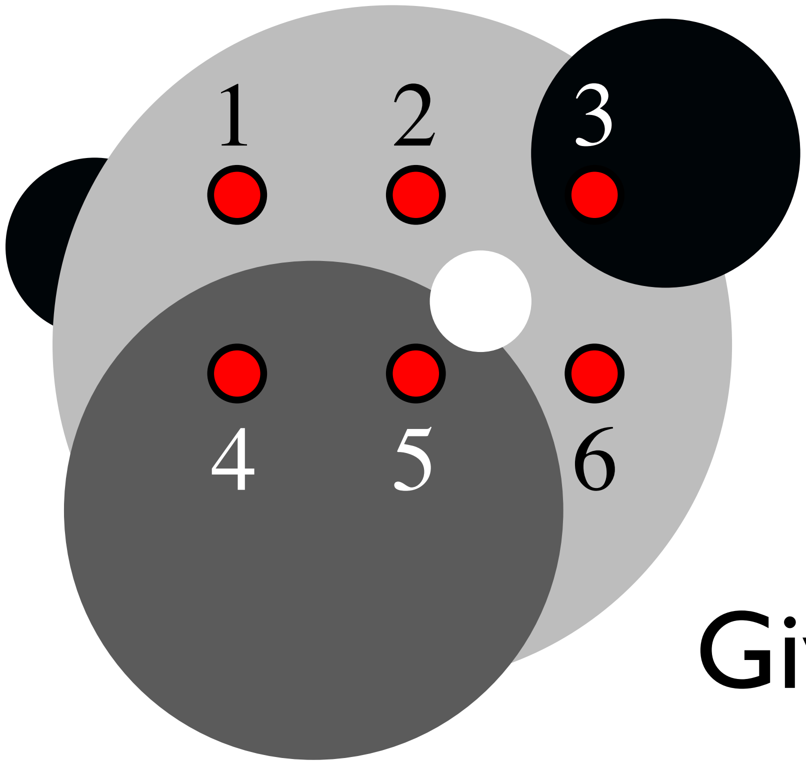

Object membership function:

which pixels are in which object?

$$מ = \{\{1,2,6\},\{3\},\{4,5\}\}$$

Given מ, the joint pixel prob factorizes

$$P(\mathbf{I}|מ) = \prod_{n=1}^{|מ|} P(\mathbf{I}_{מ_n})$$

$$= P(I_1, I_2, I_6) P(I_3) P(I_4, I_5)$$

Pitkow, in review

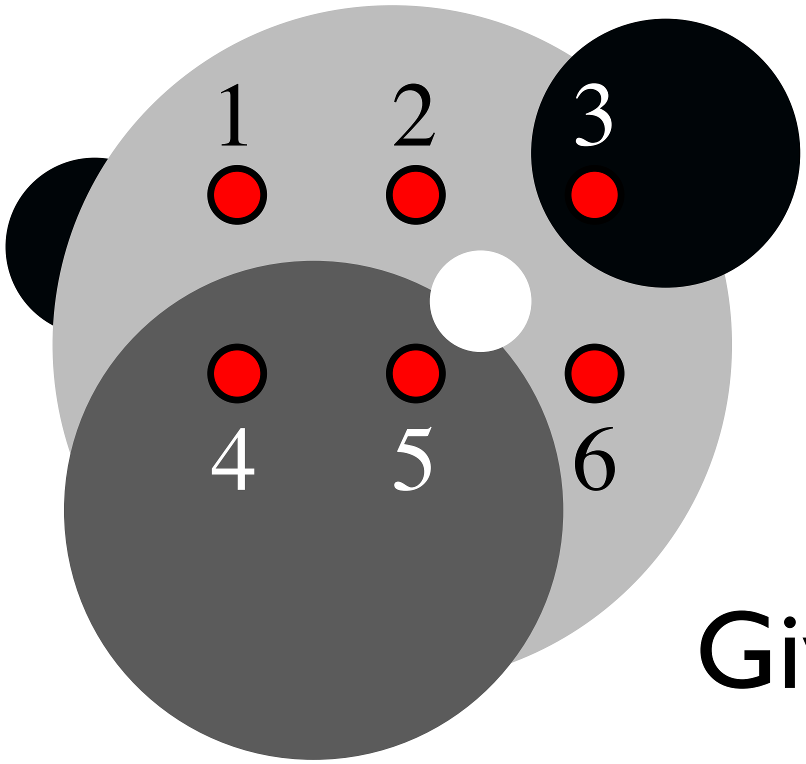

Object membership function:

which pixels are in which object?

$$מ = \{\{1, 2, 6\}, \{3\}, \{4, 5\}\}$$

Given מ, the joint pixel prob factorizes

$$P(\mathbf{I}|מ) = \prod_{n=1}^{|מ|} P(\mathbf{I}_{מ_n})$$

Full joint distribution is a mixture of different correlated distributions

$$P(\mathbf{I}) = \sum_{מ} P(\mathbf{I}|מ) P(מ)$$

Pitkow, in review

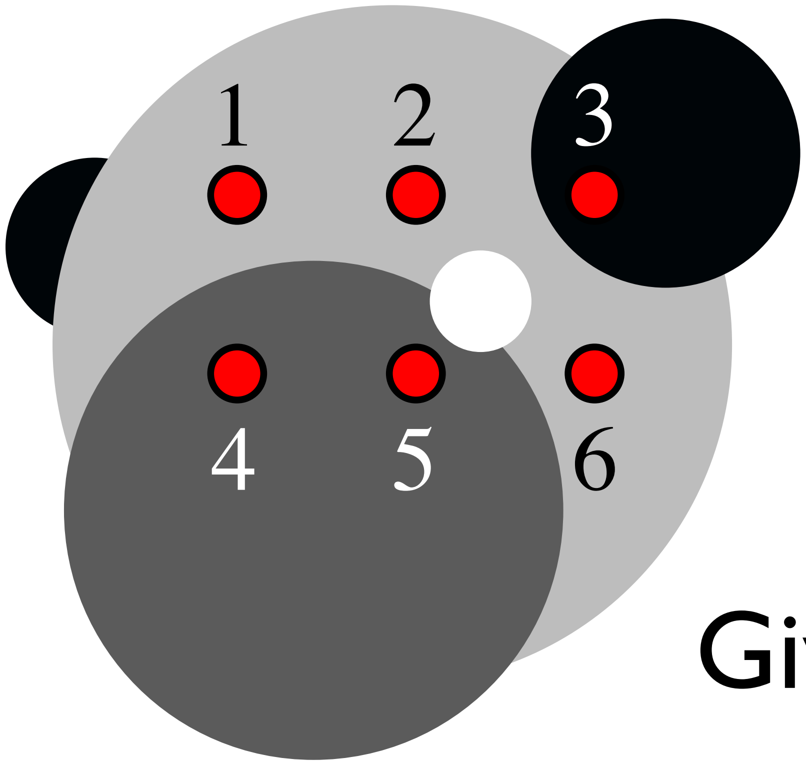

Object membership function:

which pixels are in which object?

$$\text{מ} = \{\{1,2,6\},\{3\},\{4,5\}\}$$

Given מ, the joint pixel prob factorizes

$$P(\mathbf{I}|\text{מ}) = \prod_{n=1}^{|\text{מ}|} P(\mathbf{I}_{\text{מ}_n})$$

Full joint distribution is a mixture of different correlated distributions

$$P(\mathbf{I}) = \sum_{\text{מ}} P(\mathbf{I}|\text{מ})P(\text{מ})$$

texture geometry

Pitkow, in review

Geometry

$P(\text{מ})$ can be calculated exactly for some shape ensembles

size

eccentricity

orientation

Texture

Correlated gaussians:

$$P(\mathbf{I}|\text{מ}) \propto \exp\left(-\frac{1}{2}\mathbf{I}^\top C_{\mathbf{I}|\text{מ}}^{-1}\mathbf{I}\right)$$

e.g. for uniform color + white noise texture:

$$C_{\mathbf{I}|\text{מ}} = \Xi_1 \mathbf{1} + \Xi_0 \begin{pmatrix} 1 & 1 & & & & 1 \\ 1 & 1 & & & & 1 \\ & & 1 & & & \\ & & & 1 & 1 & \\ & & & 1 & 1 & \\ 1 & 1 & & & & 1 \end{pmatrix}$$

- Applications: Natural Scene Statistics
  - marginal wavelet distributions
  - joint wavelet distributions
  - contour statistics

van hateren 1998

raw image

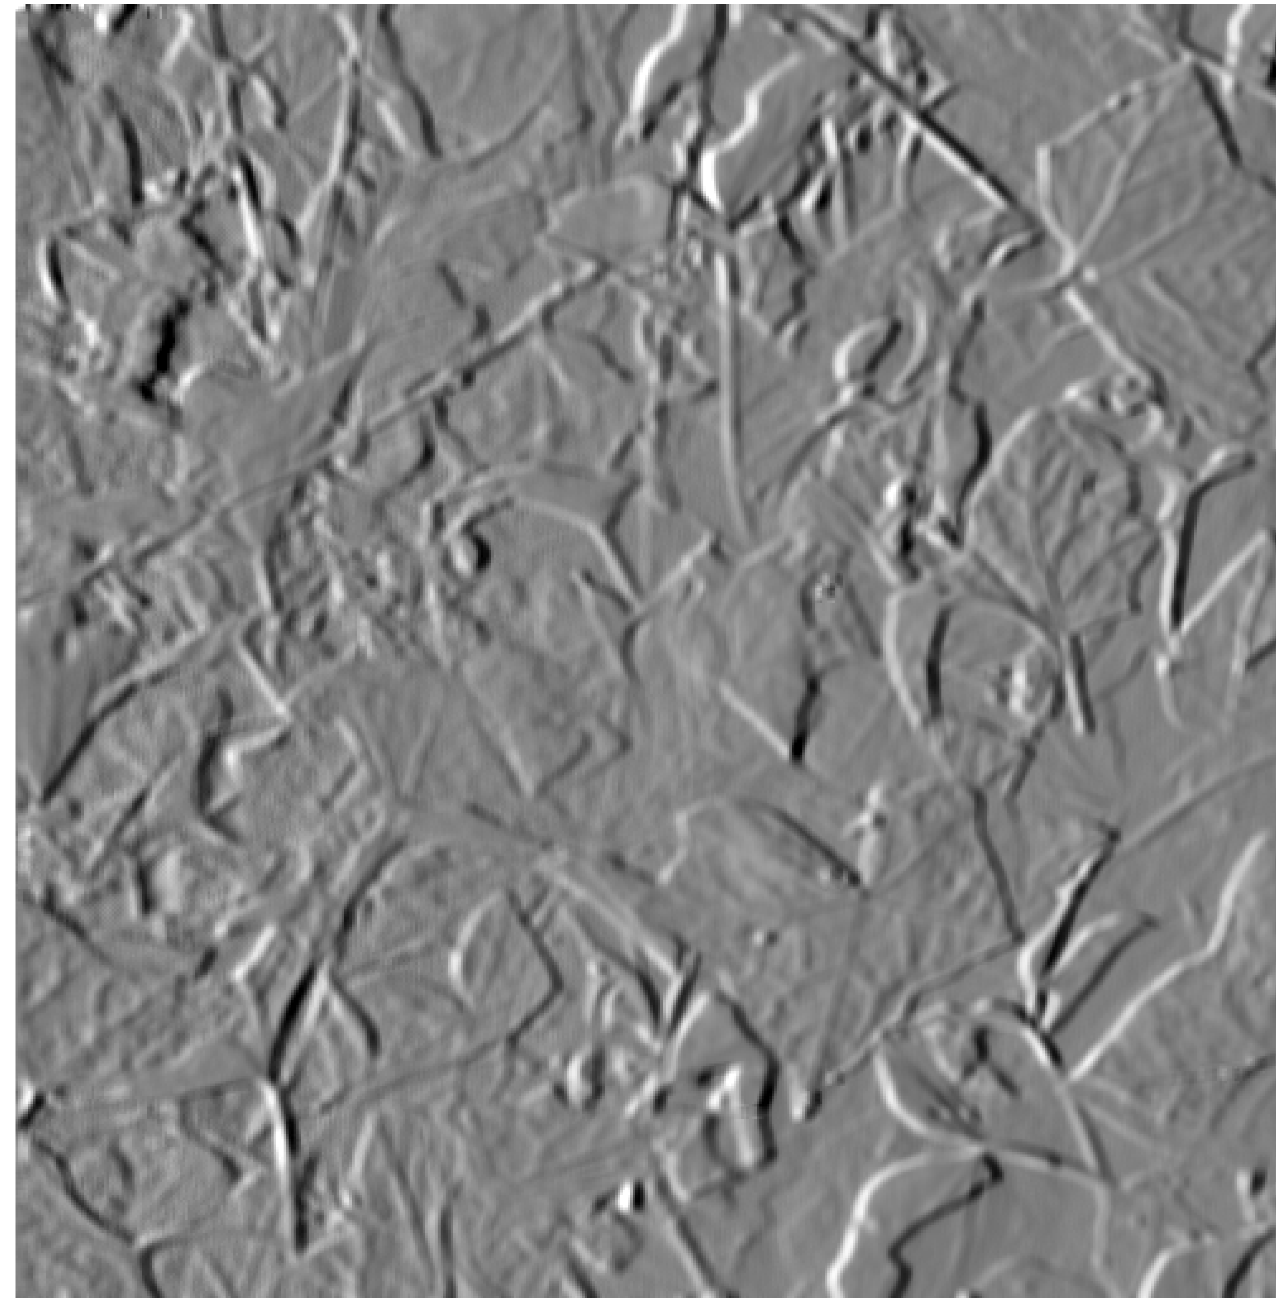

filtered image

f 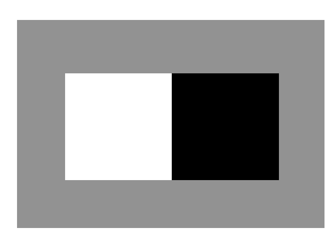 Haar wavelet features

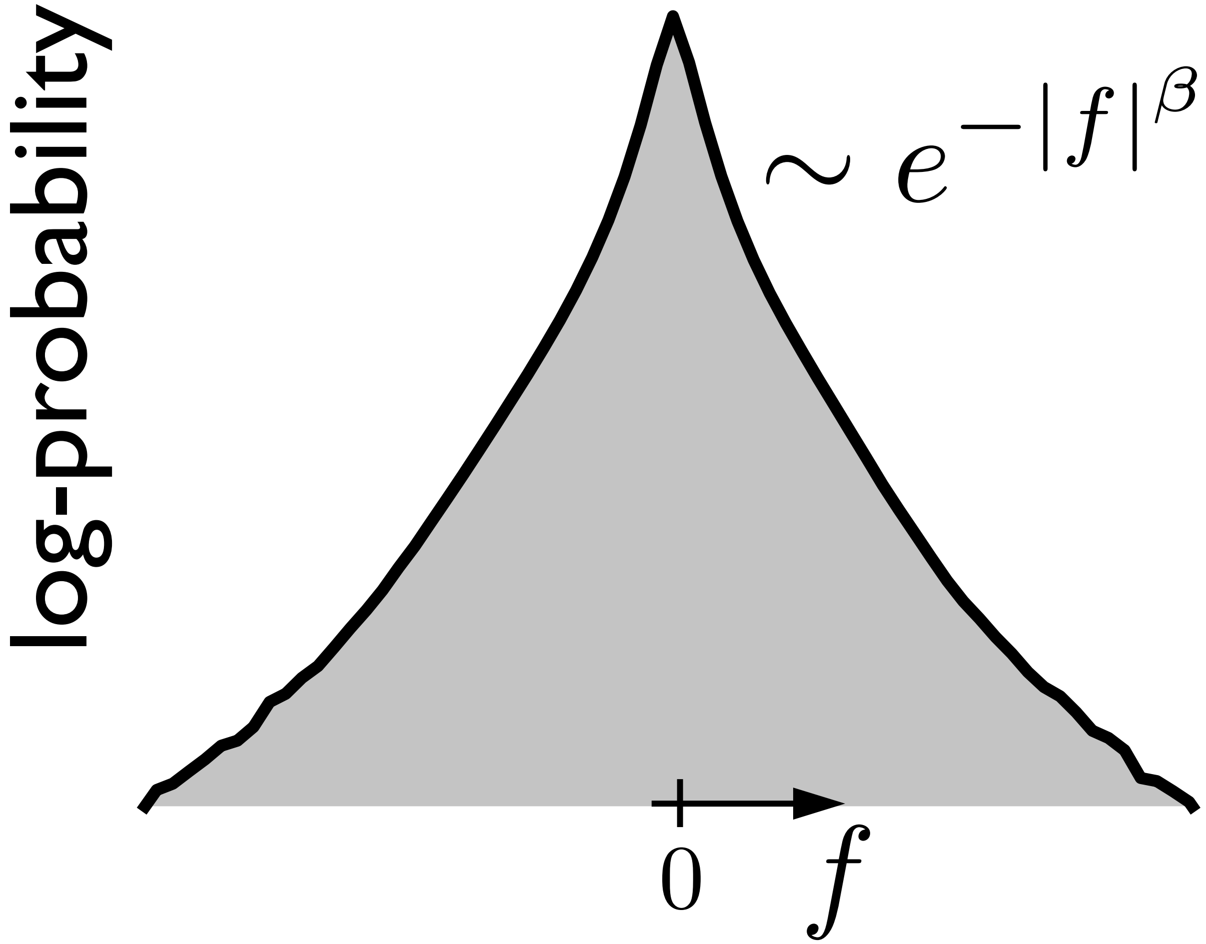

feature distribution

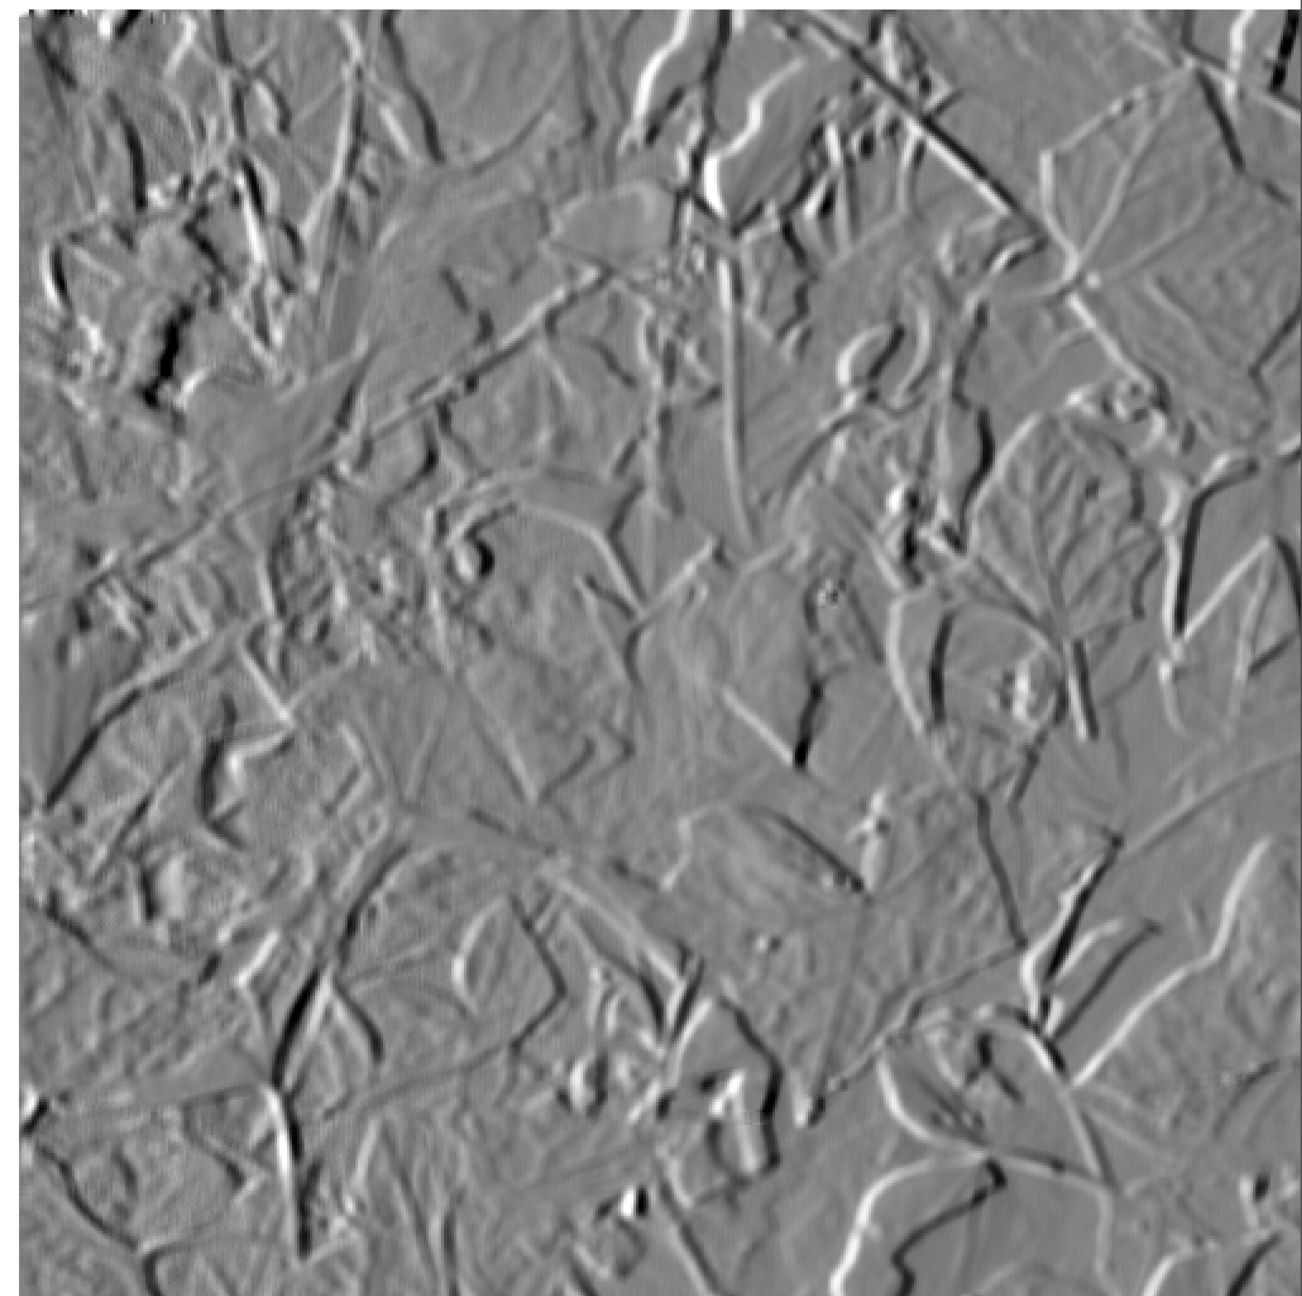

filtered image

f

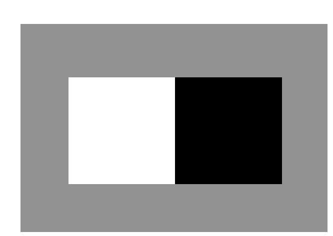

Haar wavelet features

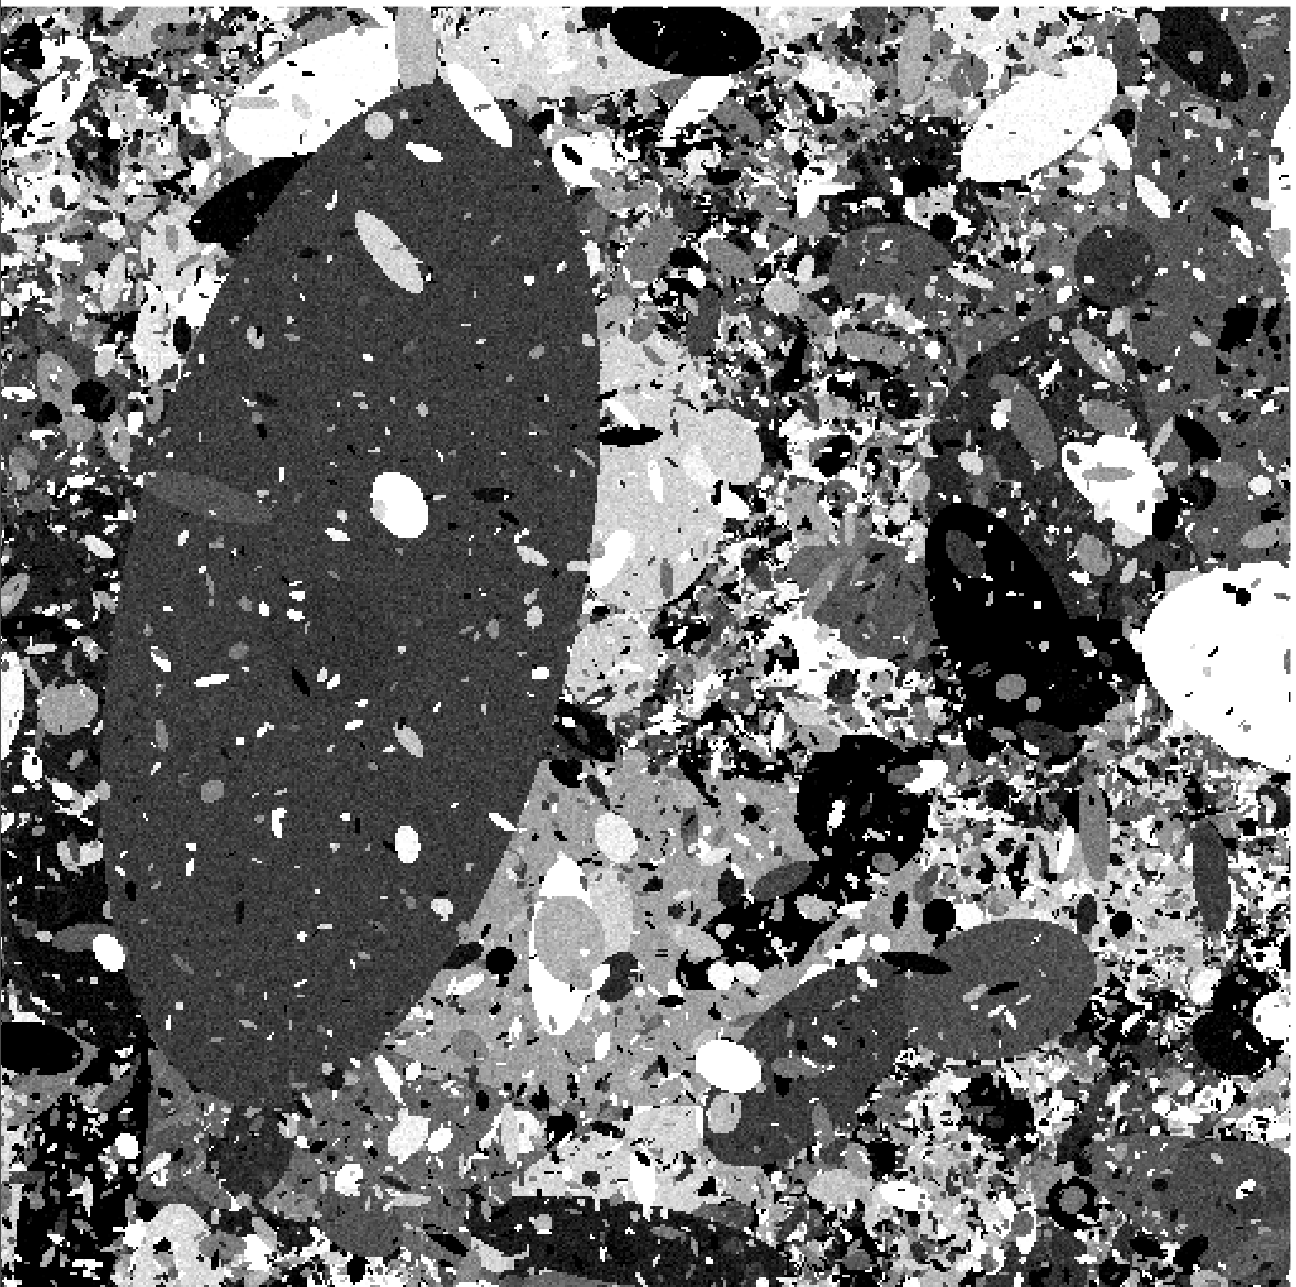

raw image

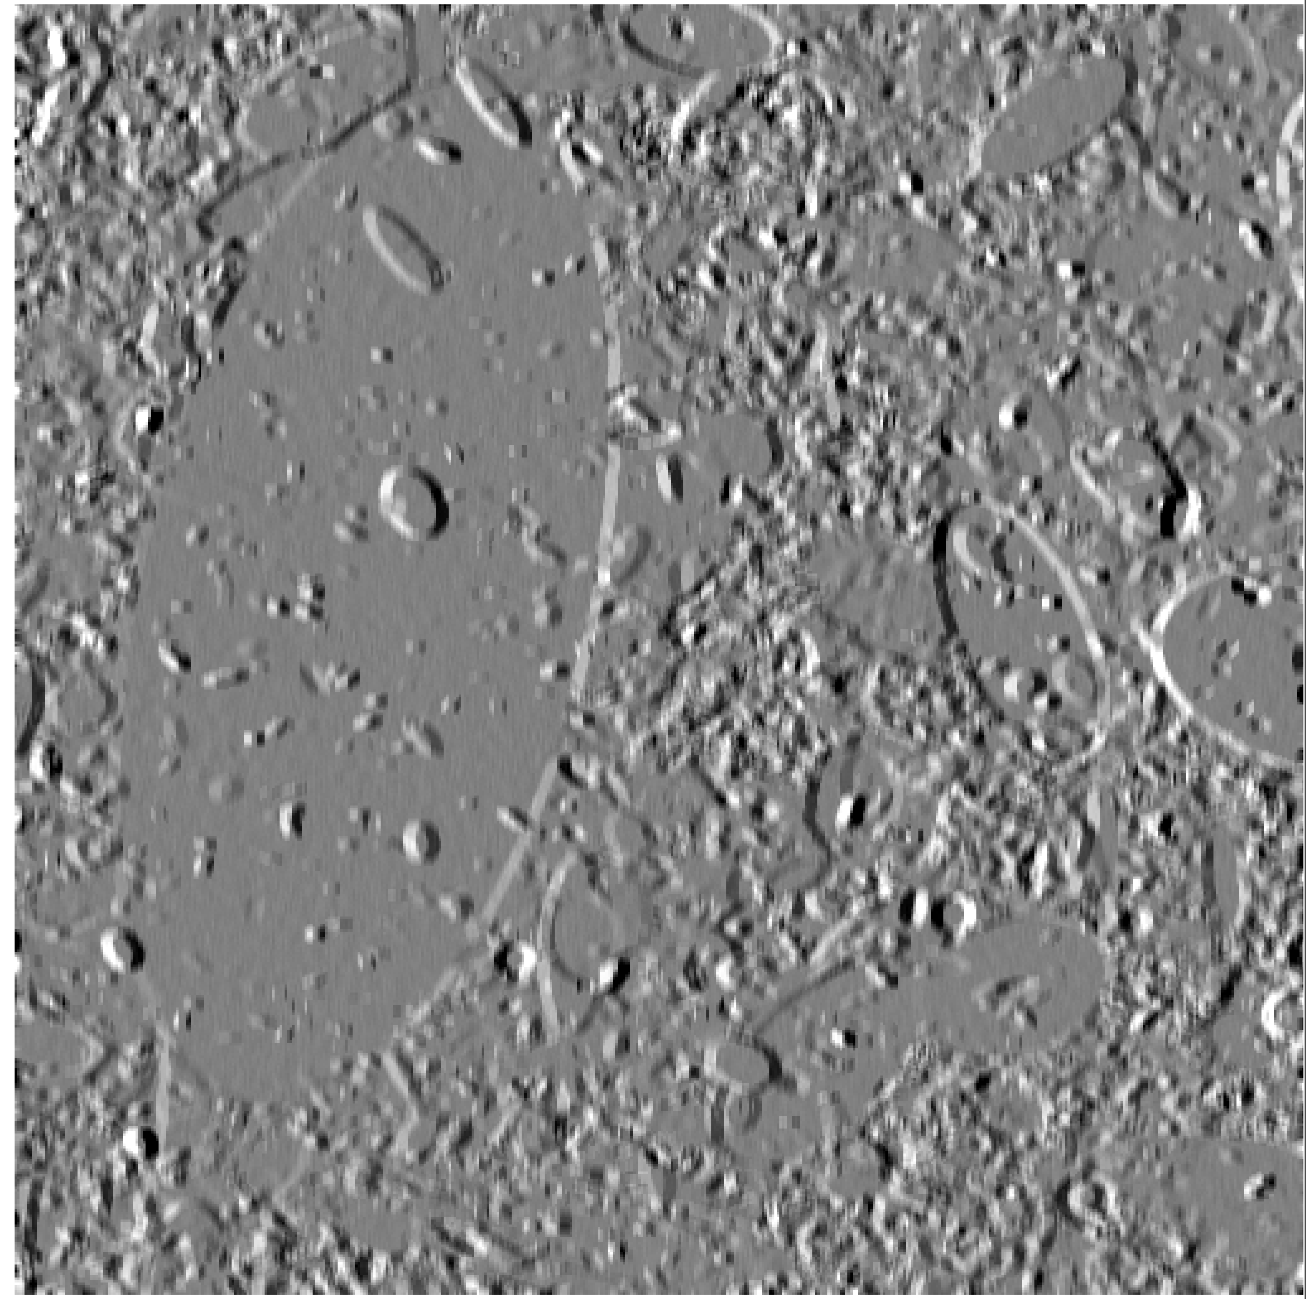

filtered image

f

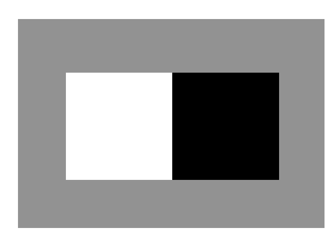

Haar wavelet features

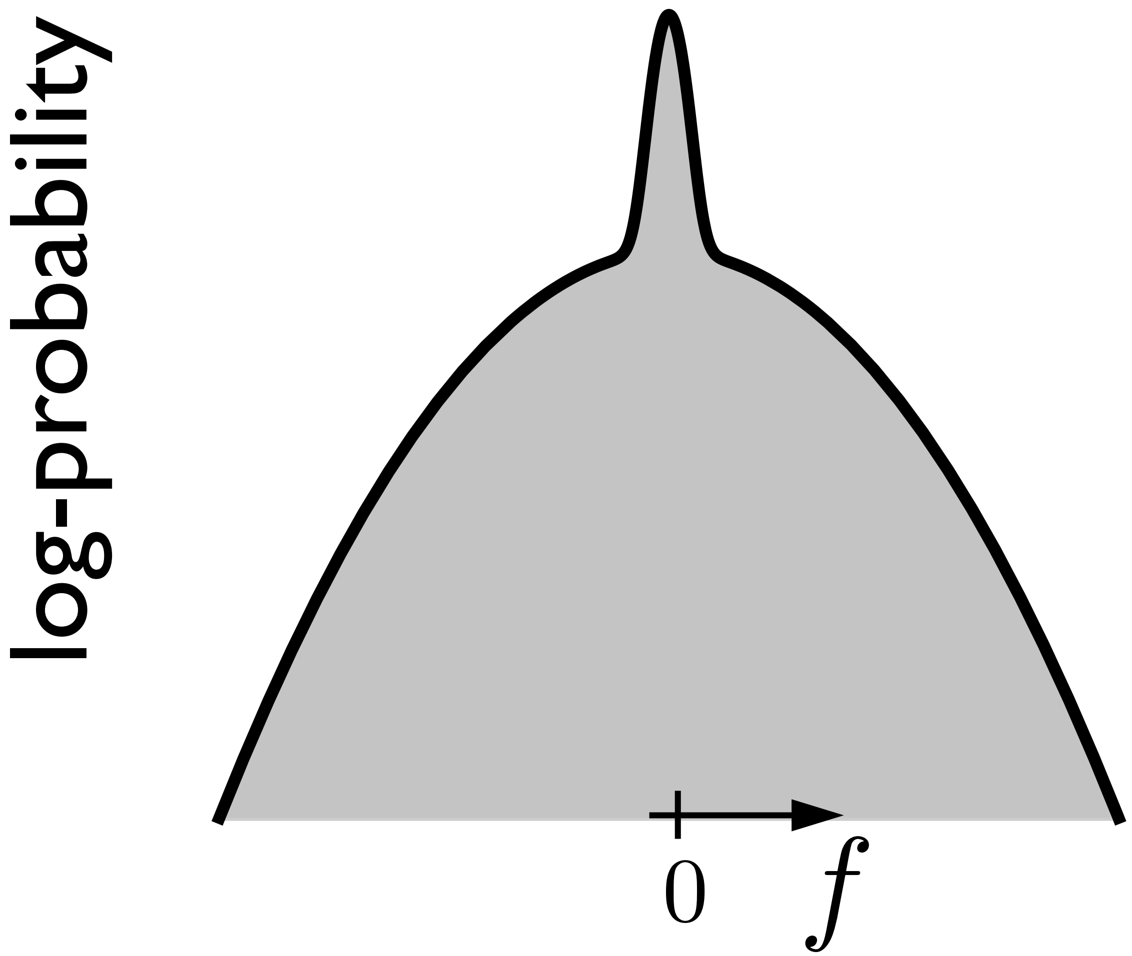

feature distribution

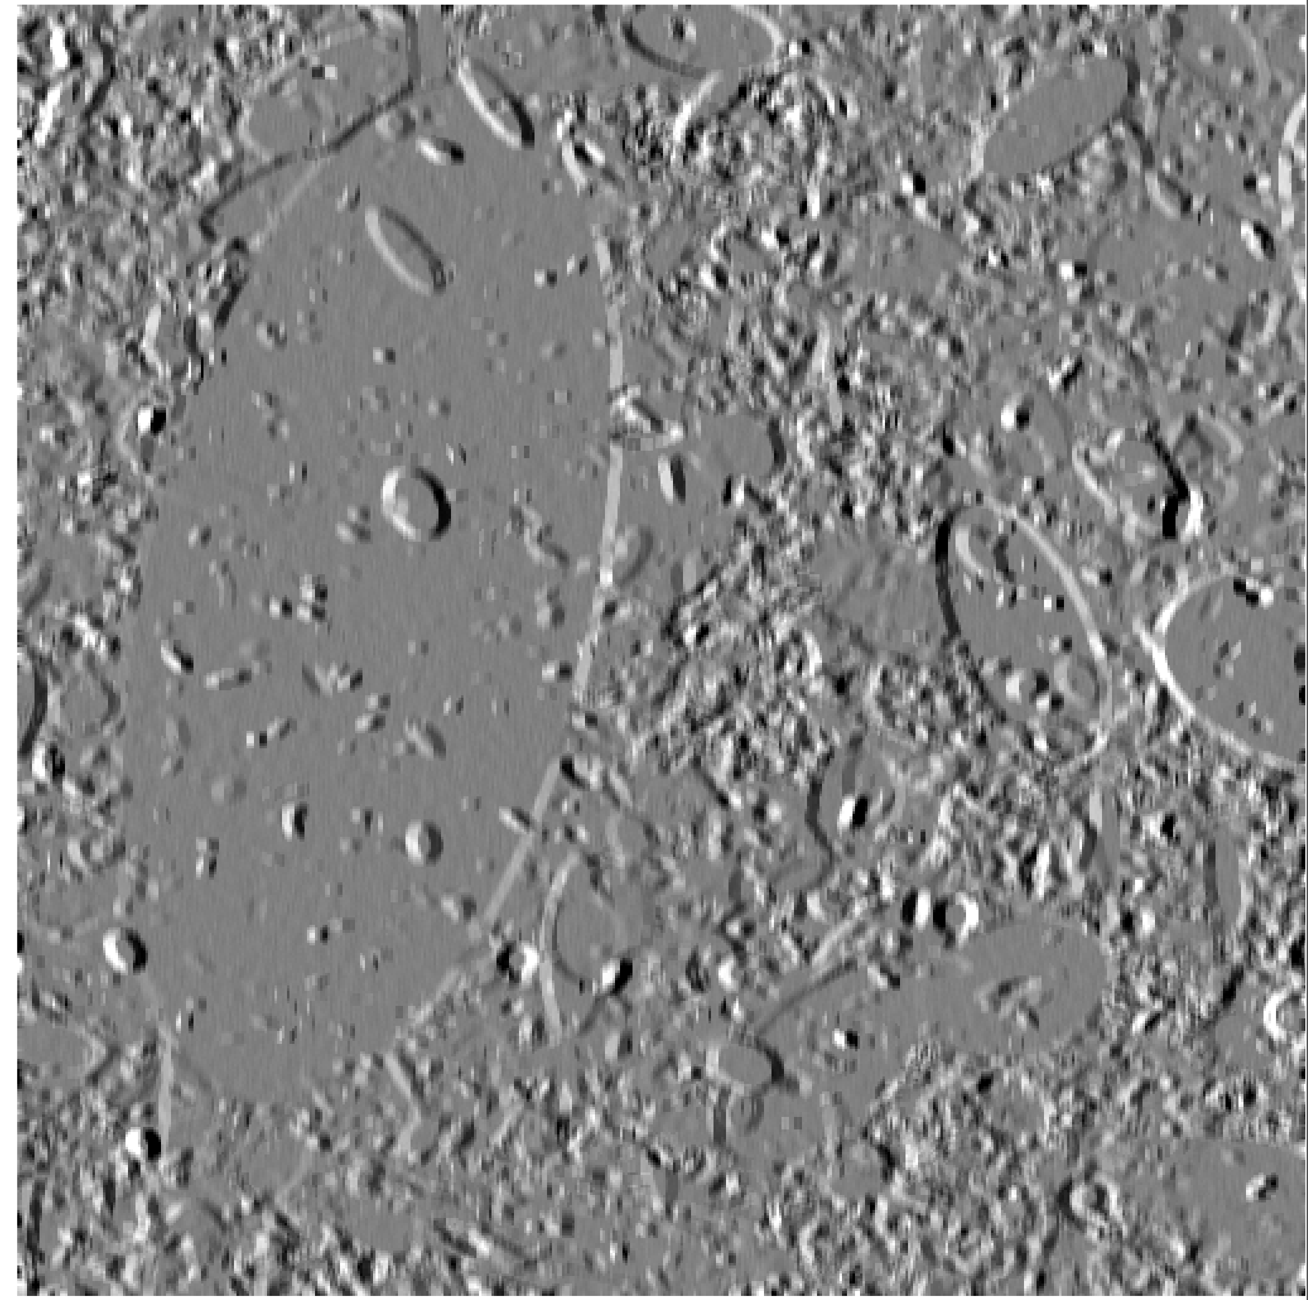

filtered image

f 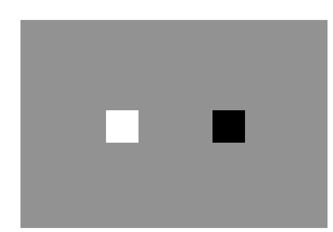 Haar wavelet features

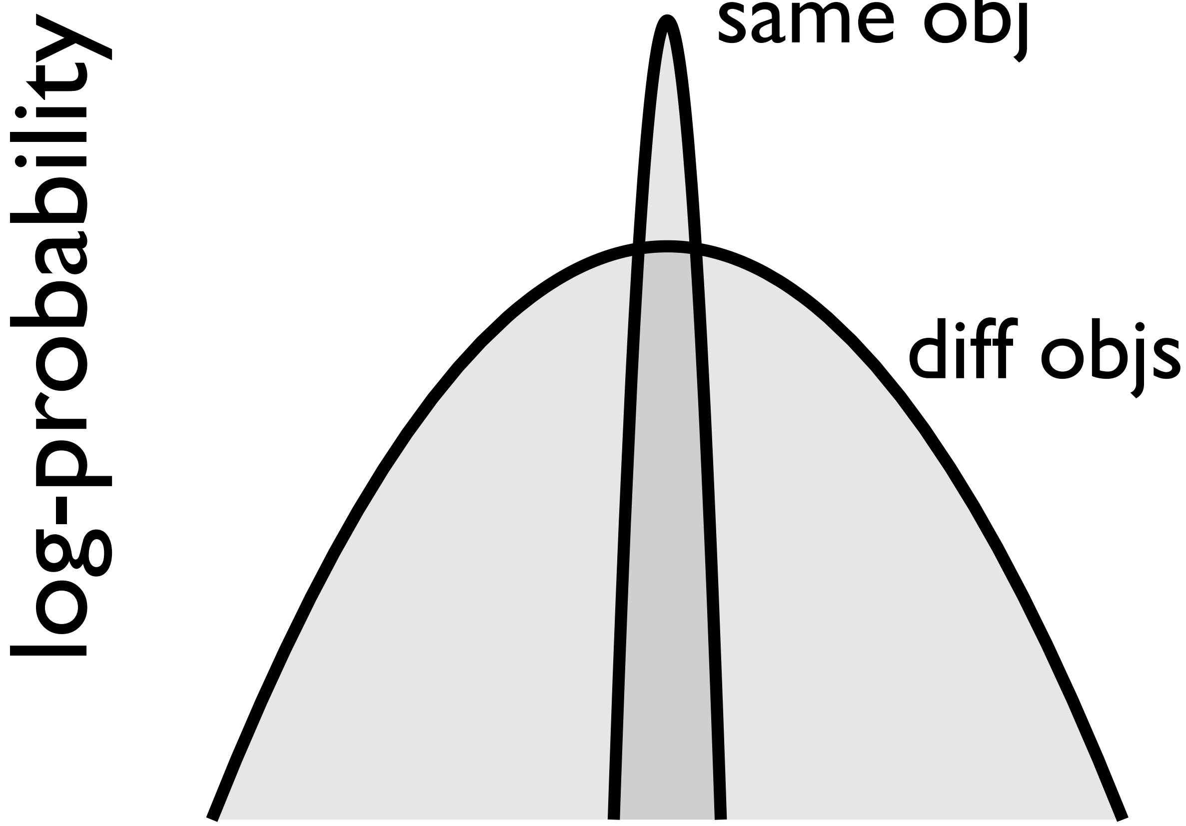

feature distribution

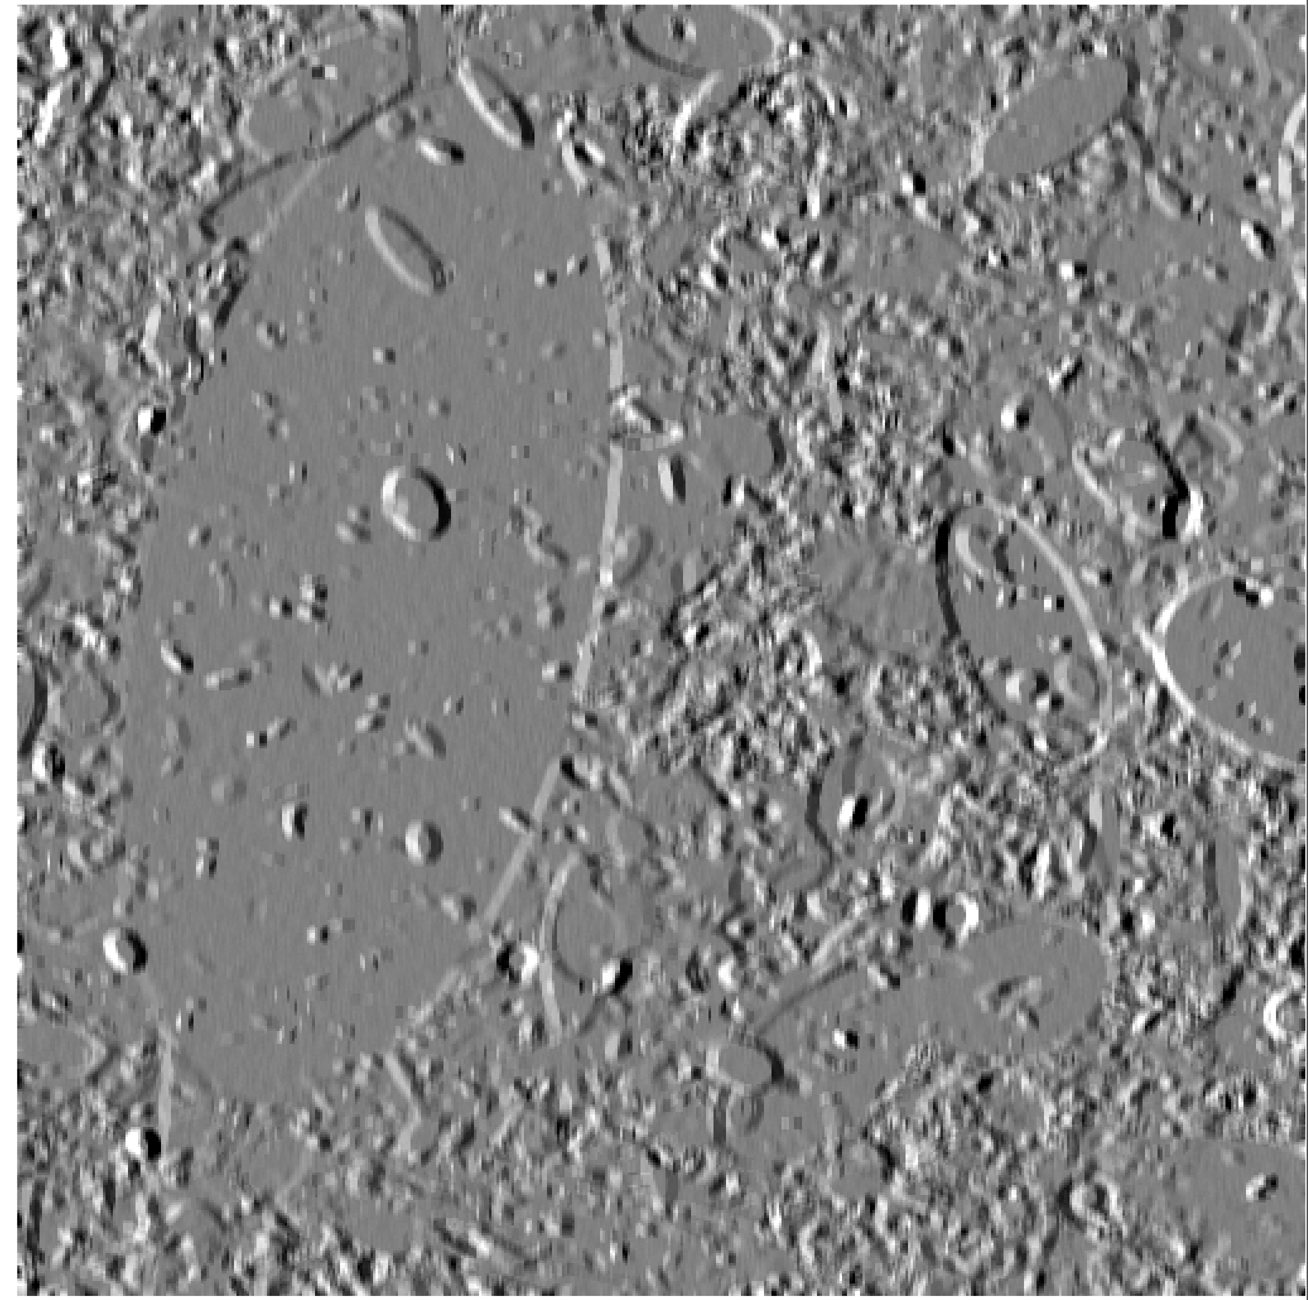

filtered image

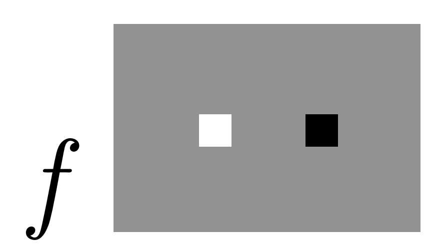

Haar wavelet features

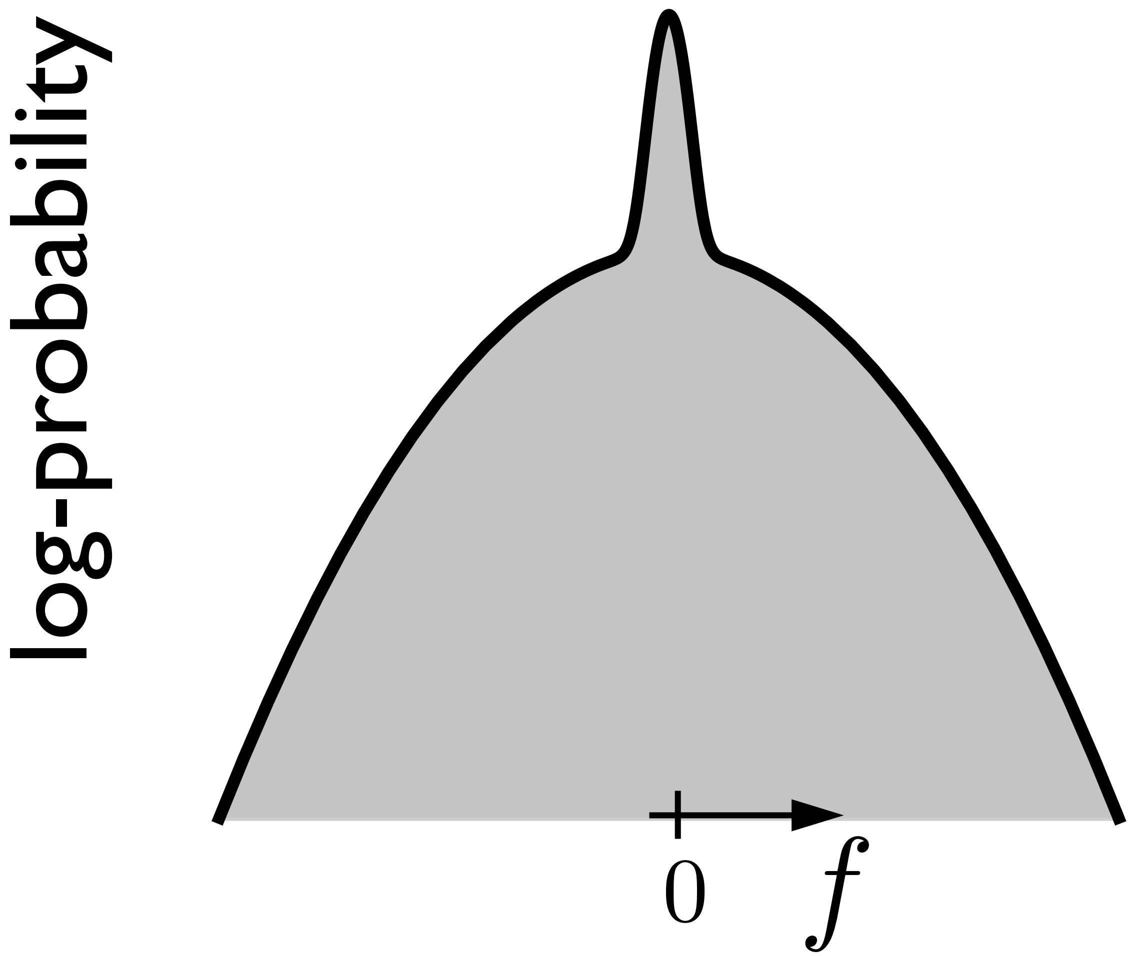

feature distribution

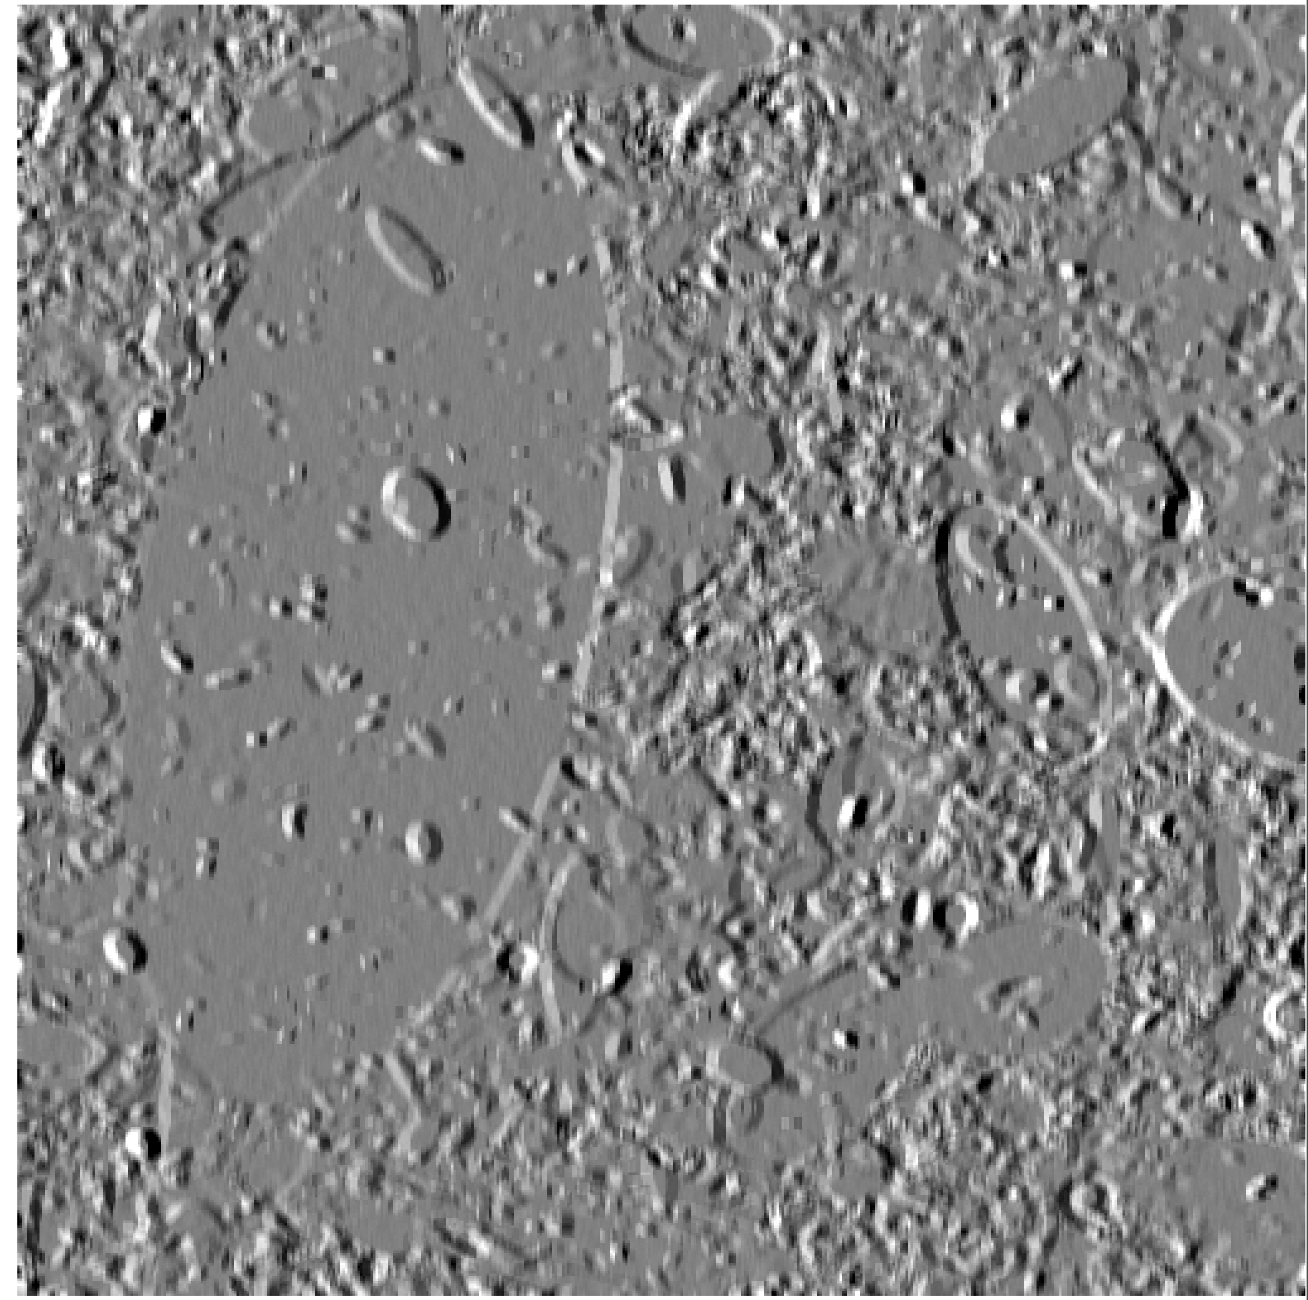

filtered image

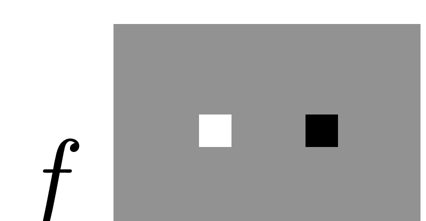

Haar wavelet features

$$\bar{I} = \sum_{j:x_j \in X} I_j$$

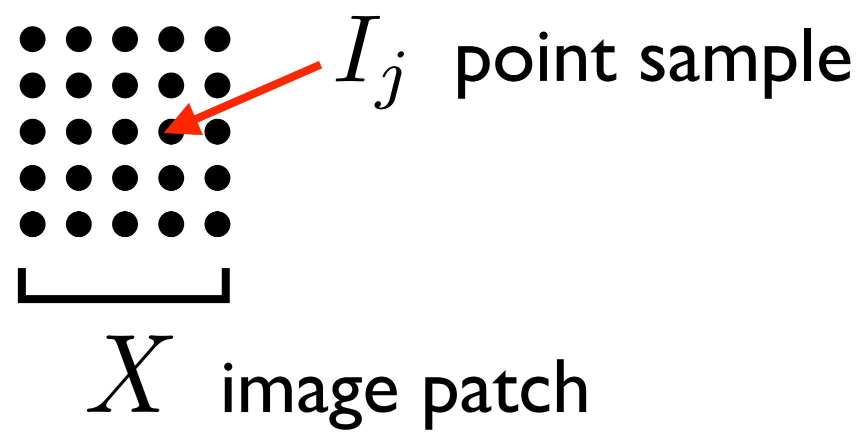

Central limit theorem:

$$\bar{I} = \sum_{j}^{N} I_j \implies P(\bar{I}) \sim \exp\left(-\frac{\bar{I}^2}{2\sigma^2}\right)$$

$$\bar{I} = \sum_{j:x_j \in X} I_j \;= \sum_{n=1}^{|\text{מ}|} |\text{מ}_n| J_n$$

for n’th object:

$|\text{מ}_n|$ area

J_n mean intensity

Central limit theorem:

$$\bar{I} = \sum_{j}^{N} I_j \implies P(\bar{I}) \sim \exp\left(-\frac{\bar{I}^2}{2\sigma^2}\right)$$

$$\bar{I} = \sum_{j:x_j \in X} I_j = \sum_{n=1}^{|\text{מ}|} |\text{מ}_n| J_n$$

for n'th object:

$|\text{מ}_n|$ area

J_n mean intensity

Central limit theorem for *random* sum:

$$\bar{I} = \sum_{j}^{\nu} I_j \implies P(\bar{I}) \sim \exp\left(-\frac{|\bar{I}|}{\sigma}\right)$$

$$P(\nu) \propto p^{\nu}$$

Kotz et al 2001
Gnedenko 1972

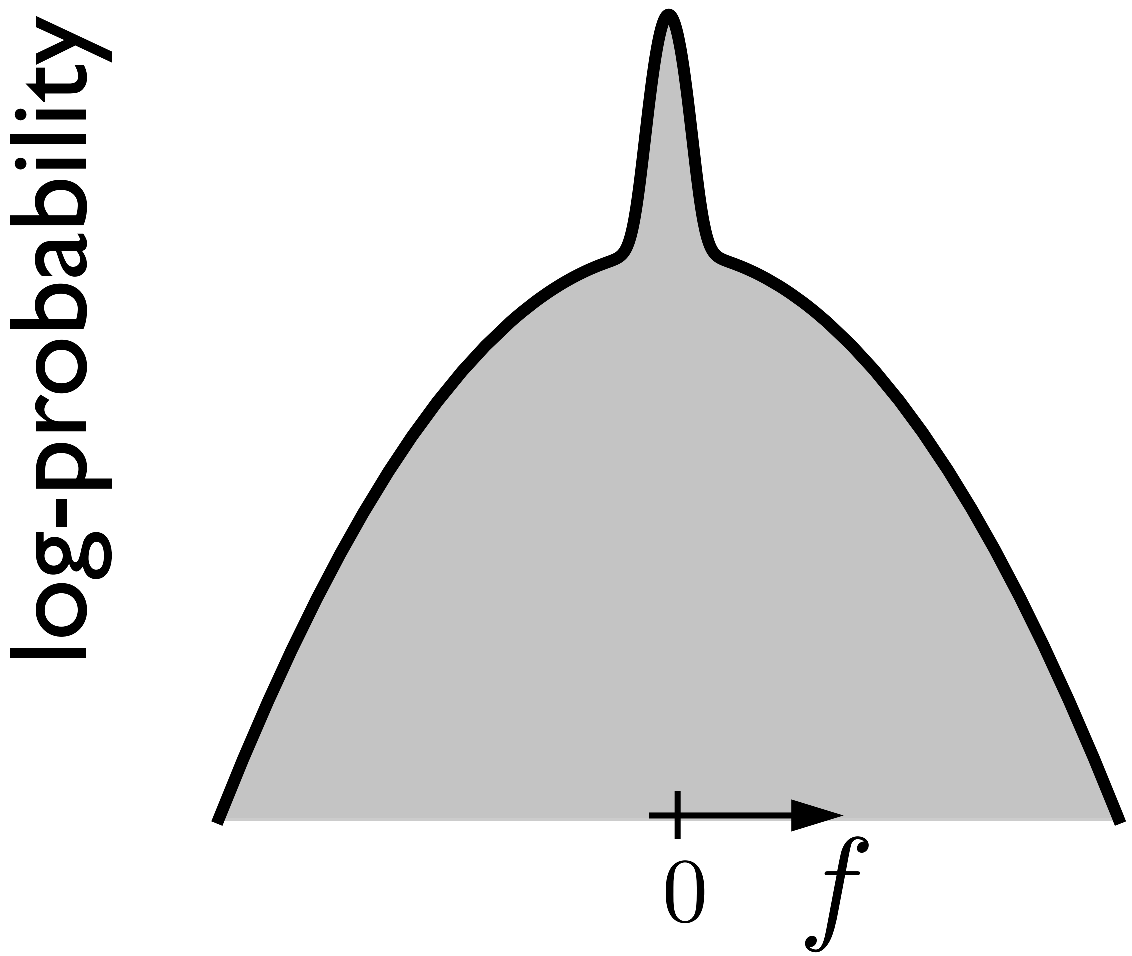

feature distribution

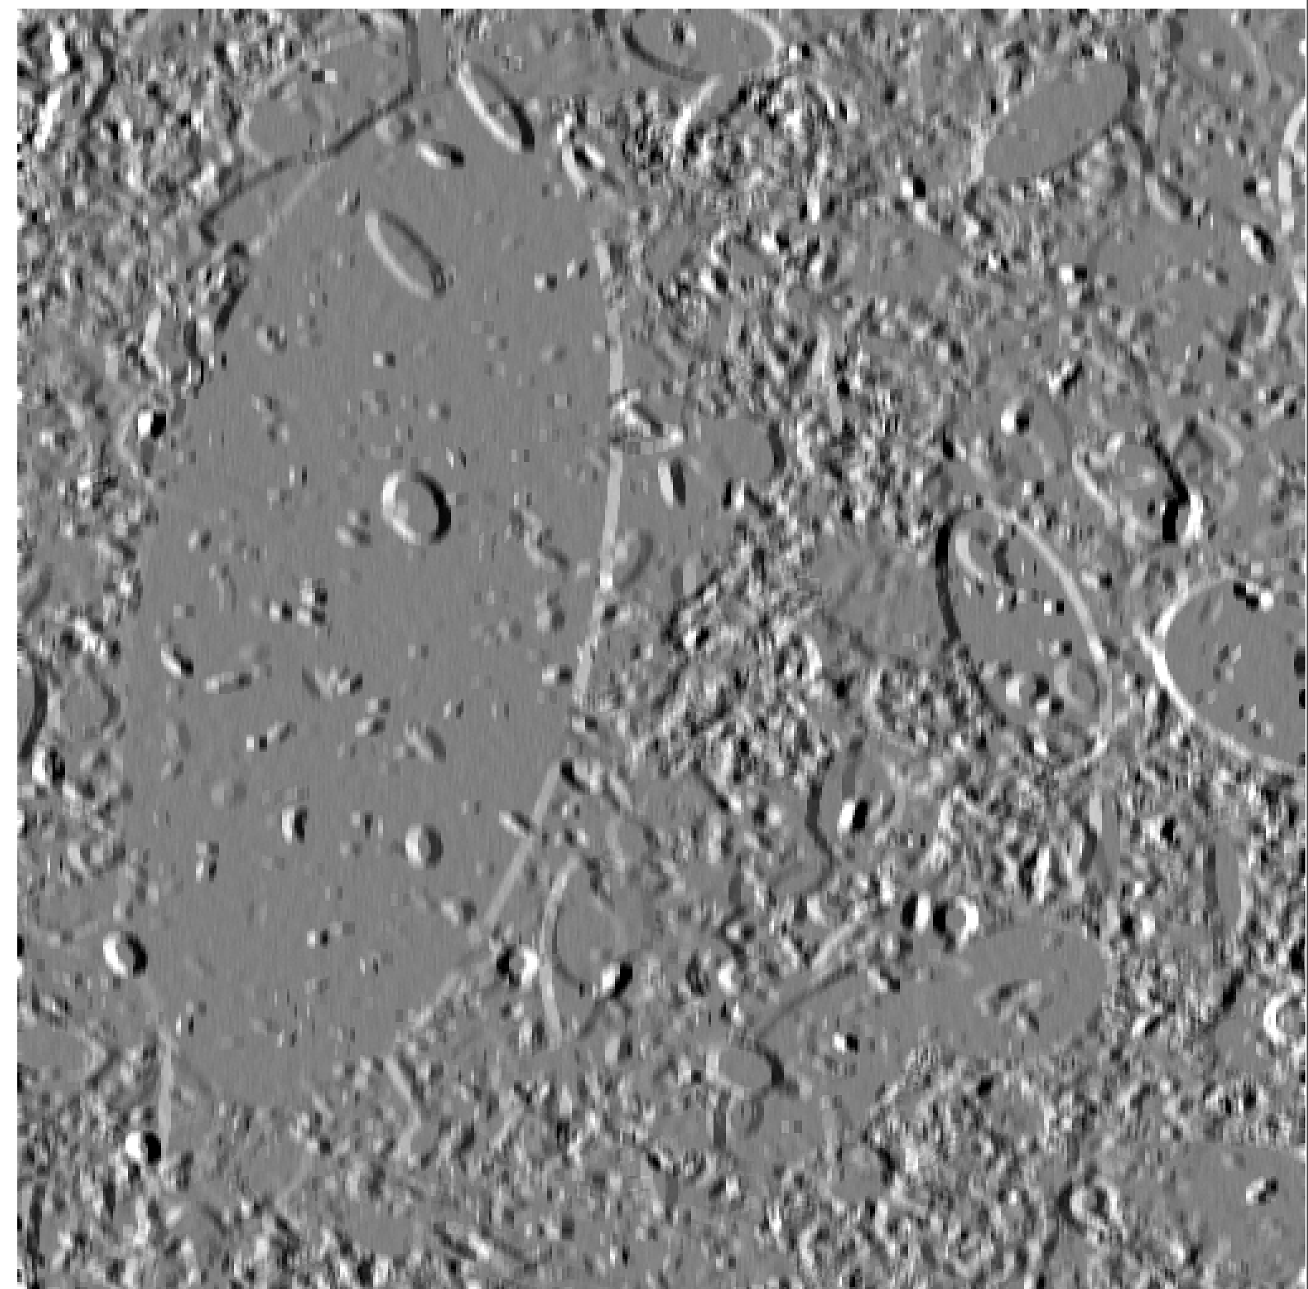

filtered image

f

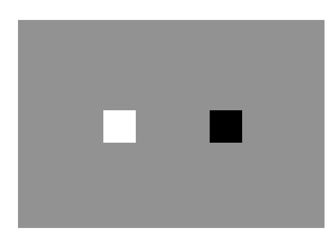

Haar wavelet features

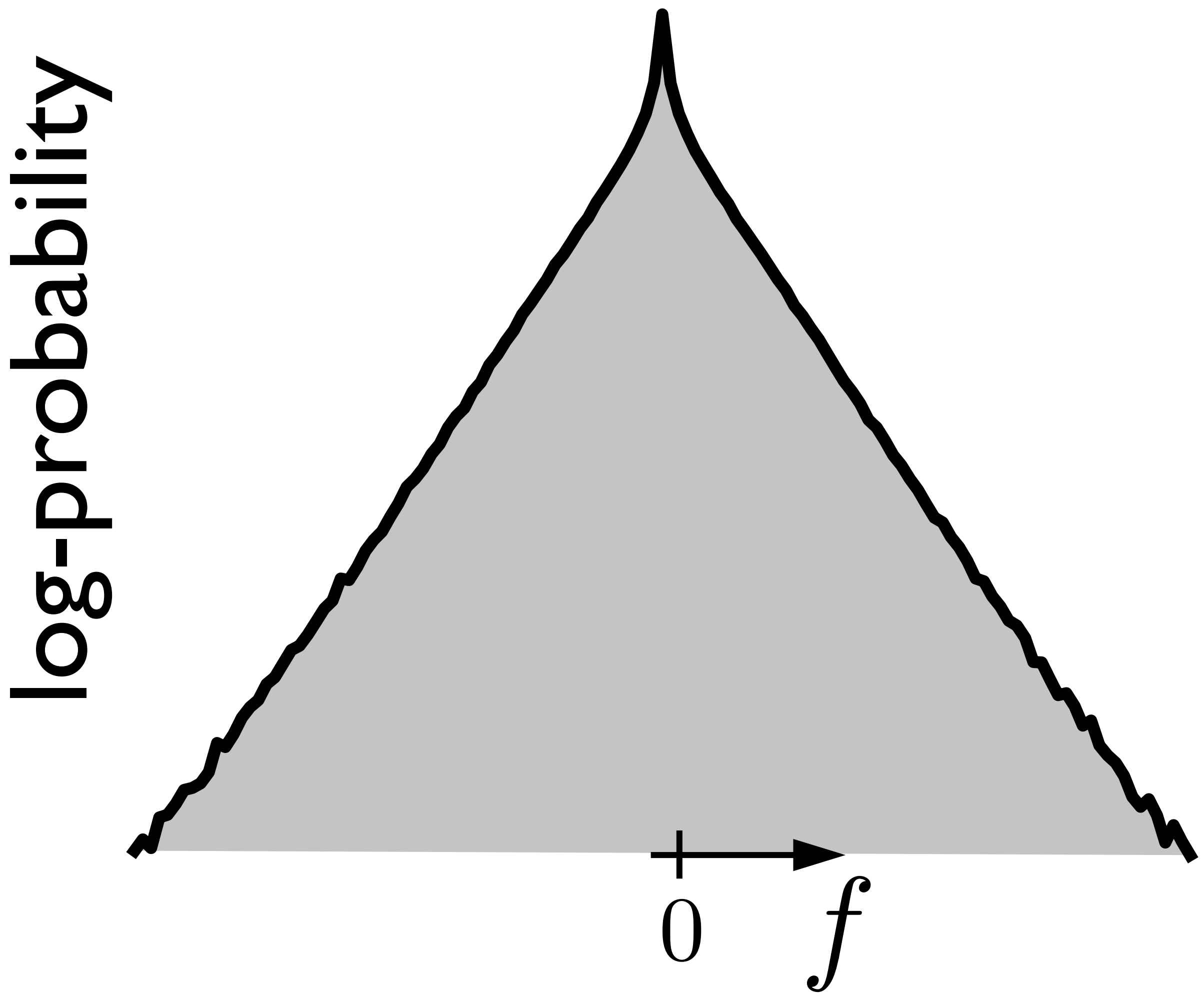

feature distribution

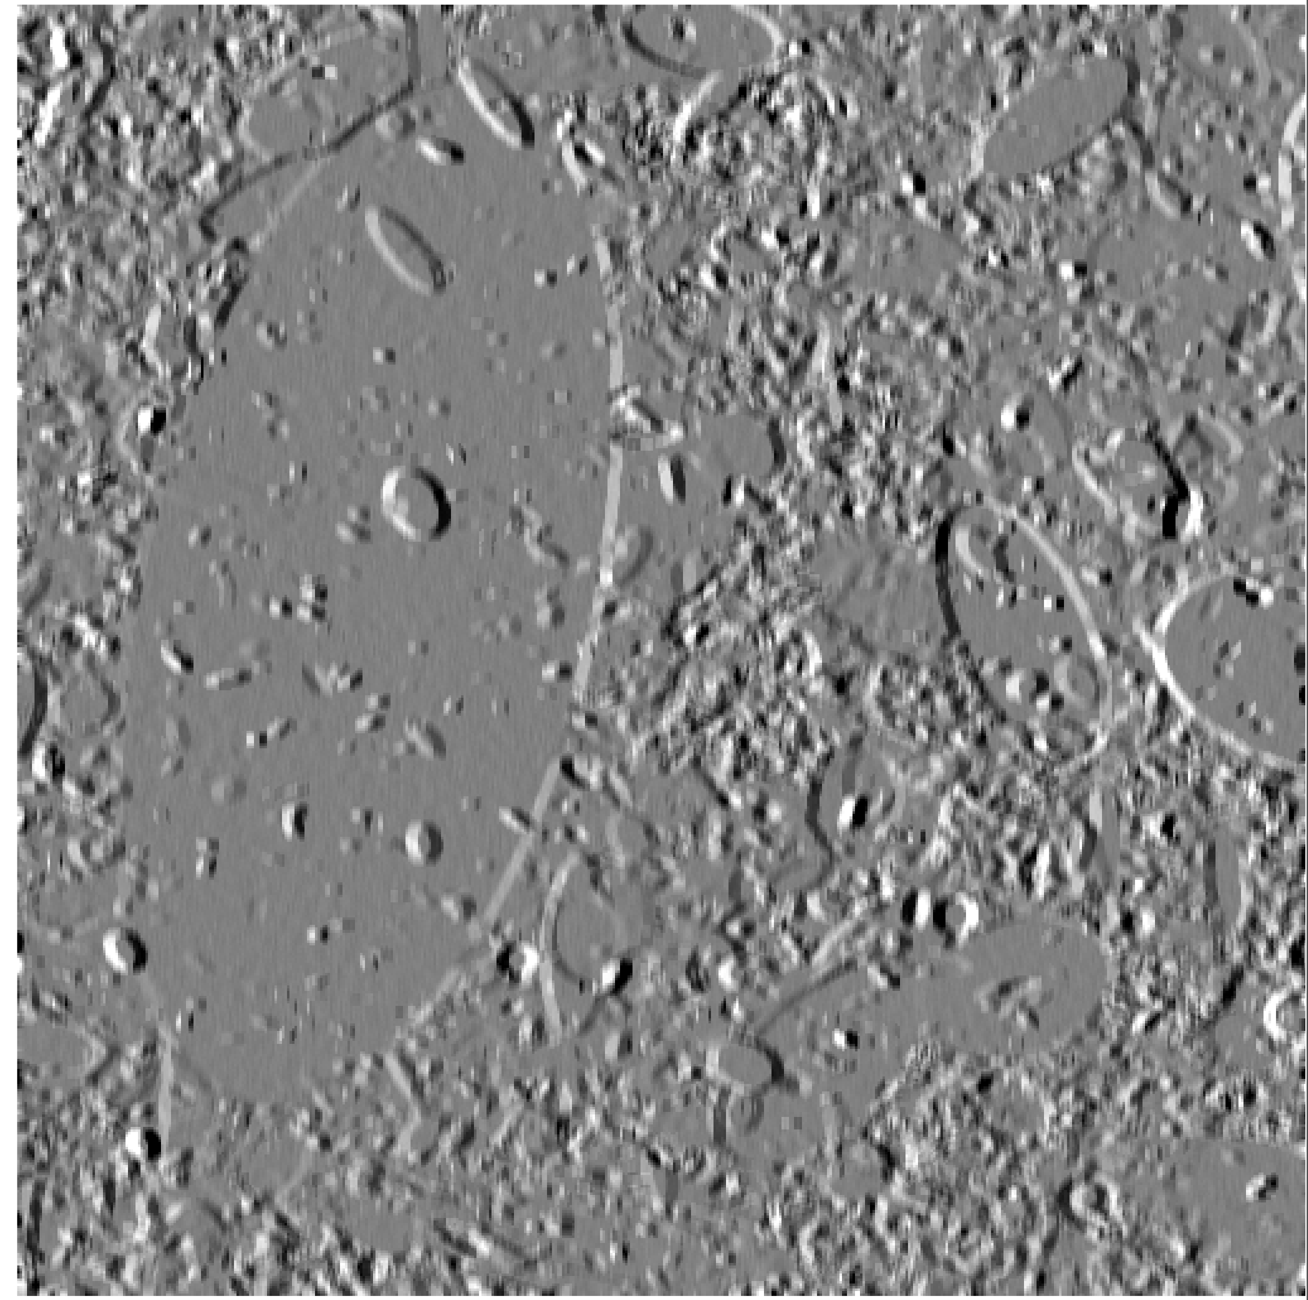

filtered image

Haar wavelet features

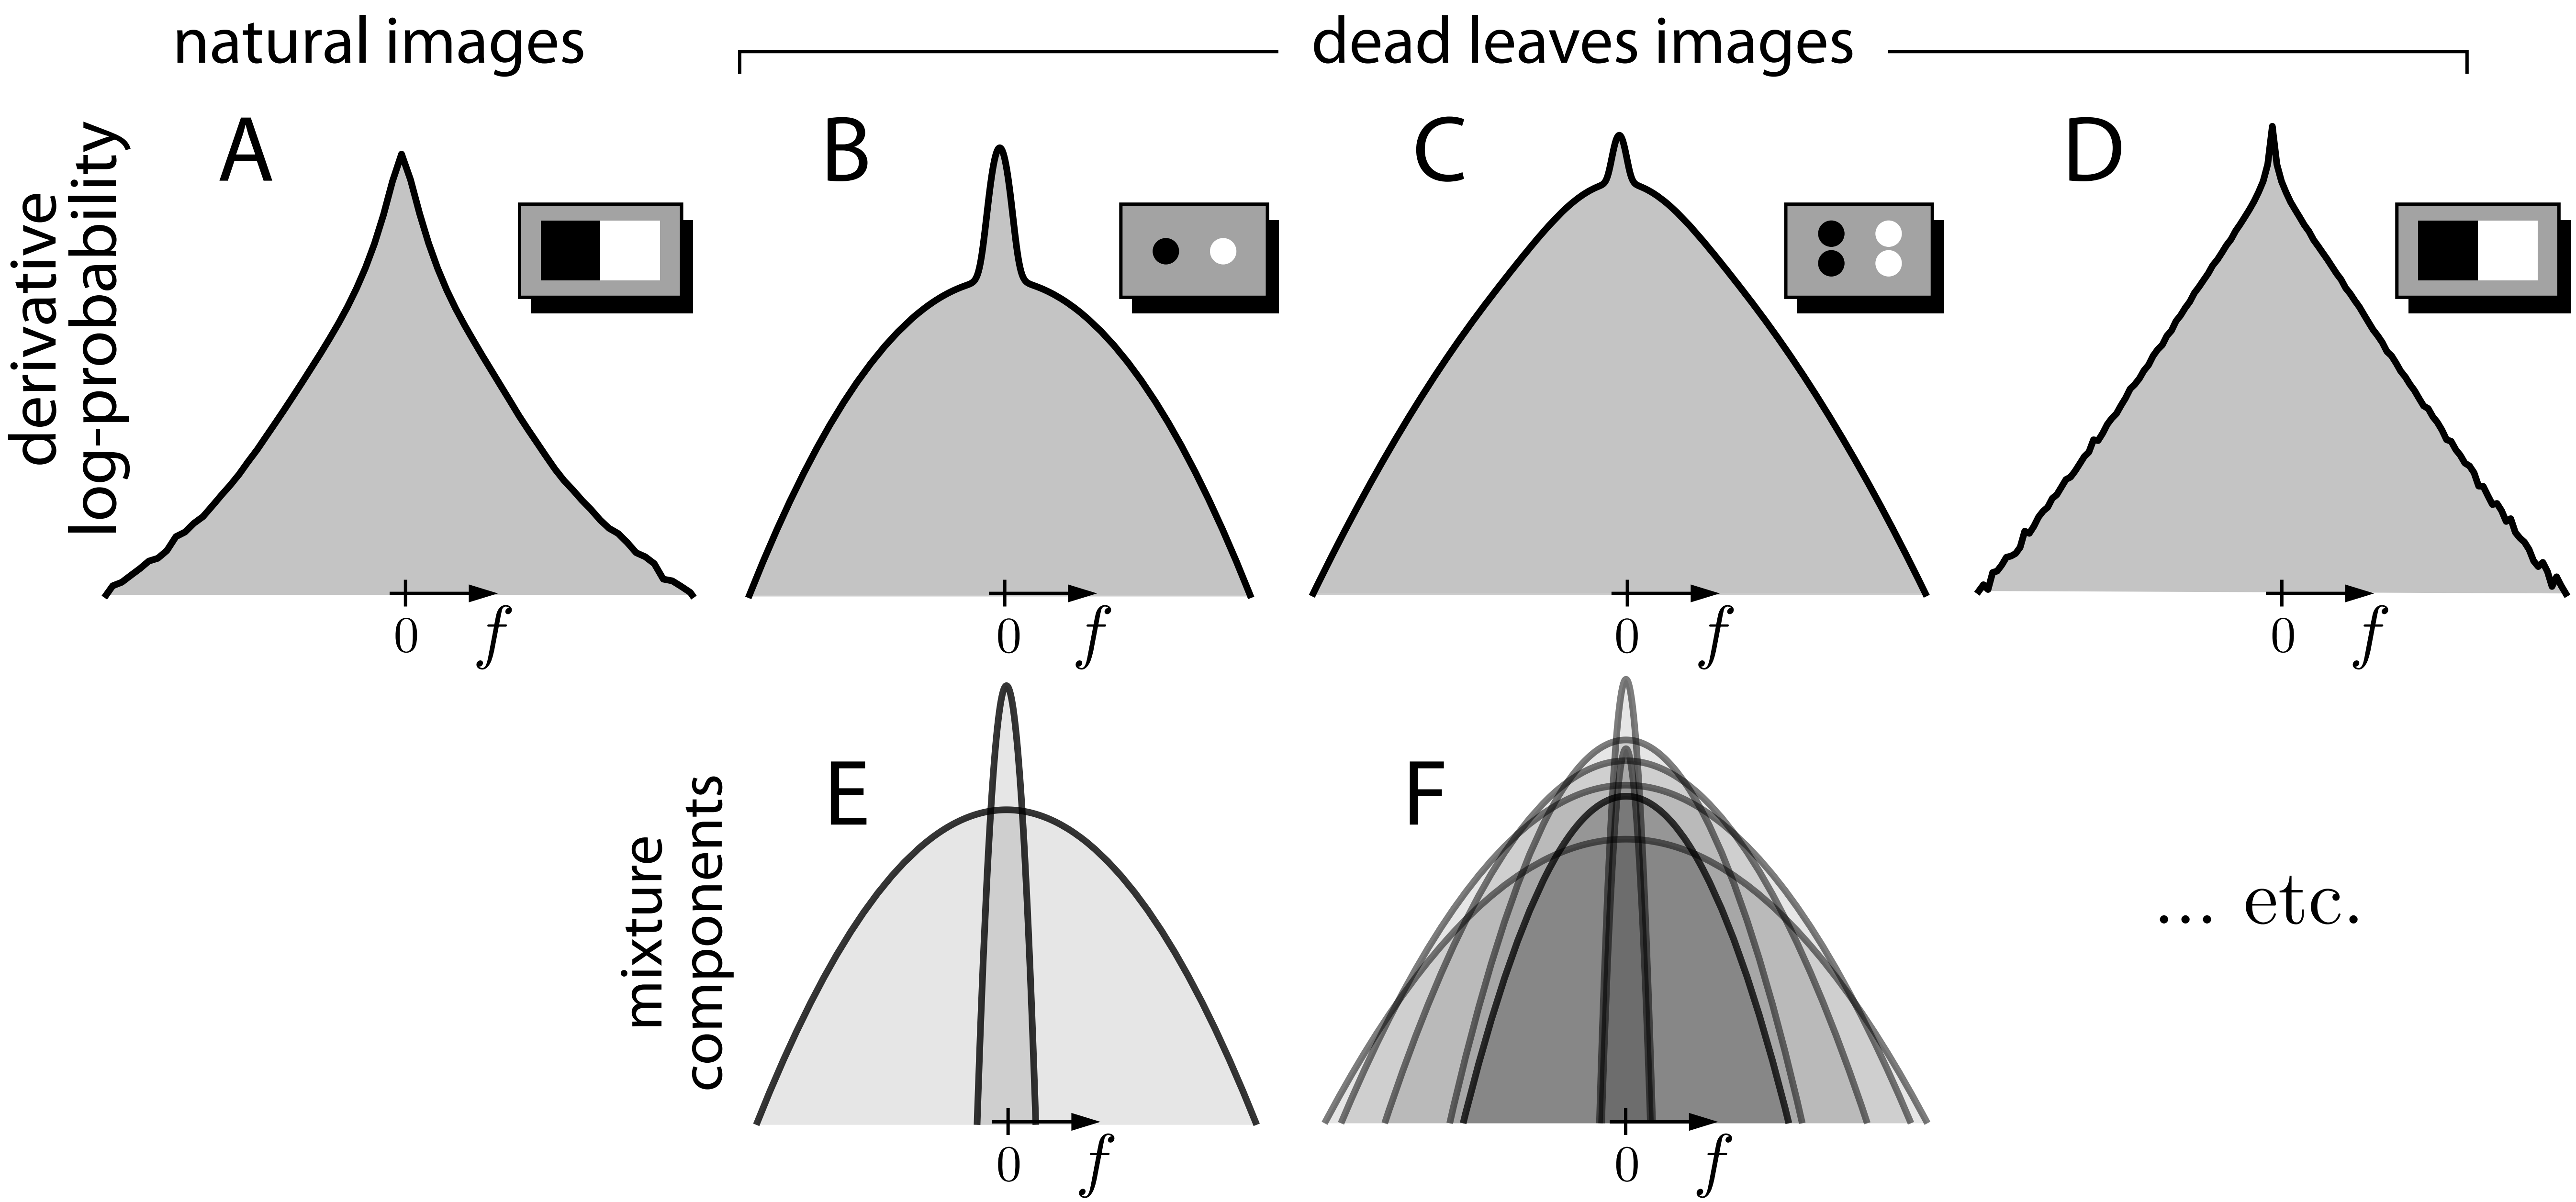

natural images
dead leaves images
A
B
C
D
derivative
log-probability
0
f
mixture
components
E
F
... etc.

joint probability

natural images

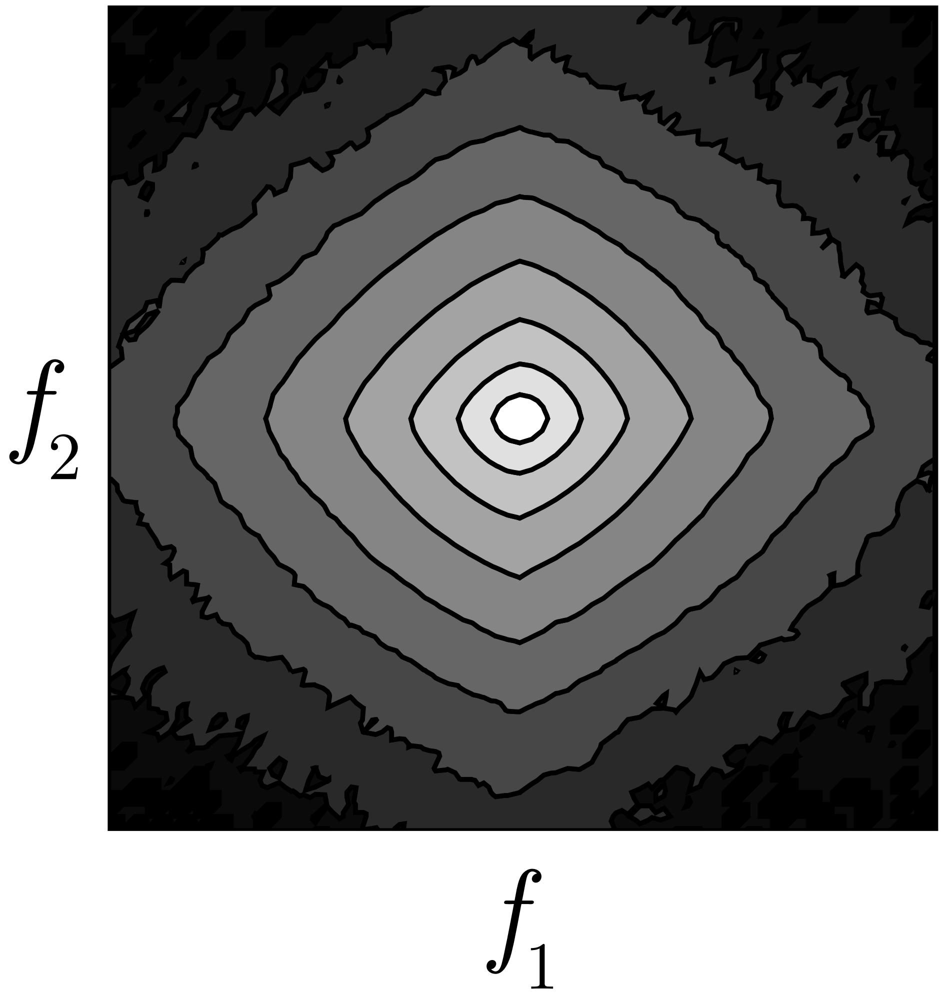

dead leaves images

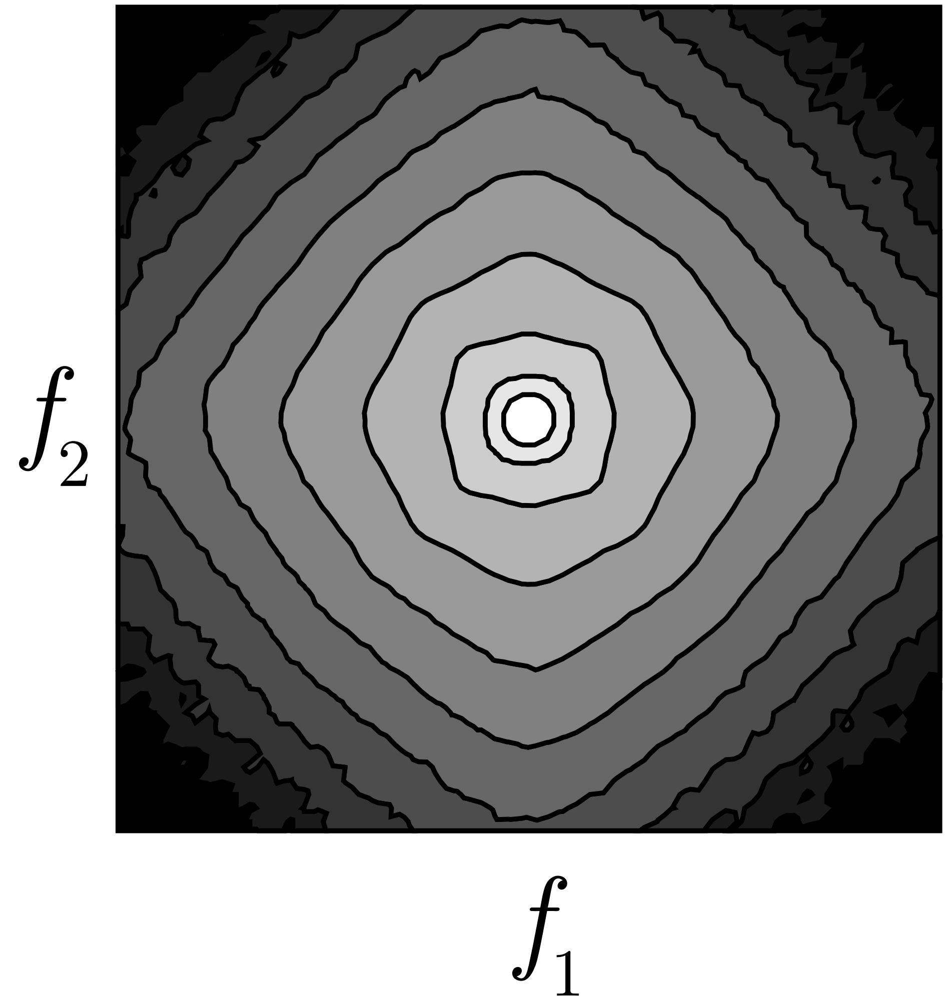

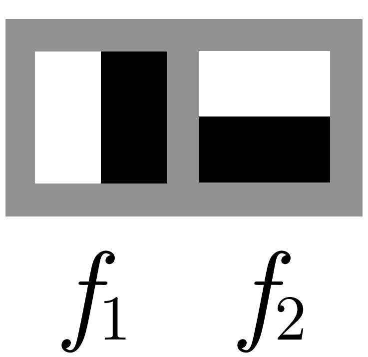

orthogonal, colocalized Haar wavelets

$$P(\mathbf{f}) = \sum_{\text{מ}} P(\mathbf{f}|\text{מ})P(\text{מ})$$

Probability at high amplitude $\mathbf{f}$ is dominated by object configurations with greatest likelihood

$$P(\mathbf{f}) = \sum_{\text{מ}} P(\mathbf{f}|\text{מ})P(\text{מ})$$

Probability at high amplitude f is dominated by object configurations with greatest likelihood

$$P(\mathbf{f}|\text{מ}) \propto \exp\left(-\frac{1}{2}\mathbf{f}^\top C_{\mathbf{f}|\text{מ}}^{-1}\mathbf{f}\right)$$

object geometry

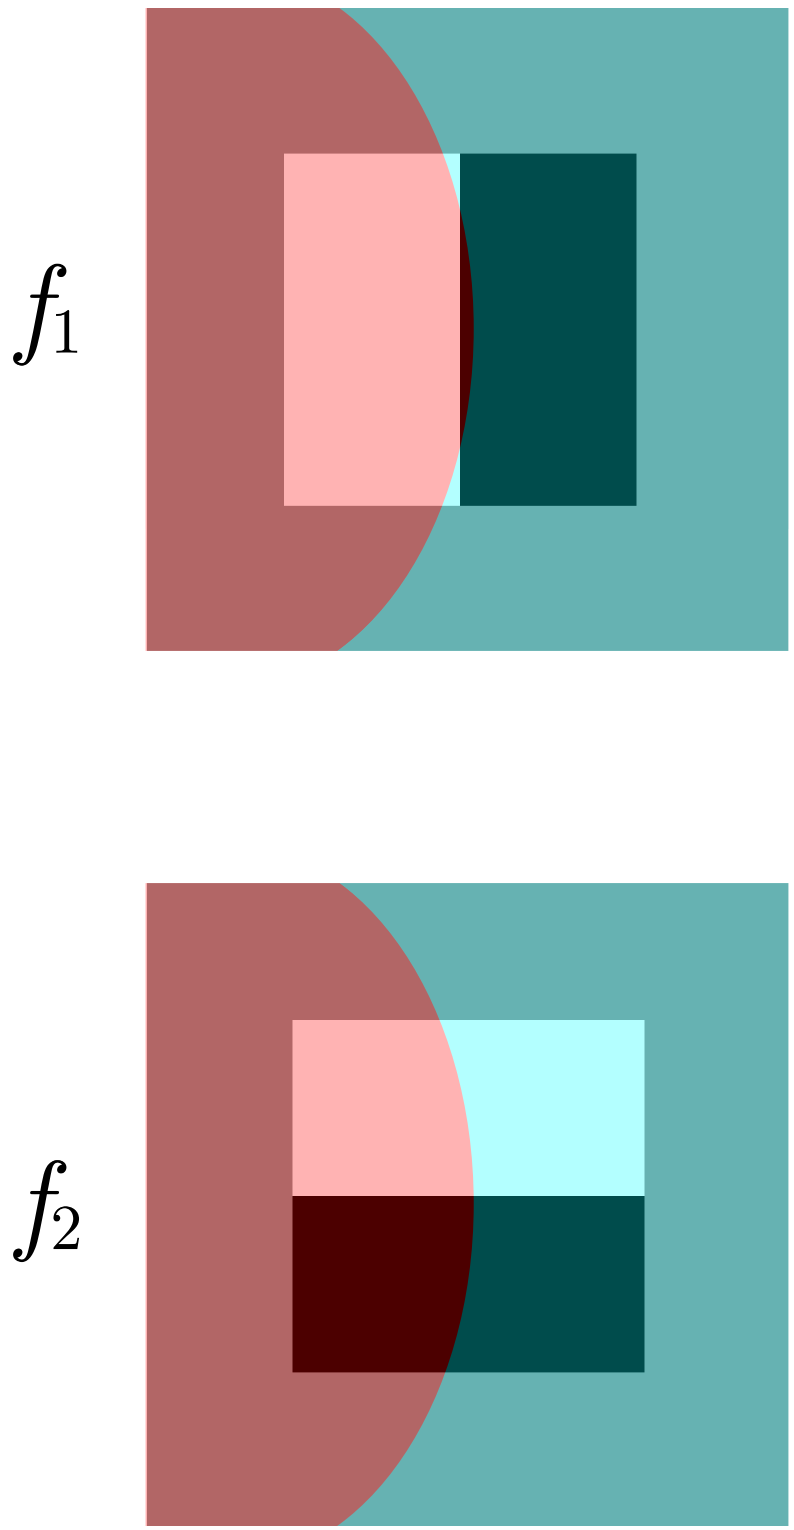

conditional probability contours

gaussian $P(f_1, f_2|\text{מ})$

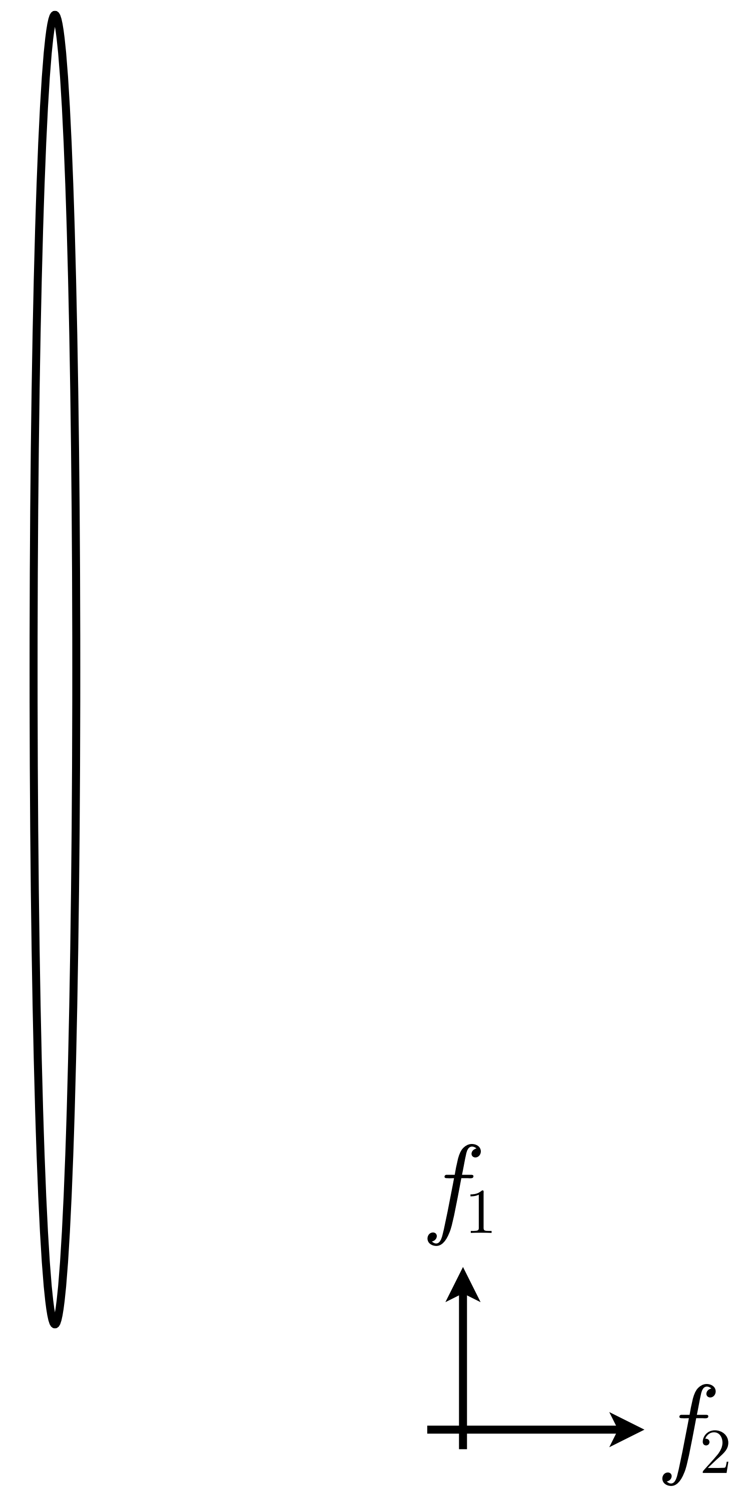

object geometry

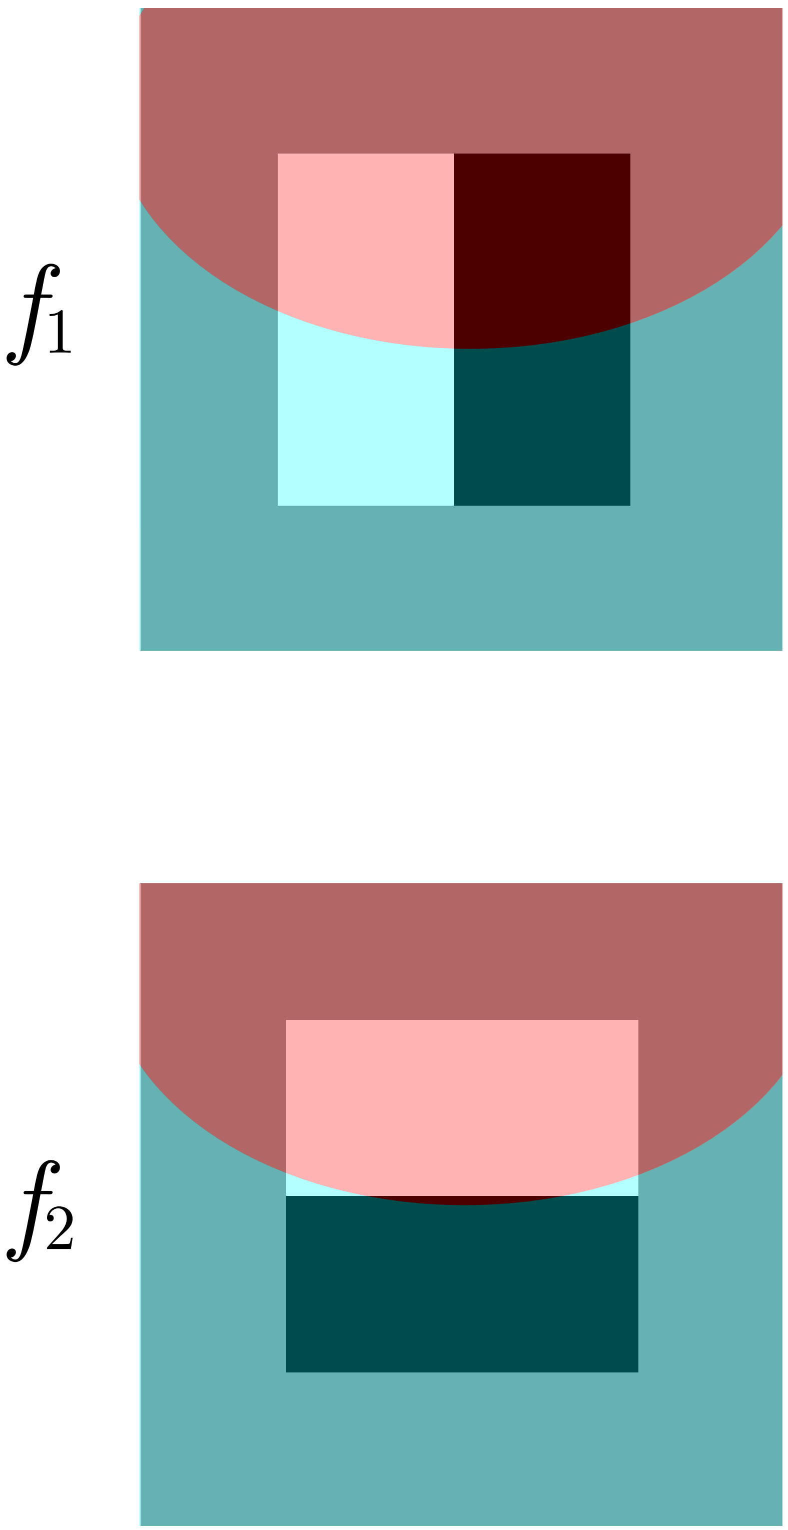

conditional probability contours
gaussian $P(f_1, f_2|מ)$

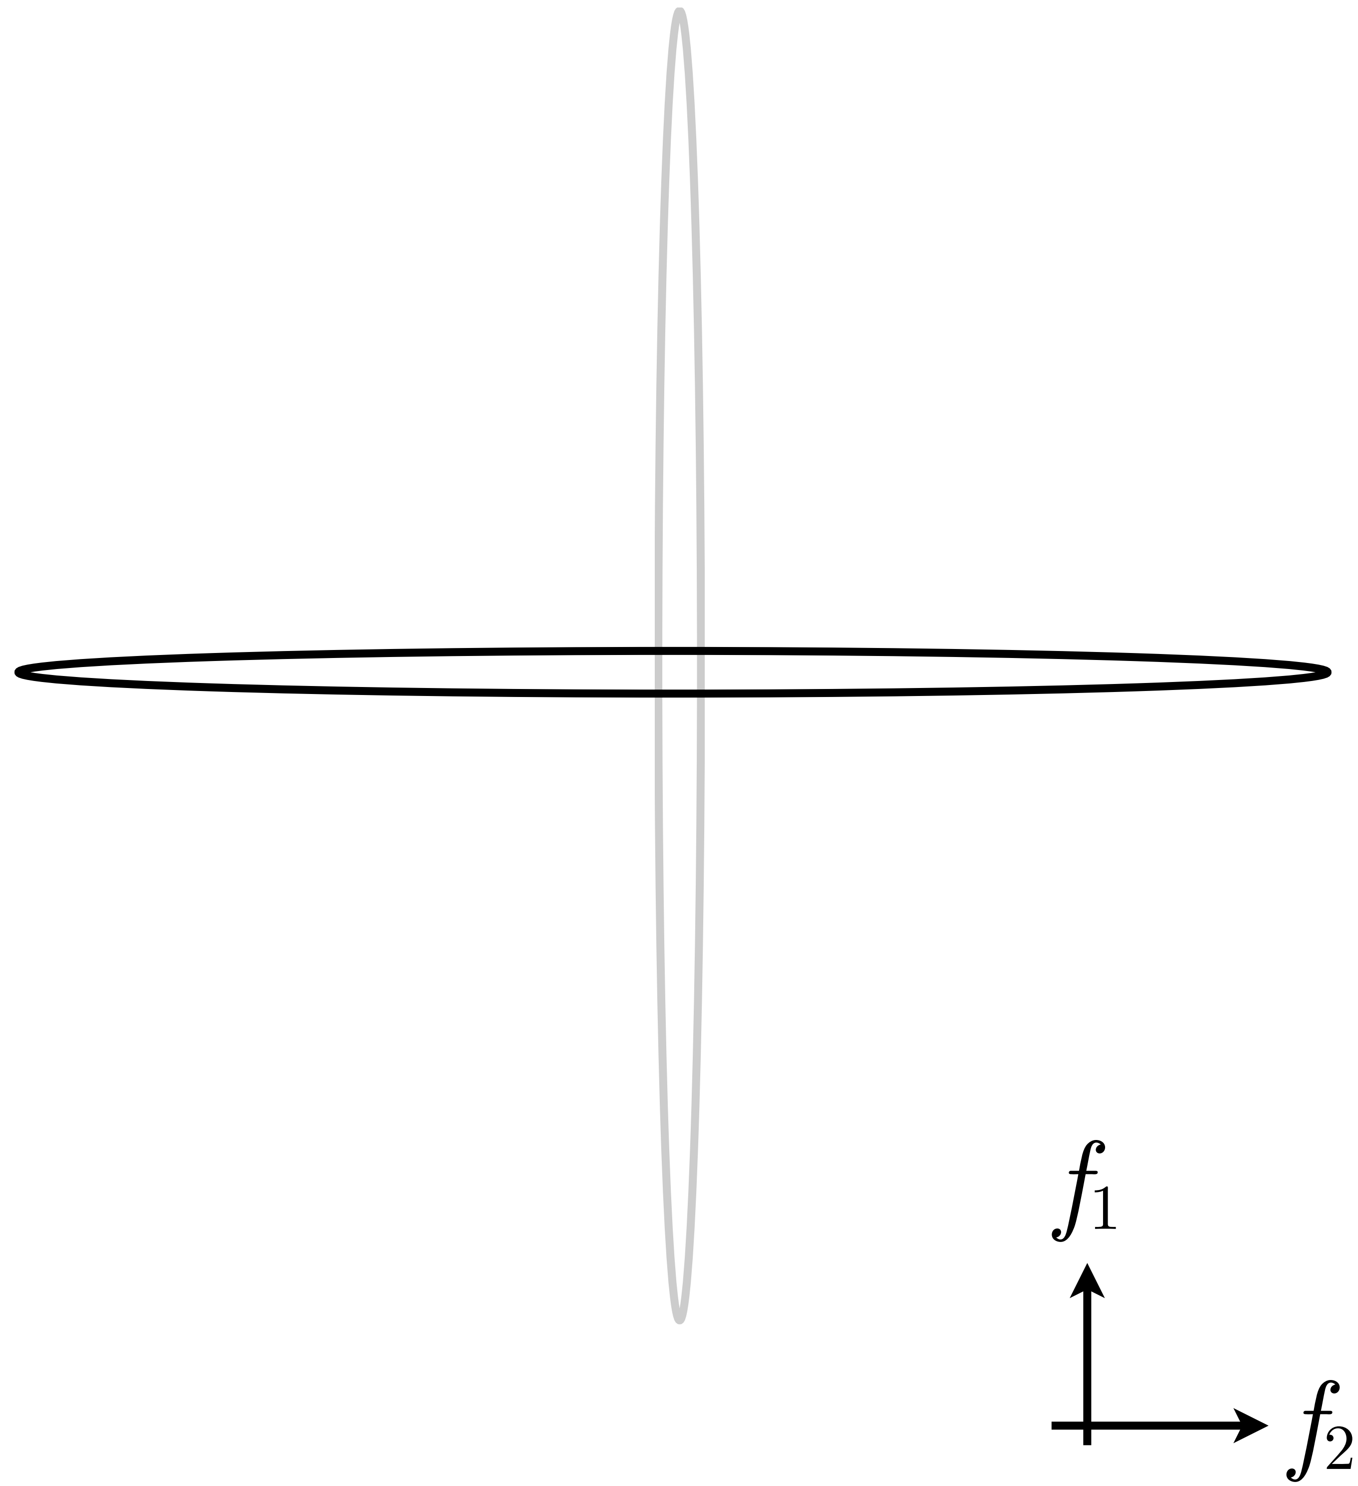

object geometry

f_1

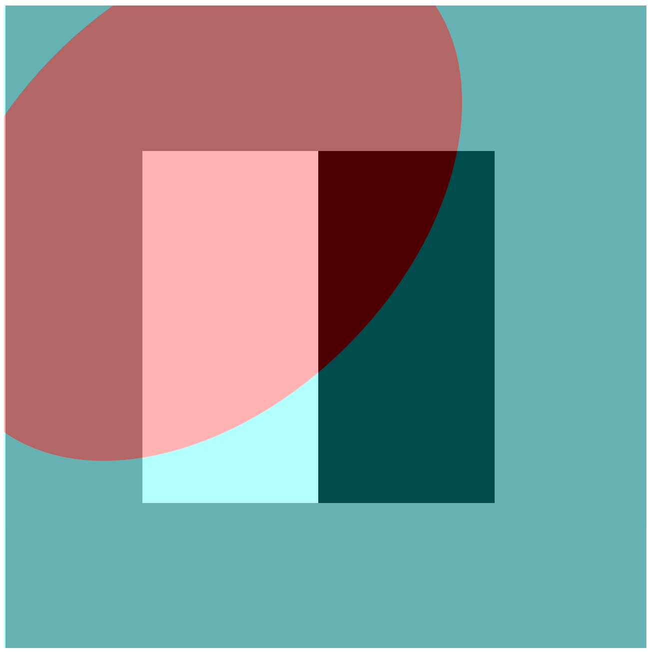

f_2

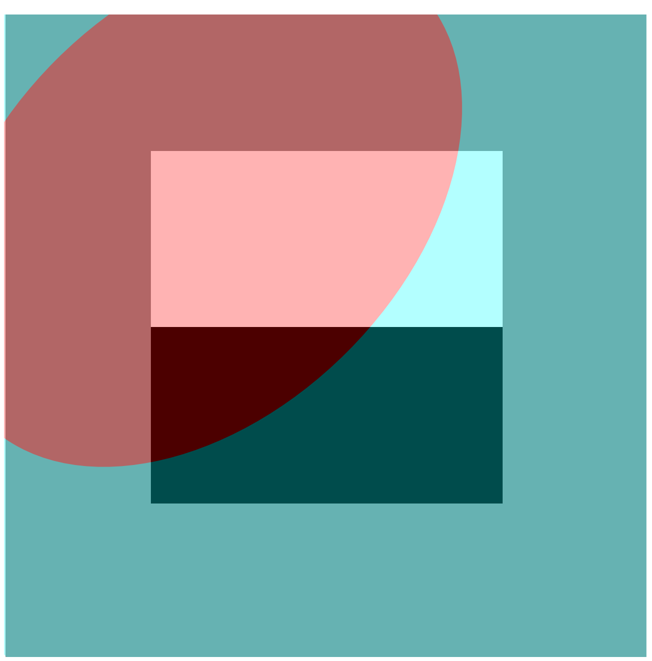

conditional probability contours

gaussian $P(f_1, f_2 | \text{מ})$

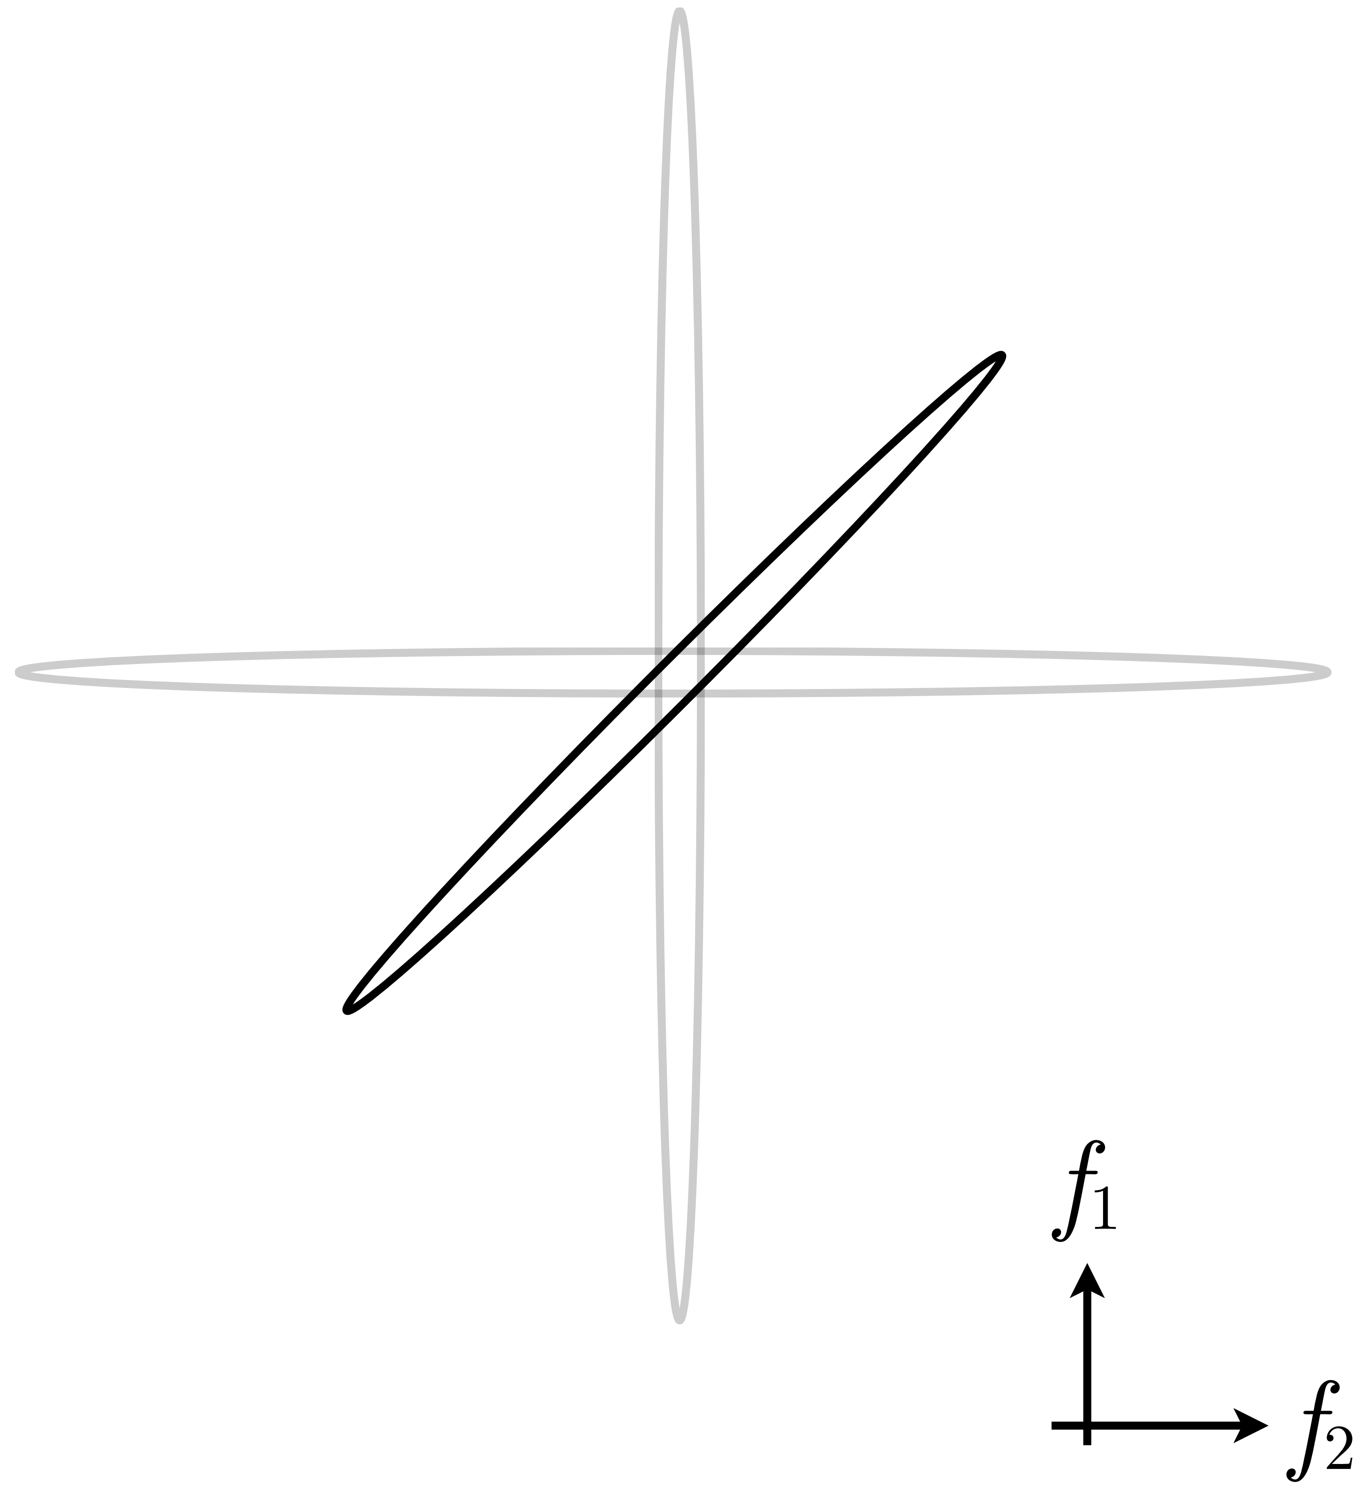

object geometry

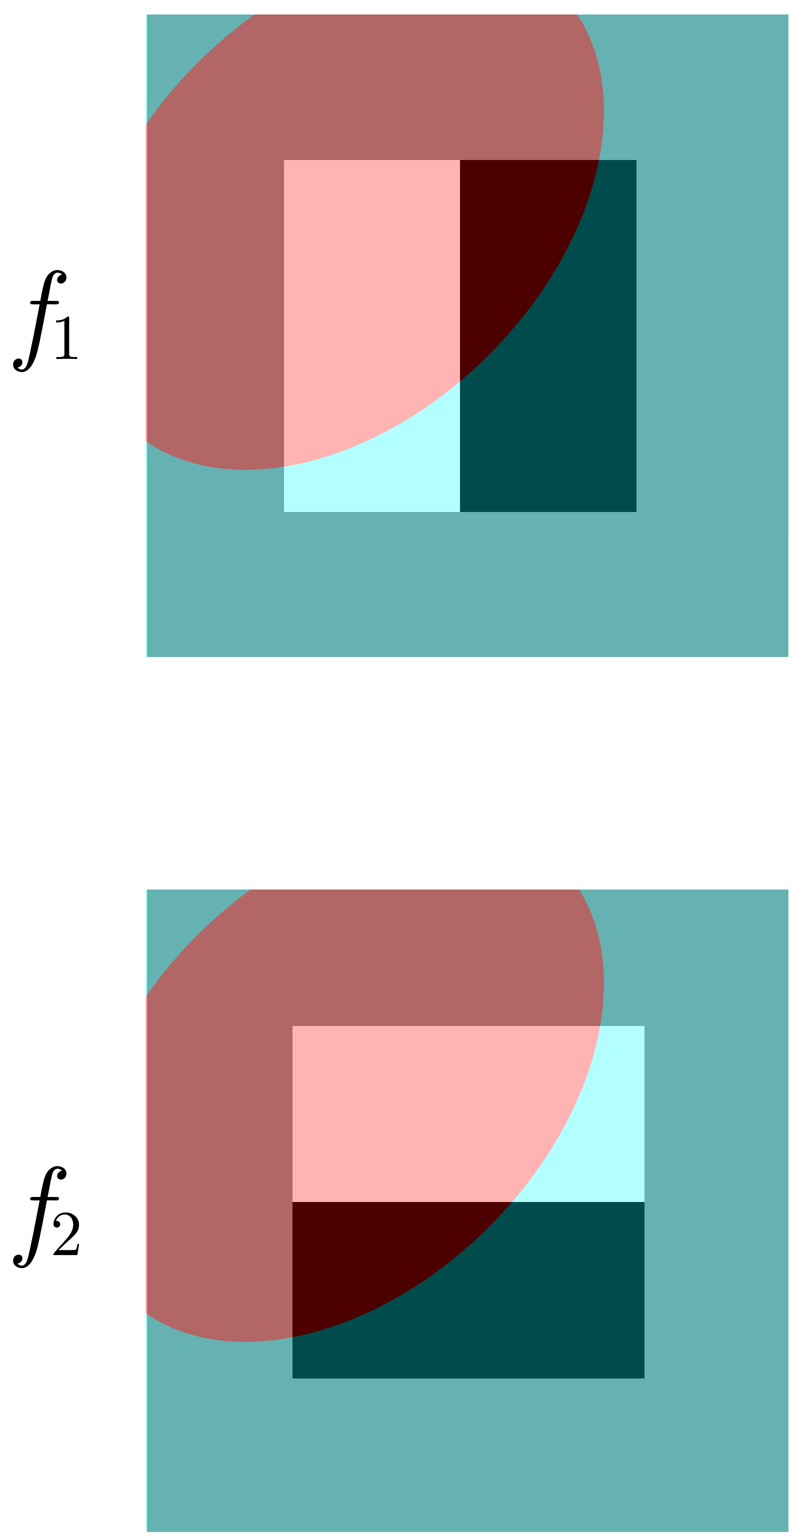

joint probability contours: envelope of $P(f_1, f_2|\text{מ})$

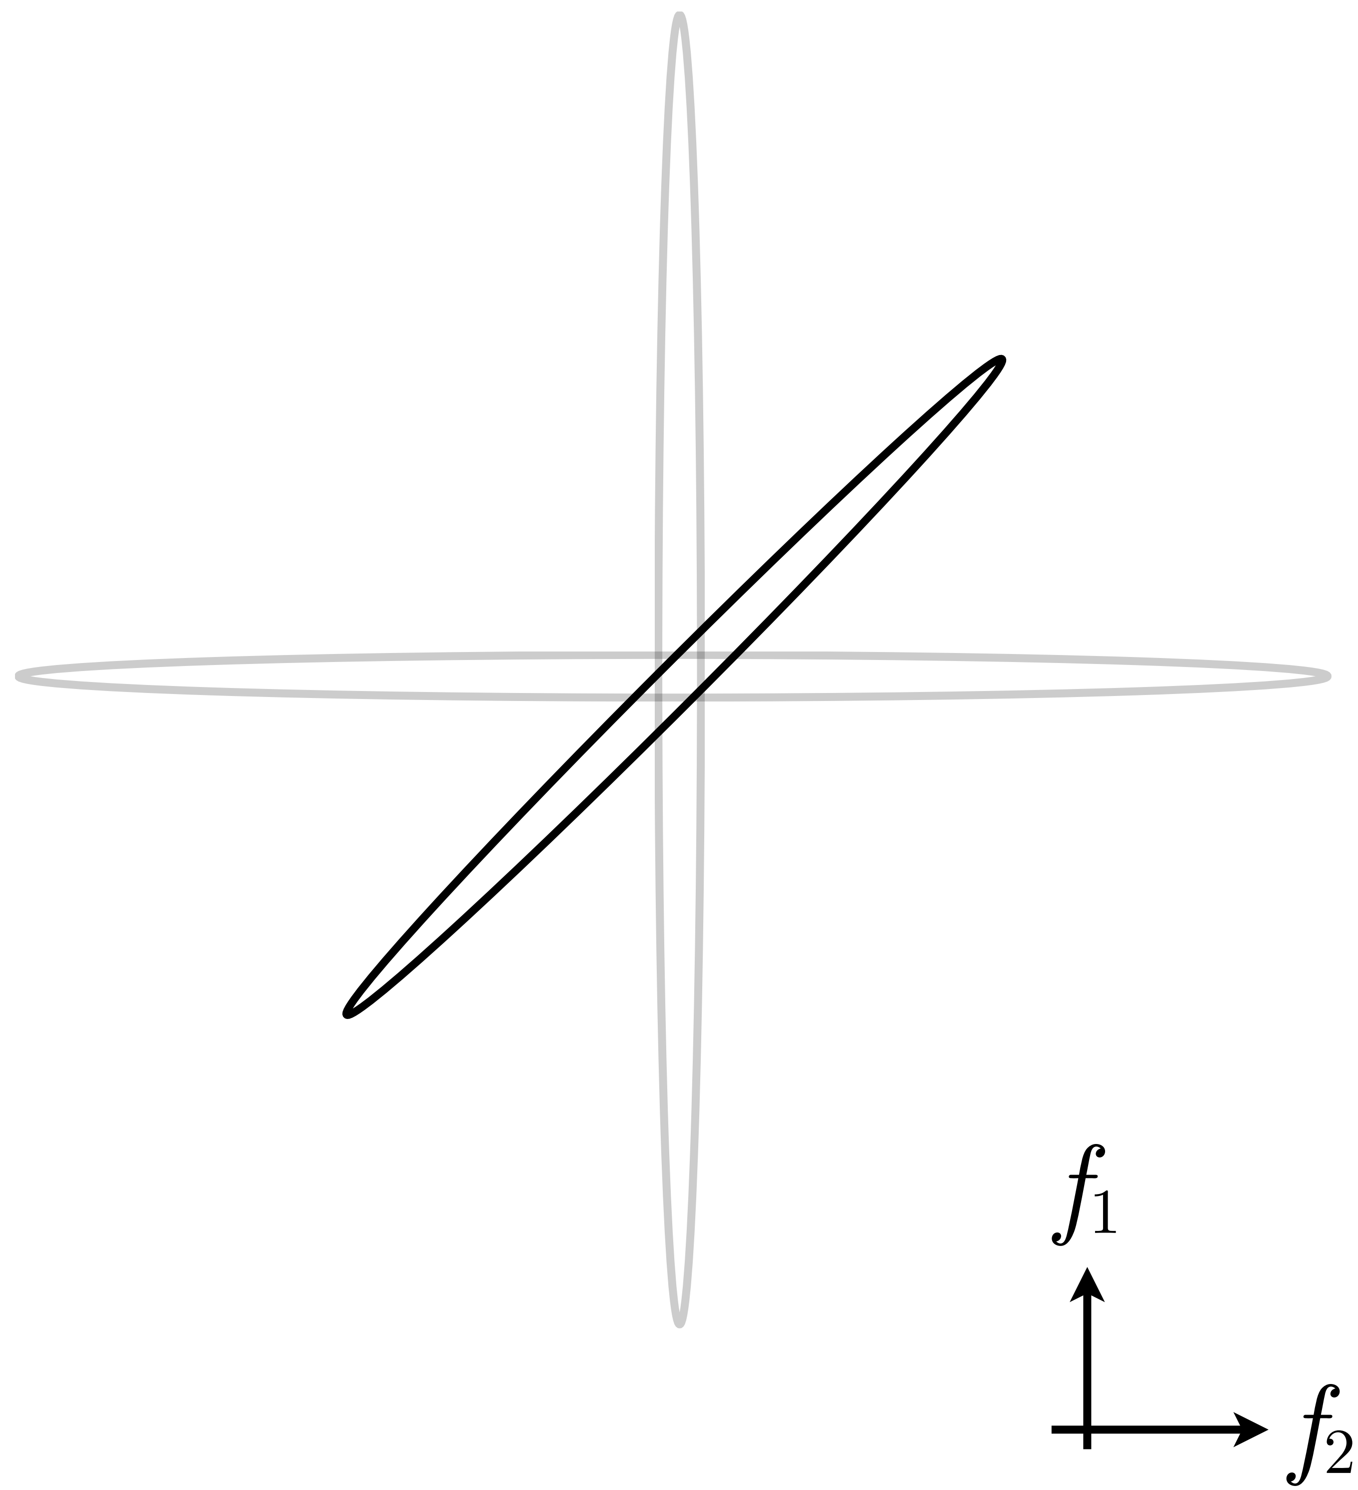

object geometry

joint probability contours: envelope of $P(f_1, f_2|מ)$

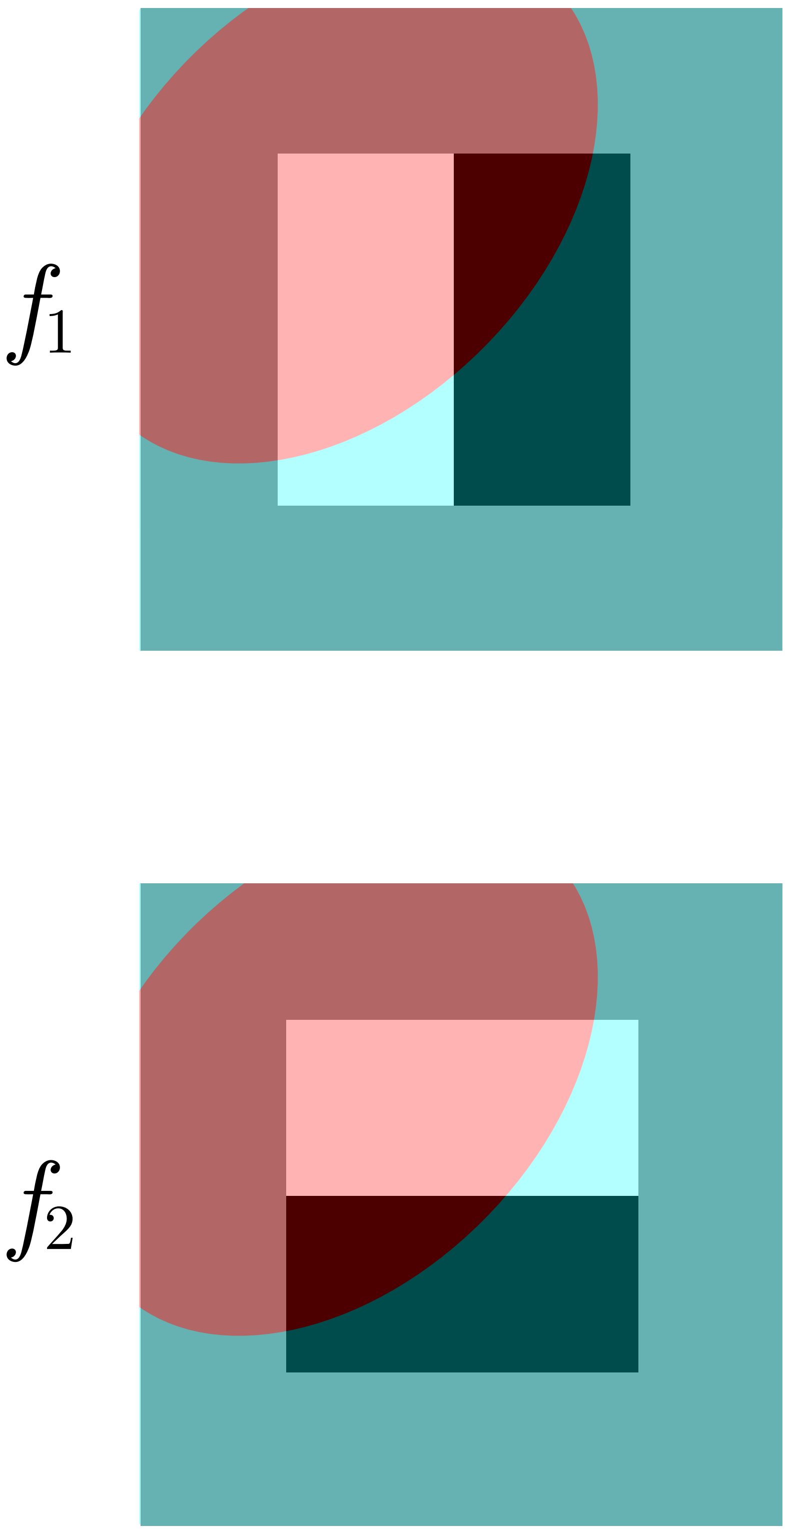

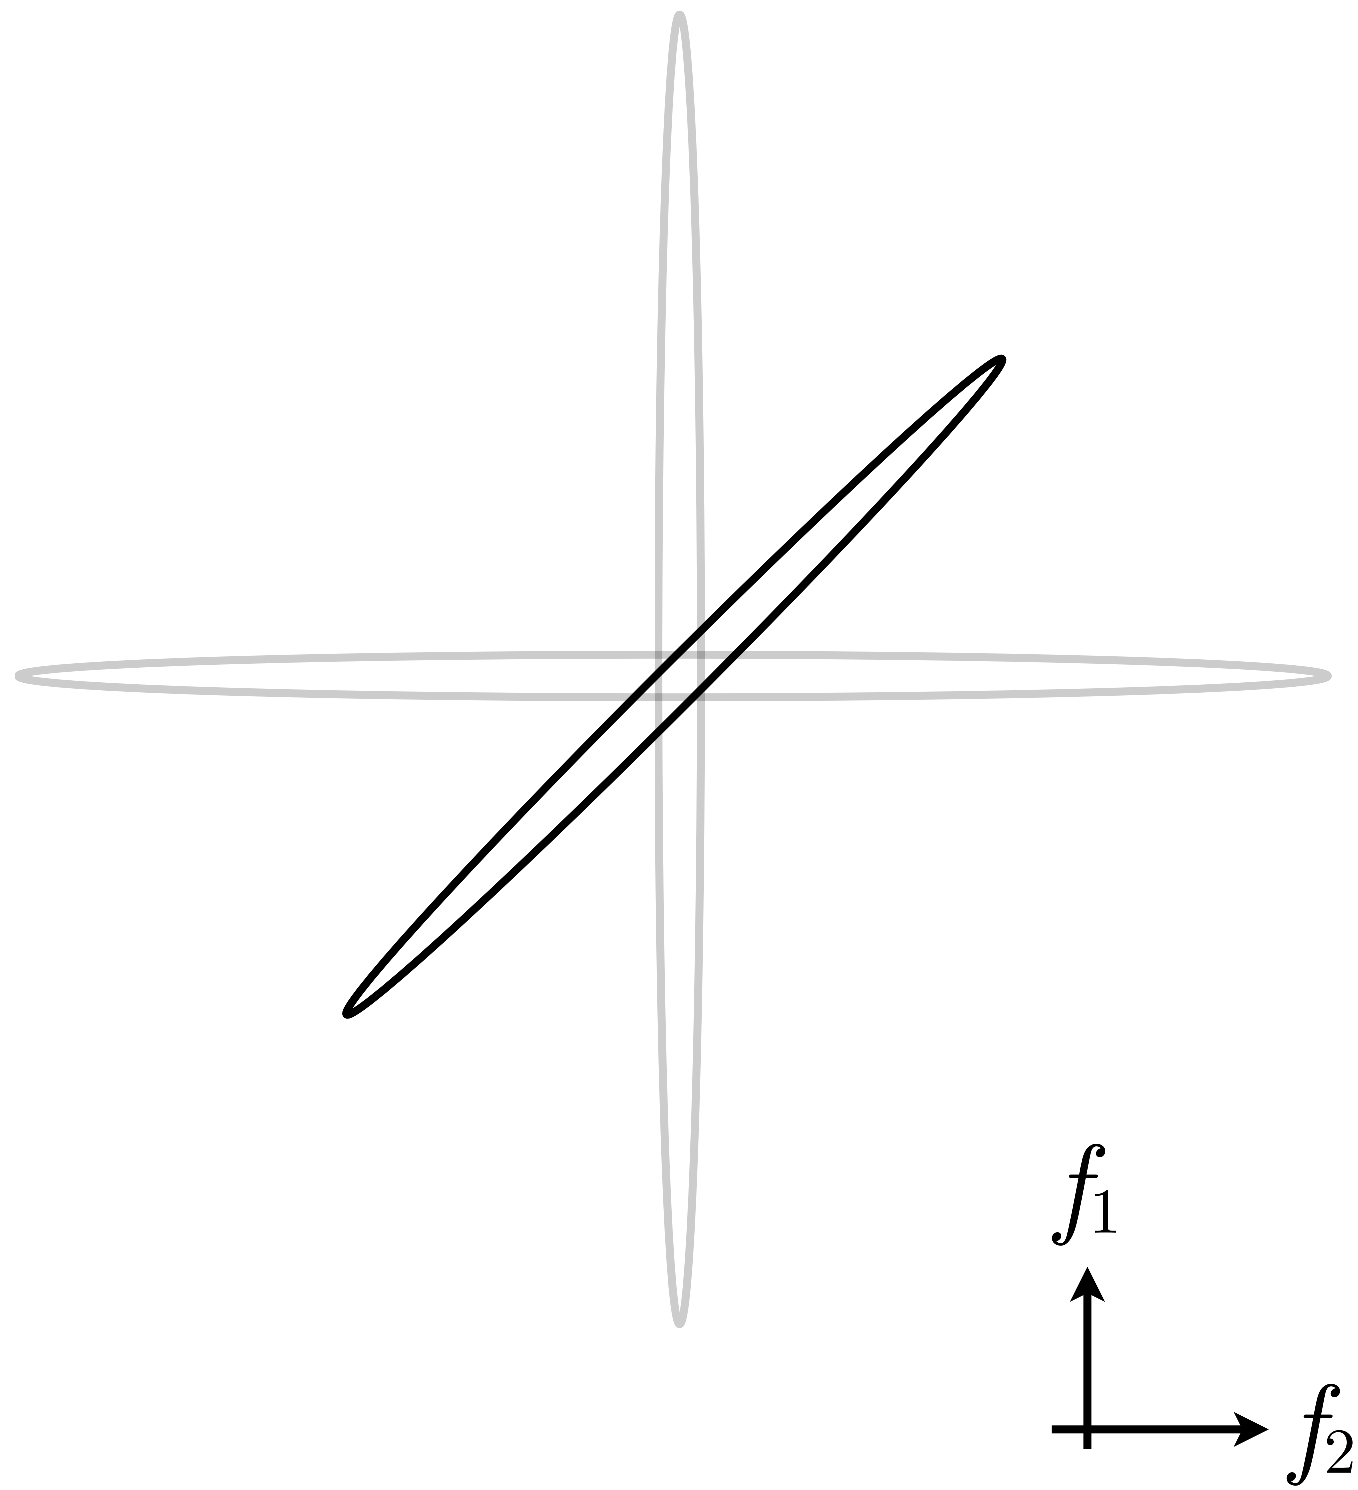

object geometry

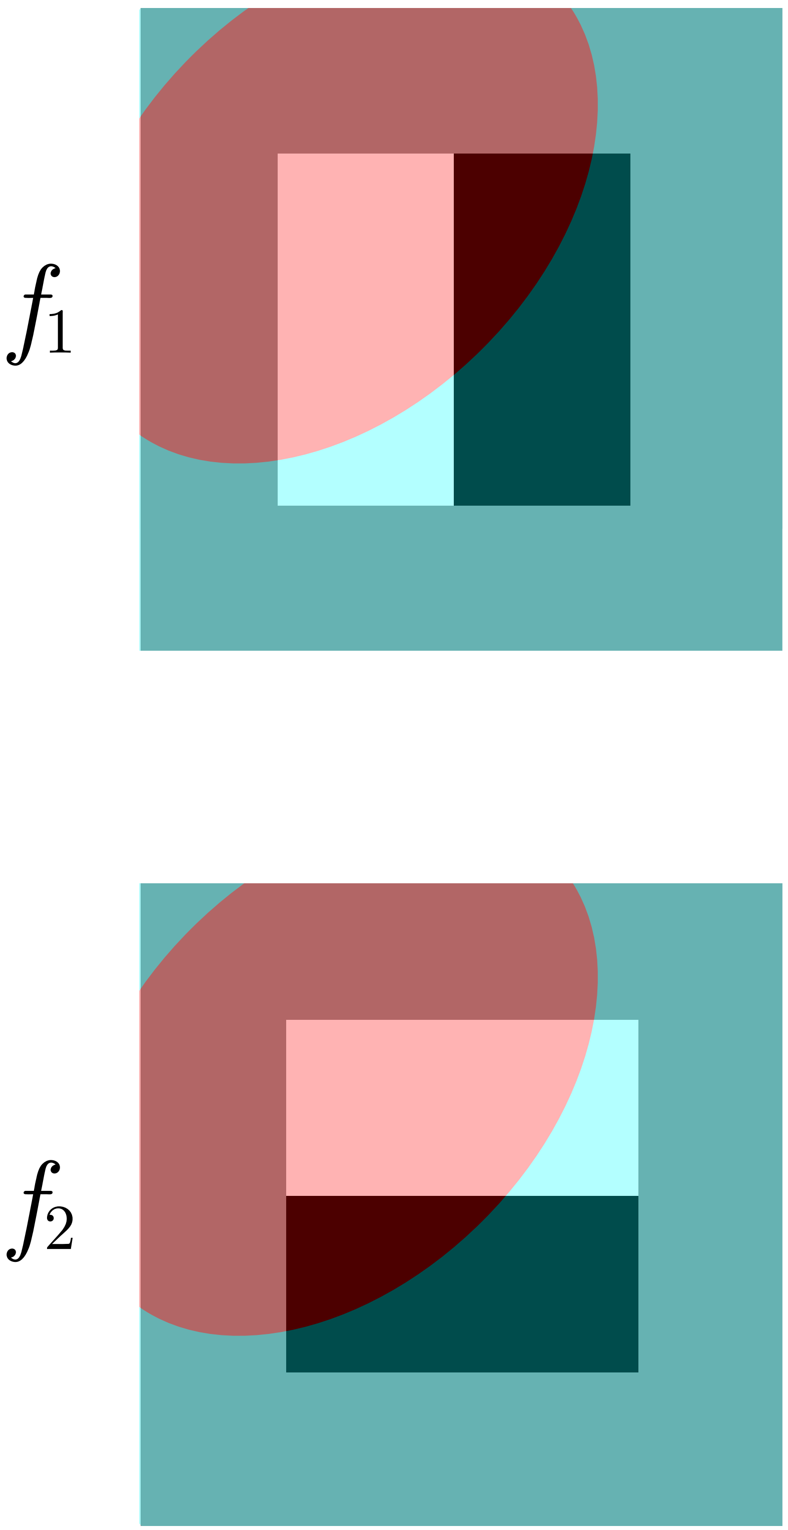

joint probability contours: envelope of $P(f_1, f_2|\text{מ})$

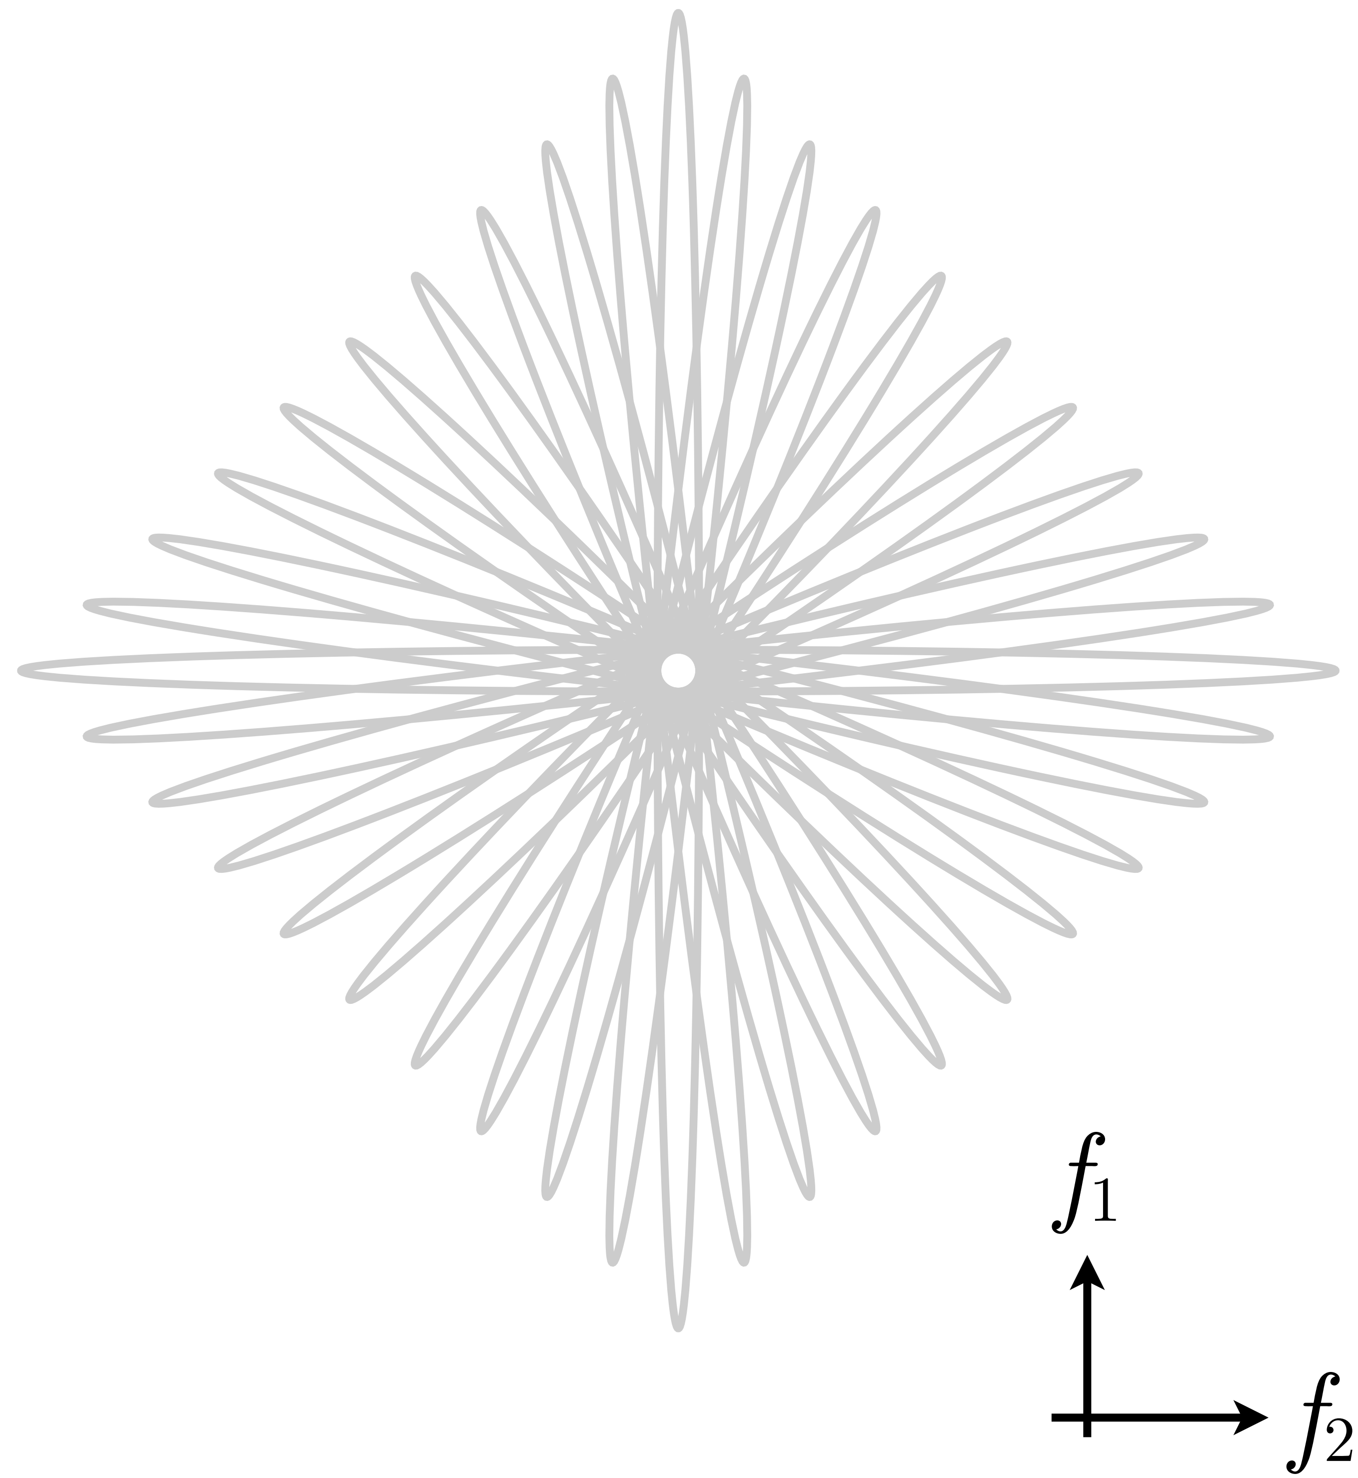

object geometry

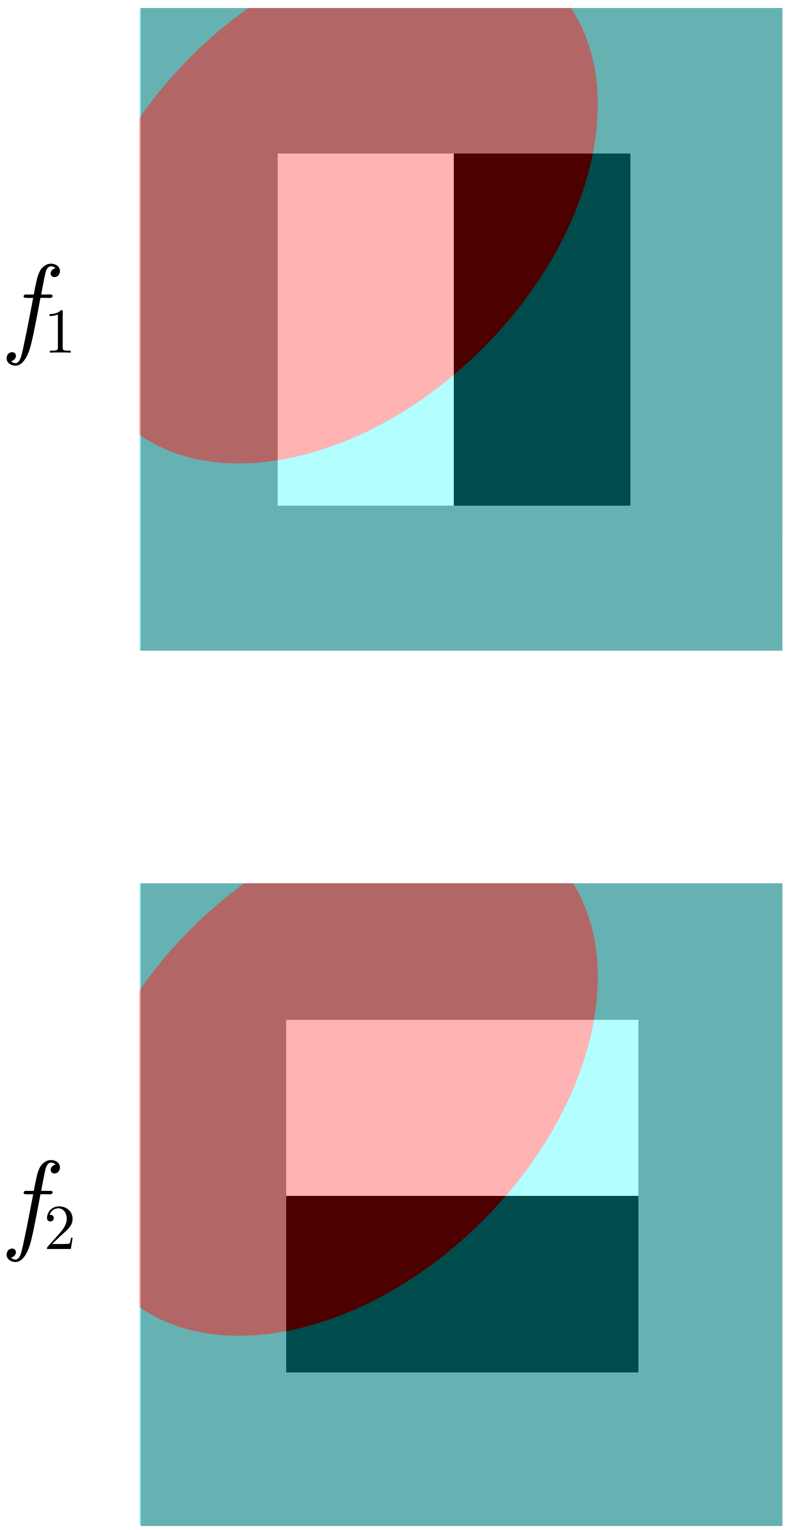

joint probability contours: envelope of $P(f_1, f_2|\text{מ})$

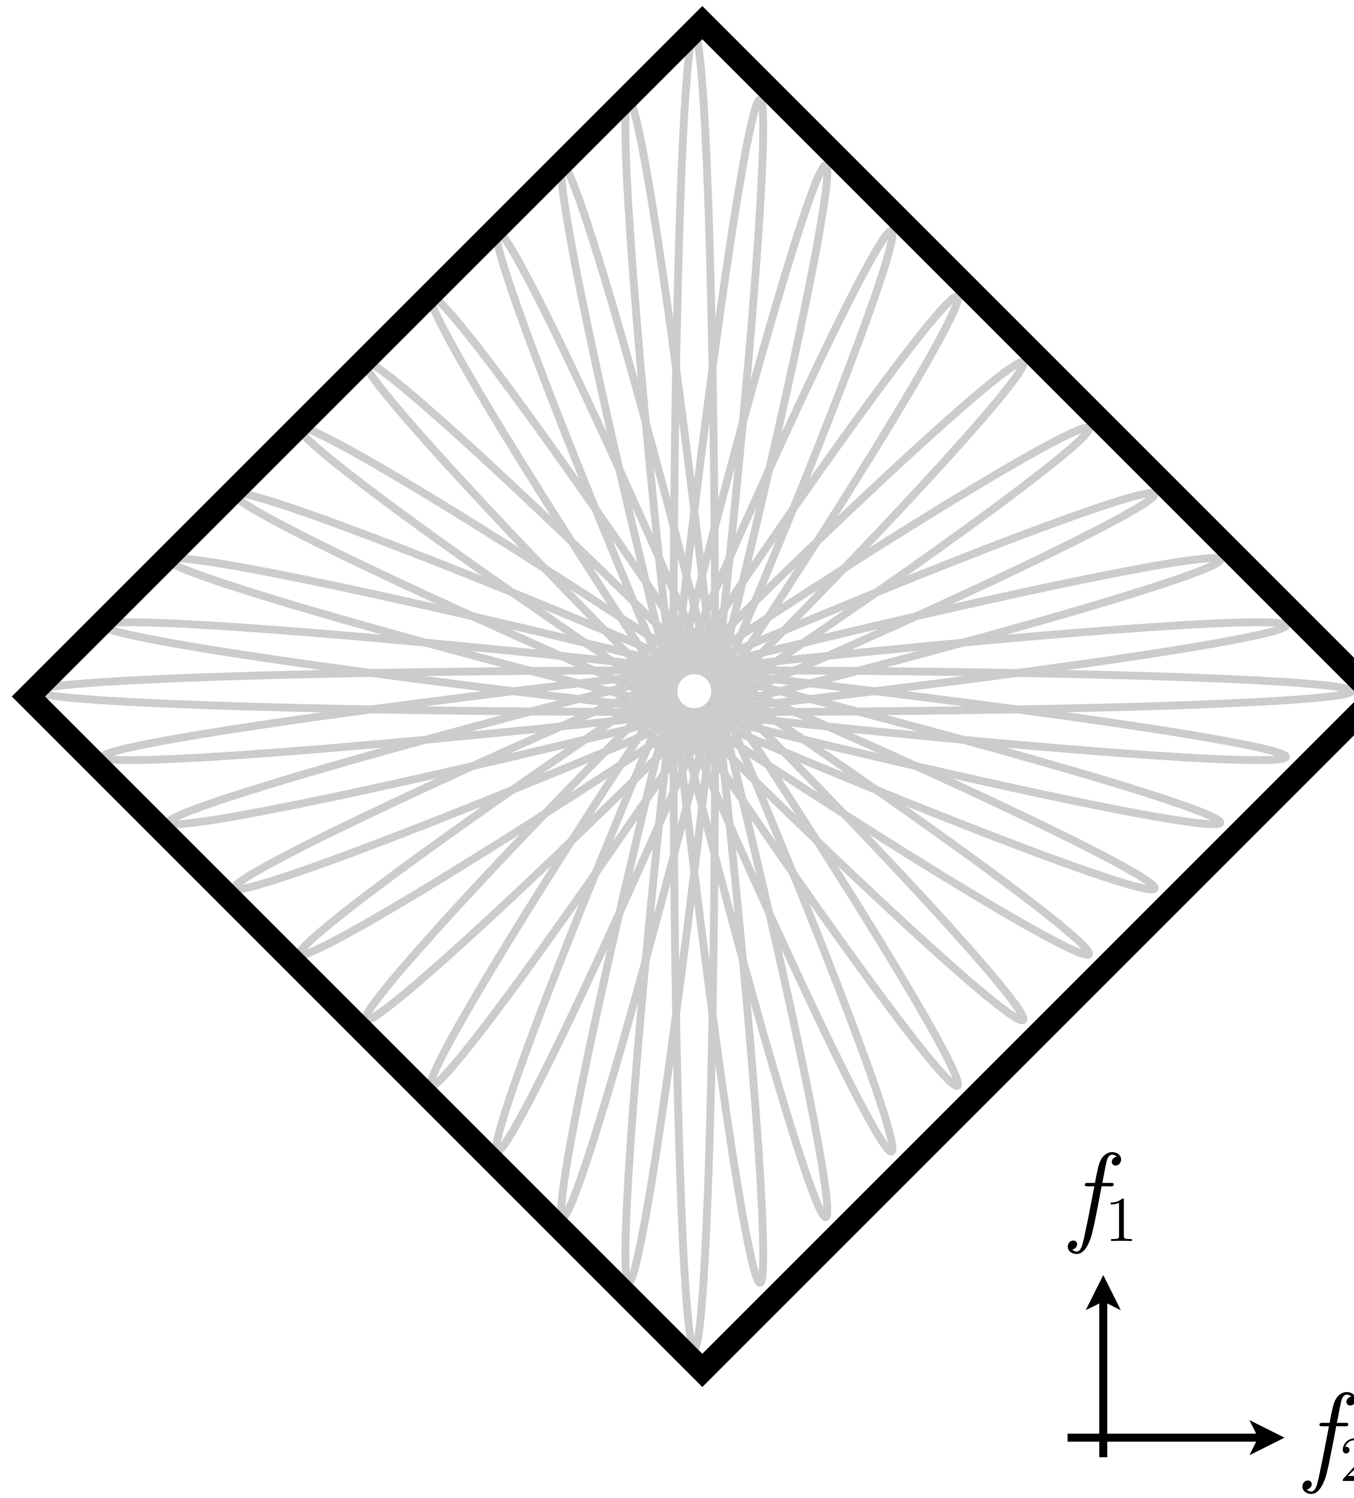

joint distributions: orthogonal, colocalized Haar wavelets

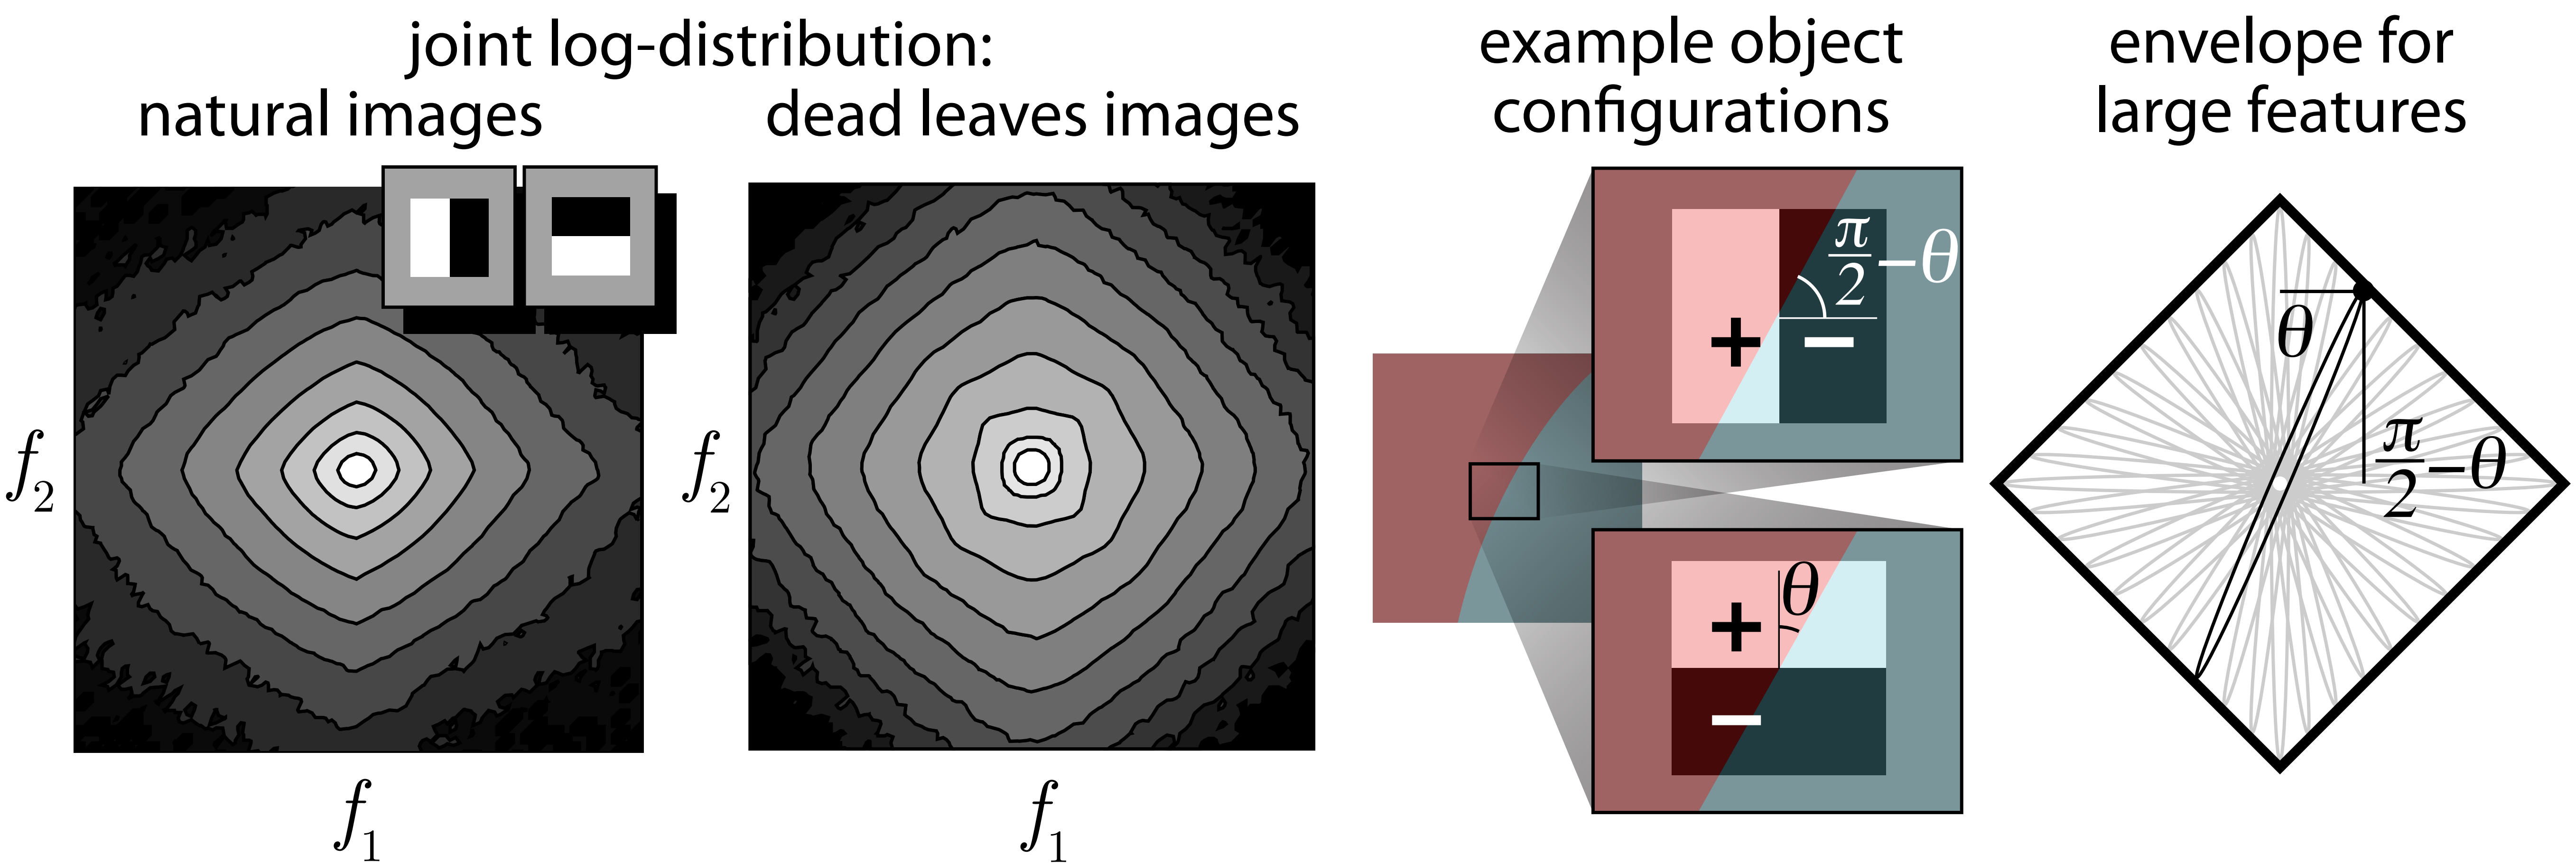

Lee, Mumford, Huang 2001

joint distributions:
parallel, neighboring Haar wavelets

joint log-distribution:

natural images | dead leaves images

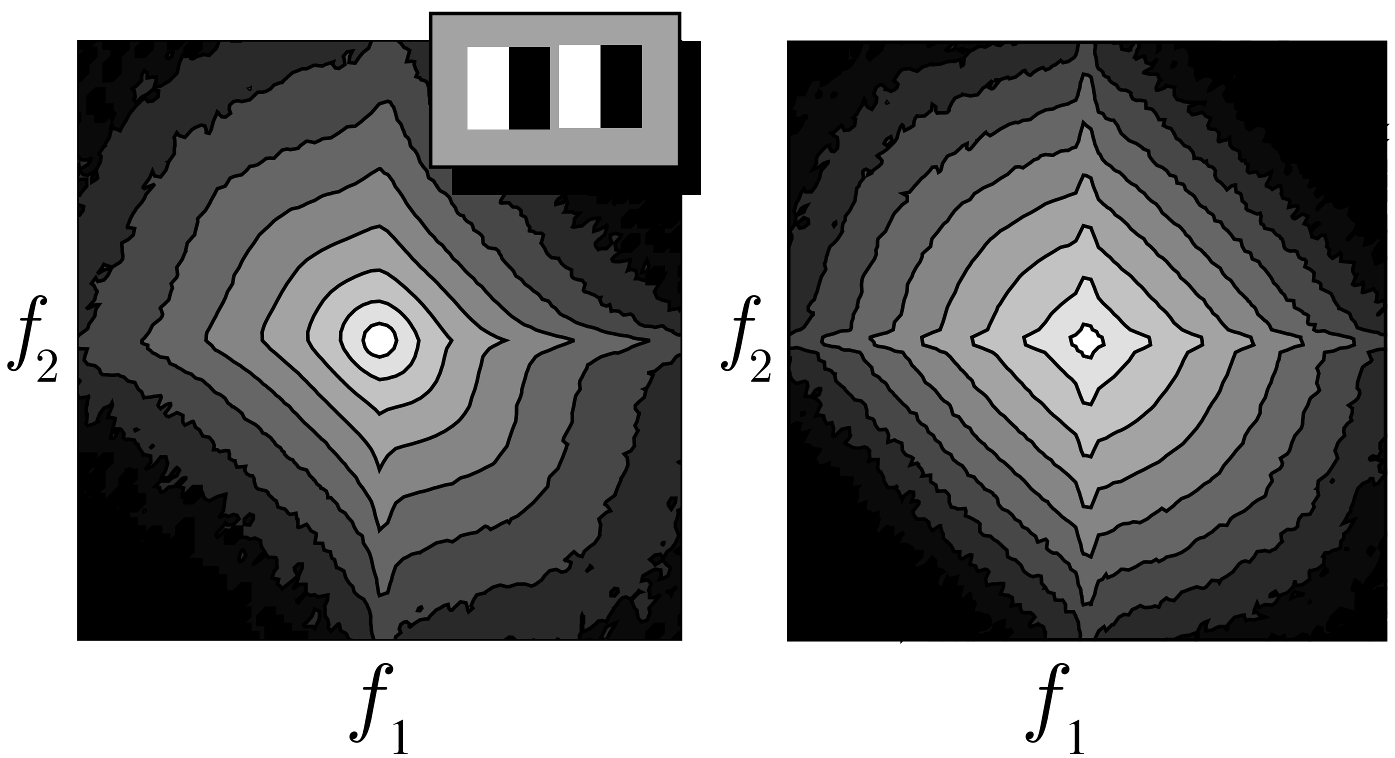

example object configurations

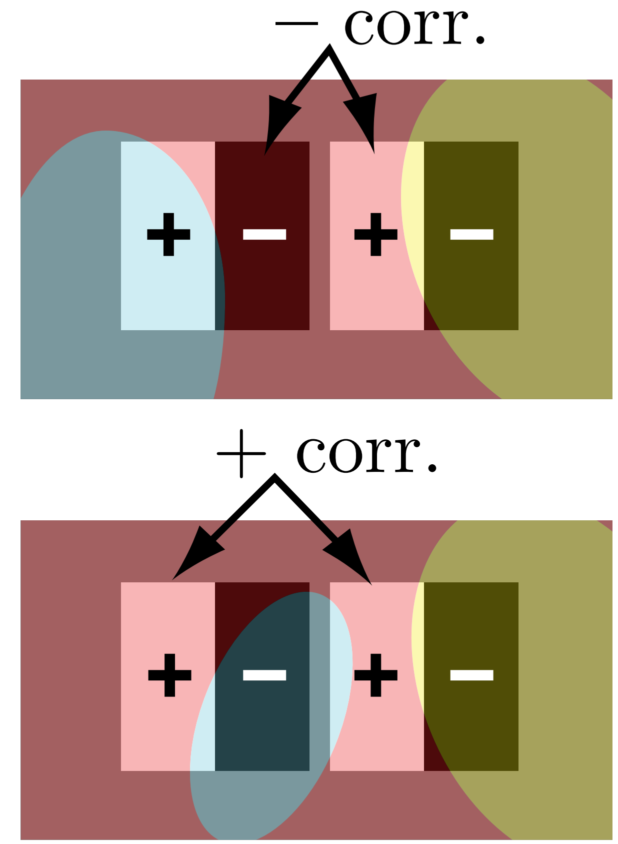

envelope for large features

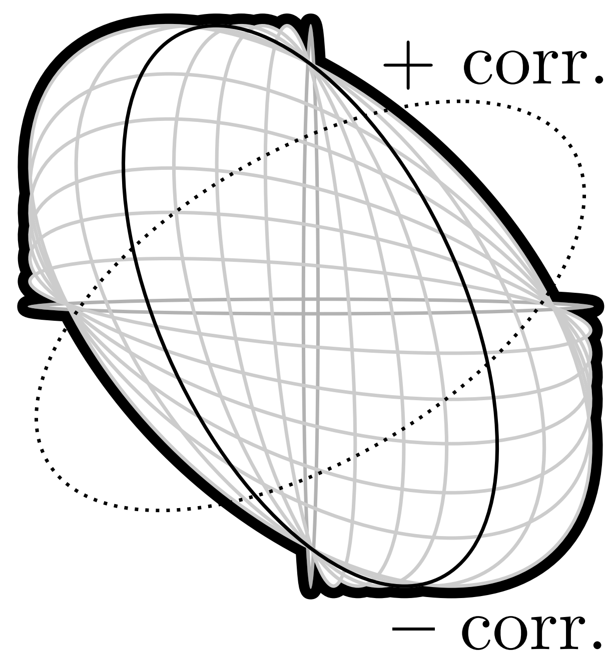

Lee, Mumford, Huang 2001

conditional distributions: orthogonal, neighboring Haar wavelets

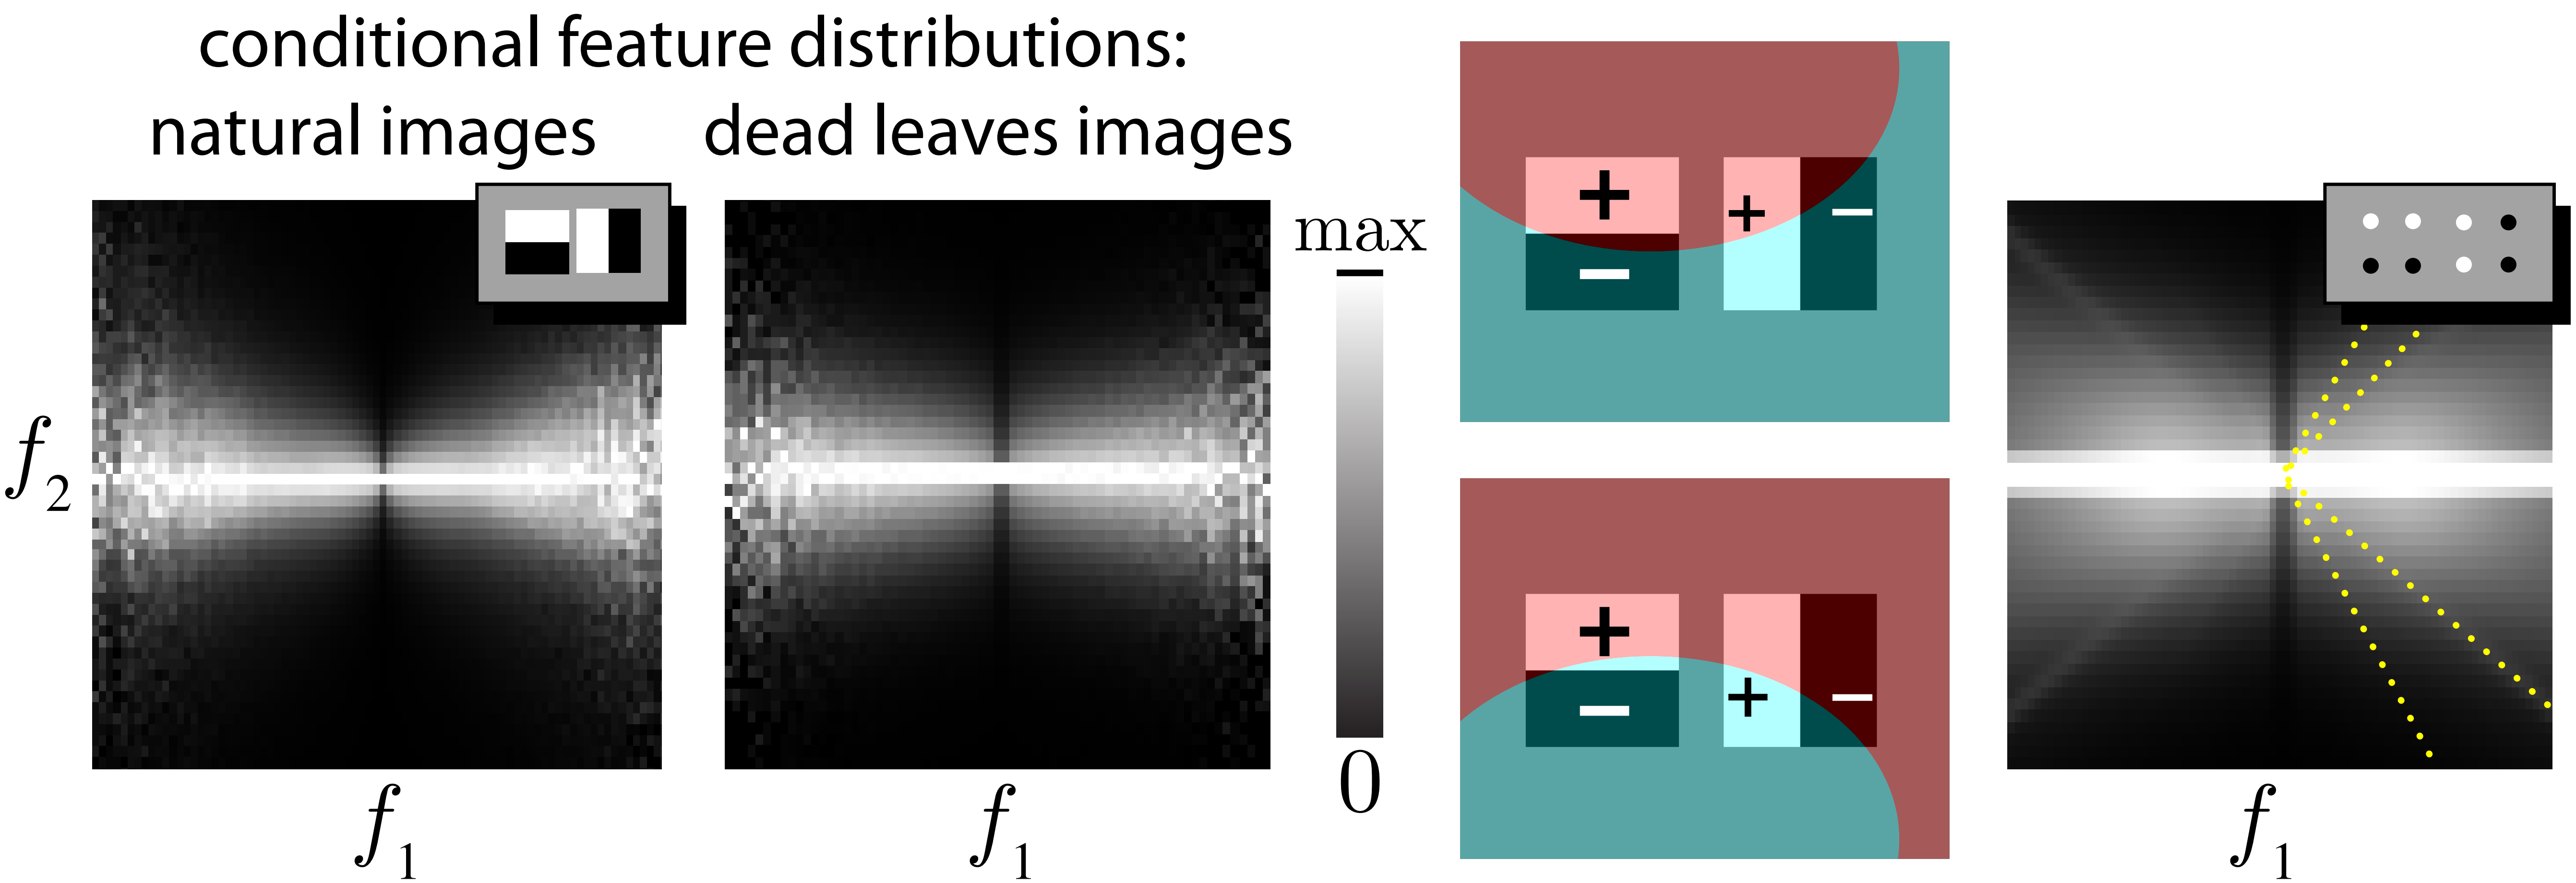

Simoncelli & Schwartz 1999
Buccigrossi & Simoncelli 1999

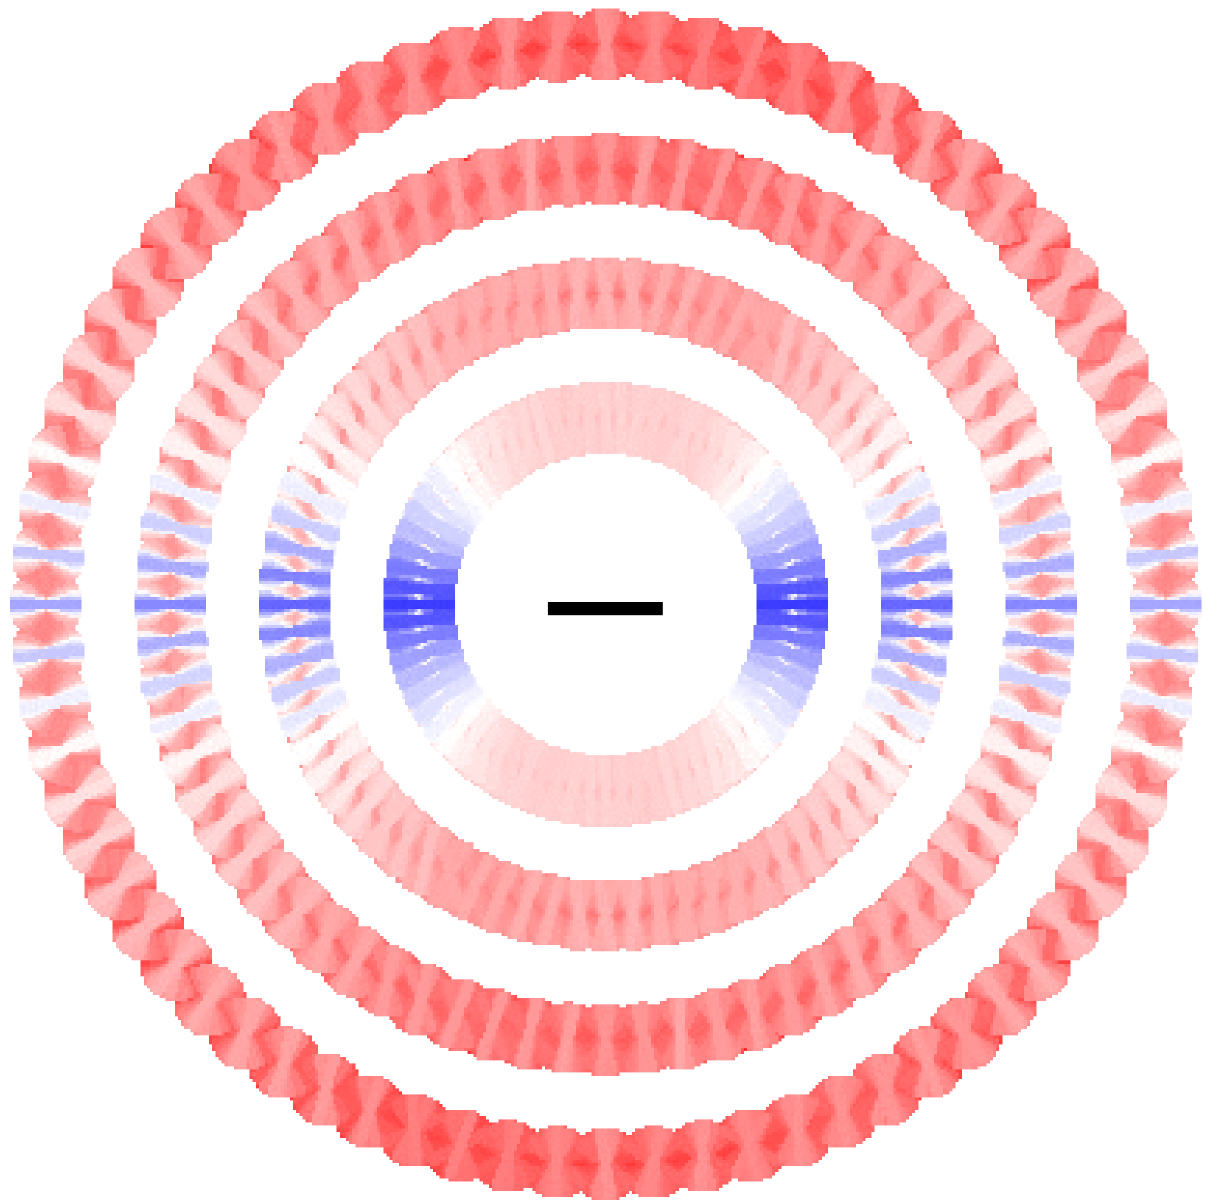

likelihood ratio:
shared cause / different causes

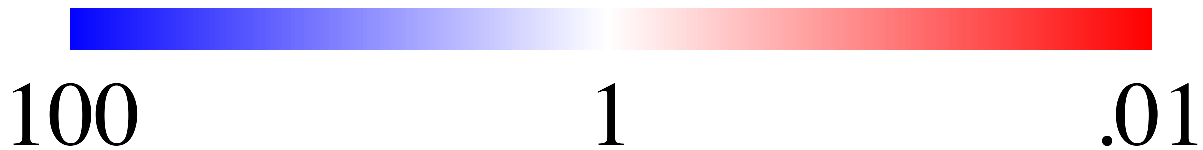

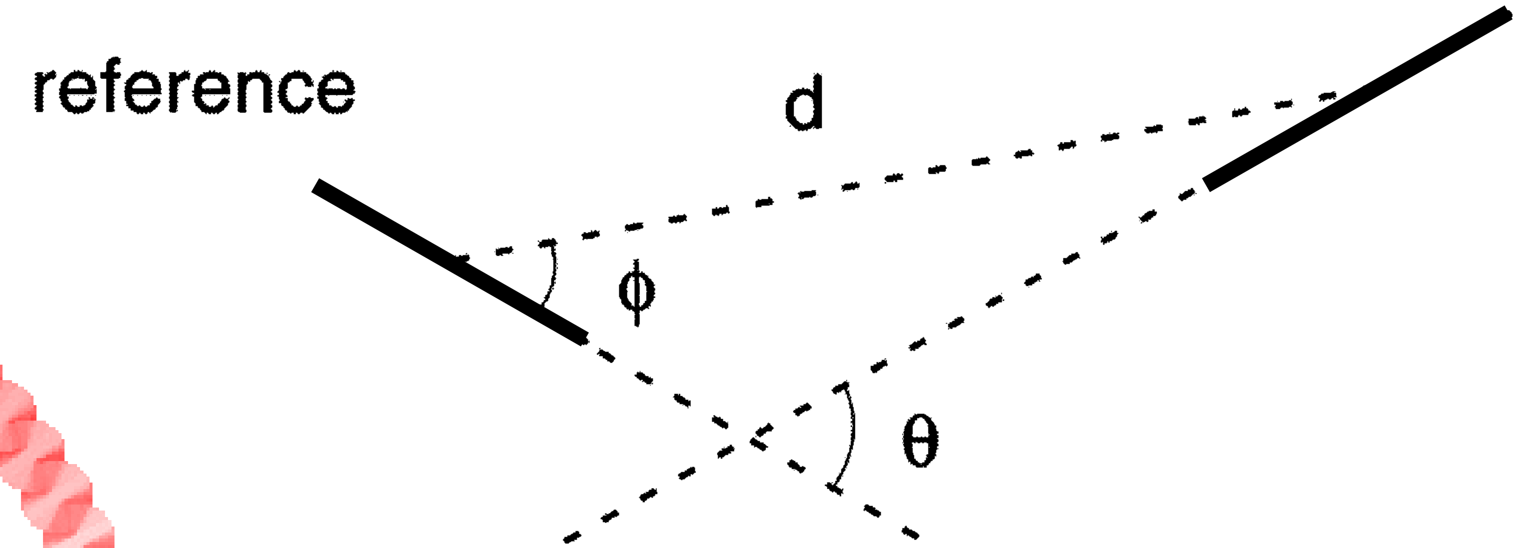

geisler et al, 2001

Likelihood of shared cause:

$$L = \frac{\sum_{מ \in S} P_{מ}}{\sum_{מ \in D} P_{מ}}$$

D : set of מ with different cause

S : set of מ with shared cause

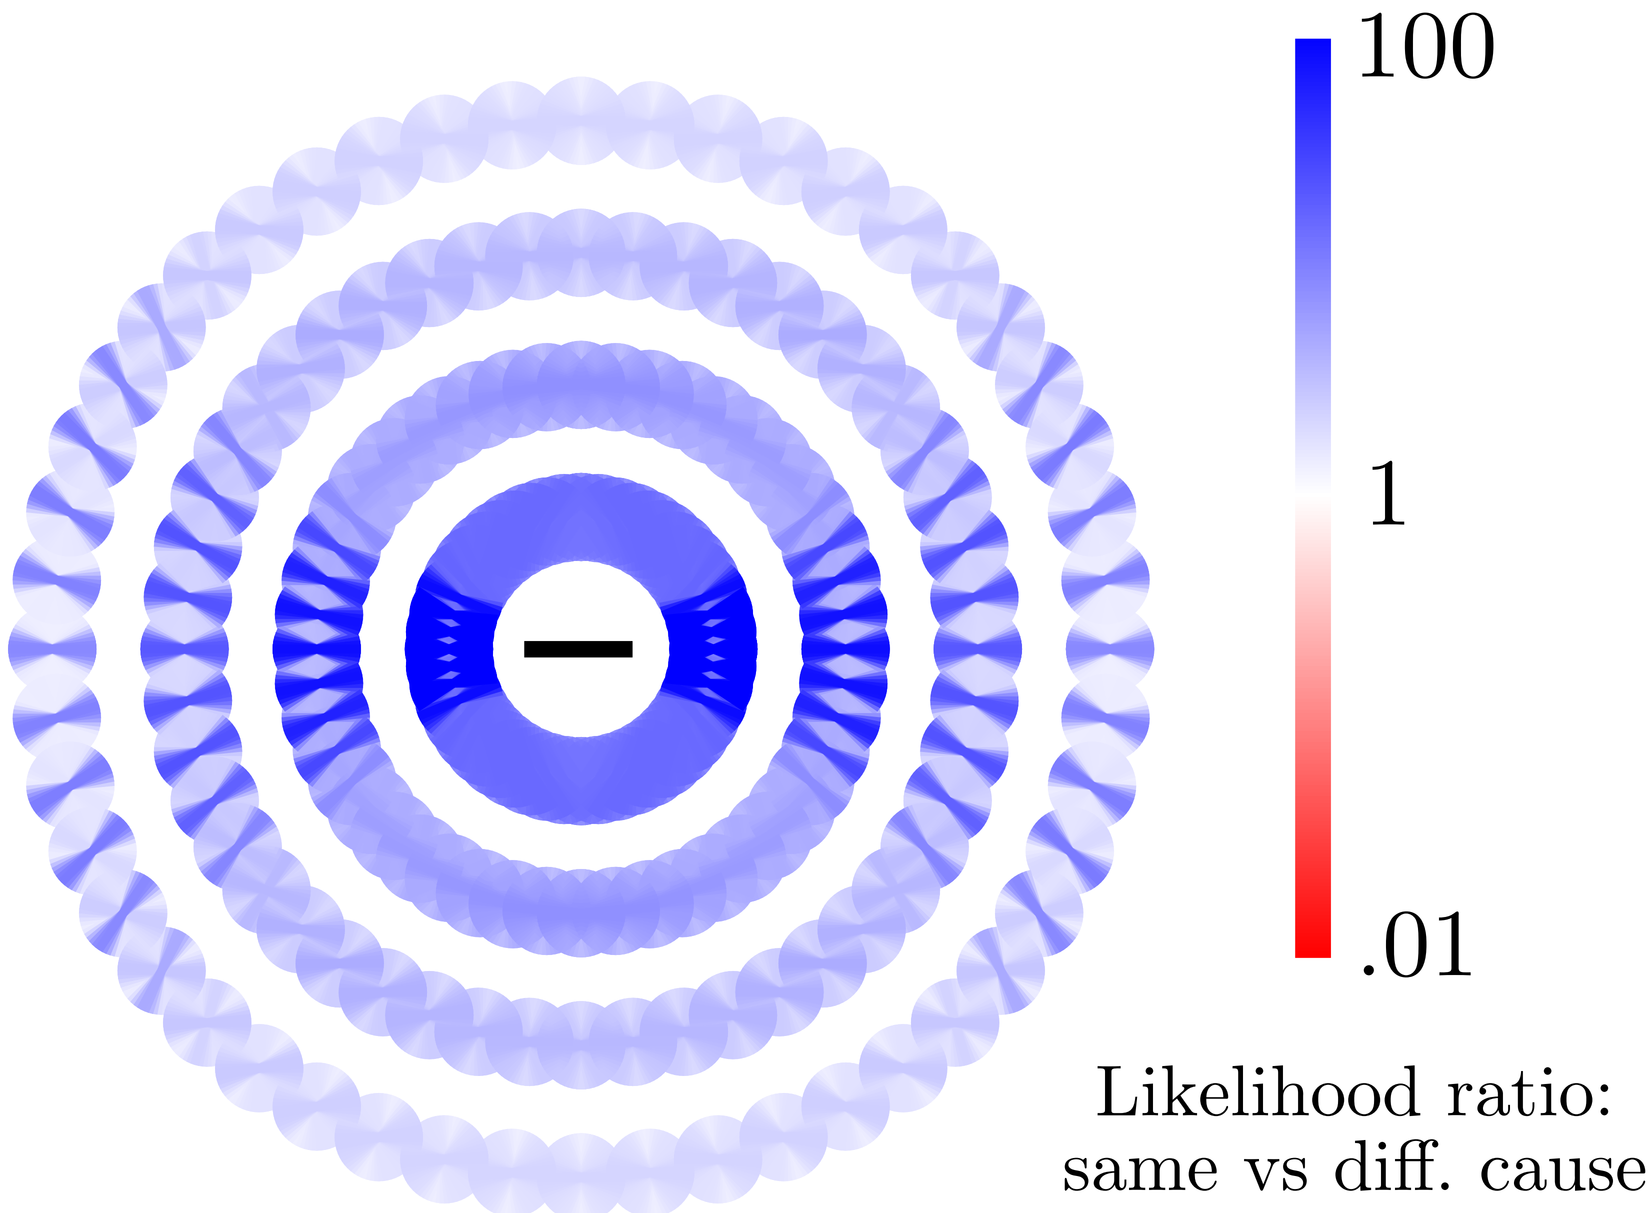
100
1
.01
Likelihood ratio:
same vs diff. cause

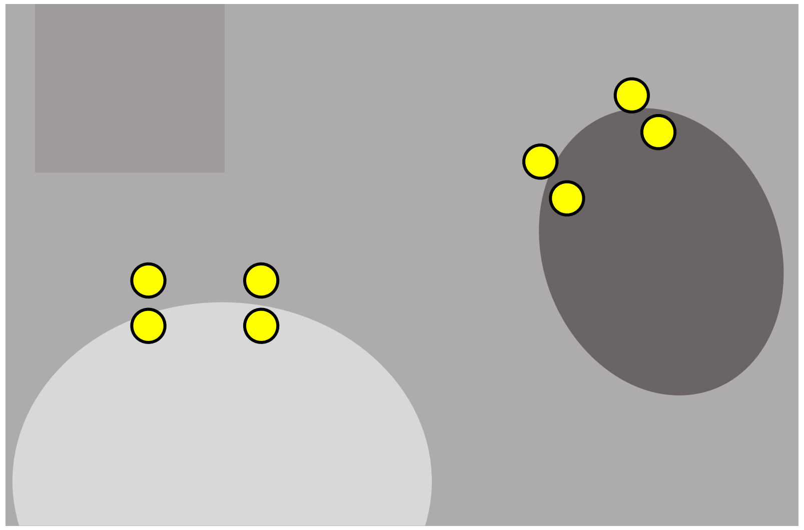

edge causes defined by object membership

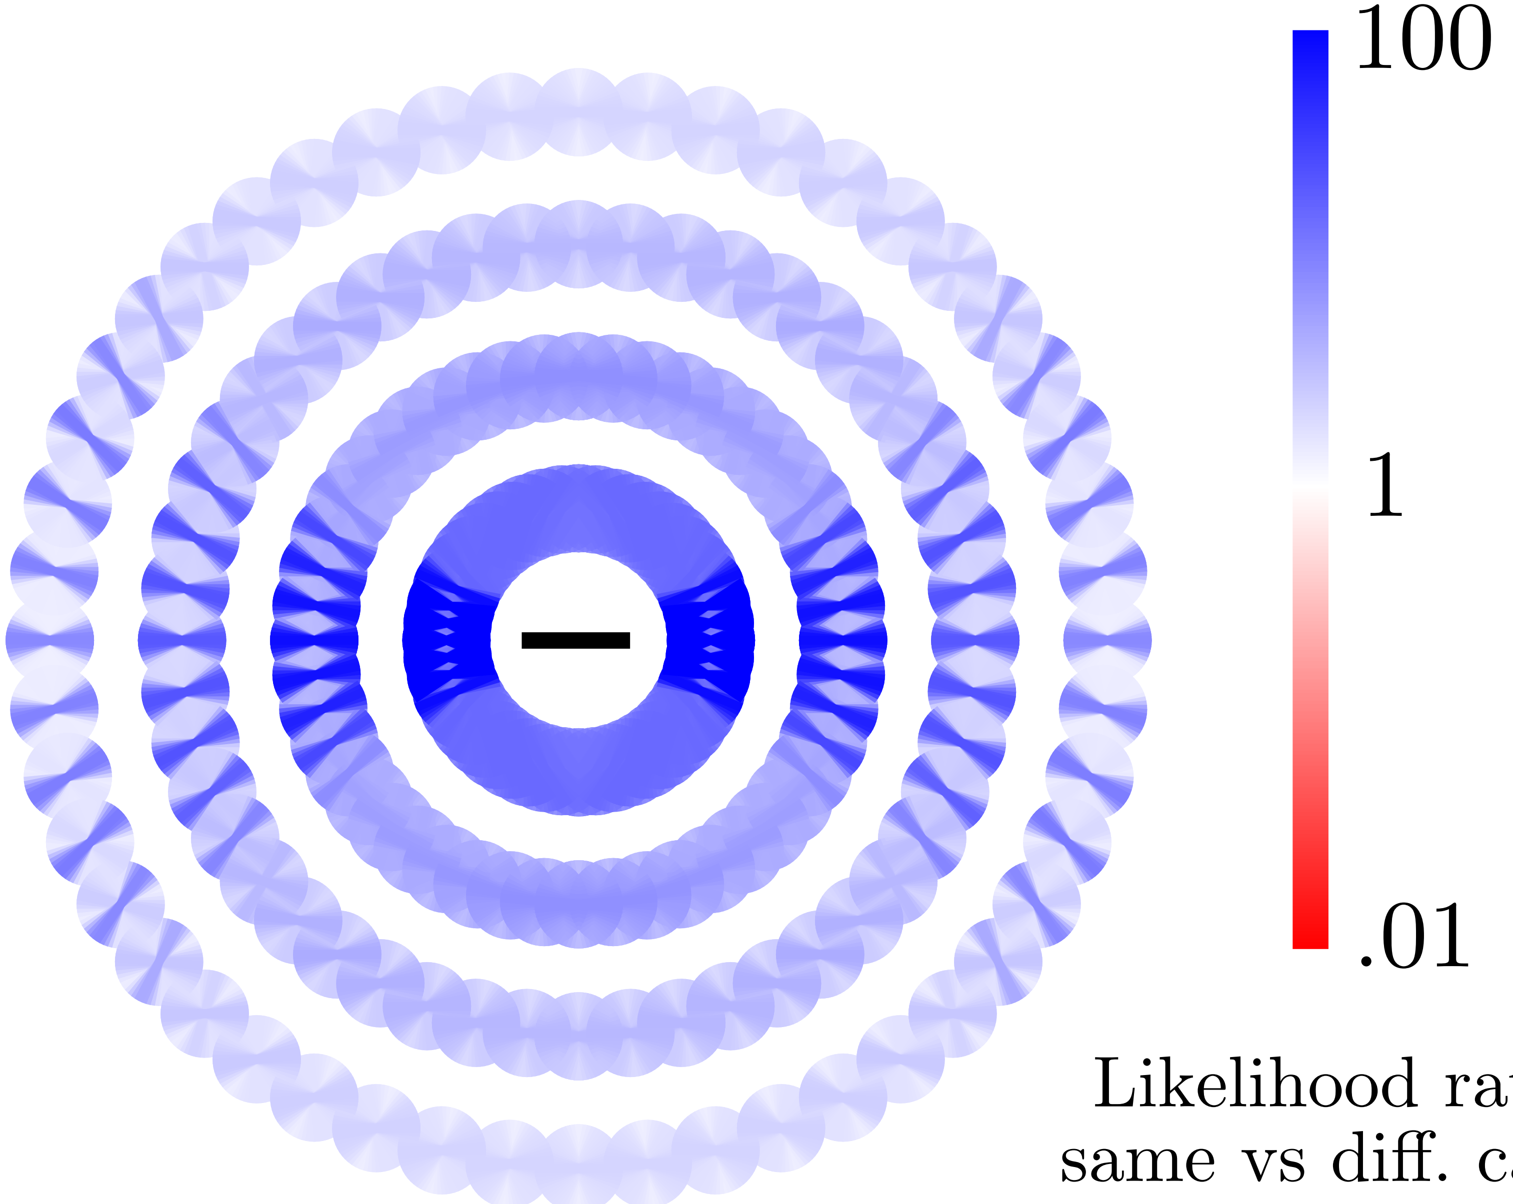

Likelihood ratio: same vs diff. cause

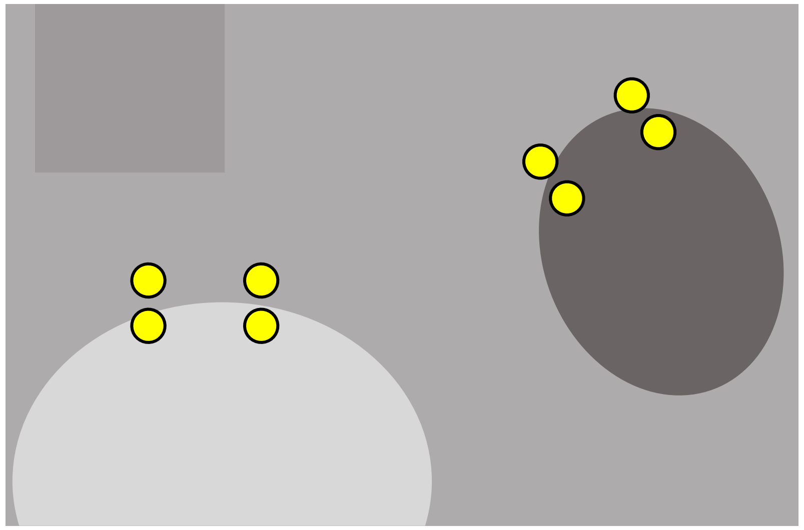

edge causes defined by object membership

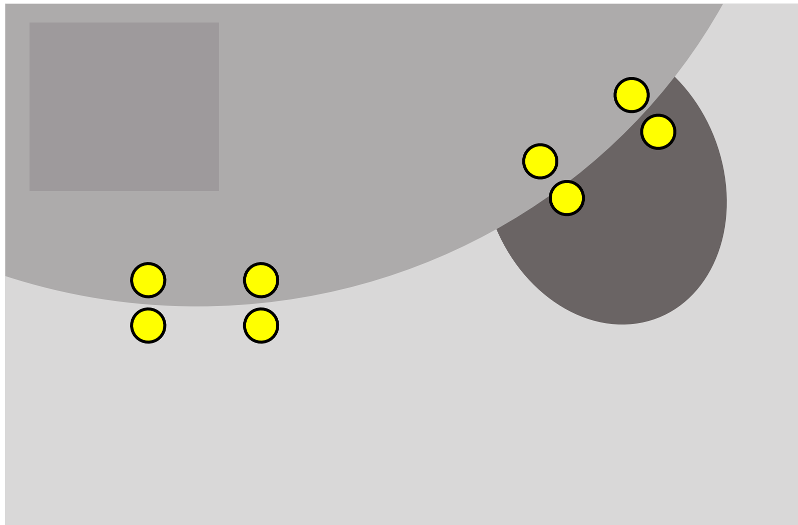

edge causes defined with depth constraints

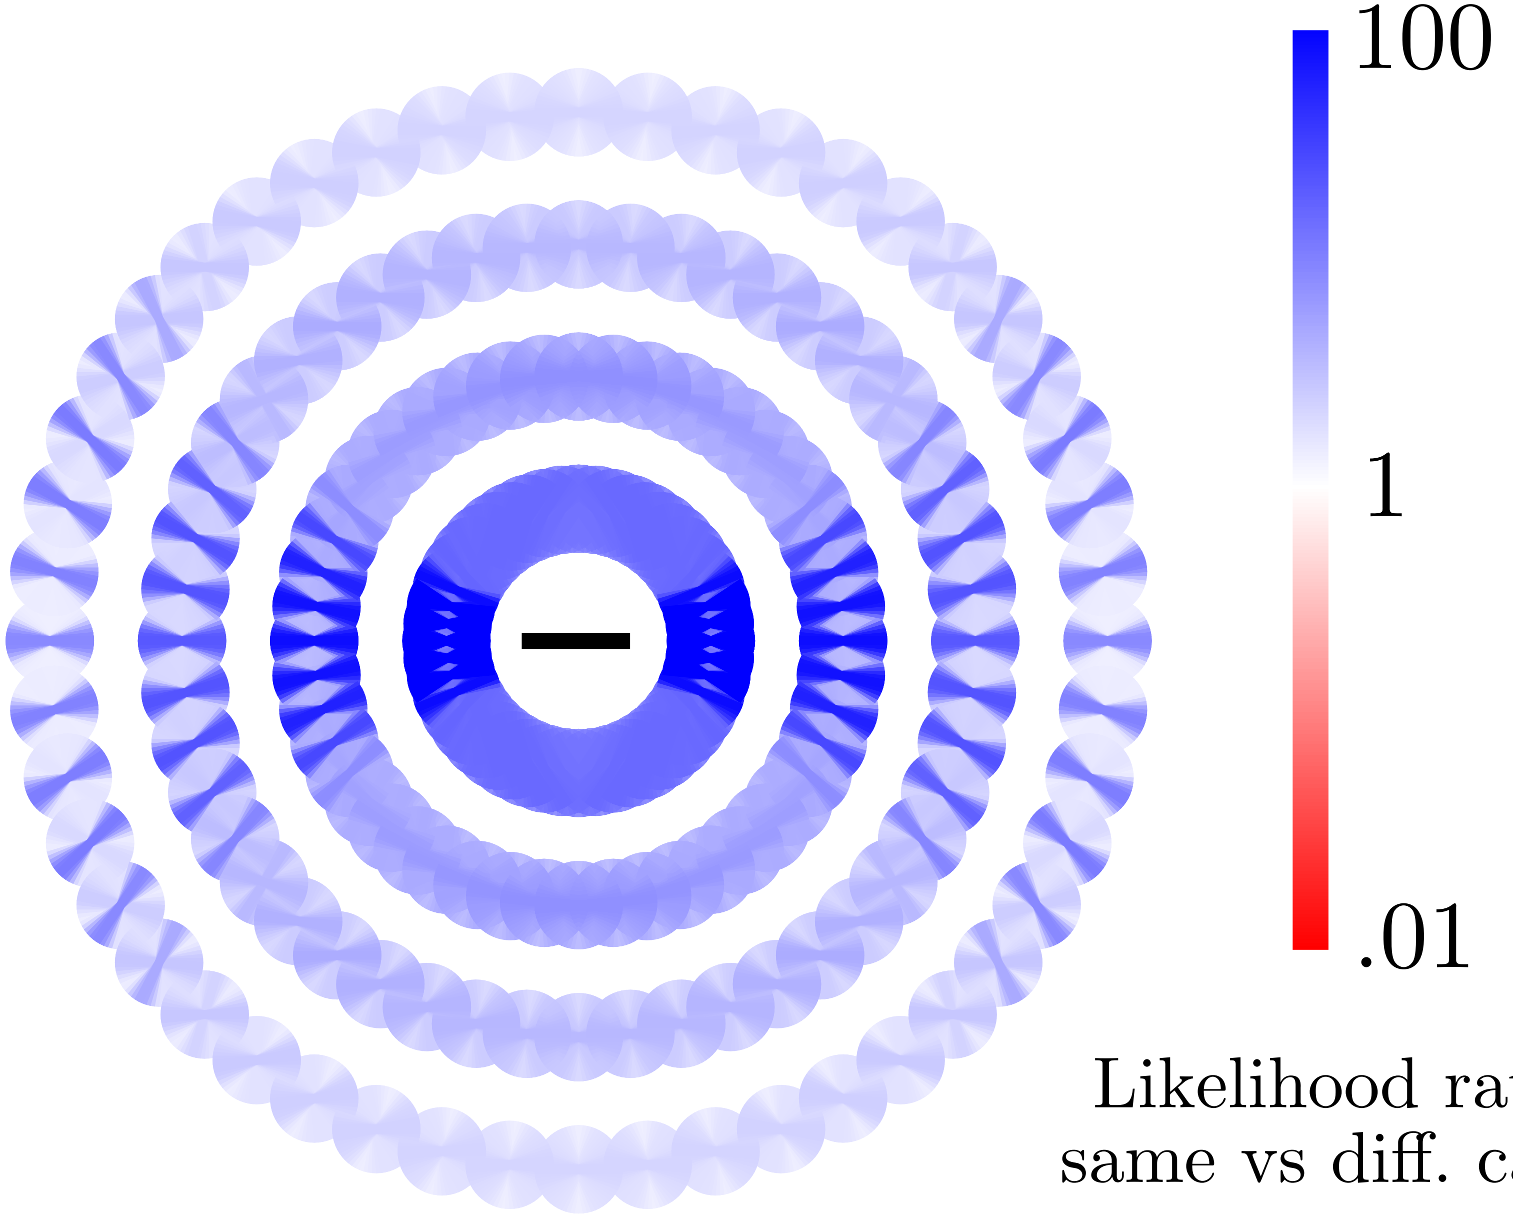

Likelihood ratio: same vs diff. cause

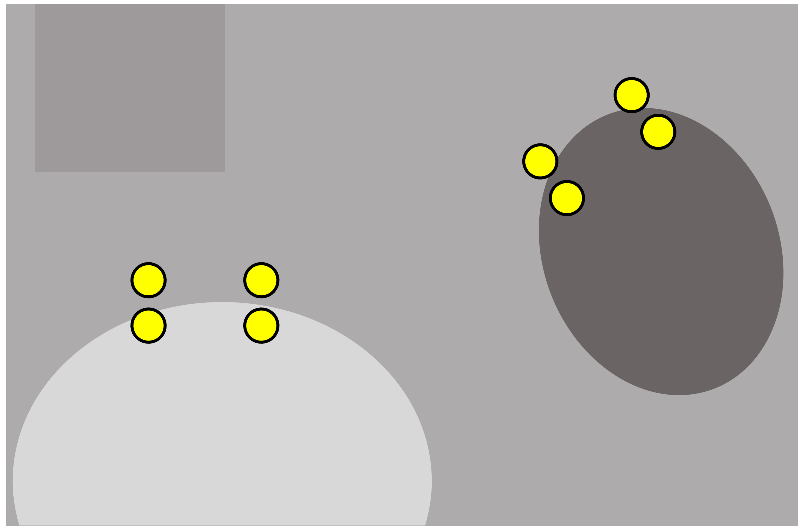

edge causes defined by object membership

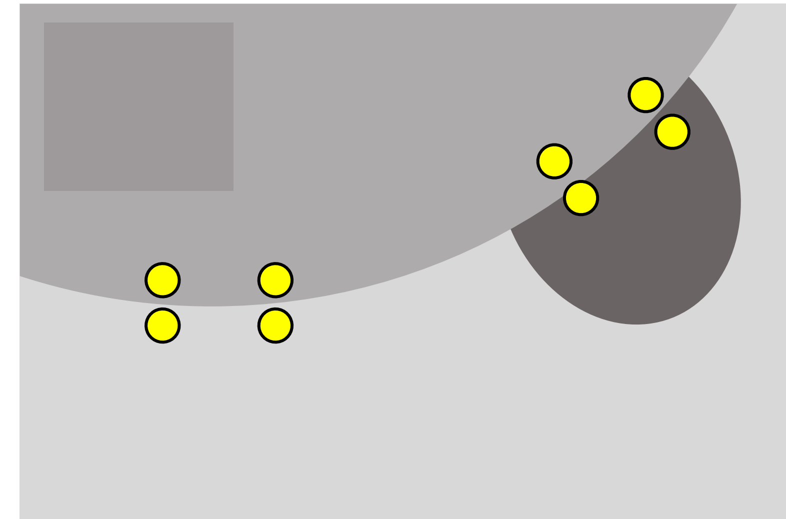

edge causes defined with depth constraints

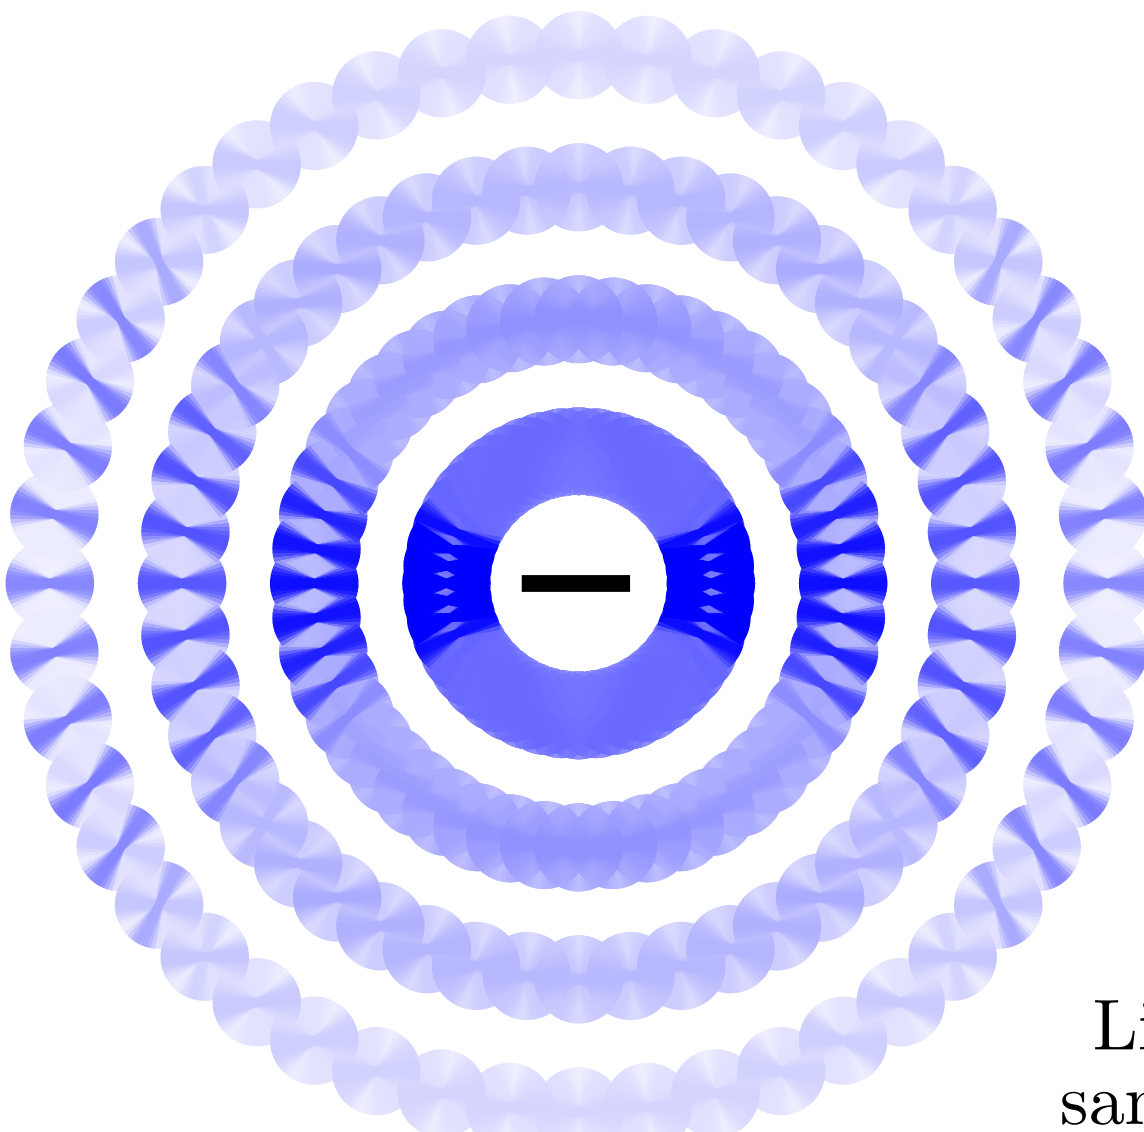

100

1

.01

Likelihood ratio: same vs diff. cause

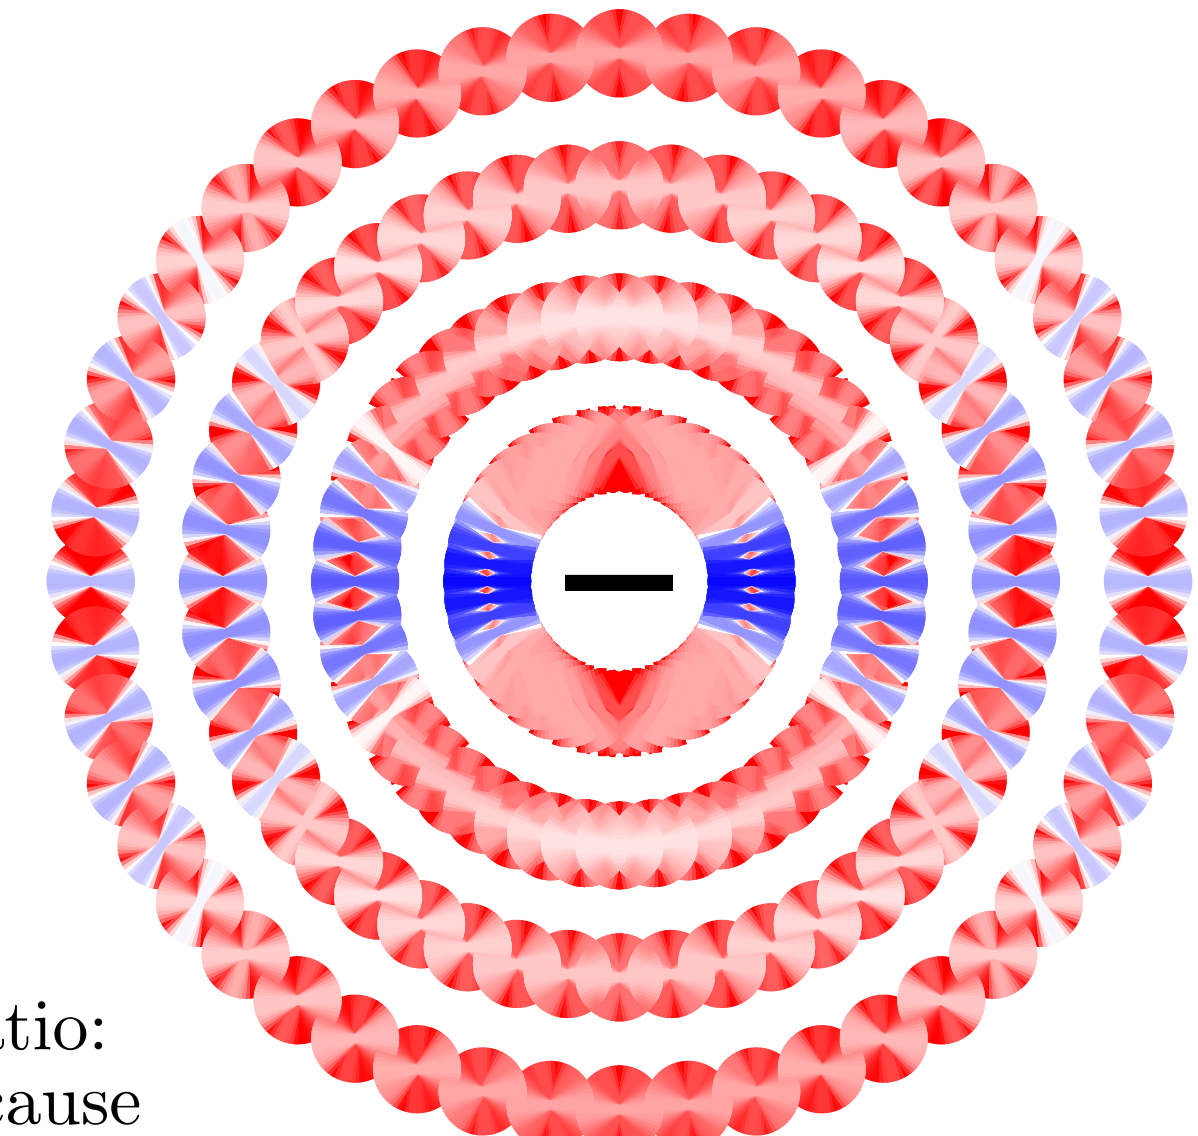

Dead Leaves Model:

- naturalistic yet tractable
- sufficient to explain many non-gaussian properties of natural images
- good stimulus for psychophysics and modeling

Thanks to:

the Swartz Foundation, NIH

the columbia center for theoretical neuroscience,

esp. Ken Miller, Larry Abbott, Taro Toyoizumi, Vladimir Itskov