Image Statistics

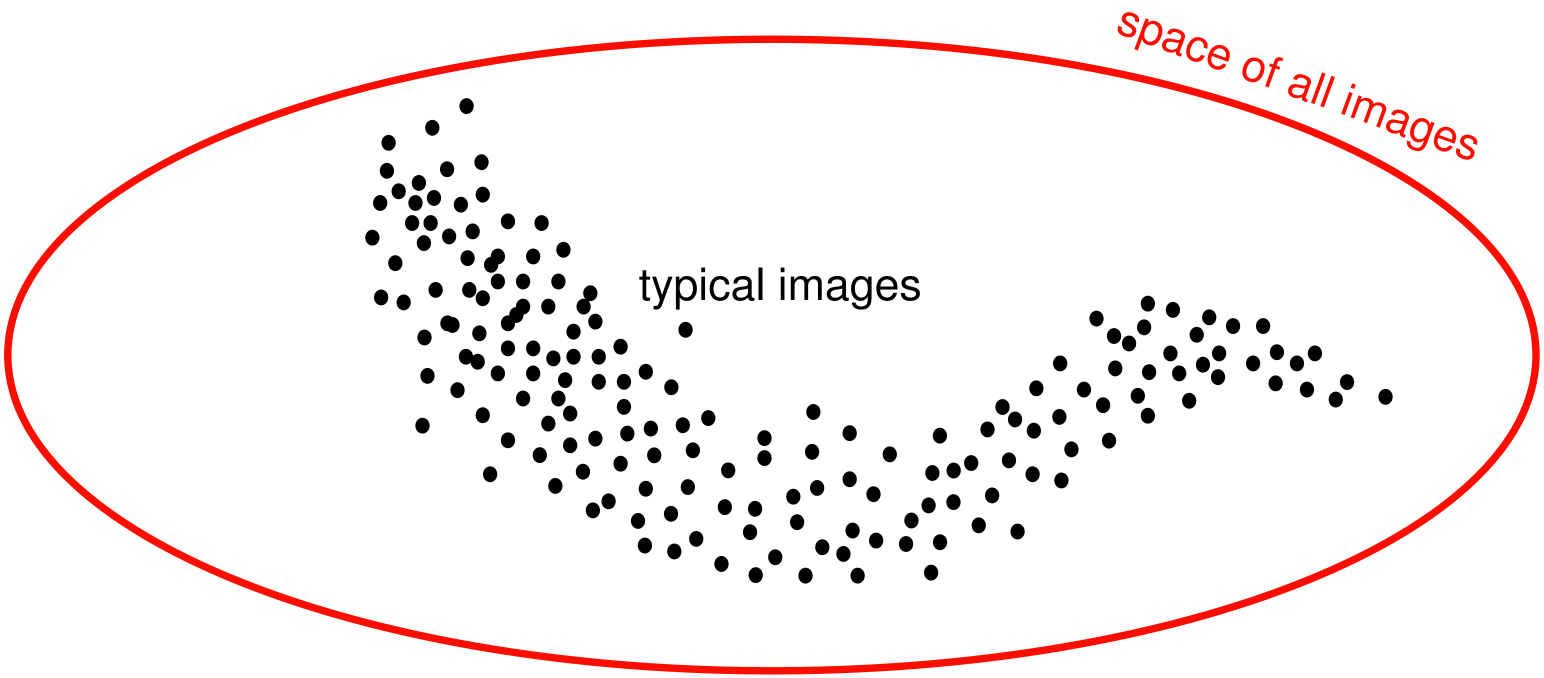

Statistical Image Models

Image Processing / Graphics:

- How compactly can we represent images?
- How easily can we detect (and remove) artifacts or distortions?
- Can we enhance images by increasing resolution or spatial extent?
- Can we synthesize realistic-looking images?

Theoretical Neurobiology:

- Do sensory neurons perform optimal decomposition of images?
- If so, how does the system learn this decomposition?

Image Statistics

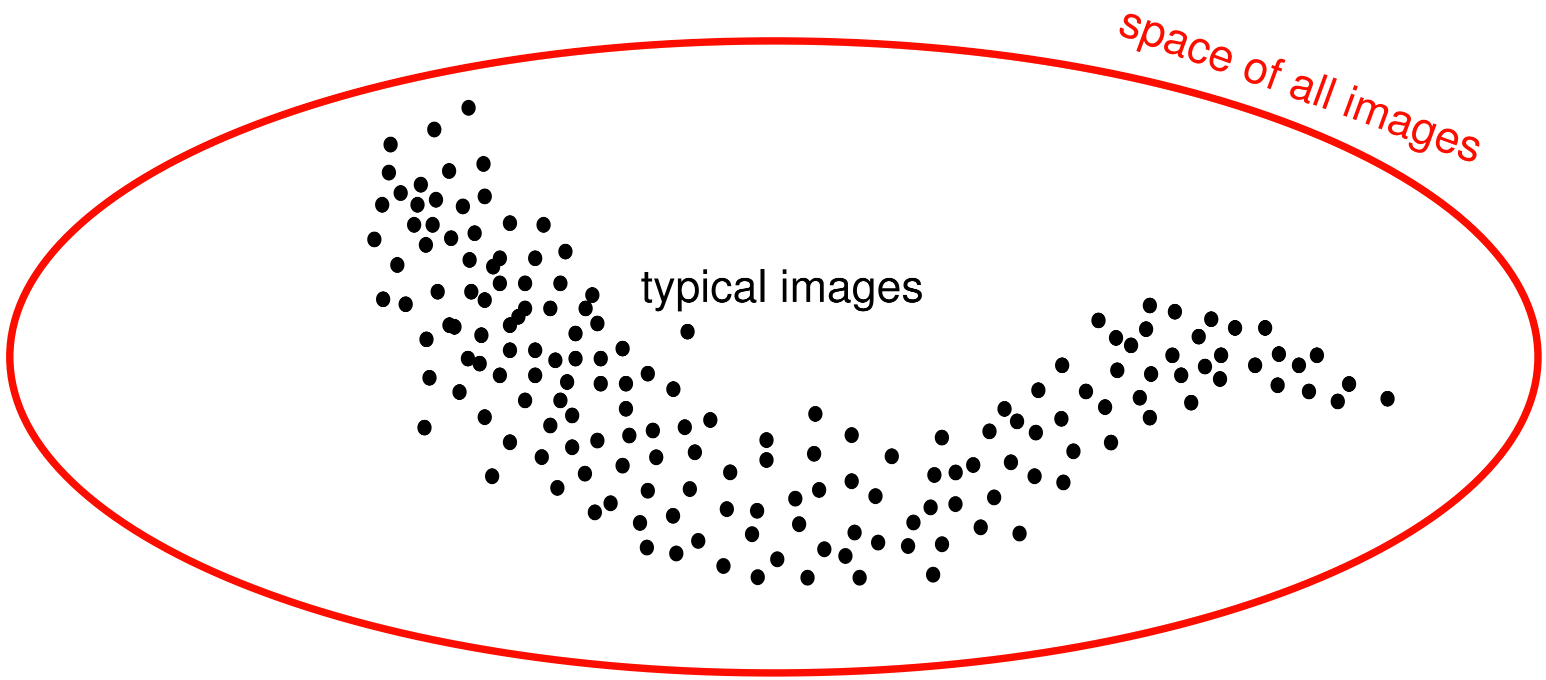

Loosely:

- Neural architecture optimized for ensemble of images
- Instantaneous pattern of responses represents relevant information about current image

Statistical Optimality

Two basic frameworks (with overlap):

- Bayes (prior)
- Efficient coding [Pam's next lecture]

How do we test these??

- Experimental
- Theoretical
- Hybrid (!)

Building a statistical model

We want to look for statistical properties that provide constraints that are both *strong* and *reliable*.

- make measurements...
  - but measurements of what?
  - beware the curse of dimensionality
- make structural assumptions (e.g., translation-invariance, locality)

Pixels

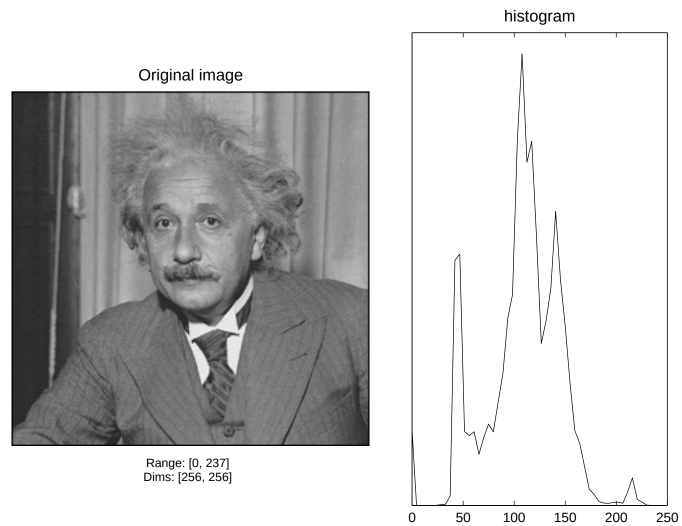
Original image
Range: [0, 237]
Dims: [256, 256]
histogram
0
50
100
150
200
250

Pixels

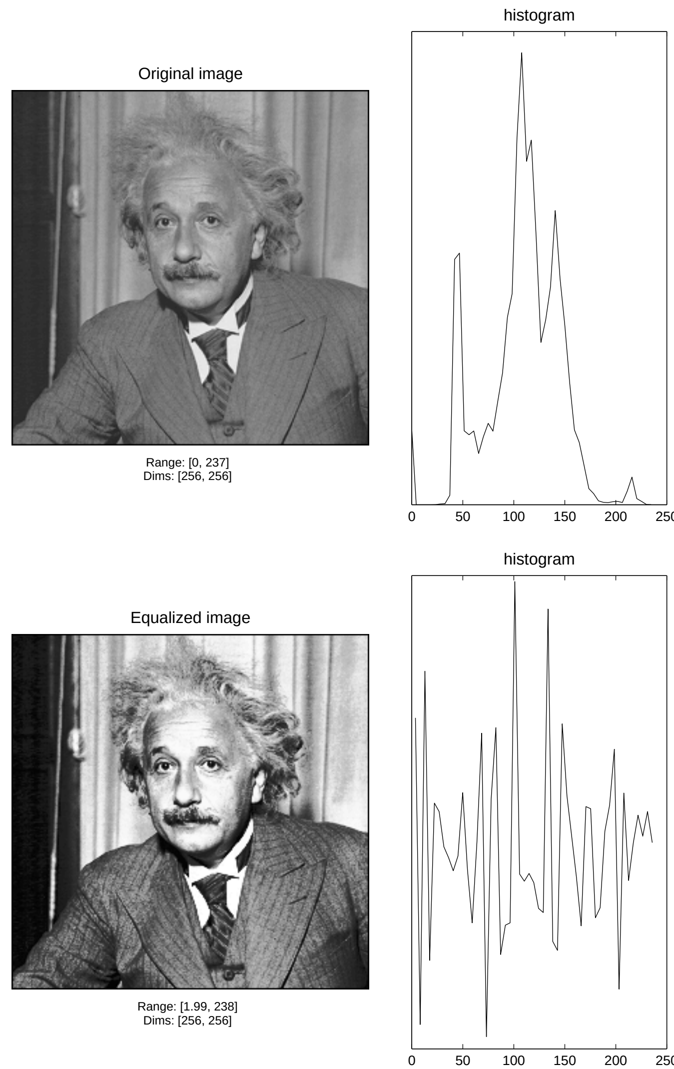
Original image
Range: [0, 237]
Dims: [256, 256]
histogram
0 50 100 150 200 250
Equalized image
Range: [1.99, 238]
Dims: [256, 256]
histogram
0 50 100 150 200 250

Pixels

Original image

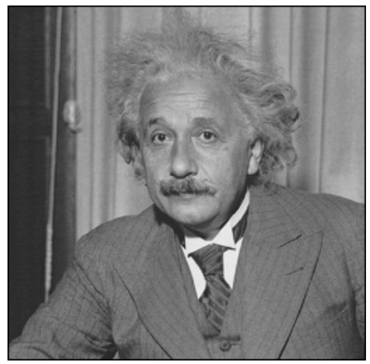

Range: [0, 237]
Dims: [256, 256]

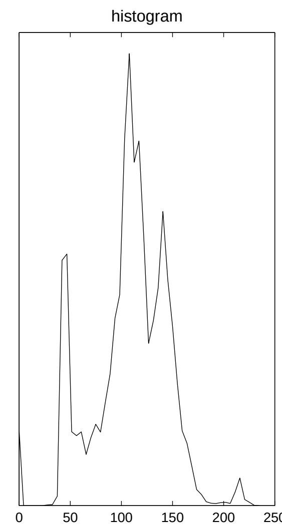

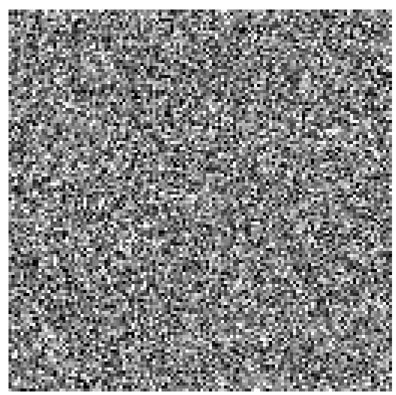

Pixel Correlations

a.

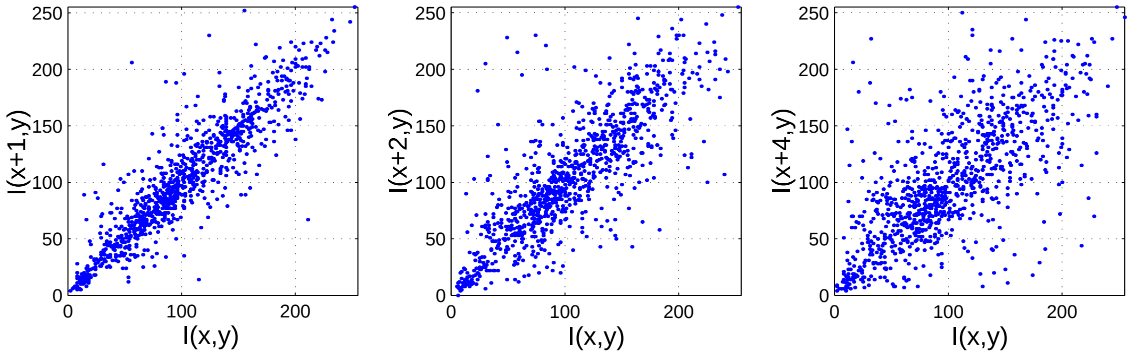

b.

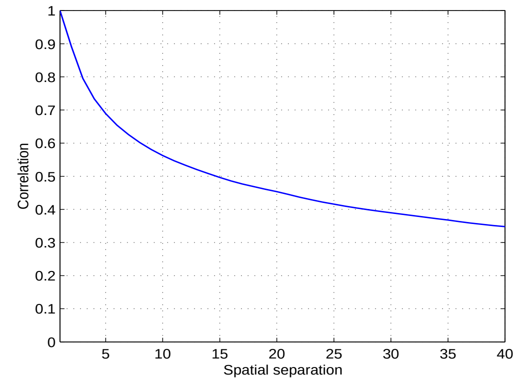

Relationship between Covariance matrix and Fourier

on board...

Principal Component Analysis (PCA)

Find linear transform (specifically, rotation and axis re-scaling) that transforms the covariance matrix to the identity.

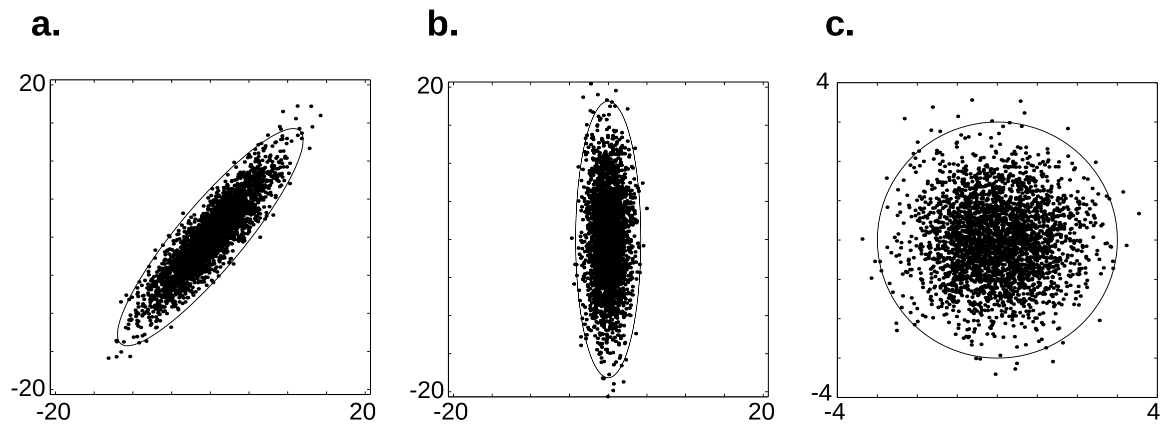

Well-known eigenvalue/eigenvector solution

Assuming translation-invariance, Fourier transform suffices

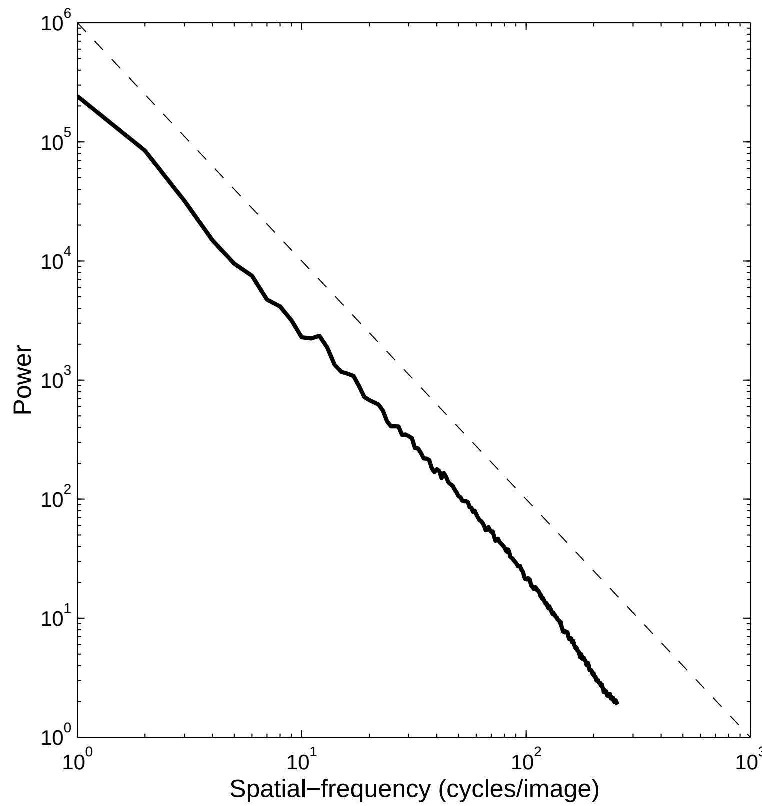

Power spectra of natural images fall as $1/f^\alpha$, $\alpha \sim 2$.

[Ritterman '52; DeRiugin '56; Field '87; Tolhurst '92; Ruderman & Bialek '94; etc]

Scale-invariance

If $\tilde{g}(\omega) = A\omega^p$, then $\tilde{g}(s\omega) = As^p\omega^p$.

i.e., for power-law spectrum, shape is preserved under scaling.

Conversely, if shape is preserved under scaling, spectrum must follow a power law!

Maximum Entropy

The distribution over response r with maximum entropy subject to a constraint of the form:

$$\mathcal{E}(f(r)) = c$$

is

$$\mathcal{P}(r) \propto \exp(-\lambda(c) f(r))$$

Examples:

- $f(r) = r^2$
- $f(r) = |r|$

Multi-dimensional Gaussians

- Characterized by mean (vector) and covariance (matrix)
- Remains Gaussian under linear transformation of space
- Conditionals (slices) and marginals (projections) are Gaussian
- Unique property: separable products are spherically symmetric
- Central limit theorem: sums of i.i.d. random variables become Gaussian
- Heisenberg: Fourier transform is Gaussian, and minimizes variance product

Model I: Fourier+Gaussian

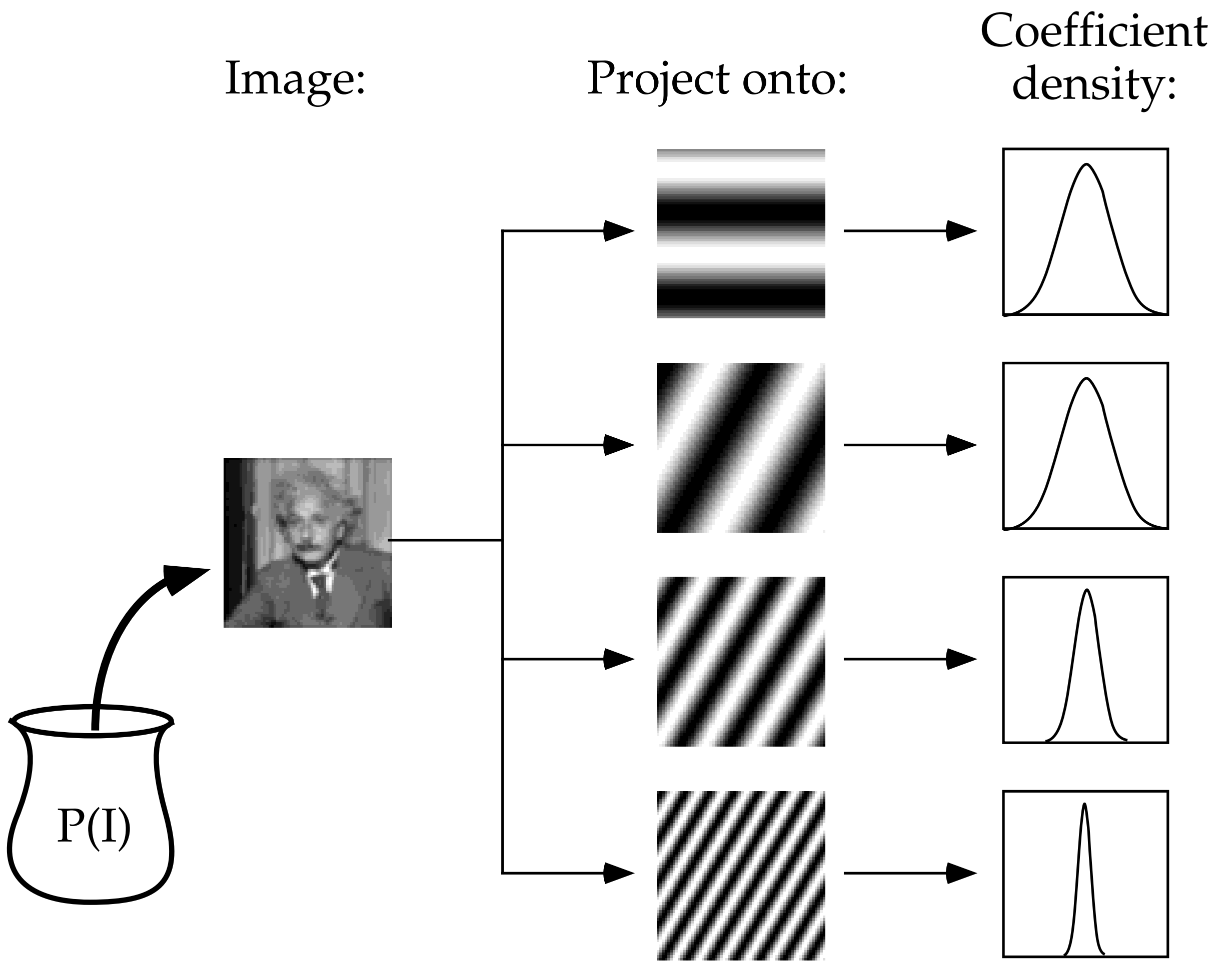

Most image processing engineering is based on this “classic” model

Synthesis

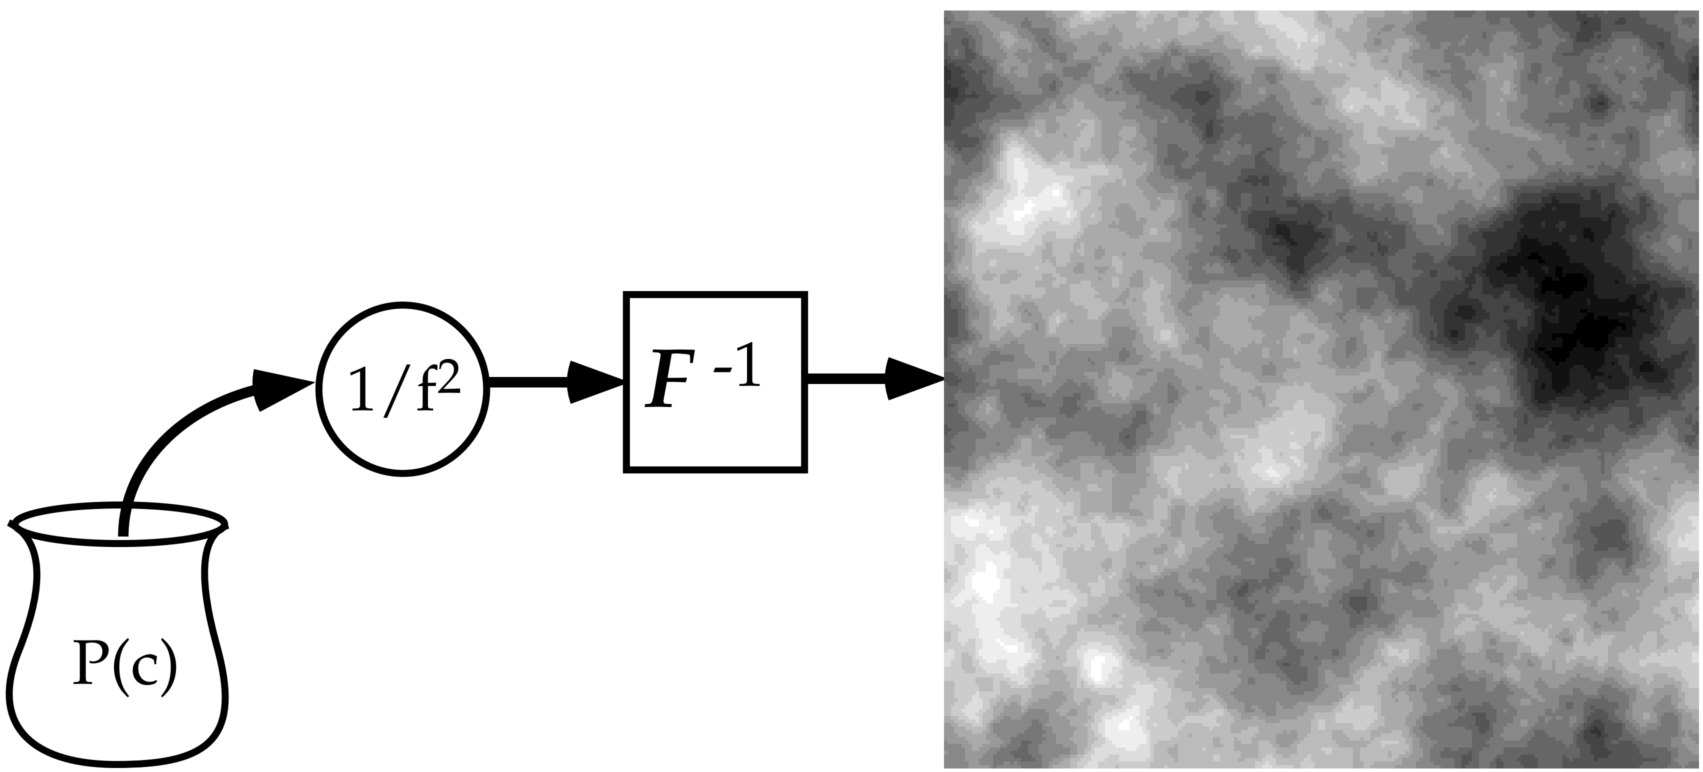

Whitening

a.

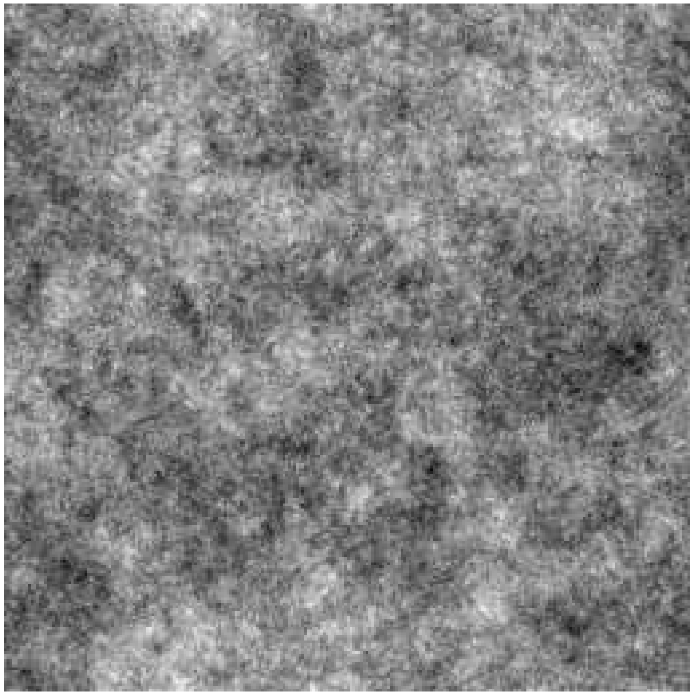

b.

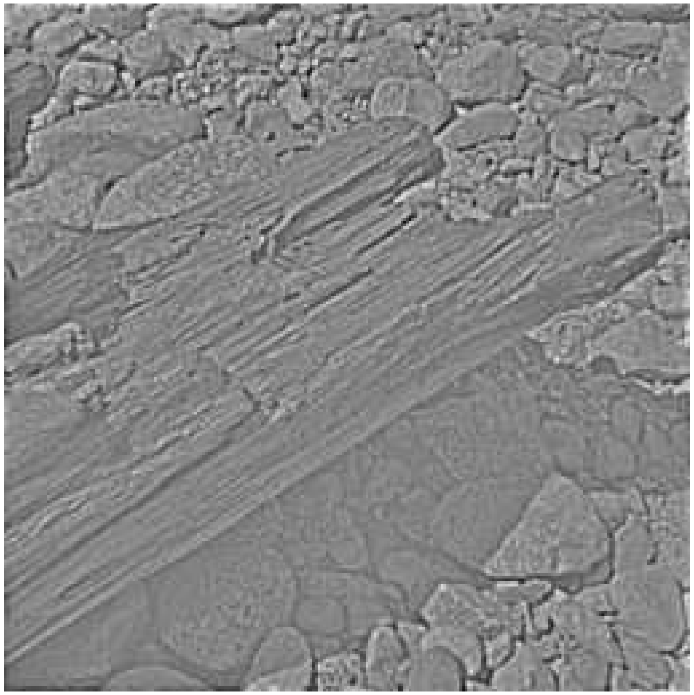

Windowed PCA

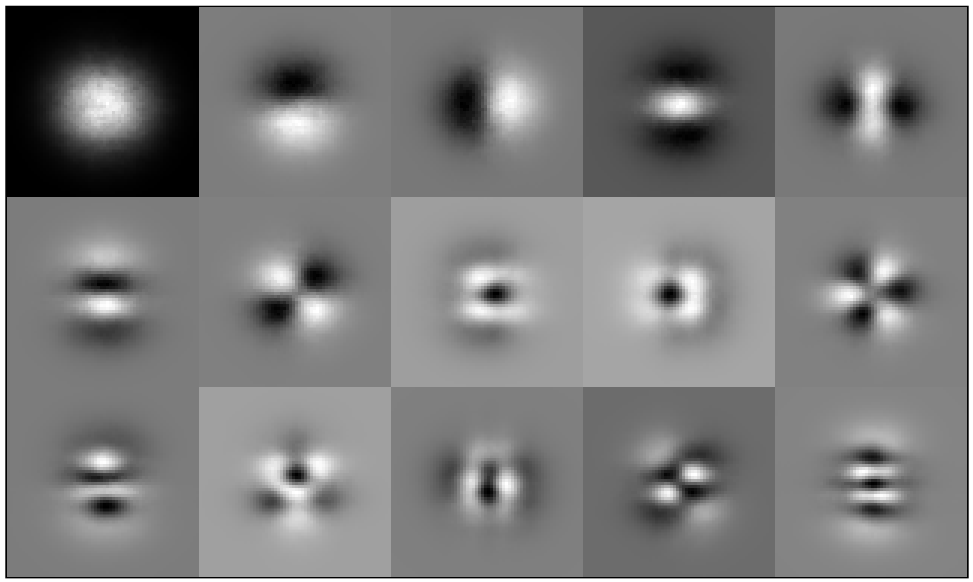

[Hancock etal, ’91]