Image Statistics

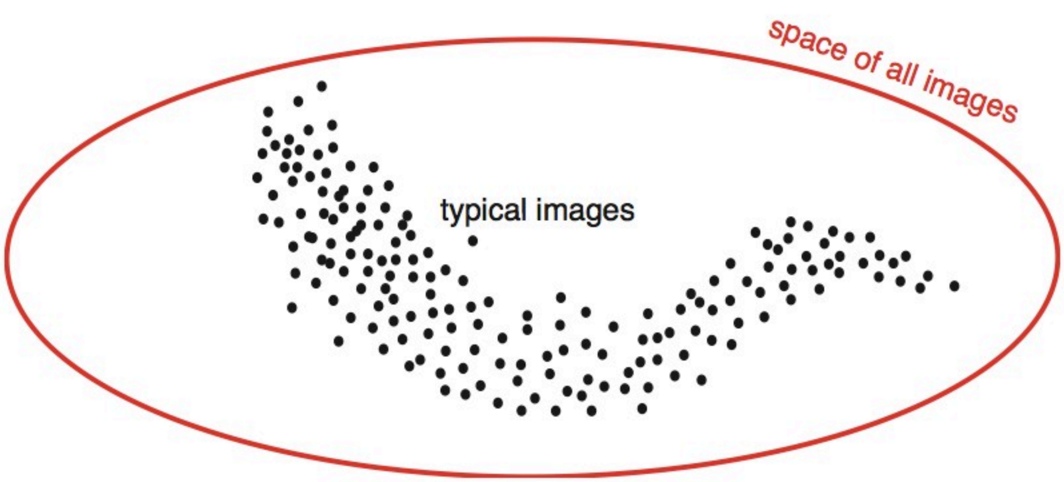

One could describe this set as a deterministic manifold....

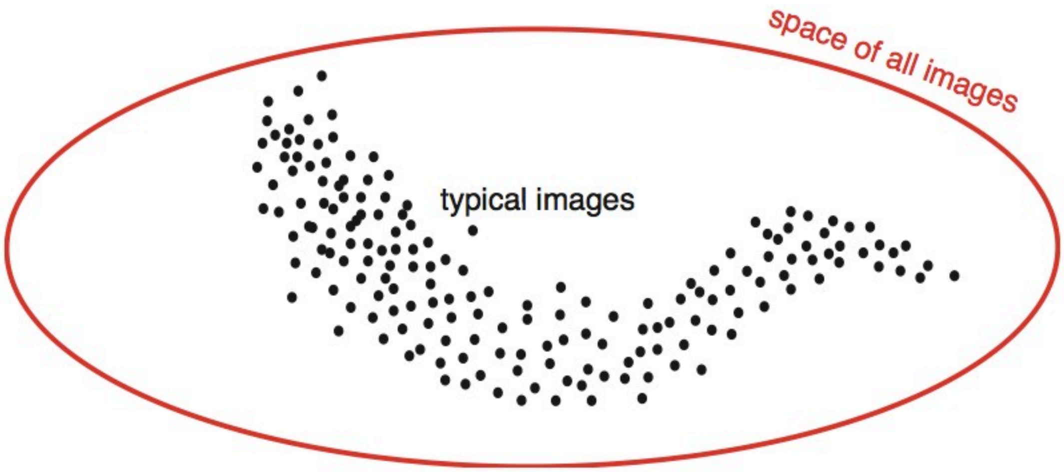

One could describe this set as a deterministic manifold....

But membership is “soft”, so seems more natural to use probability

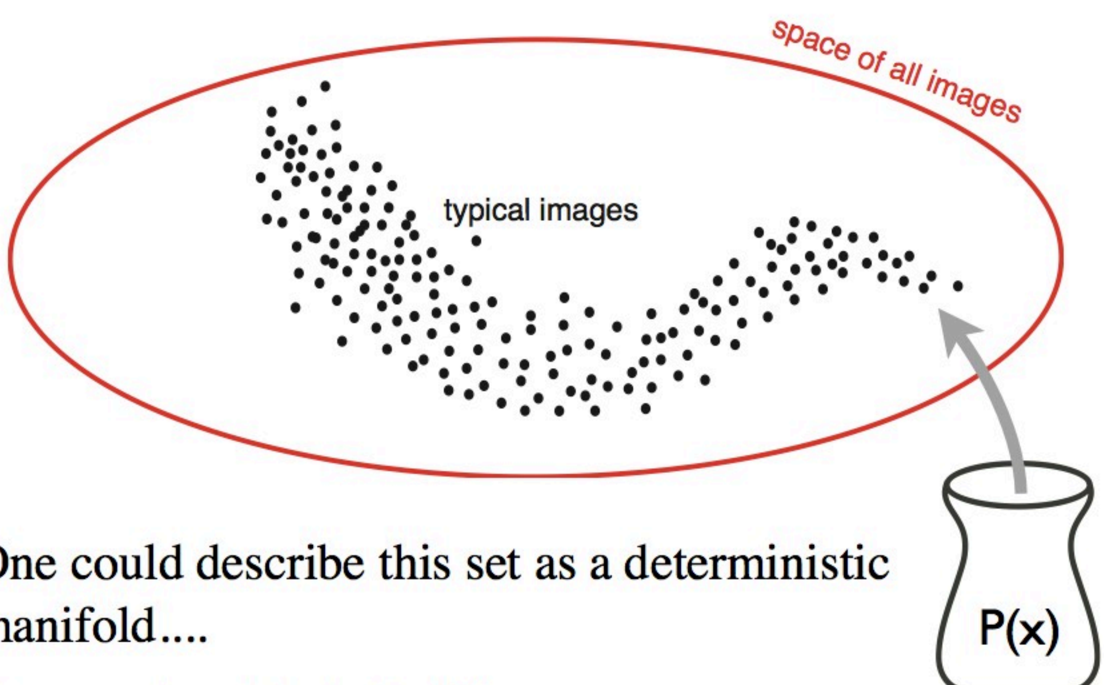

One could describe this set as a deterministic manifold....

But membership is “soft”, so seems more natural to use probability

Evolution of image models

I. (1950's): Fourier + Gaussian

II. (mid 80's - late 90's): Wavelets + kurtotic marginals

III. (mid 90's - mid 00's): Wavelets + local context

- local amplitude (contrast)
- local orientation

IV. (recent): Hierarchical models

General methodology

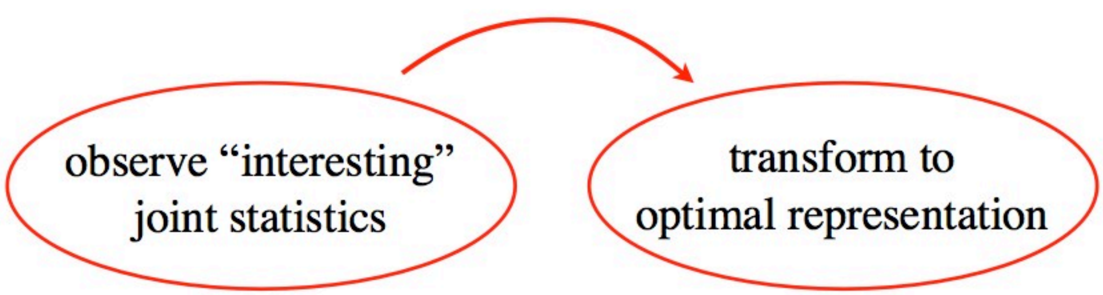

General methodology

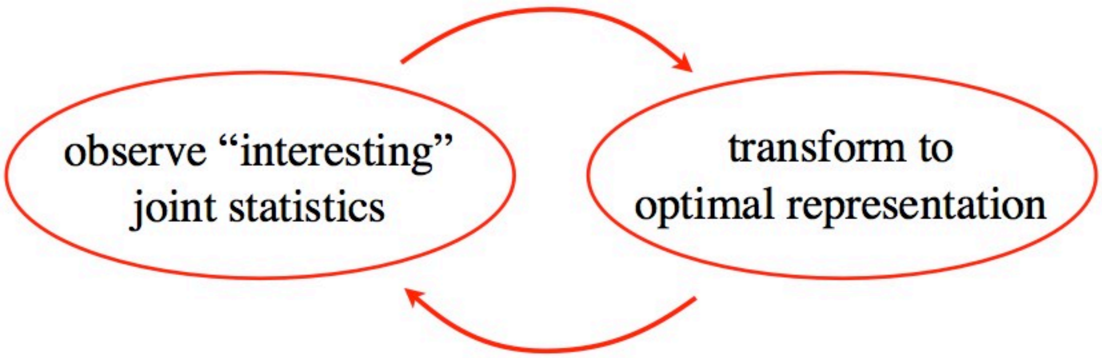

General methodology

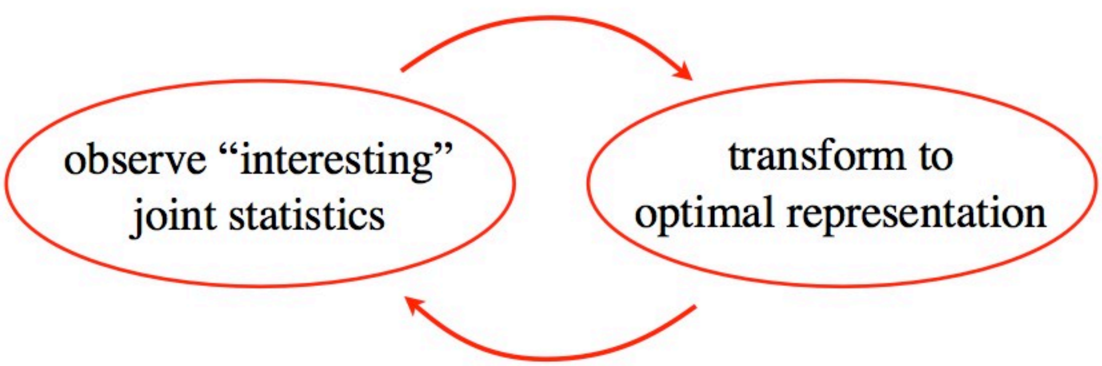

"Onion peeling"

Pixel histograms carry very little information about image content...

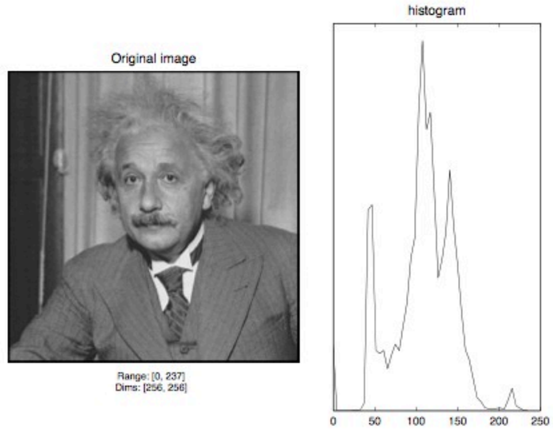

Pixel histograms carry very little information about image content...

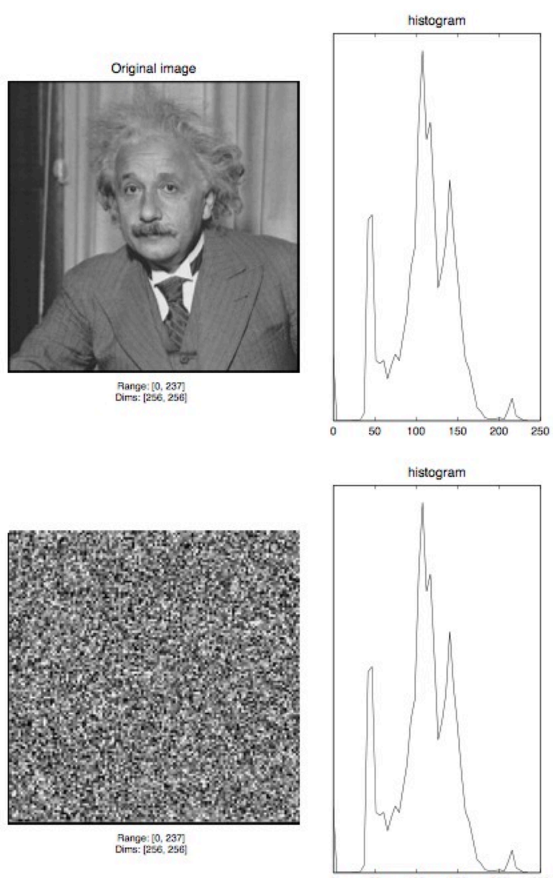

Pixel histograms carry very little information about image content...

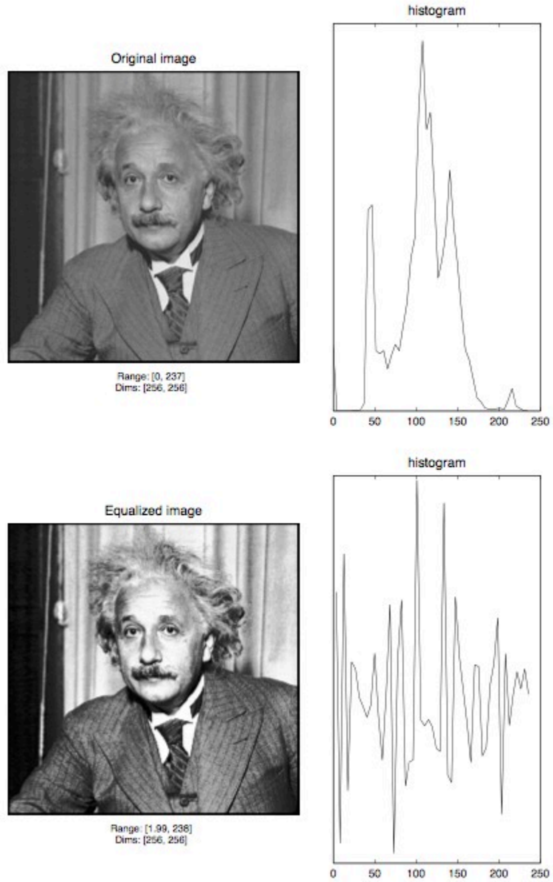

Pixel correlation

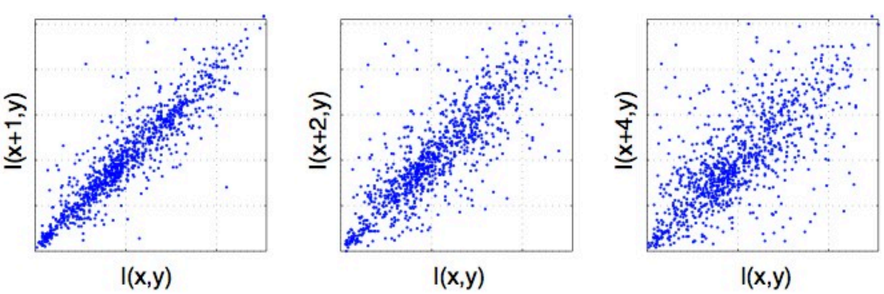

Pixel correlation

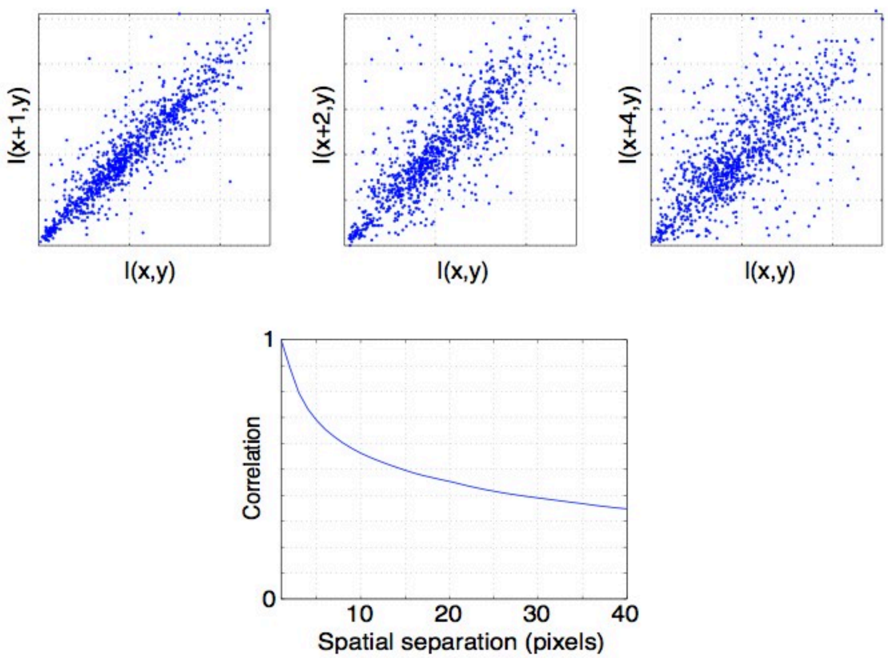

Translation invariance

Assuming translation invariance,

Translation invariance

Assuming translation invariance,

=> covariance matrix is Toeplitz (convolutional)

Translation invariance

Assuming translation invariance,

=> covariance matrix is Toeplitz (convolutional)

=> eigenvectors are sinusoids (complex exp.)

Translation invariance

Assuming translation invariance,

=> covariance matrix is Toeplitz (convolutional)

=> eigenvectors are sinusoids (complex exp.)

=> can diagonalize (decorrelate) with F.T.

Translation invariance

Assuming translation invariance,

=> covariance matrix is Toeplitz (convolutional)

=> eigenvectors are sinusoids (complex exp.)

=> can diagonalize (decorrelate) with F.T.

The variances of the frequency components (known as the "Power Spectrum") capture the full covariance structure!

im

log(abs(fftshift(fft2(im))))

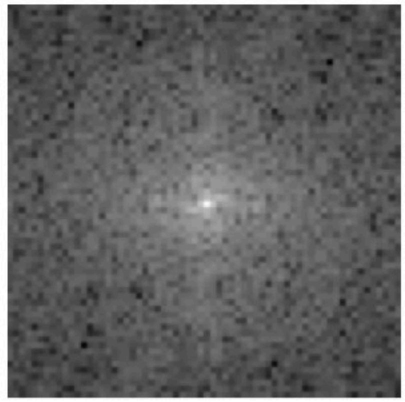

im

log(abs(fftshift(fft2(im))))

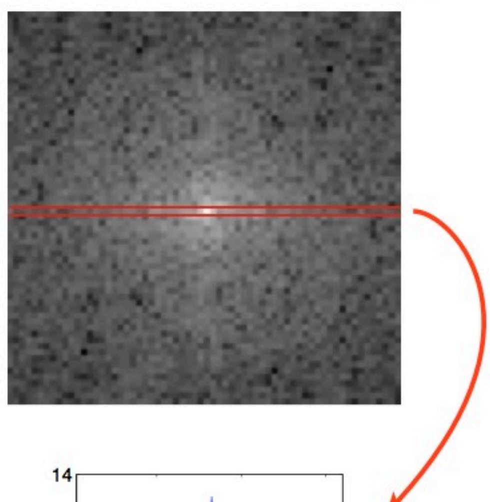

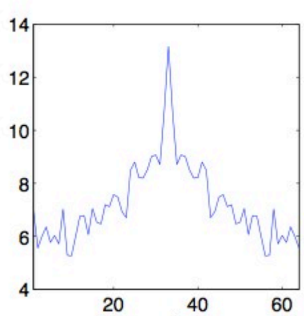

Spectral power

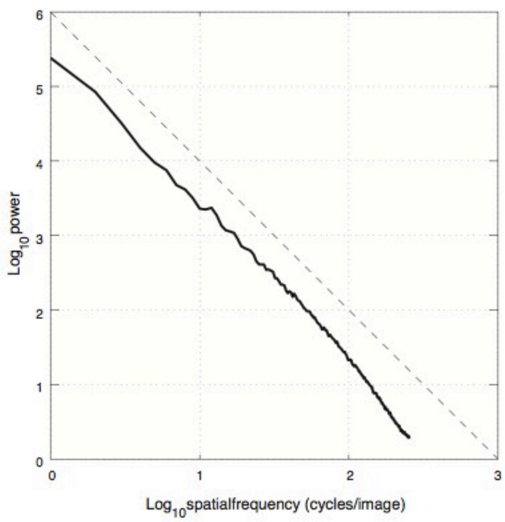

$$F(\omega) \propto \frac{1}{\omega^2}$$

(averaged over many images)

[Ritterman 52; DeRiugin 56; Field 87; Tolhurst 92; Ruderman/Bialek 94; ...]

Spectral power

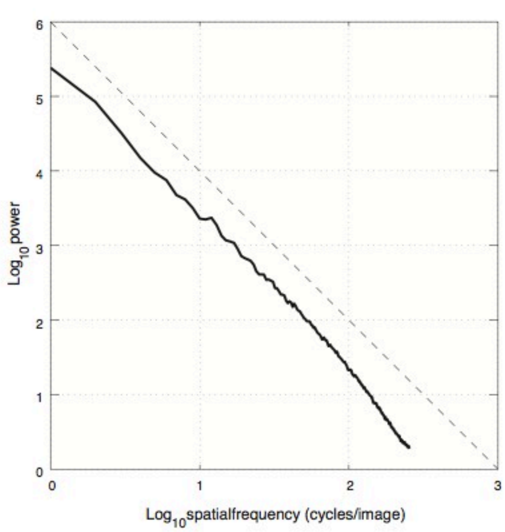

$$F(\omega) \propto \frac{1}{\omega^2}$$

(averaged over many images)

Note: this implies scale-invariance:

$$F(s\omega) \propto F(\omega), \qquad \forall s$$

[Ritterman 52; DeRiugin 56; Field 87; Tolhurst 92; Ruderman/Bialek 94; ...]

Maximum entropy (maxEnt)

The density with maximal entropy satisfying

$$E\left(f(x)\right) = c$$

is of the form

$$p_{\mathrm{ME}}(x) \propto \exp\left(-\lambda f(x)\right)$$

where λ depends on c

Examples: $f(x) = x^2$ $f(x) = |x|$

Model I (Fourier/Gaussian)

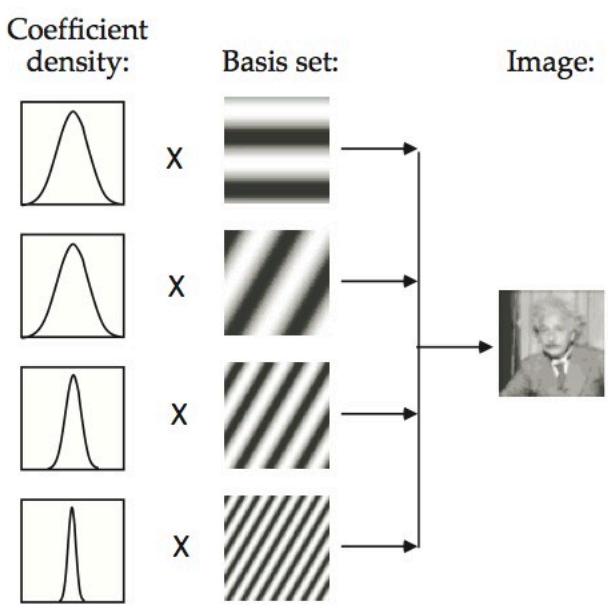

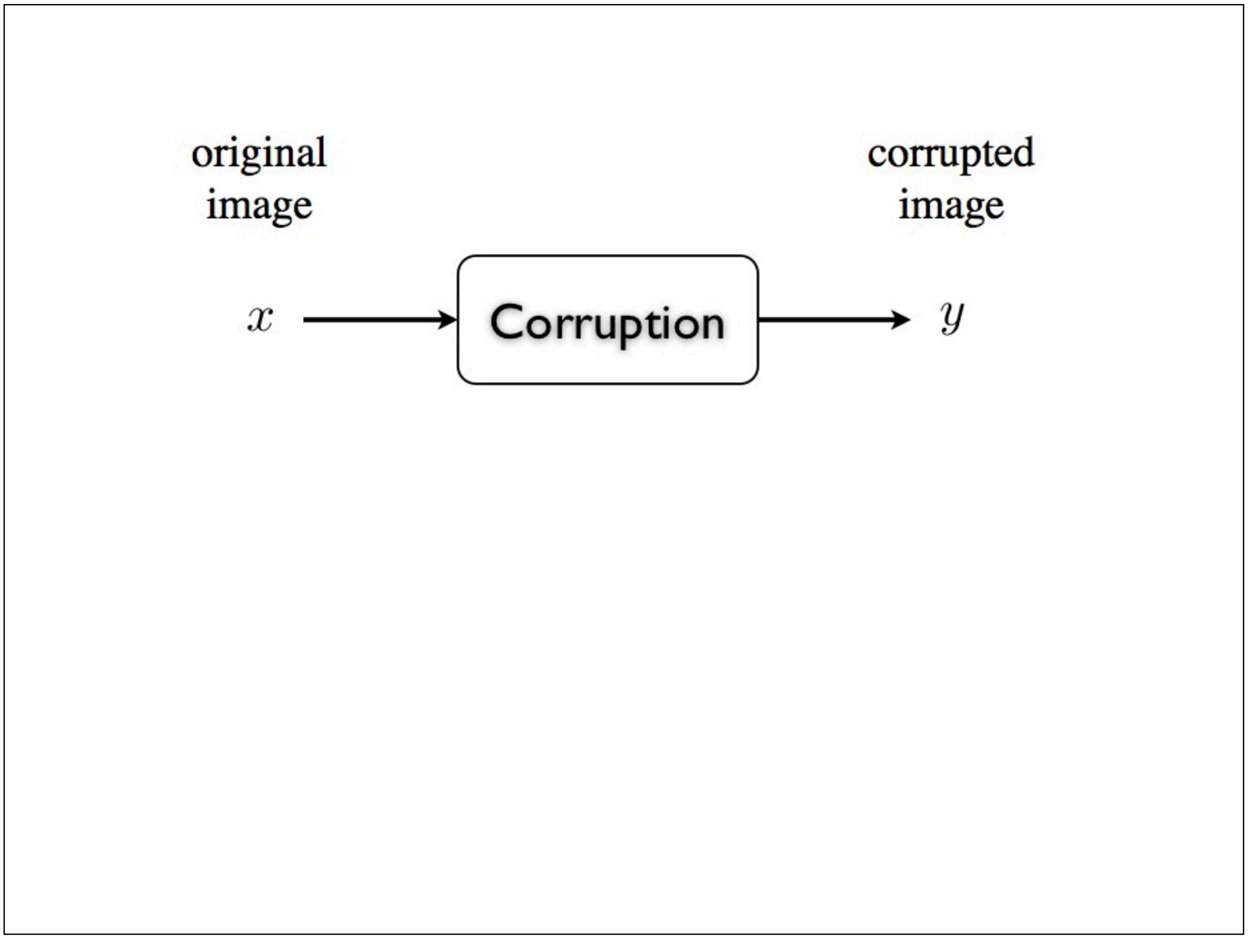
original
image
corrupted
image
x
Corruption
y

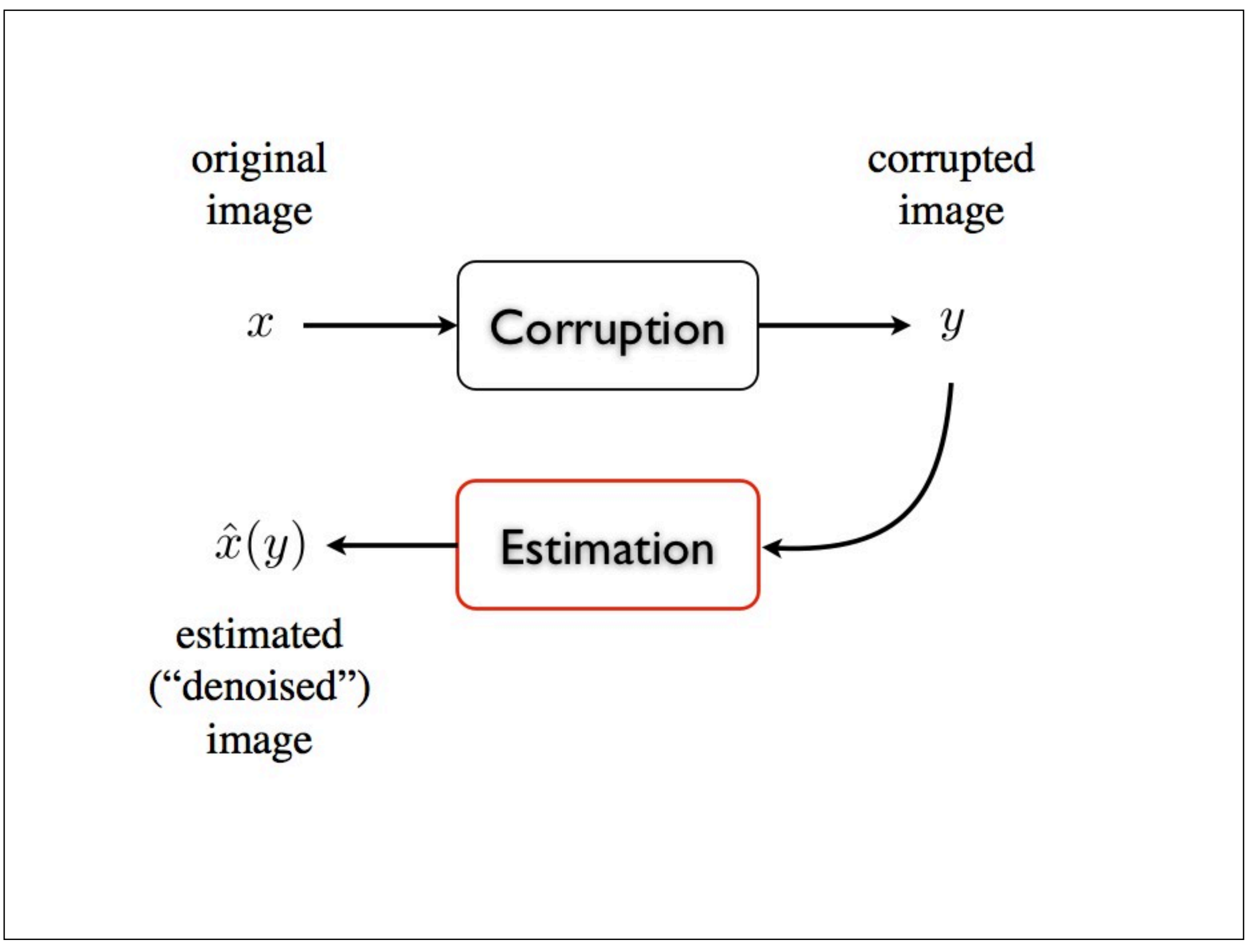
original
image
corrupted
image
x
Corruption
y
$\hat{x}(y)$
Estimation
estimated
("denoised")
image

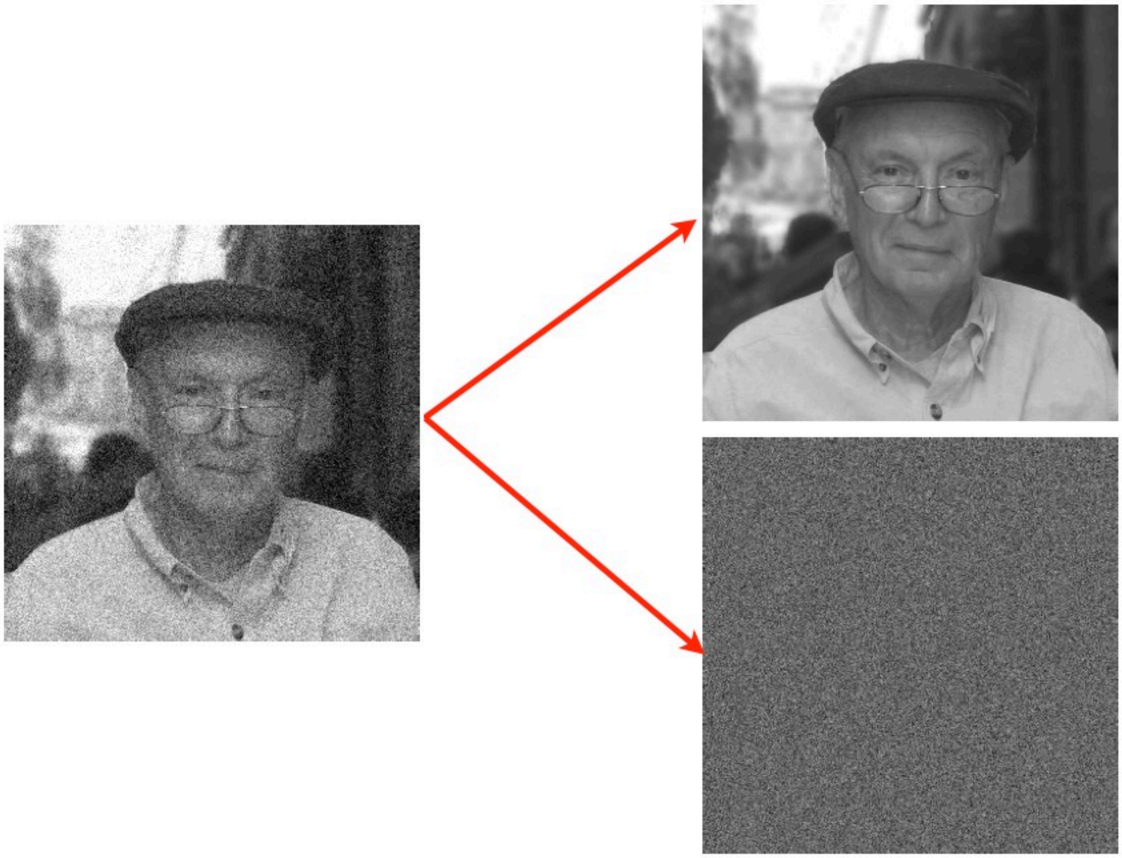

Common statistical estimators:

- maximum likelihood (ML)

$$x_{\mathrm{ML}} = \arg\max_{x} p(y|x)$$

- maximum a posterior (MAP)

$$x_{\mathrm{MAP}} = \operatorname*{argmax}_{x} p(x|y) = \operatorname*{argmax}_{x} p(y|x)p(x)$$

- minimum mean squares error (MMSE)

$$x_{\mathrm{MMSE}} = \arg\min_{x'} \int ||x - x'||^2 \; p(x|y) \; dx$$

$$= \int x \; p(x|y) \; dx = E(x|y)$$

Bayesian estimation

$$x_{\text{Bayes}}(y) = \arg\min_{x'} \int L(x, x')\ p(x|y)\ dx$$
$$\propto \arg\min_{x'} \int L(x, x')\ p(y|x)\ p(x)\ dx$$

Given observation y, choose estimate that minimizes the average loss

Denoising: classical

If signal is Gaussian, MAP/MMSE estimate is linear:

$$\mathbb{E}(\vec{x}|\vec{y}) = C_x(C_x + C_w)^{-1}\vec{y}$$

$$\hat{X}(\omega) = \frac{A/\omega^p}{A/\omega^p + \sigma^2} \cdot Y(\omega)$$

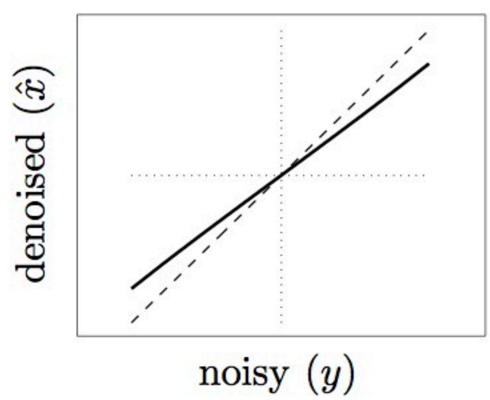

[proof on board]

Denoising: classical

If signal is Gaussian, MAP/MMSE estimate is linear:

$$\mathbb{E}(\vec{x}|\vec{y}) = C_x(C_x + C_w)^{-1}\vec{y}$$

$$\hat{X}(\omega) = \frac{A/\omega^p}{A/\omega^p + \sigma^2} \cdot Y(\omega)$$

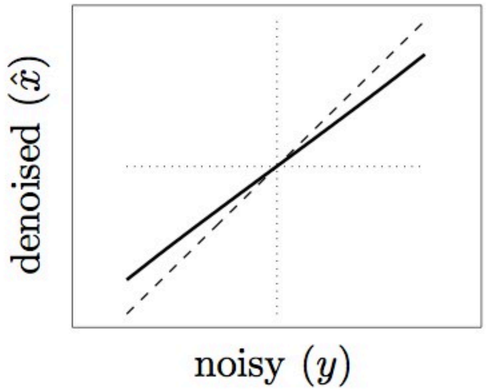

[proof on board]

Bottom line: Suppress fine scales, retain coarse scales

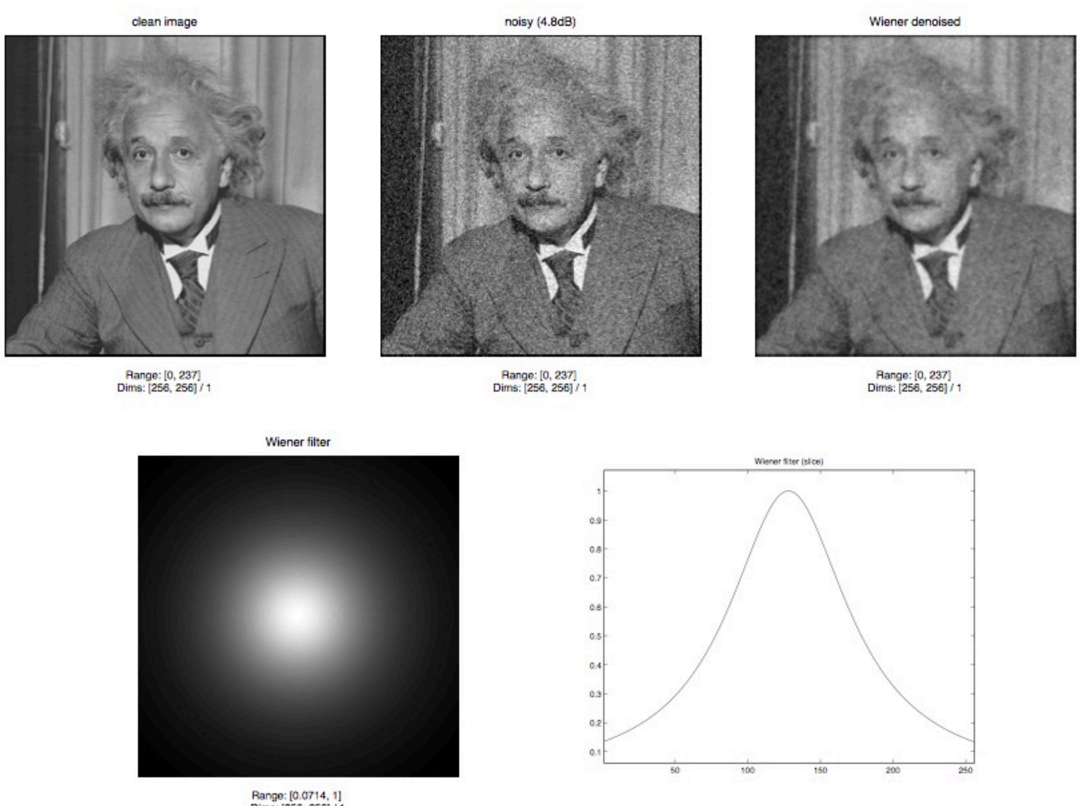

Whitening

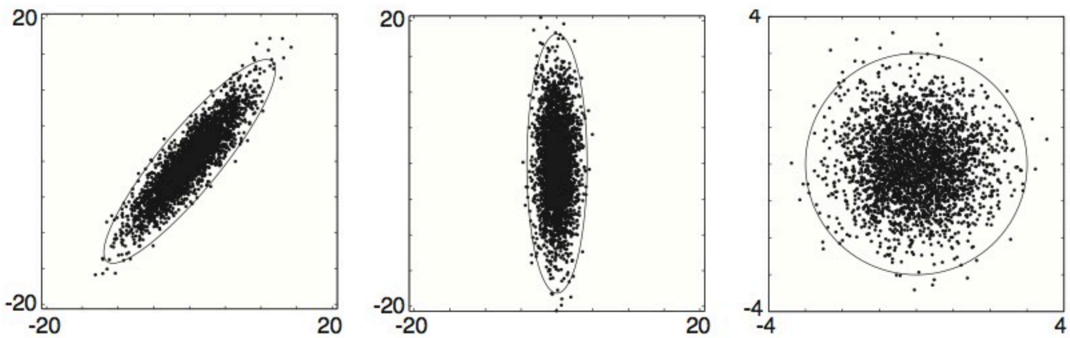

Fourier/Gaussian model is weak

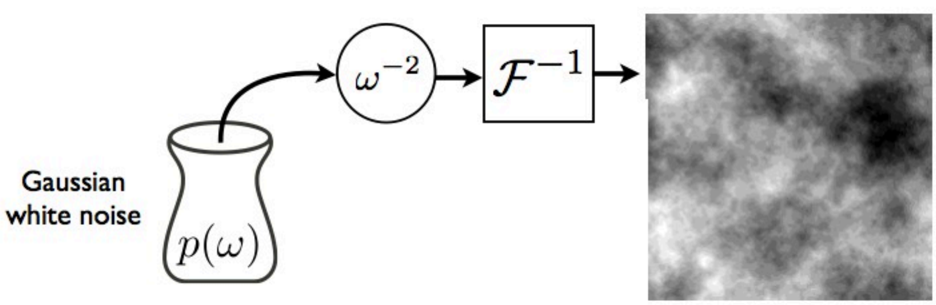

Fourier/Gaussian model is weak

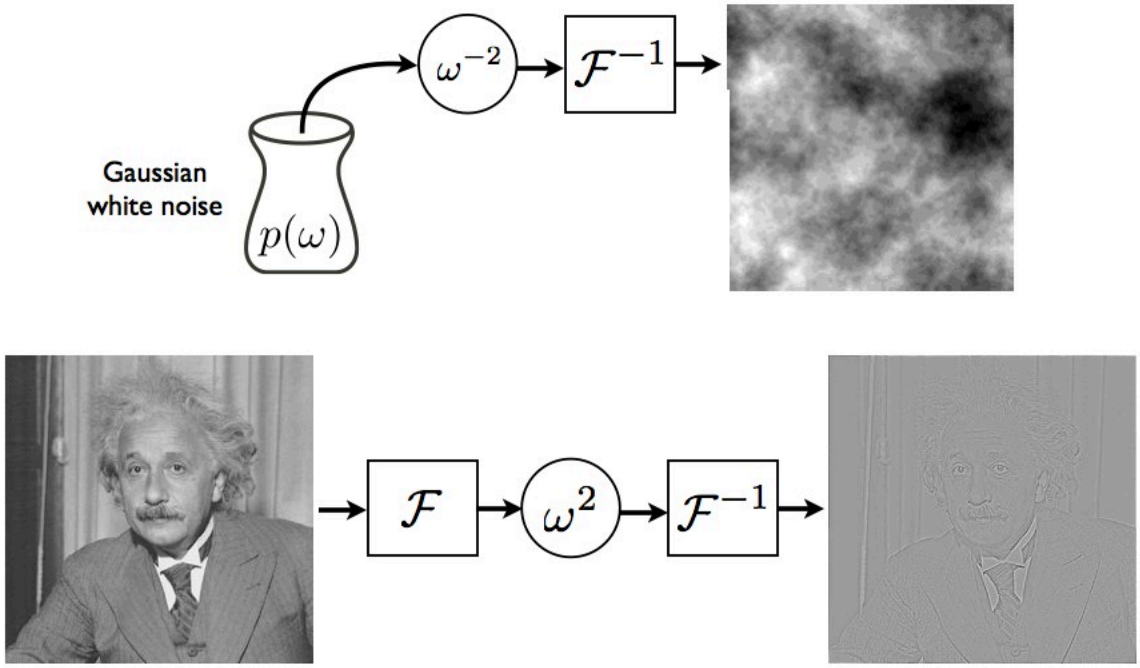

Fourier/Gaussian model is weak

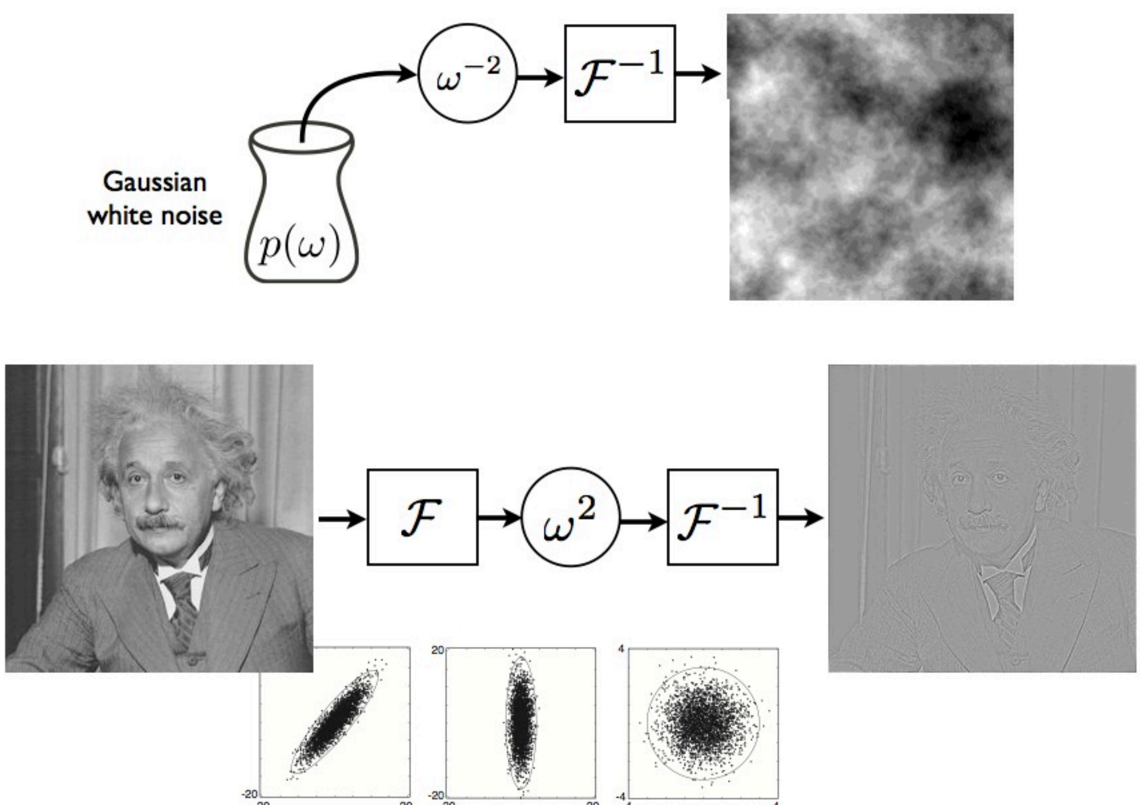

PCA

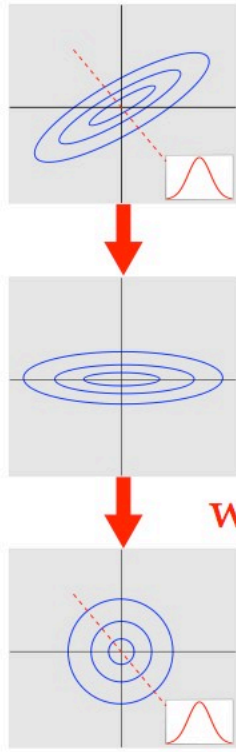

$\vec{x}$

$\Sigma = U \Lambda U^T$

$U^T \vec{x}$

whitening

$\Lambda^{-\frac{1}{2}} U^T \vec{x}$

not unique! $V \Lambda^{-\frac{1}{2}} U^T \vec{x}$

Fourier/Gaussian image model

Fourier/Gaussian image model

+ Simple/tractable/well-understood

Fourier/Gaussian image model

+ Simple/tractable/well-understood

+ Pseudo-biological (Hebbian) whitening algorithms

Fourier/Gaussian image model

+ Simple/tractable/well-understood
+ Pseudo-biological (Hebbian) whitening algorithms
- Fourier basis is **global**

Fourier/Gaussian image model

+ Simple/tractable/well-understood
+ Pseudo-biological (Hebbian) whitening algorithms
- Fourier basis is **global**
- Whitening transform is not unique

Fourier/Gaussian image model

+ Simple/tractable/well-understood
+ Pseudo-biological (Hebbian) whitening algorithms
- Fourier basis is **global**
- Whitening transform is not unique
- Important dependencies/structures remain