Inference

- Wiener filter
- coring/shrinkage

Simple example

$$y = x + n$$

You observe y, what is x?

$$P(x|y) \propto \underbrace{P(y|x)}_{\text{likelihood}} \underbrace{P(x)}_{\text{prior}}$$

$$P(x) = \frac{1}{Z_x} e^{-\frac{(x-\mu_x)^2}{2\sigma_x^2}}$$

$$P(y|x) = \frac{1}{Z_n} e^{-\frac{(y-x)^2}{2\sigma_n^2}}$$

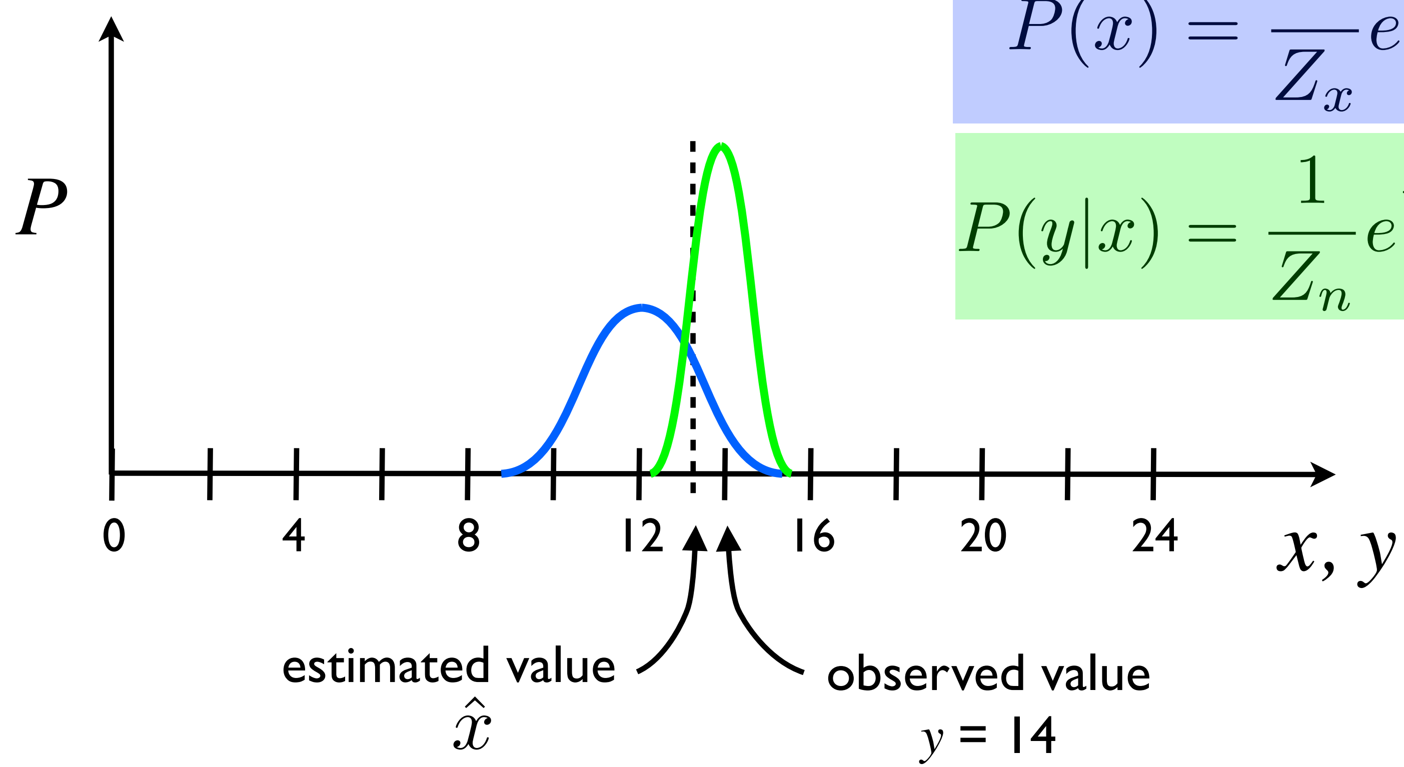

How to compute $\hat{x}$?

$$P(x|y) \propto P(y|x)\, P(x)$$

$$= \frac{1}{Z_n} e^{-\frac{(y-x)^2}{2\sigma_n^2}} \; \frac{1}{Z_x} e^{-\frac{(x-\mu_x)^2}{2\sigma_x^2}}$$

$$-\log P(x|y) = \frac{(y-x)^2}{2\sigma_n^2} + \frac{(x-\mu_x)^2}{2\sigma_x^2} + \text{const.}$$

$$-\frac{\partial}{\partial x} \log P(x|y) = -\frac{(y-x)}{\sigma_n^2} + \frac{(x-\mu_x)}{\sigma_x^2} = 0$$

$$\Rightarrow \boxed{\hat{x} = \frac{\sigma_x^2\, y + \sigma_n^2\, \mu_x}{\sigma_x^2 + \sigma_n^2}}$$

NOISE REMOVAL VIA BAYESIAN WAVELET CORING

Eero P. Simoncelli

Computer and Information Science Dept.
University of Pennsylvania
Philadelphia, PA 19104

Edward H. Adelson

Brain and Cognitive Science Dept.
Massachusetts Institute of Technology
Cambridge, MA 02139

The classical solution to the noise removal problem is the Wiener filter, which utilizes the second-order statistics of the Fourier decomposition. Subband decompositions of natural images have significantly non-Gaussian higher-order point statistics; these statistics capture image properties that elude Fourier-based techniques. We develop a Bayesian estimator that is a natural extension of the Wiener solution, and that exploits these higher-order statistics. The resulting nonlinear estimator performs a "coring" operation. We provide a simple model for the subband statistics, and use it to develop a semi-blind noise-removal algorithm based on a steerable wavelet pyramid.

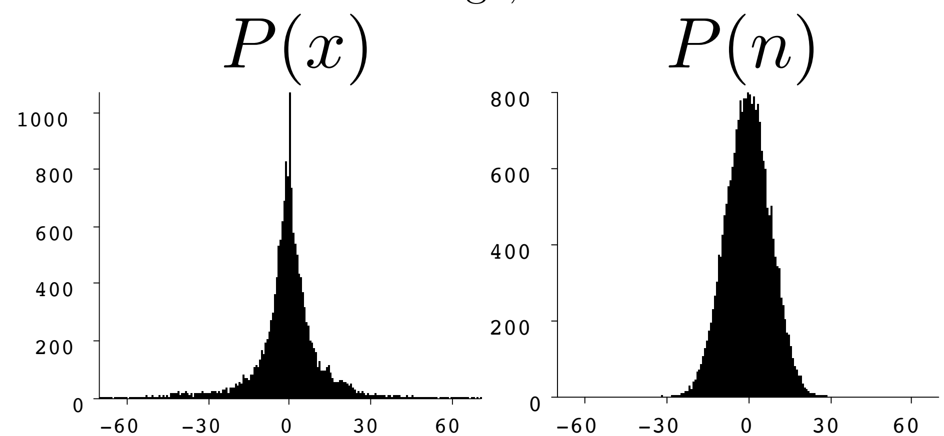

Figure 1: Histograms of a mid-frequency subband in an octave-bandwidth wavelet decomposition for two different images. Left: The "Einstein" image. Right: A white noise image with uniform pdf.

$$y = x + n$$

$$y = x + n$$

$$P(x) = \frac{1}{Z_s} e^{-|\frac{x}{s}|^p}$$

$$P(x|y) \propto P(y|x)\, P(x)$$

MAP estimate:

$$\hat{x} = \arg\min_x \left[\frac{|y-x|^2}{2\sigma_n^2} + |\frac{x}{s}|^p \right]$$

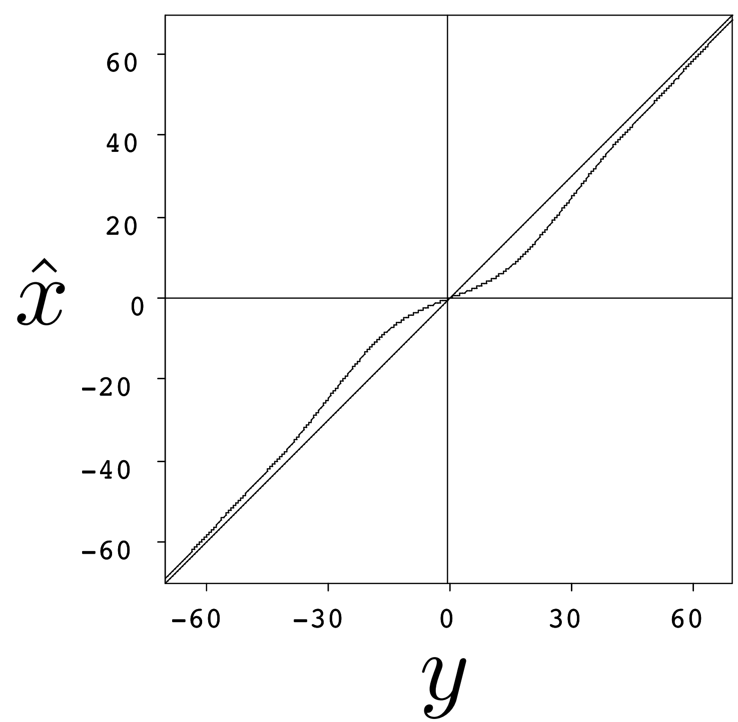

original image image + noise

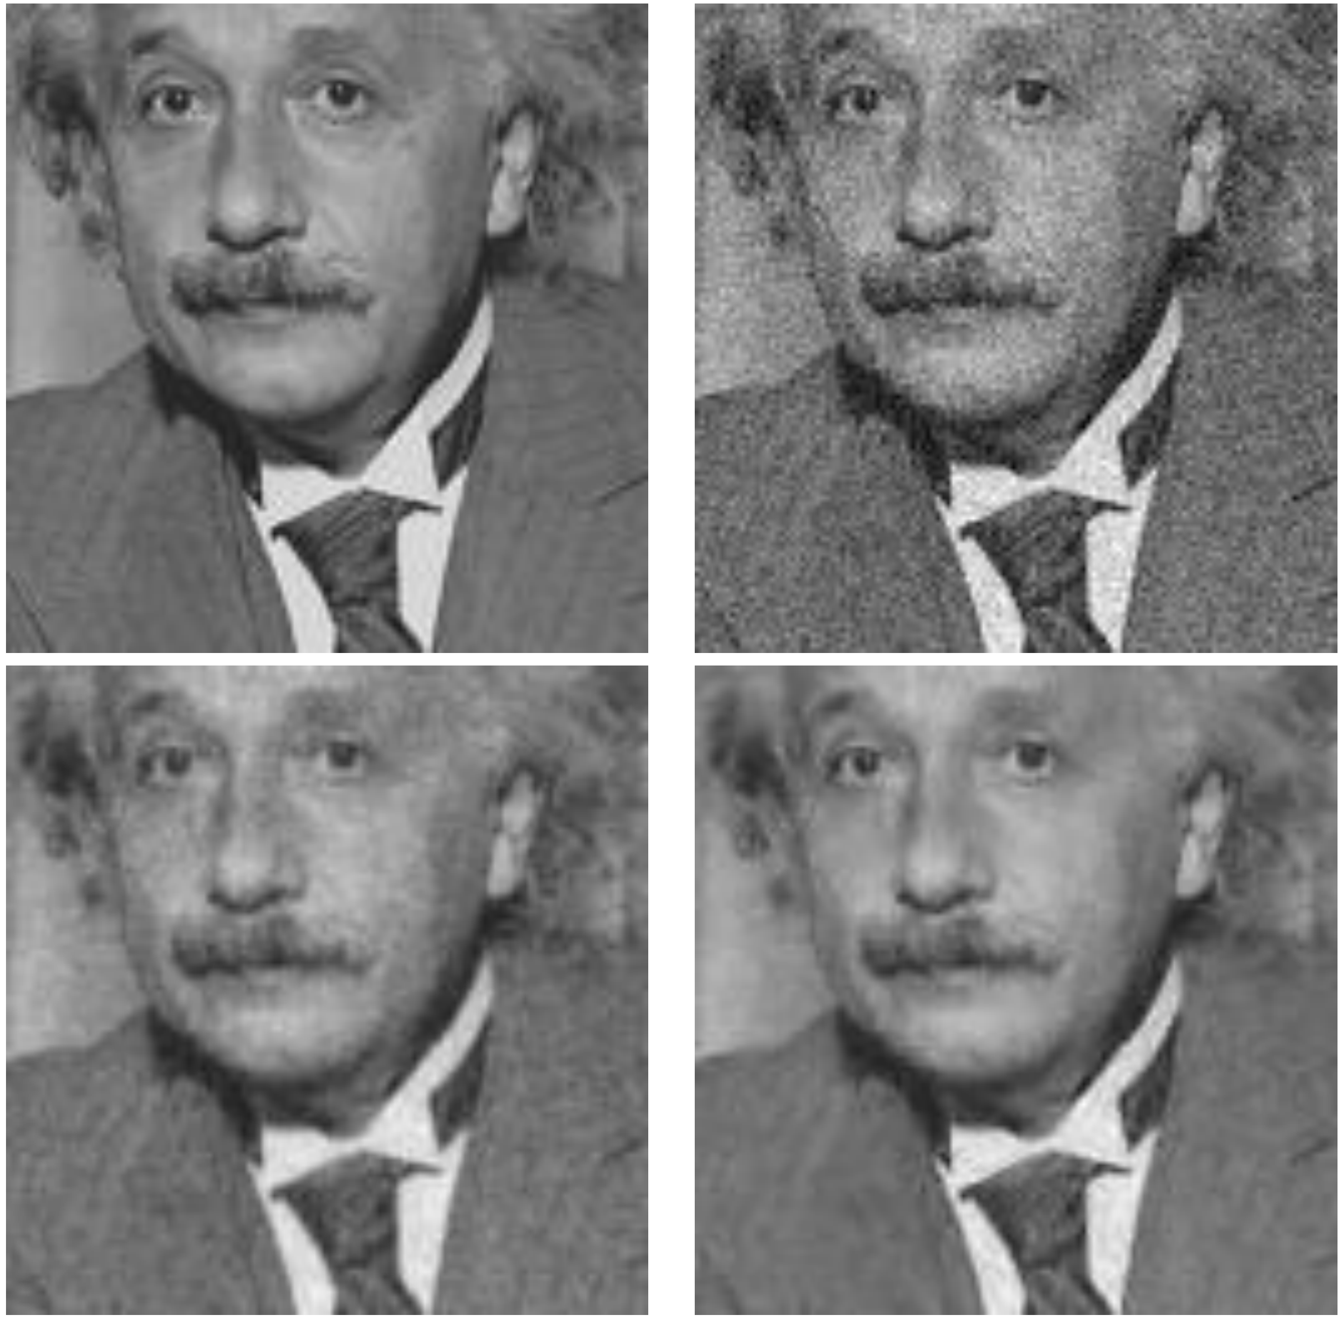

Wiener filter wavelet coring

Learning Horizontal Connections in a Sparse Coding Model of Natural Images

Pierre J. Garrigues
Department of EECS
Redwood Center for Theoretical Neuroscience
Univ. of California, Berkeley
Berkeley, CA 94720
garrigue@eecs.berkeley.edu

Bruno A. Olshausen
Helen Wills Neuroscience Inst.
School of Optometry
Redwood Center for Theoretical Neuroscience
Univ. of California, Berkeley
Berkeley, CA 94720
baolshausen@berkeley.edu

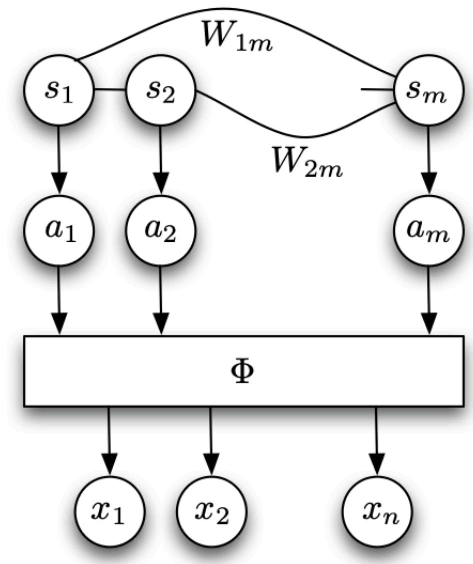

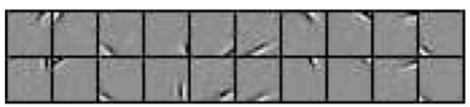

(a) 10 most positive weights

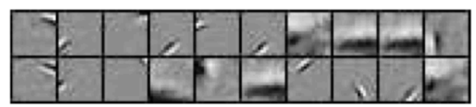

(b) 10 most negative weights

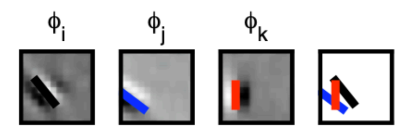

(c) Weights visualization

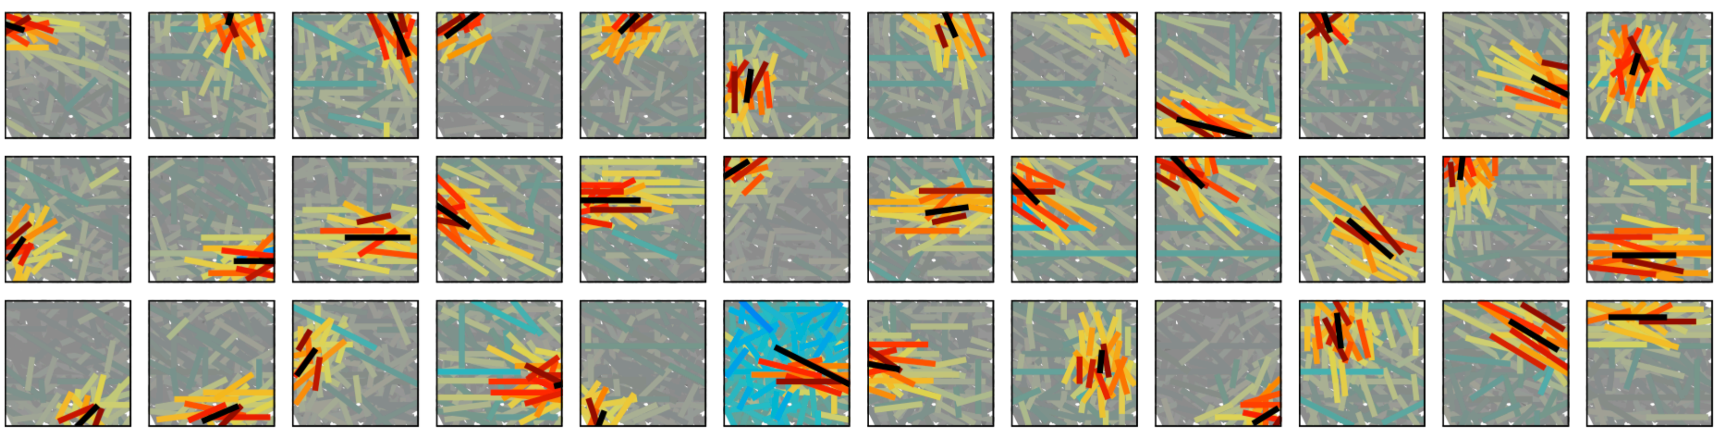

(d) Association fields

Dynamical models
(Kalman filter)

First-order Markov process

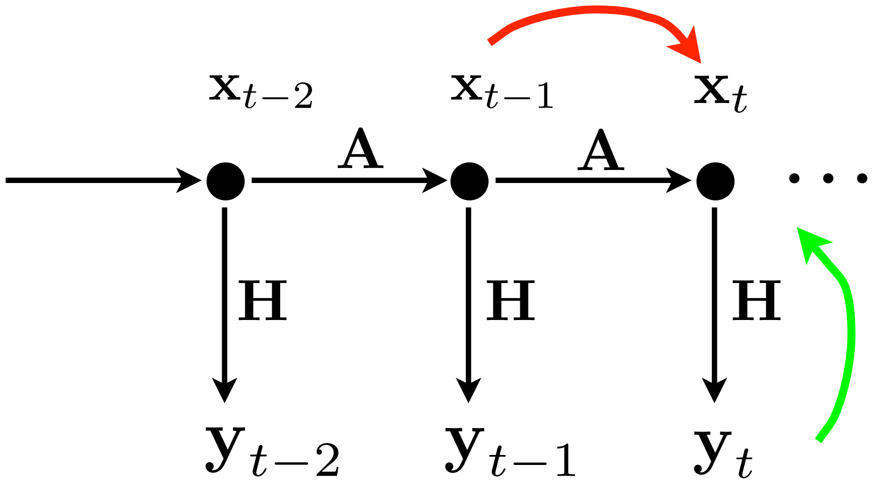

Linear generative model:

$$\begin{aligned}\mathbf{x}_t &= \mathbf{A}\,\mathbf{x}_{t-1} + \mathbf{w}_{t-1}\\ \mathbf{y}_t &= \mathbf{H}\,\mathbf{x}_t + \mathbf{n}_t\end{aligned}$$

Prediction:

$$P(\mathbf{x}_t|\mathbf{y}_0 \ldots \mathbf{y}_{t-1}) = \int_{-\infty}^{\infty} P(\mathbf{x}_t|\mathbf{x}_{t-1})\, P(\mathbf{x}_{t-1}|\mathbf{y}_0 \ldots \mathbf{y}_{t-1})\, d\mathbf{x}_{t-1}$$

Update:

$t \leftarrow t+1$

$$P(\mathbf{x}_t|\mathbf{y}_0 \ldots \mathbf{y}_t) \propto P(\mathbf{y}_t|\mathbf{x}_t)\, P(\mathbf{x}_t|\mathbf{y}_0 \ldots \mathbf{y}_{t-1})$$

Sparse coding of time-varying images

$$I(x, y, t) = \sum_i a_i(t) * \phi_i(x, y, t) + \nu(x, y, t)$$

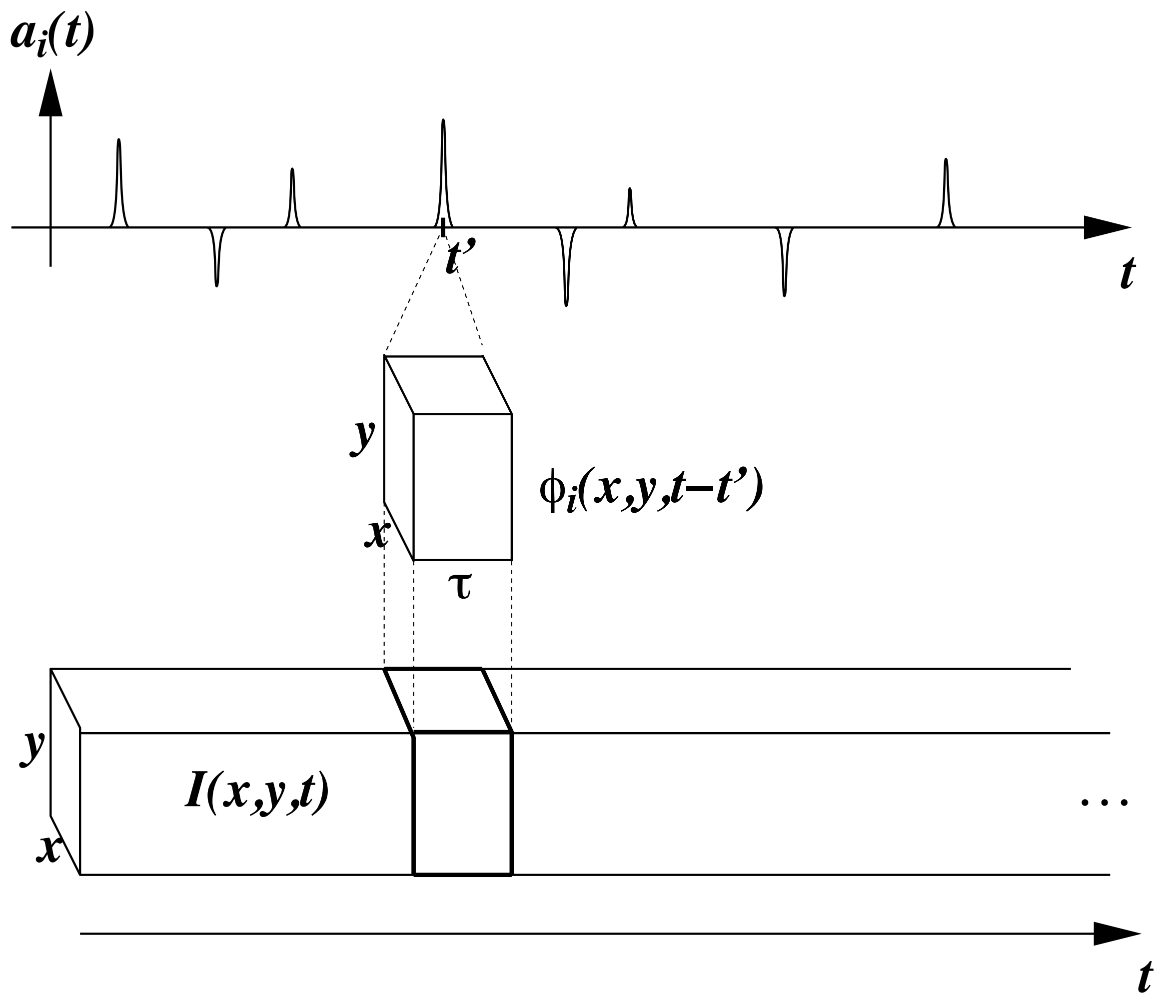

Learned basis space-time basis functions
(200 bfs, 12 x12 x 7)

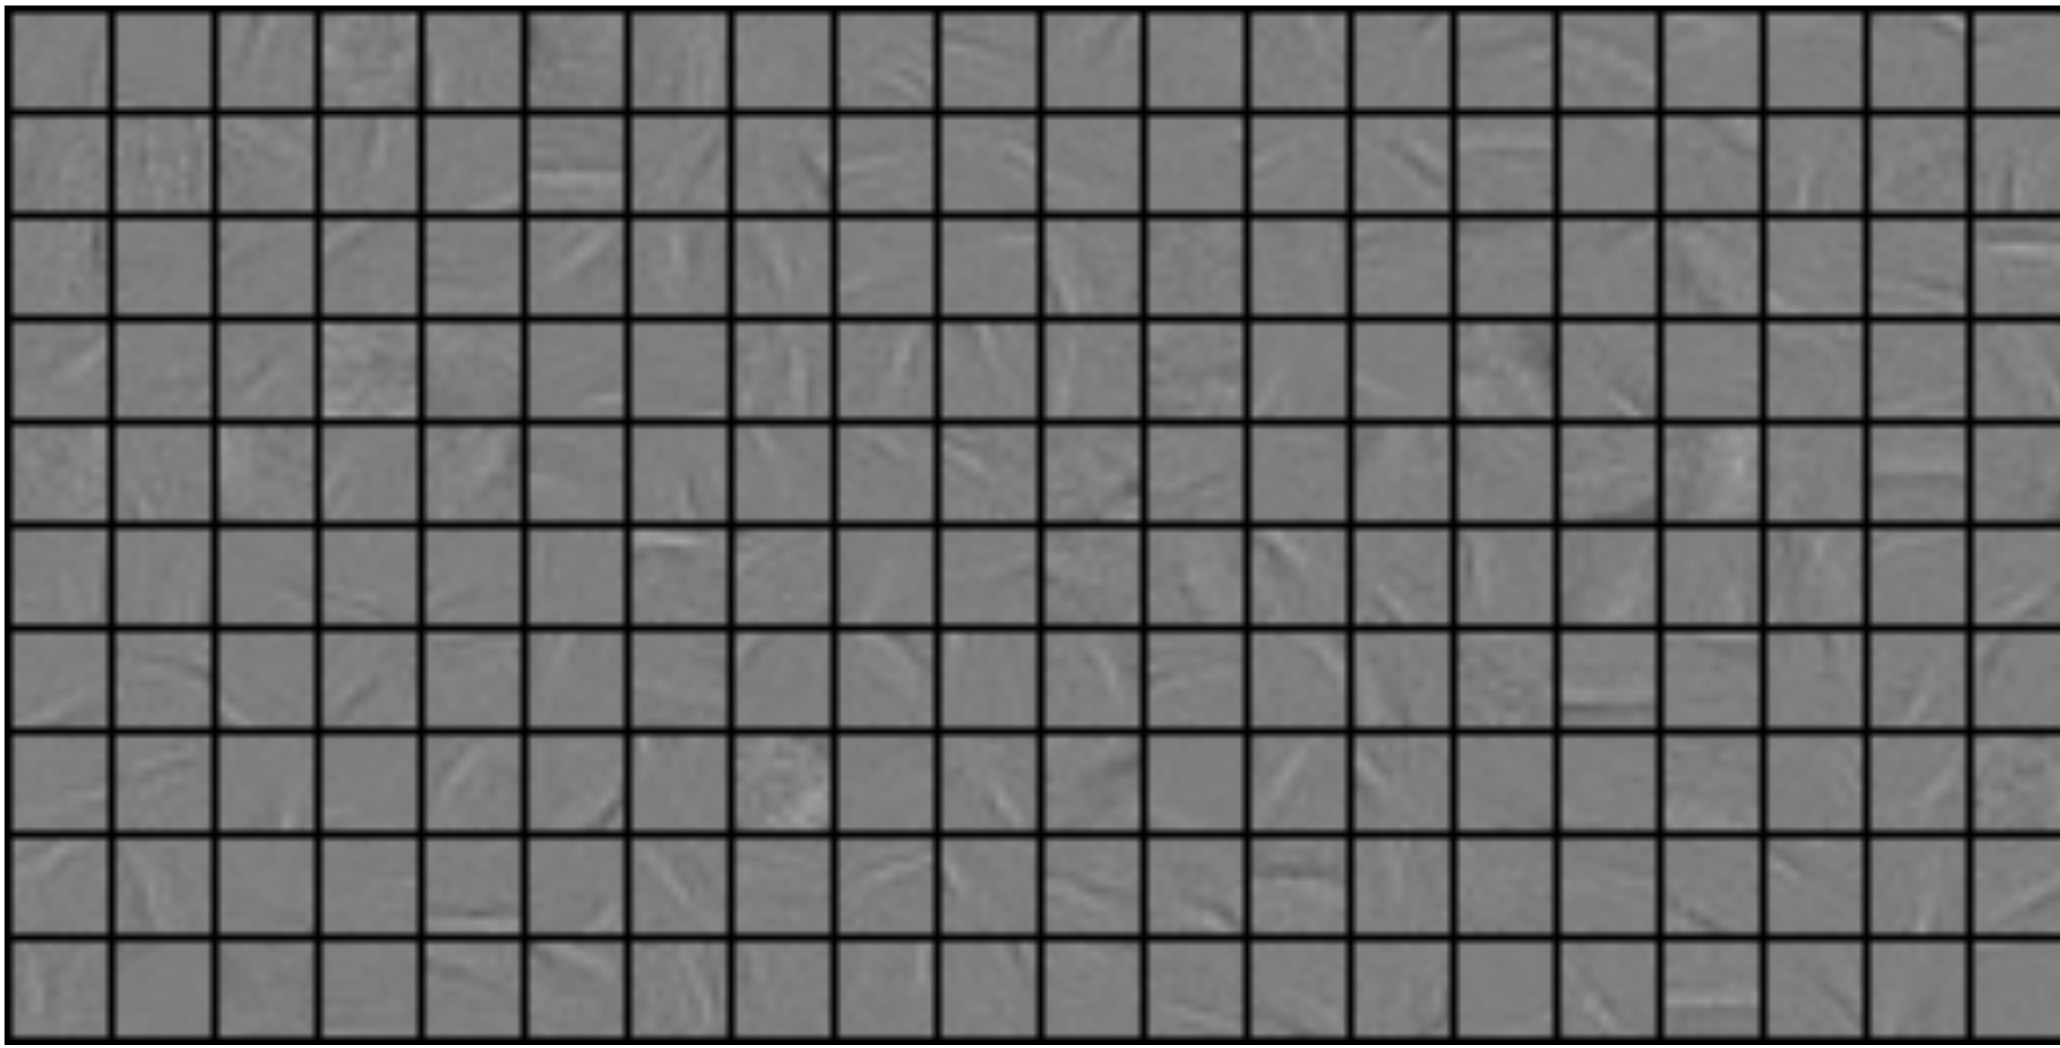

Sparse coding and reconstruction