Deblurring Noisy Gaussian Blurred Signals

In this example, we compare several methods for recovering a signal that has been blurred and then corrupted by noise. Let h, shown in Figure 1a, be the original clean signal. It is a discrete signal of length $N = 1000$ that we can think of as a one dimensional array with N elements or as a vector in $\mathbb{R}^N$. Note that h is sparse, meaning most of its entries $h(n)$ are zero. In fact, only 22 of its 1000 elements are nonzero.

The blurry signal shown in Figure 1c is obtained by convolving h with a normalized Gaussian blur kernel k having a standard deviation of 5. The blur kernel k is shown centered at 0 in Figure 1b. Since k is symmetric, the convolution $k * h$ is equivalent to computing the correlation $k \star h$. The effect is that each element of h is replaced by a weighted average of neighboring elements where the weights are defined by the Gaussian plotted in 1b. For simplicity, h is considered to be periodic, so $h(N+n) = h(n)$.

After h is blurred, it is further corrupted by the addition of white Gaussian noise, η, shown in Figure 1d. Each $\eta(n)$ is sampled independently from a zero mean normal distribution of standard deviation $\sigma = .005$. Although its amplitude seems small, it is still significant relative to the amplitude of the blurred signal. This can be seen in Figure 1e, which shows the noisy, blurry signal $f = k * h + \eta$.

Suppose we are given f and k and want to recover the original signal h. This is a deconvolution problem. Sometimes we might also not know k, in which case it is a much more difficult blind deconvolution problem. Here, we assume k is known. It is also reasonable in some cases to assume that the variance of the Gaussian noise is known. It can often be estimated from f by computing the mean and variance of some part of f that looks like it should correspond to a constant portion of h.

A naive approach to deconvolving f might be to assume the noise is negligible and attempt to solve $k * u$ for u. This is not a good approach even when there is very little noise. Since convolution is linear, we can define a matrix K so that $k * u = Ku$ and solving $k * u = f$ amounts to solving a linear system $Ku = f$. However, the matrix K is so ill-conditioned that even very small changes in f caused by things like noise or even round off error can result in very large changes in u.

Since $f = Kh + \eta$, we can attempt to find u such that $Ku - f$ has the same variance as η, namely

$$\frac{1}{N}\sum_n (Ku - f)_n^2 = \sigma^2.$$

We would like the vector u to be as close to h as possible, so among all vectors u satisfying the variance constraint, we might additionally look for the one with the smallest Euclidean norm $\|u\| = (\sum_n u_n^2)^{\frac{1}{2}}$. As long as $0 \leq \sigma\sqrt{N} \leq \|f\|$, there exists some $\lambda \geq 0$ such that the problem of minimizing $\|u\|$ over all u satisfying $\|Ku - f\|^2 = N\sigma^2$ is equivalent to minimizing

$$\frac{1}{2}\|u\|^2 + \frac{\lambda}{2}\|Ku - f\|^2.$$

The formula for the minimizer is given by

$$u = (I + \lambda K^T K)^{-1}\lambda K^T f,$$

which can be easily computed using the fast Fourier transform. The correct value of λ can be determined experimentally. The result for $\lambda = 40$ is shown in Figure 1f. Although it's an improvement over f, there are many spurious oscillations.

We might try to eliminate these oscillations by instead minimizing

$$\frac{1}{2}\|\nabla u\|^2 + \frac{\lambda}{2}\|Ku - f\|^2,$$

with

$$\nabla u(n) = u(n+1) - u(n)$$

and u treated as being periodic. There is a similar formula for the minimizer,

$$u = (\triangle + \lambda K^T K)^{-1}\lambda K^T f,$$

where $\triangle$ denotes the discrete Laplacian matrix defined by

$$\triangle u(n) = u(n+1) - 2u(n) + u(n-1).$$

The formula for the minimizer can again be computed using the fast Fourier transform. The result for $\lambda = 2$ is shown in Figure 1g. Oscillations are slightly reduced, but there is no significant improvement.

A much better result can be obtained by minimizing

$$\|u\|_1 \qquad \text{subject to} \qquad \|Ku - f\|^2 \leq N\sigma^2,$$

where $\|u\|_1 = \sum_n |u_n|$. There is no simple formula for the minimizer in this case, but it can still be found using an iterative procedure. The result is shown in Figure 1h. Minimizing $\|u\|_1$ has done a much better job of approximating our original sparse signal.

MATLAB: The accompanying MATLAB code can be run by typing sparse_deblur. It should automatically compute and plot the signals shown in Figure 1. The noise and blur parameters can be modified by editing sparse_deblur.m.

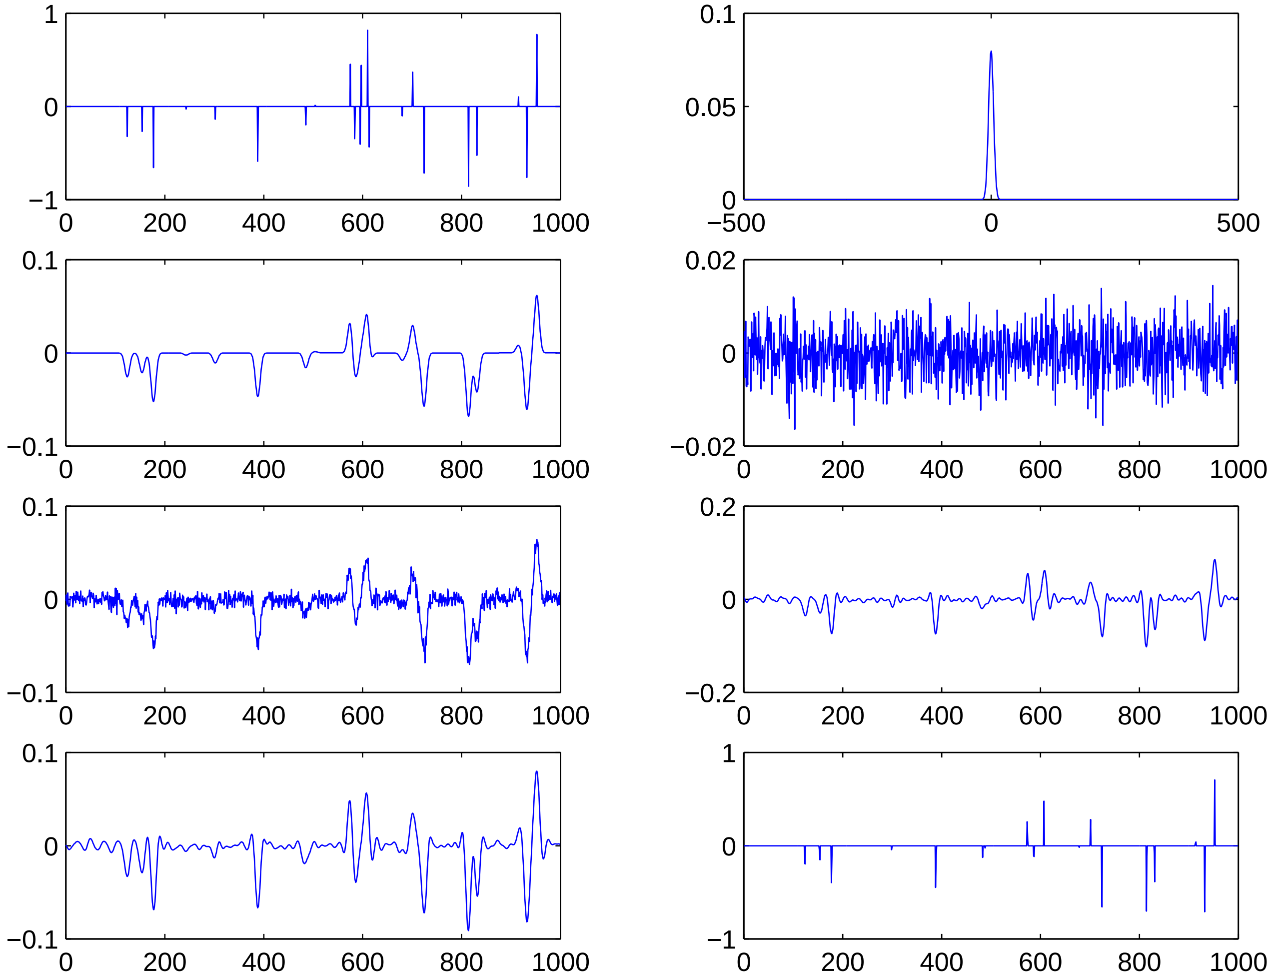

Figure 1: Comparison of several deblurring methods

a. clean signal h

b. Gaussian blur kernel k

c. blurry signal $k * h$

d. Gaussian noise η

e. noisy blurry signal $f = k * h + \eta$

f. $\arg\min \frac{1}{2}\|u\|^2 + \frac{\lambda}{2}\|k * u - f\|^2$

g. $\arg\min \frac{1}{2}\|\nabla u\|^2 + \frac{\lambda}{2}\|k * u - f\|^2$

h. $\arg\min \|u\|_1$ st $\|k * u - f\|^2 \leq N\sigma^2$