Lecture 16: Image Deblurring by Frequency Domain Operations

Harvey Rhody
Chester F. Carlson Center for Imaging Science
Rochester Institute of Technology
rhody@cis.rit.edu

November 8, 2005

Abstract

Image restoration by reduction of blurring is an important application of linear filter techniques. These filtering techniques are most easily understood in the frequency domain. Wiener, Constrained Least-Squares deblurring and blind deconvolution are presented.

Image Blur Model

Image blur is a common problem. It may be due to the point spread function of the sensor, sensor motion, or other reasons.

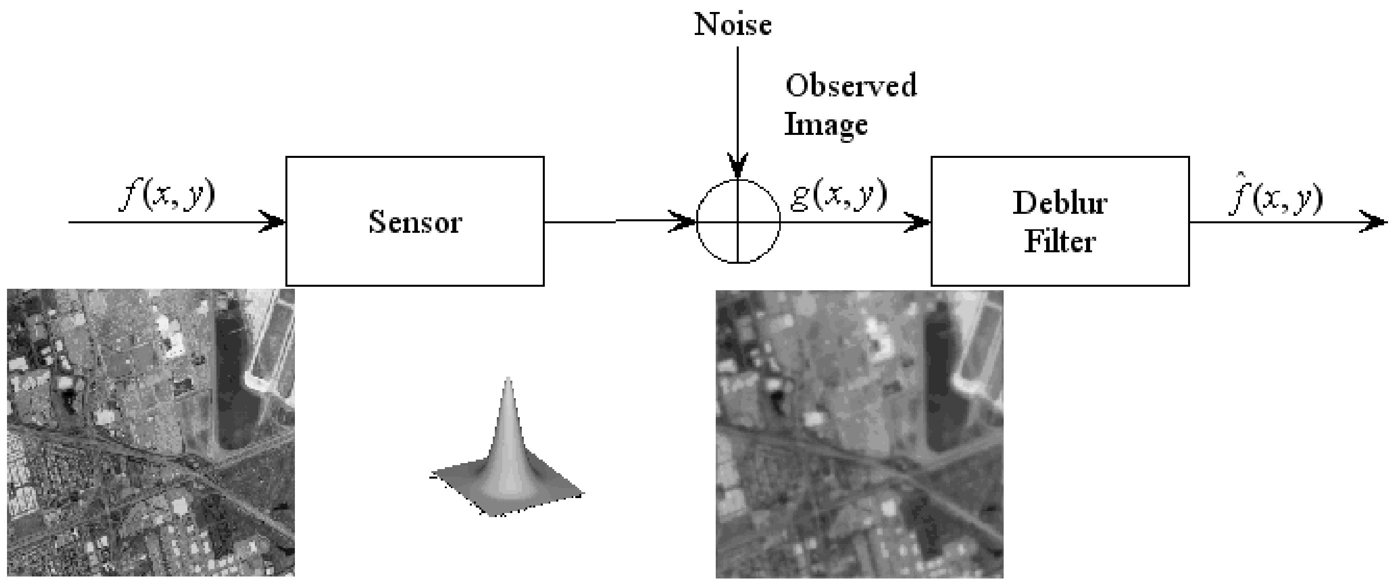

Linear model of observation system

$$g(x,y) = f(x,y) \star h(x,y) + \eta(x,y)$$

What deblurring filter will construct the best estimate $\hat{f}(x,y)$?

Frequency Domain Model

The observation equation can also be expressed in the frequency domain as

$$G(u,v) = F(u,v)H(u,v) + \mathcal{N}(u,v)$$

We can construct an estimate of $F(u,v)$ by filtering the observation $G(u,v)$. Let $T(u,v)$ be a linear shift-invariant reconstruction filter.

$$\hat{F}(u,v) = G(u,v)T(u,v)$$

Our task is to find a filter $T(u,v)$ that provides a good estimate of the original image.

The solution must balance noise reduction and sharpening of the image. These are conflicting goals.

Inverse Filtering

As a first attempt at a solution we can try the inverse filter. Assume that we know the sensor function $H(u, v)$. Try

$$T(u, v) = H^{-1}(u, v)$$

$$\hat{F}(u, v) = G(u, v)H^{-1}(u, v) = F(u, v) + \mathcal{N}(u, v)H^{-1}(u, v)$$

The result will be filtered noise added to the desired image.

The problem is that the inverse filter typically has very high gain at certain frequencies so that the noise term completely dominates the result.

Inverse Filter

Original Image

Blurred Image

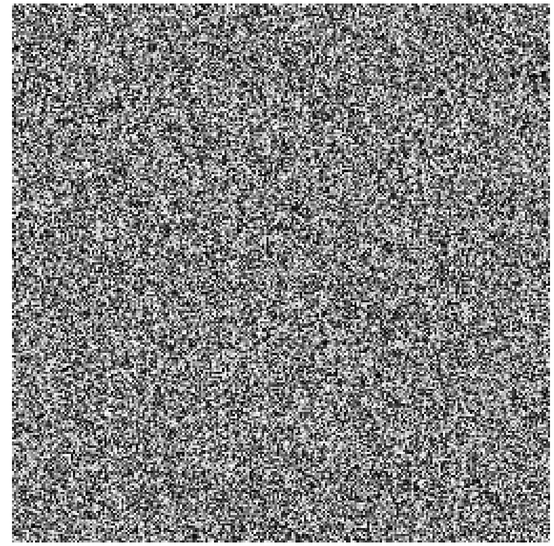

Restored with $H^{-1}(u,v)$

A small amount of noise saturates the inverse filter.

Inverse Filter

The inverse filter has very high gain at frequencies where $H(u,v)$ is small. The amplified noise at these frequencies will dominate the output.

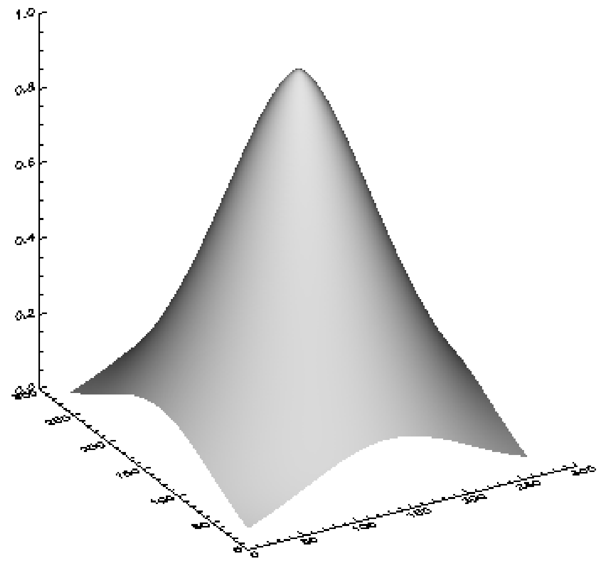

Sensor MTF

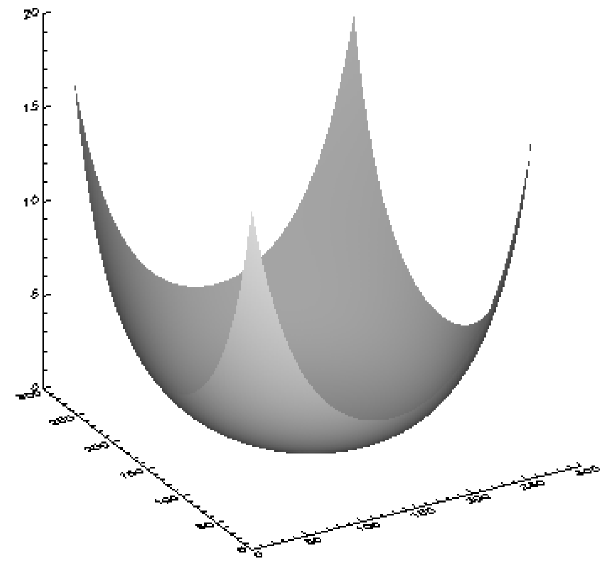

Inverse Filter

The frequency response of a very low-quality sensor and the inverse filter for it are shown above.

Modified Inverse Filter

One can attempt to reduce the noise amplification by modifying the inverse filter to lower high-frequency gain. Consider weighting H^{-1} with a Butterworth response.

$$B(u,v) = \frac{1}{1 + \left(\frac{u^2+v^2}{D^2}\right)^n}$$

$$T(u,v) = \frac{B(u,v)}{H(u,v)}$$

$$\begin{aligned} \hat{F}(u,v) &= (F(u,v)H(u,v) + \mathcal{N}(u,v))T(u,v) \\ &= F(u,v)B(u,v) + \frac{\mathcal{N}(u,v)B(u,v)}{H(u,v)} \end{aligned}$$

The goal is to select $B(u,v)$ so that the image is not distorted and the noise is still suppressed.

Modified Inverse Filter

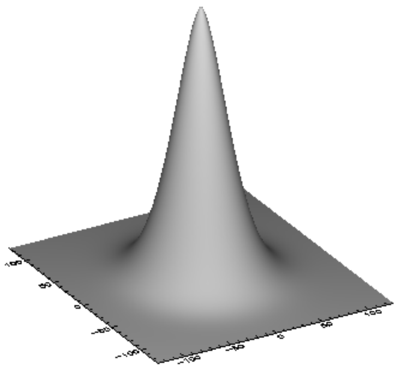

$H(u,v)$

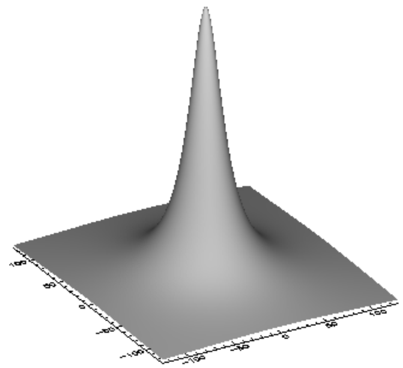

$B(u,v)$: $R = 20, n = 1$

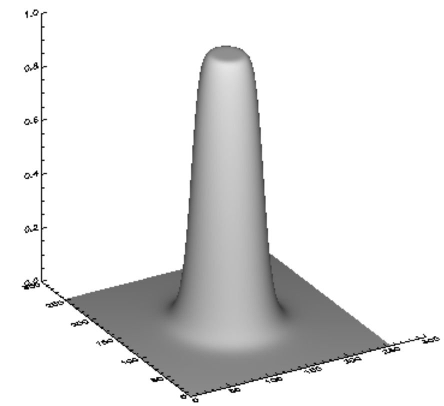

$R = 40, n = 1$

Blurred Image $G(u,v)$

Restored using $R = 20$

Restored using $R = 40$

Modified Inverse Filter

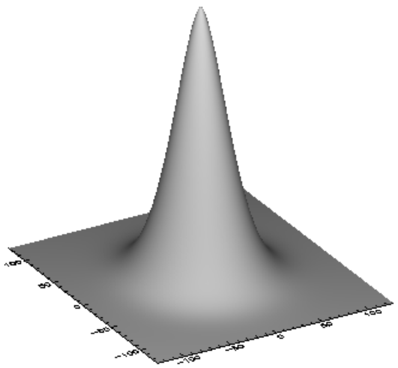

$H(u,v)$

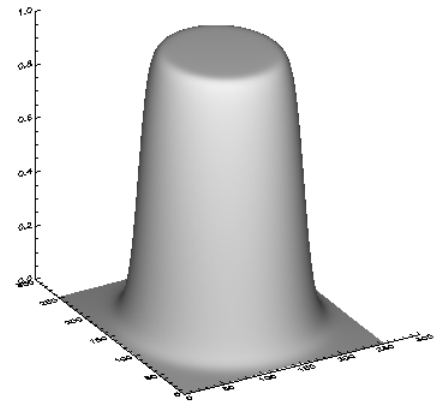

$B(u,v)$: $R = 90, n = 8$

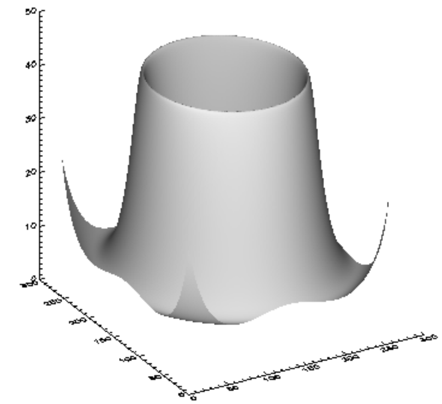

Inverse B/H

Original Image $G(u,v)$

Blurred using $R = 20$

Restored

Wiener Filter

The Wiener filter minimizes the mean-squared error

$$e^2 = E\left[(f(x,y) - \hat{f}(x,y))^2\right]$$

The frequency-domain expression for the Wiener filter is

$$T(u,v) = \frac{H^*(u,v)}{|H(u,v)|^2 + S_\eta(u,v)/S_f(u,v)}$$

where $S_\eta(u,v)$ and $S_f(u,v)$ are the power spectra of the noise and the original image.

When the power spectra are not known it is common to use a Wiener inspired approximation

$$T(u,v) = \frac{H^*(u,v)}{|H(u,v)|^2 + K}$$

where K has a small positive value.

Norbert Wiener 1894-1964

- Born in 1894 in Columbia, Missouri
- 1918 – Ph.D. in Philosophy, Harvard, age 18
- Instructor of Mathematics at MIT in 1919, Assistant Professor in 1929 and Professor in 1931
- 1942 – "The Yellow Peril" (classified)
- 1948 – "Cybernetics or Control and Communication in the Animal and the Machine"
- 1949 – "Extrapolation, Interpolation and Smoothing of Stationary Time Series with Engineering Applications"

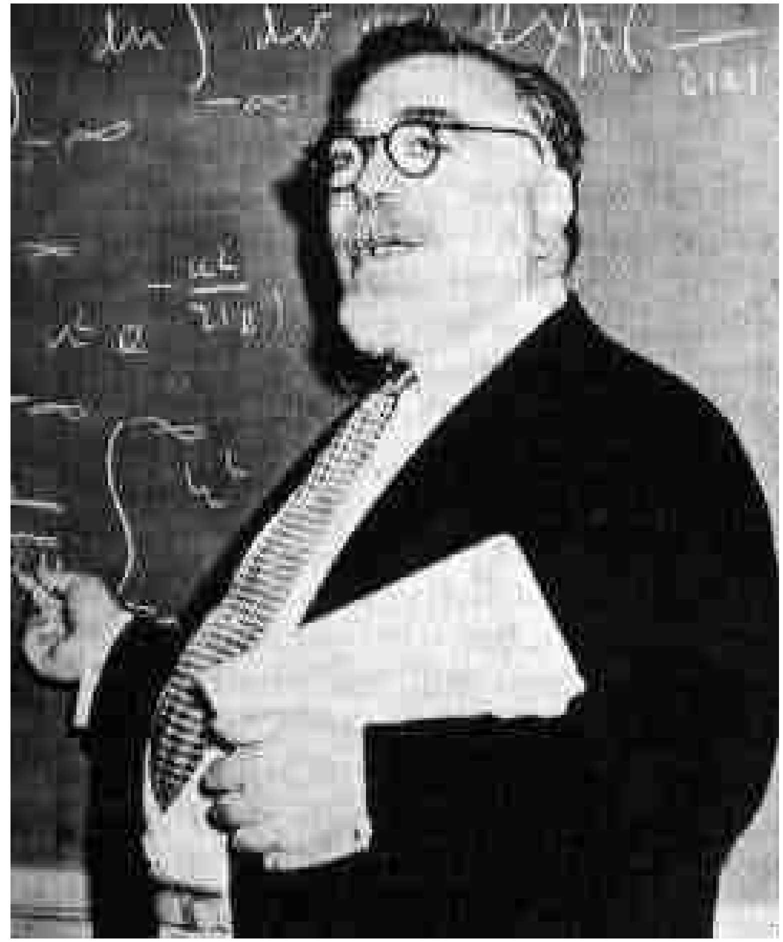

Norbert Wiener

Wiener Filter

Original image (left), blurred image (right)

Restored with Wiener filters with $K = 0.01$ (left) and $K = 0.0001$ (right)

Wiener Filter Response

The frequency response of Wiener filters with $K = 0.01$ and $K = 0.0001$ are shown below.

Note that they do not suffer from catastrophic gain at any frequency.

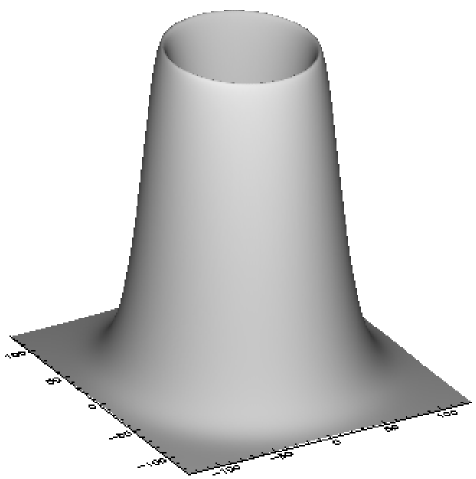

$K = 0.01$

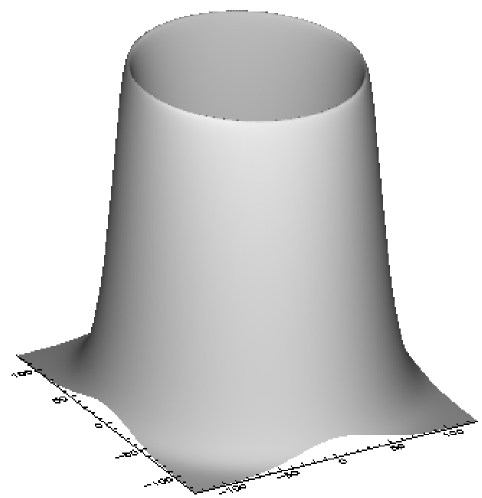

$K = 0.0001$

Comparison

Restoration with the Wiener approximation and the Butterworth modified inverse filter are similar.



Original — Wiener $K = 0.0001$ — Inverse Butterworth $[90, 8]$

Motion Blur Model

An image can be blurred by motion while the camera shutter is open. If an image $f(x, y)$ is moving with constant velocity and in T seconds moves to $f(x - \alpha, y - \beta)$ then the image registered by a sensor is

$$g(x, y) = \int_0^T f(x - \frac{\alpha t}{T}, y - \frac{\beta t}{T})dt$$
$$G(u, v) = F(u, v)H(u, v)$$

where the system transfer function is

$$H(u, v) = T\frac{\sin \pi(\alpha u + \beta v)}{\pi(\alpha u + \beta v)}e^{-i\pi(\alpha u+\beta v)}$$

Motion Blur Example

An image that has motion blur with parameters $\alpha = \beta = 0.1$ and $T = 1$ is shown below. (Original left, blurred right).

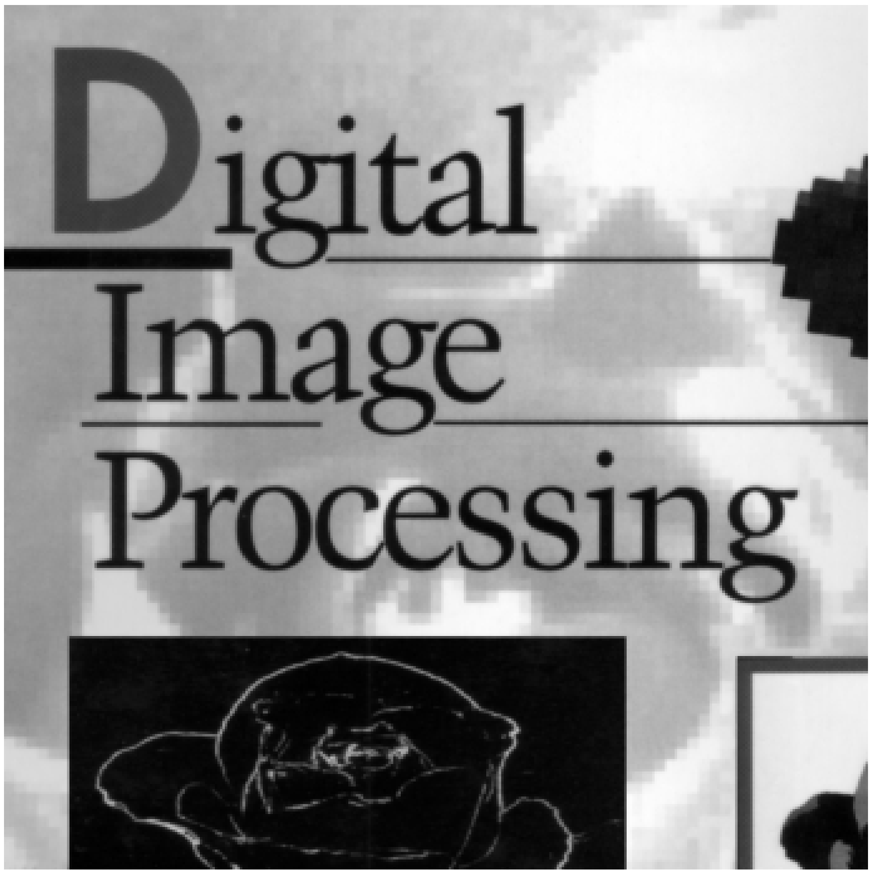

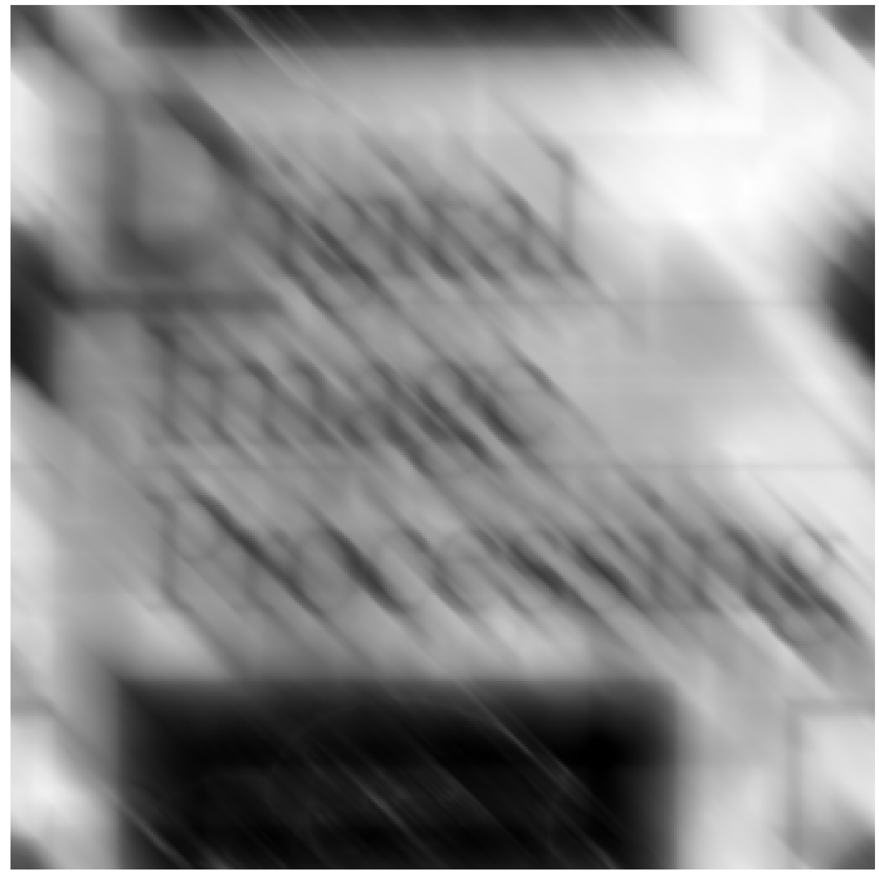

Wiener Filter Restoration of Motion Blur

Shown below are three blurred images with different noise levels and the images recovered by Wiener filtering.

The restoration filter was the Wiener filter approximation

$$T(u,v) = \frac{H^*(u,v)}{|H(u,v)|^2 + K}$$

with K chosen for best visual quality for each image by trial and error.

Wiener Restoration Results

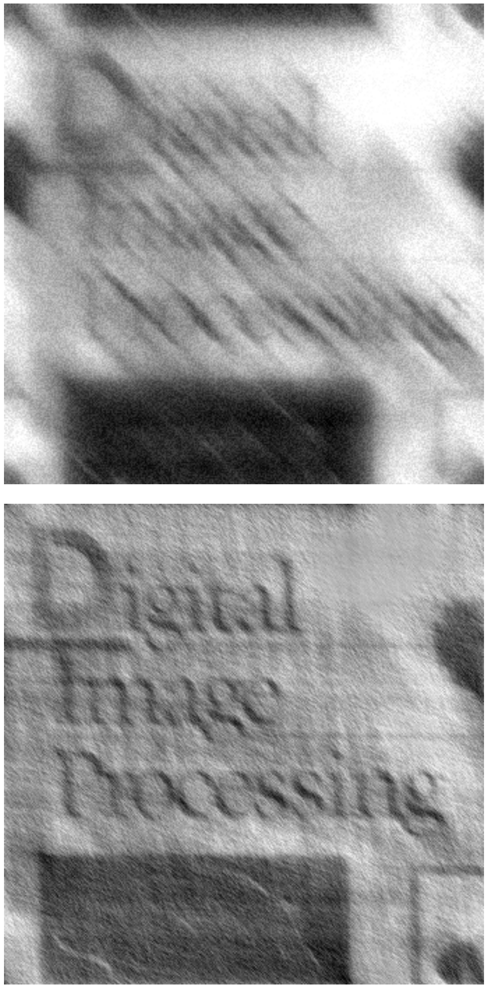

High Noise

Medium Noise

Low Noise

Constrained Least Squares Filtering

The Wiener filter uses the power spectrua of the actual image and noise. The CLS method seeks to constrain the variation in the image due to noise without actual knowledge of either power spectrum.

The CLS reconstruction is

$$\hat{F}(u,v) = \left[\frac{H^*(u,v)}{|H(u,v)|^2 + \gamma |P(u,v)|^2} \right] G(u,v)$$

$P(u,v)$ is the Fourier transform of the Laplacian filter

$$p(x,y) = \begin{bmatrix} 0 & -1 & 0 \\ -1 & 4 & -1 \\ 0 & -1 & 0 \end{bmatrix}$$

The filter $P(u,v)$ has a large amplitude at high frequencies, where the noise tends to be dominant. It modifies the denominator to reduce the noise effects at high frequencies.

CLS-Wiener Comparison

The reconstruction results are shown for CLS (top) and Wiener (bottom).

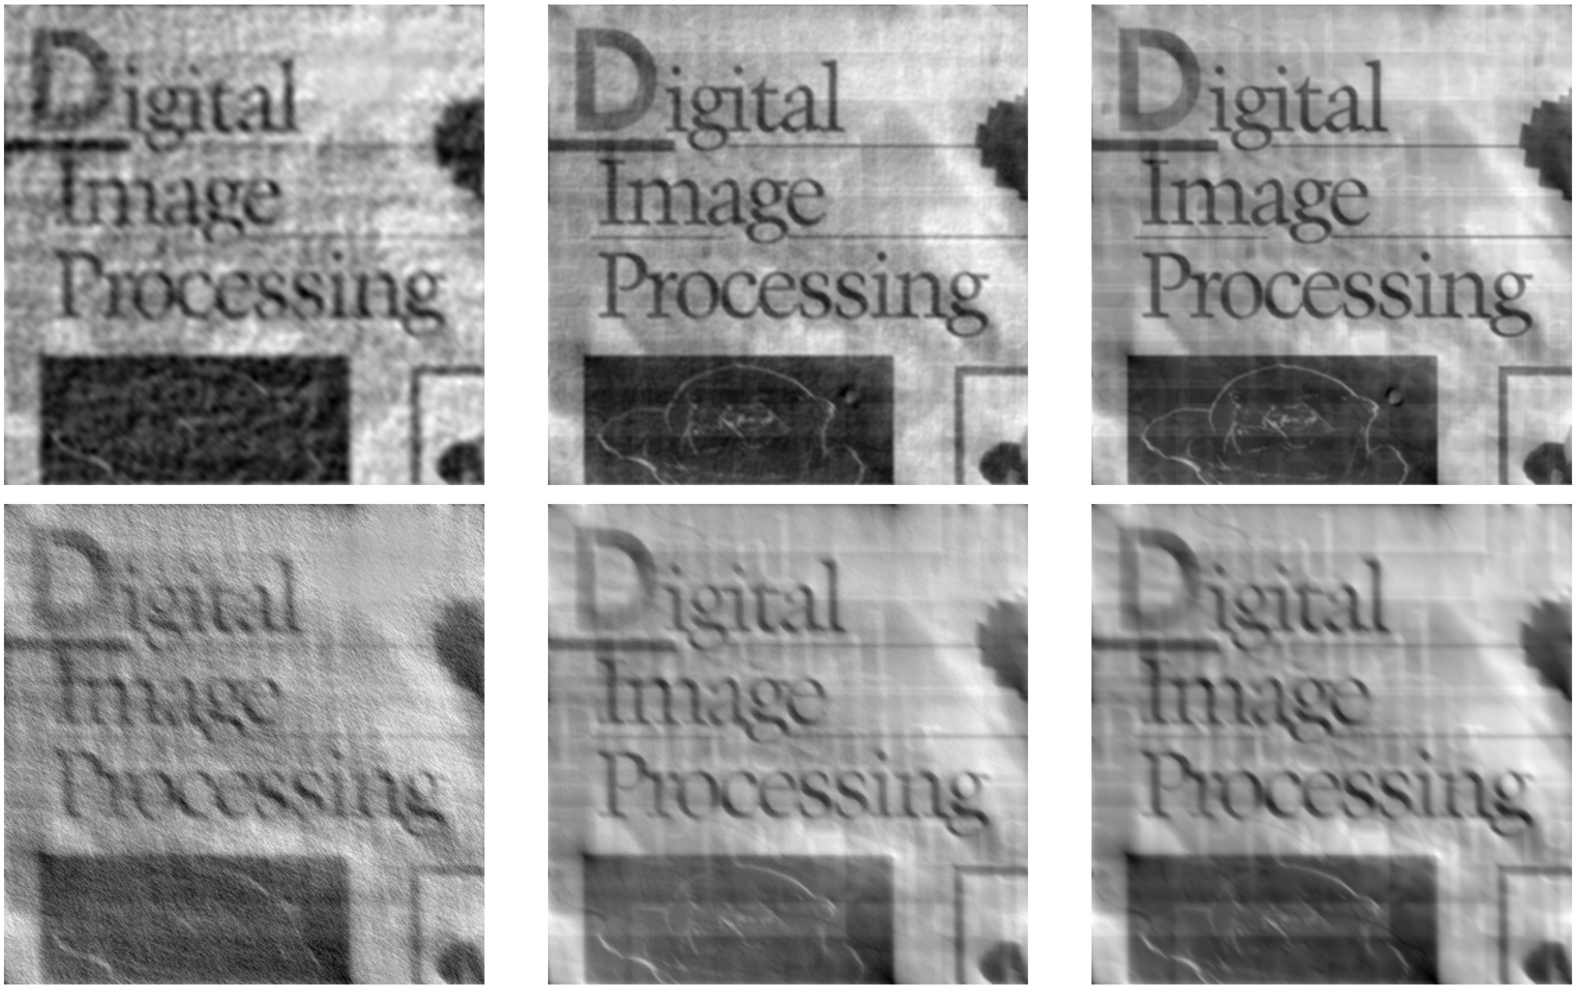

High Noise | Medium Noise | Low Noise

Blind Deconvolution

If the blur function $H(u, v)$ is not known then recovery of $F(u, v)$ from

$$G(u, v) = F(u, v)H(u, v) + \mathcal{N}(u, v)$$

is more difficult.

Blind deconvolution is a description of a class of techniques that estimate $H(u, v)$ from one or more images and then use that estimate to recover an estimate of $F(u, v)$.

If nothing is known about the image or blur function then it is not possible to solve the blind deconvolution problem.

Blind deconvolution algorithms differ in what is assumed about $F(u, v)$ and $H(u, v)$.

In developing a blind deconvolution by heuristics we will ignore the additive noise. It may be possible to modify the process to reduce noise effects where that is necessary.

Blind Deconvolution

One practical algorithm[1] is based on the assumption that $|H(u,v)|$ is a smooth function. This is reasonable for many sensor systems.

The concept for the algorithm is to separate the image and blur functions by using smoothness. If we ignore noise, we have

$$|G(u,v)| = |F(u,v)||H(u,v)|$$

Express $|F(u,v)|$ as the sum of a slowly and rapidly varying components.

$$|F(u,v)| = |F(u,v)|_L + |F(u,v)|_H$$

Then

$$|G(u,v)| = |H(u,v)||F(u,v)|_L + |H(u,v)||F(u,v)|_H$$

[1] Jae S. Lim, Two-Dimensional Signal and Image Processing, Prentice-Hall, 1990. Section 9.3.2, "Algorithms for Blind Deconvolution"

Blind Deconvolution

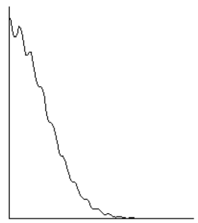

$\bullet$

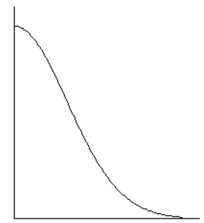

$=$

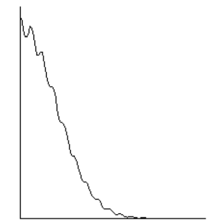

$|F(u,v)|$ $\qquad |H(u,v)|$ $\qquad |G(u,v)|$

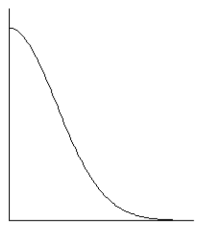

$+$

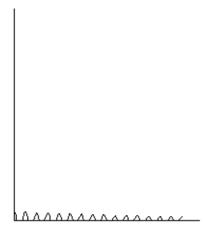

$= |F(u,v)|$

$|F(u,v)|_L$ $\qquad |F(u,v)|_H$

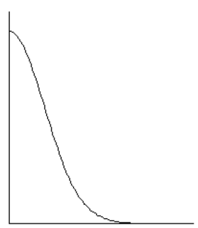

$+$

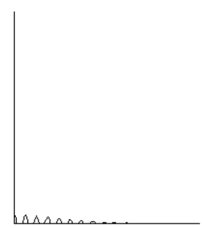

$= |G(u,v)|$

$|F(u,v)|_L \cdot |H(u,v)|$ $\qquad |F(u,v)|_H \cdot |H(u,v)|$

Blind Deconvolution

The low-frequency term in the expression

$$|G(u,v)| = |H(u,v)||F(u,v)|_L + |H(u,v)||F(u,v)|_H$$

can be extracted with a suitable smoothing filter, $S(u,v)$.

$$\begin{aligned} S(u,v)|G(u,v)| &\approx S(u,v)|H(u,v)||F(u,v)|_L + S(u,v)|H(u,v)||F(u,v)|_H \\ &\approx |H(u,v)||F(u,v)|_L \end{aligned}$$

We can then estimate the blur function from

$$|\hat{H}(u,v)| \approx \frac{S(u,v)|G(u,v)|}{|F(u,v)|_L}$$

The numerator can be constructed from the image. The denominator can be estimated from similar images.

Blind Deconvolution for Atmospheric Blur Removal

A 1.6 meter ground-based telescope sensing green light, for example, should resolve features on the order of 10 inches when observing the Hubble Space Telescope in its 600 km orbit. Time-varying changes in the refractive index of Earth's atmosphere can, however, swell this resolution by a factor of 15 making meaningful inference about the satellite from ground-based pictures nearly impossible.[2]

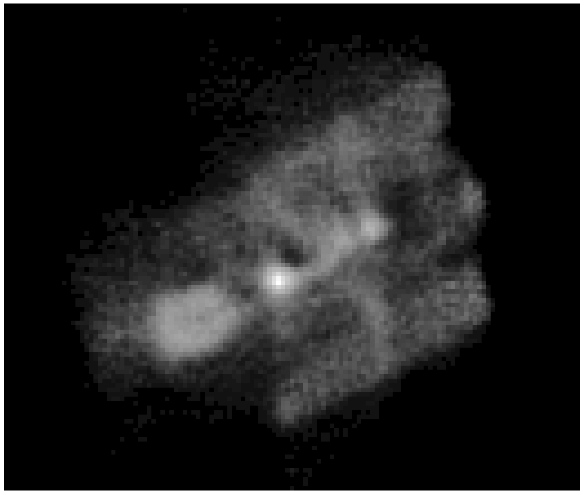

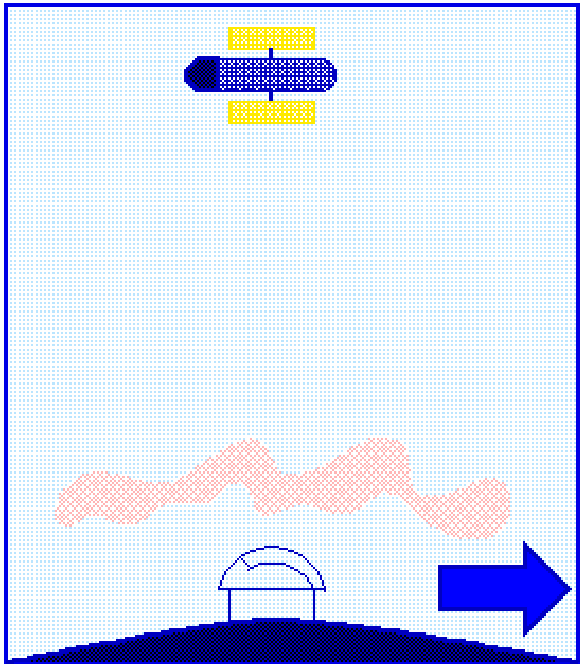

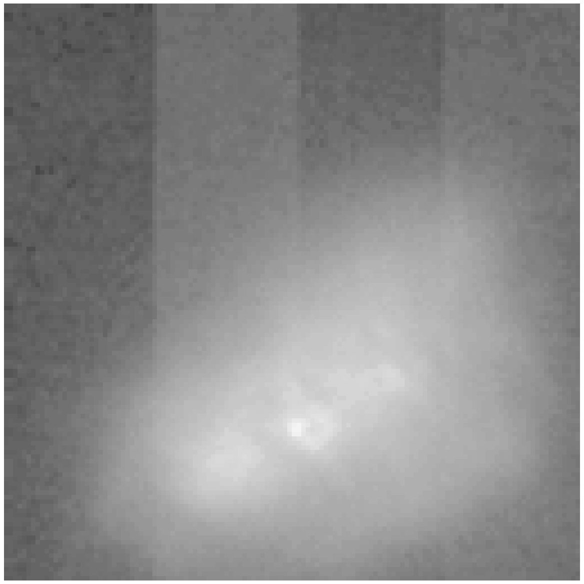

Hubble Space Telescope as acquired by a 1.6 m telescope at Air Force Maui Optical Station

[2] Timothy J. Schulz, Michigan Technical University, http://www.ece.mtu.edu/faculty/schulz/research/blind_deconvolution.html

Blind Deconvolution for Atmospheric Blur Removal

A maximum-likelihood estimation method has been considered for recovering fine-resolution imagery from a sequence of noisy, blurred images. A numerical technique based on the expectation-maximization (EM) procedure has been developed for solving this multiframe blind deblurring problem. A parallel implementation of this algorithm on an IBM SP2 computer at the Maui High Performance Computing Center has been used to restore the resolution of ground-based telescope imagery of the Hubble Space Telescope.

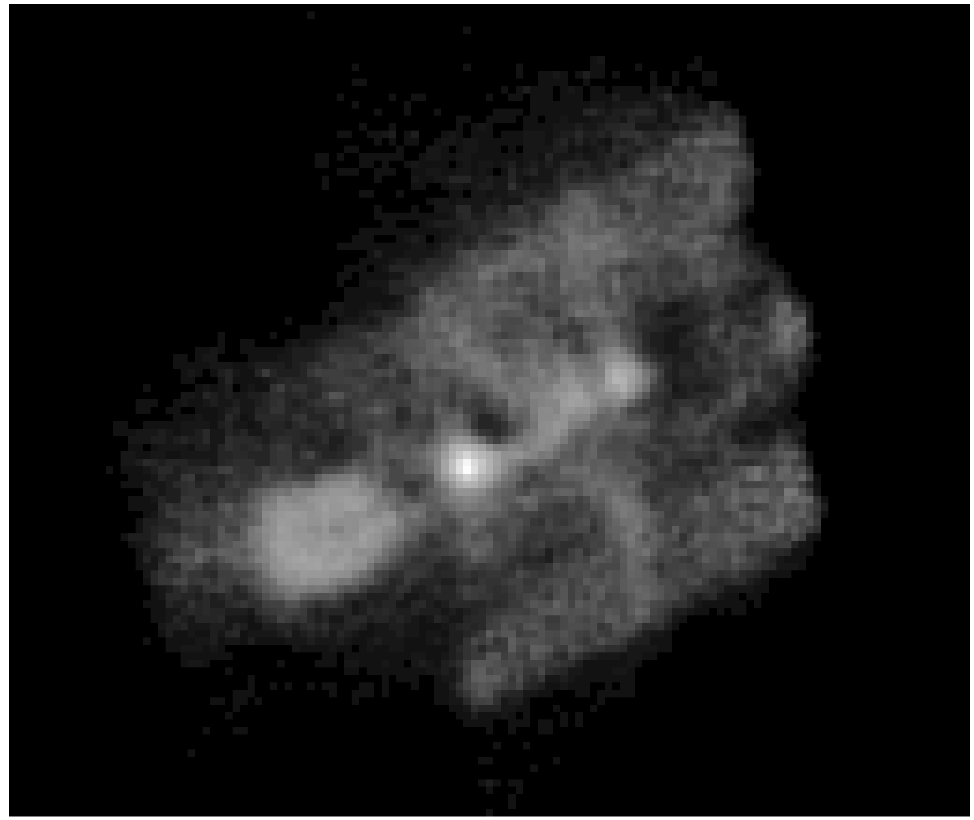

Phase vs Magnitude

The Fourier transform of an image $f(x, y)$ can be expressed in magnitude-phase form as

$$F(u, v) = |F(u, v)|e^{i\theta(u,v)}$$

A question is sometimes asked "What is the most important, the magnitude spectrum or the phase spectrum?"

The magnitude spectrum describes the energy distribution as a function of frequency. We have seen that the mean-squared value of $f(x, y)$ can be found by summing over $|F(u, v)|^2$.

However, the phase spectrum contains important information about image detail. One cannot expect to reconstruct $f(x, y)$ from $|F(u, v)|$ because the phase-shift for each component is completely ambiguous. However, as shown on the next page, one can reconstruct a reasonable approximation from the phase spectrum.

Phase Substitution

The magnitude spectrum and phase spectrum are both important, but which is more powerful? Let A and B be image arrays, and let

$$
\begin{aligned}
AF &= \mathcal{F}(A) \\
BF &= \mathcal{F}(B) \\
CF &= |AF| e^{i\phi_B} \\
DF &= |BF| e^{i\phi_A} \\
C &= \mathcal{F}^{-1}(CF) \\
D &= \mathcal{F}^{-1}(DF)
\end{aligned}
$$

The effect of the interchange of phase and magnitude spectra is illustrated here.