ELEG 404/604 - Digital Image and Audio Signal Processing

Restoration

Gonzalo R. Arce

Charles Black Evans Professor

Department of Electrical and Computer Engineering
University of Delaware

Spring 2014

What is Restoration (vs. Enhancement)

Restoration attempts to recover an image that has been degraded by using a priori knowledge of the degradation phenomenon.

Enhancement	Restoration
• Making pleasing images • Often no specific model of the degradation • Ad hoc procedures	• Undoing (inverting) and unwanted effect • Model-based approach • Optimality criteria

Image Restoration Problem

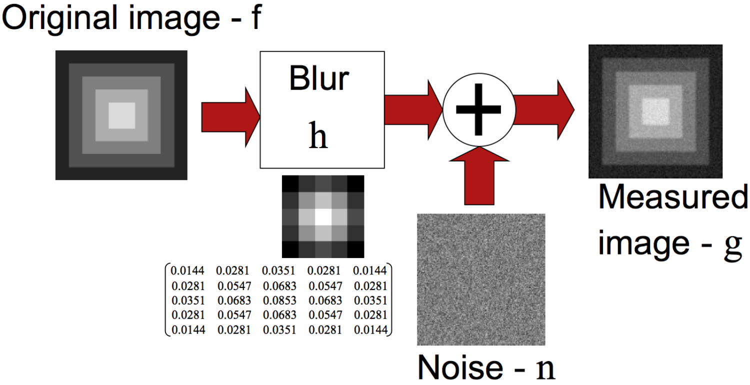

$$\begin{pmatrix} 0.0144 & 0.0281 & 0.0351 & 0.0281 & 0.0144 \\ 0.0281 & 0.0547 & 0.0683 & 0.0547 & 0.0281 \\ 0.0351 & 0.0683 & 0.0853 & 0.0683 & 0.0351 \\ 0.0281 & 0.0547 & 0.0683 & 0.0547 & 0.0281 \\ 0.0144 & 0.0281 & 0.0351 & 0.0281 & 0.0144 \end{pmatrix}$$

f,g,n of size $M \times M$

Image Restoration Problem

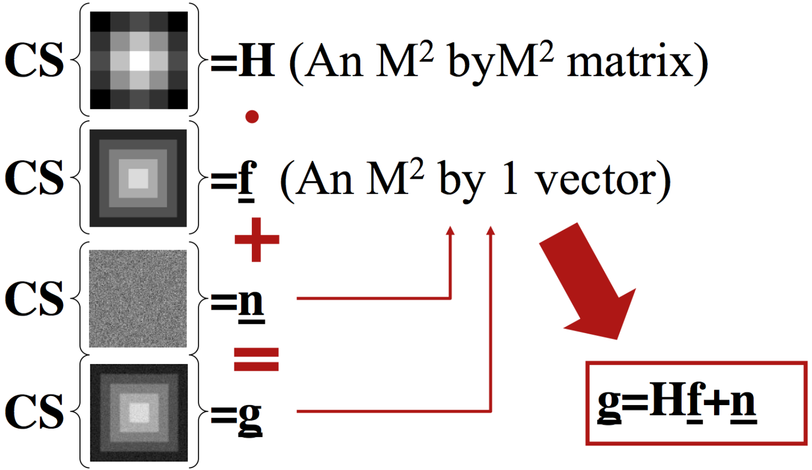

Image Restoration Problem

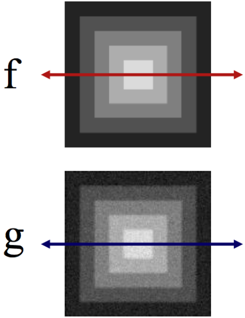

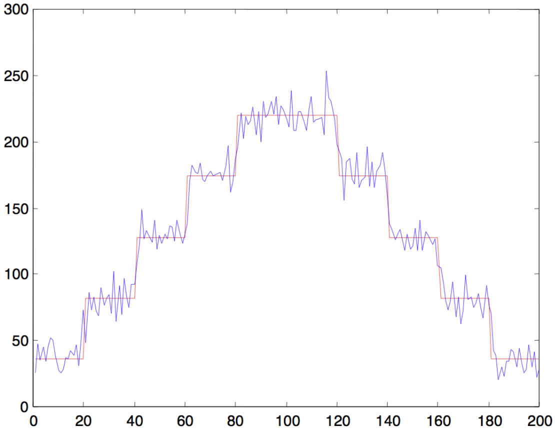

A Model of the Image Degradation/Restoration Process

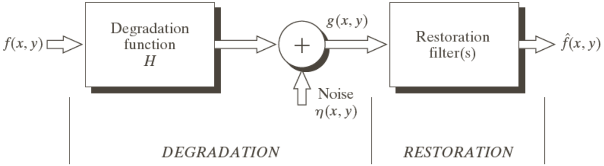

$f(x, y)$: an input image

$g(x, y)$: a degraded image

$h(x, y)$: the degradation function

$\eta(x, y)$: the additive noise

$\hat{f}(x, y)$: an estimate of the original image

The more we know about h and η, the closer $\hat{f}(x, y)$ will be to $f(x, y)$

Degradation and Restoration Model

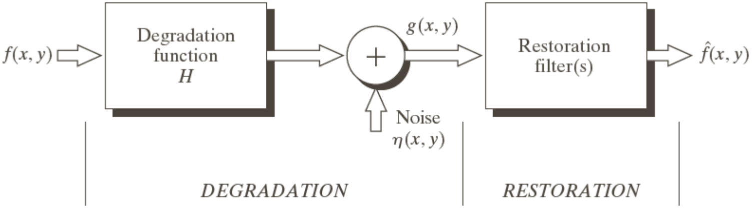

- Degradation is taken to be a linear spatially invariant operator

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

$$G(u,v) = H(u,v)F(u,v) + N(u,v)$$

Noise Properties

- Arises in acquisition, digitization, and transmission/storage processes
- CCD cameras are affected by:
  - Light levels
  - Sensor temperature
  - Bad sensors
- Transmission noise can be due to interference
  - Wireless transmission interference
  - Lost networking packets

Noise Probability Density Functions

- Gaussian (normal) PDF:

$$p(z) = \frac{1}{\sqrt{2\pi\sigma}} e^{-(z-\bar{z})^2/2\sigma^2}$$

where z represents intensity, $\bar{z}$ is the mean value of z, and σ is its standard deviation.

  - Typically models electronics and sensor noise
  - Central Limit Theorem justification

- Rayleigh PDF:

$$p(z) = \begin{cases} \frac{2}{b}(z-a)e^{-(z-a)^2/b} & \text{for } z \geq a \\ 0 & \text{for } z < a \end{cases}$$

$$\bar{z} = a + \sqrt{\pi b/4}, \quad \sigma^2 = \frac{b(4-\pi)}{4}$$

  - Skewed distribution typically models range imaging noise

Noise Probability Density Functions

- Gamma PDF:

$$p(z) = \begin{cases} \frac{a^b z^{b-1}}{(b-1)!}(z-a)e^{-az} & \text{for } z \geq 0 \\ 0 & \text{for } z < 0 \end{cases}$$

$$\bar{z} = \frac{b}{a}, \quad \sigma^2 = \frac{b}{a^2}$$

- Exponential PDF:

$$p(z) = \begin{cases} ae^{-az} & \text{for } z \geq 0 \\ 0 & \text{for } z < 0 \end{cases}$$

$$\bar{z} = \frac{1}{a}, \quad \sigma^2 = \frac{1}{a^2}$$

  - Both appropriate for laser imaging
  - Heavy-tailed distributios
    - samples contain frequent outliers

Noise Probability Density Functions

- Uniform PDF:

$$p(z) = \begin{cases} \dfrac{1}{b-1} & \text{for } a \leqslant z \leqslant b \\ 0 & \text{otherwise} \end{cases}$$

$$\bar{z} = \frac{a+b}{2}, \quad \sigma^2 = \frac{(b-a)^2}{12}$$

- Impusive (salt and pepper) PDF:

$$p(z) = \begin{cases} P_a & \text{for } z = a \\ P_b & \text{for } z = b \\ 0 & \text{otherwise} \end{cases}$$

  - Shot or spike notice appropriate for faulty sensor or electrnics, transmission error/drop.

Noise Probability Density Functions

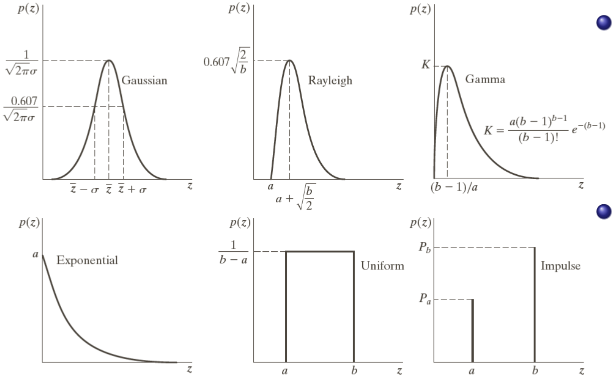

- Gaussian distribution is most widely used
  - Desirable properties
  - Independence and correlated
- Other distributions appropriate dor specific cases
  - Simplicity of uniform enables derivation of results

Test Pattern Corruption Example

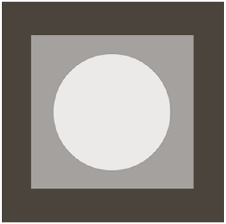

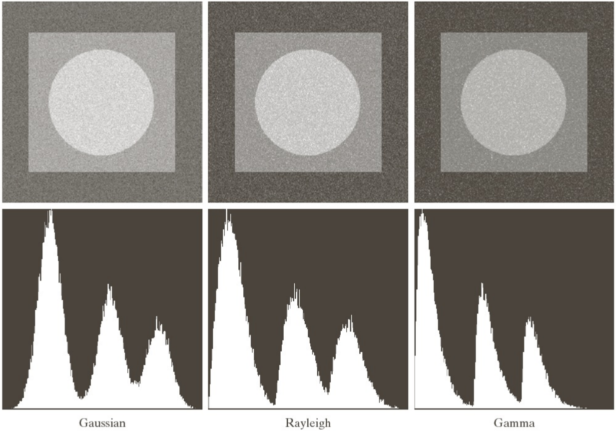

Test Pattern Corruption Example

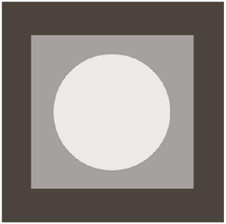

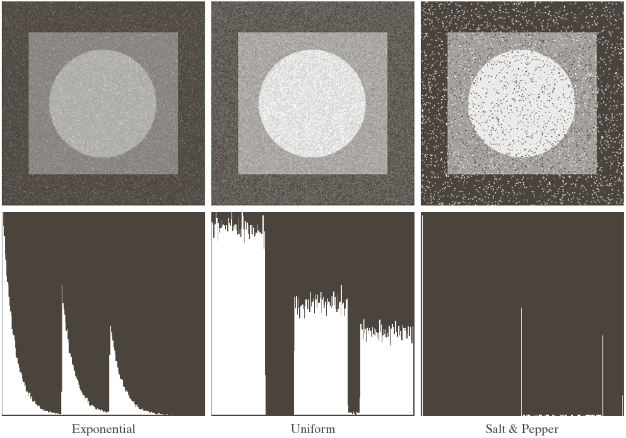

Noise Parameter Estimation

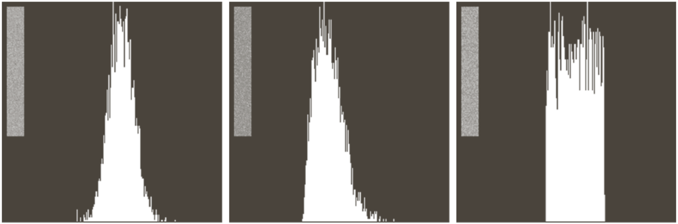

- Generate statistics from uniform region
  - Histogram matching
  - Quantile-Quantile (Q-Q) plot
- Determined defining statistics, e.g.

$$\bar{z} = \sum_{i=0}^{L-1} z_i p_S(z_i), \quad \sigma^2 = \sum_{i=0}^{L-1} (z_i - \bar{z})^2 p_S(z_i)$$

where S is a subimage (strip).

Q-Q Plot Example

- Are samples $x_1, x_2, \ldots, x_N$ and $y_1, y_2, \ldots, y_N$ governed by the same distribution?
  - Order the samples: $x_{(1)}, x_{(2)}, \ldots, x_{(N)}$ and $y_{(1)}, y_{(2)}, \ldots, y_{(N)}$
  - Plot $(x_{(1)}, y_{(1)}), (x_{(2)}, y_{(2)}), , (x_{(N)}, y_{(N)})$
  - Samples governed by the same distribution will lie (approximately) along a line

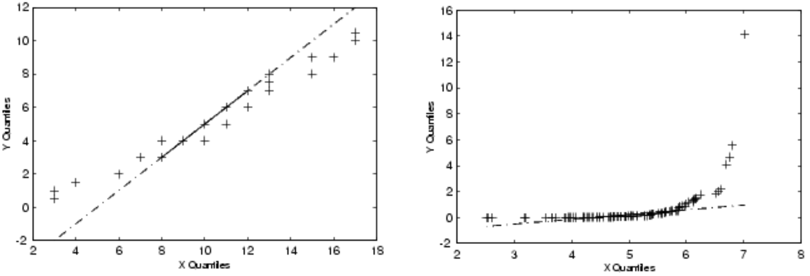

Noise-only degradation - Spatial filtering

$$g(x,y) = f(x,y) + \eta(x,y); \qquad G(u,v) = F(u,v) + N(u,v)$$

We are looking for the best compromise between noise attenuation and detail preservation.
Let S_{xy} be an $m \times n$ neighborhood of (x,y):

Arithmetic mean

$$\hat{f} = \frac{1}{mn} \sum_{(s,t) \in S_{xy}} g(s,t)$$

Geometric mean

$$\hat{f} = \left[\prod_{(s,t)] \in S_{xy}} g(s,t) \right]^{\frac{1}{mn}}$$

Contraharmonic mean

$$\hat{f} = \frac{\sum_{(s,t) \in S_{xy}} g(s,t)^{Q+1}}{\sum_{(s,t) \in S_{xy}} g(s,t)^{Q}}$$

$Q = 0 \rightarrow$ arithmetic mean
$Q > 0 \rightarrow$ good for dark impulse noise
$Q < 0 \rightarrow$ good for light impulse noise

Mean in Geometric Mean Example

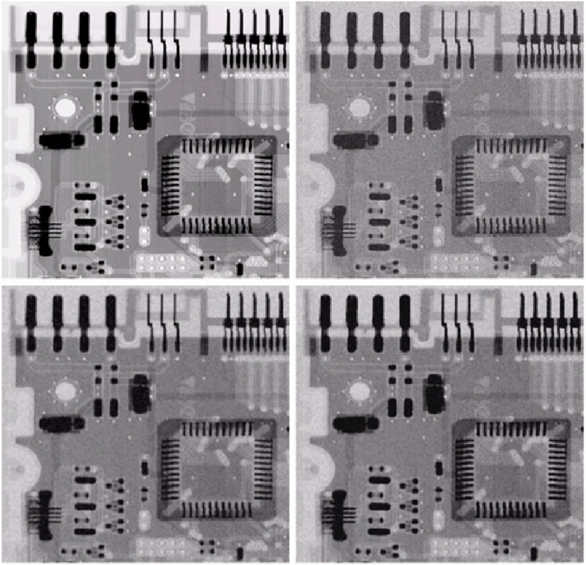

FIGURE 5.7 (a) X-ray image. (b) Image corrupted by additive Gaussian noise. (c) Result of filtering with an arithmetic mean filter of size 3×3. (d) Result of filtering with a geometric mean filter of the same size.

Multiple Applications of the Median

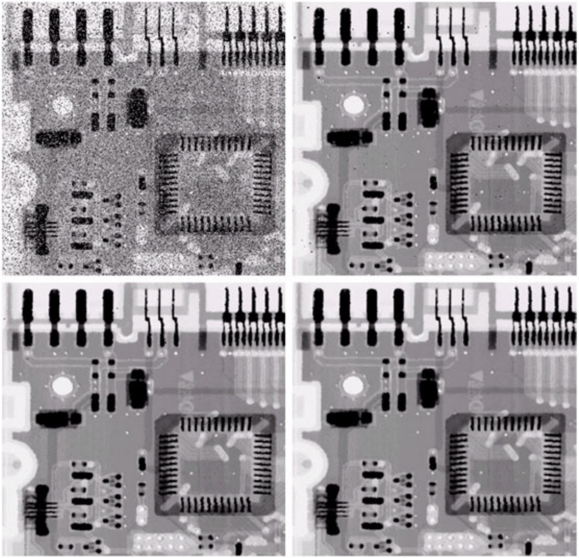

a b
c d

FIGURE 5.10 (a) Image corrupted by salt-and-pepper noise with probabilities $P_a = P_b = 0.1$. (b) Result of one pass with a median filter of size 3×3. (c) Result of processing (b) with this filter. (d) Result of processing (c) with the same filter.

Mixed Noise Example

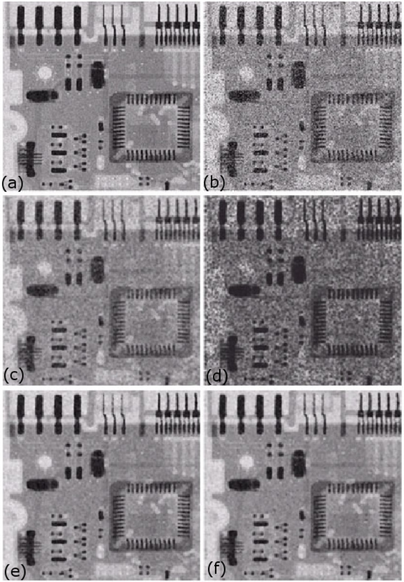

- Image corrupted by
  - Uniform noise
  - Uniform and salt and pepper noise
- Filtering methods
  - Arithmetic mean
  - Geometric mean
  - Median
  - Alpha-trimmed mean

Mean, Arithmetic Mean, and Adaptive Approach

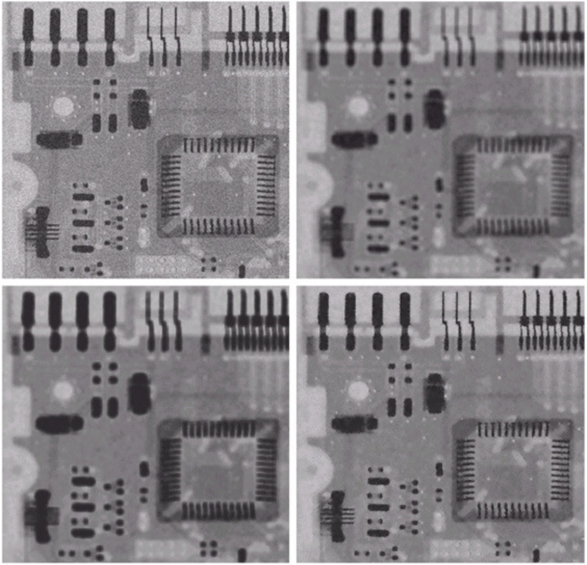

FIGURE 5.13 (a) Image corrupted by additive Gaussian noise of zero mean and variance 1000. (b) Result of arithmetic mean filtering. (c) Result of geometric mean filtering. (d) Result of adaptive noise reduction filtering. All filters were of size 7×7.

Median and Adaptive Median

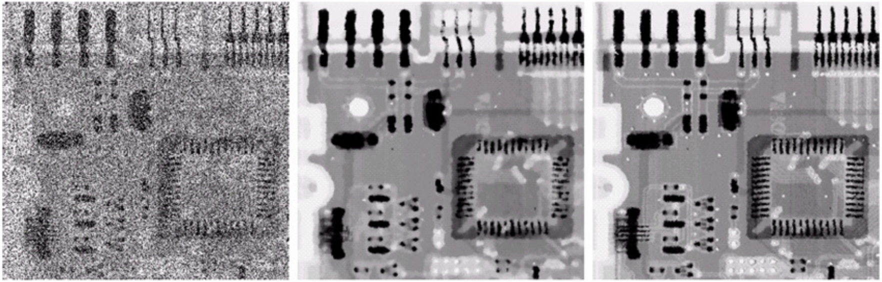

a b c

FIGURE 5.14 (a) Image corrupted by salt-and-pepper noise with probabilities $P_a = P_b = 0.25$. (b) Result of filtering with a 7×7 median filter. (c) Result of adaptive median filtering with $S_{max} = 7$.

Frequency Domain Filtering - Periodic Interference

FIGURE 5.15 From left to right, perspective plots of ideal, Butterworth (of order 1), and Gaussian bandreject filters.

- Band reject filters
  - Straightforward extension from a high/low pass case
    - Ideal, Butterworth, Gaussian

Sinusoidal Corruption Example

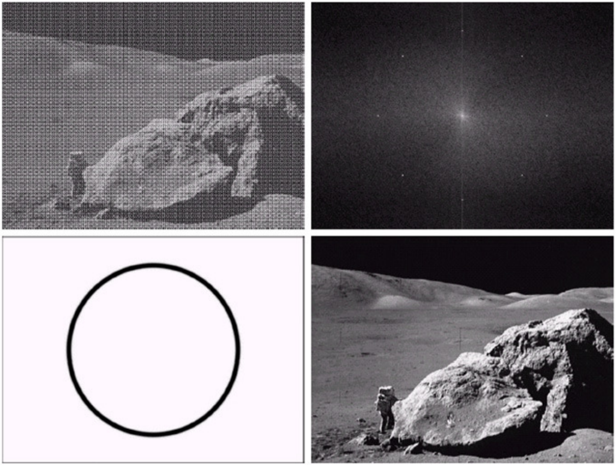

FIGURE 5.16 (a) Image corrupted by sinusoidal noise. (b) Spectrum of (a). (c) Butterworth bandreject filter (white represents 1). (d) Result of filtering. (Original image courtesy of NASA.)

NASA Image, Spectrum, and Filtering Result

a b

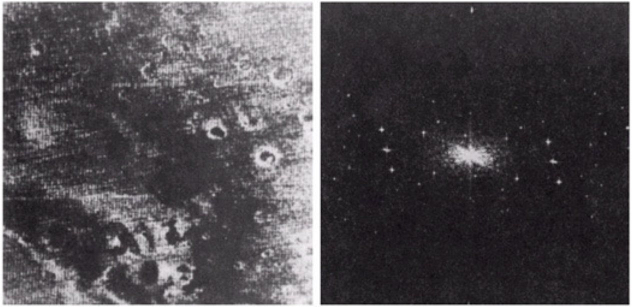

FIGURE 5.20 (a) Image of the Martian terrain taken by ***Mariner 6.*** (b) Fourier spectrum showing periodic interference. (Courtesy of NASA.)

Image Enhancement and Restoration

The image acquisition process can be modeled by

$$g(x,y) = \int_0^T \int \int_{-\infty}^{\infty} h(x,y,x',y',t) f(x',y',t) dx' dy' dt + n(x,y)$$

where T is the exposure time, $n(x,y)$ is some additive noise, and $h(x,y,x',y',t)$ characterizes the distortion introduced by the imaging system.

- limited aperture
- out of focus
- atmospheric turbulence
- relative motion
- additive noise (shot noise)

Image enhancement typically refers to noise removal or point processing. Image Restoration exploits a-priori knowledge of image degradation.

If the imaging system is ideal, spatial and time invariant, and noise-free, i.e.,

$$h(x,y,x',y',t) = \delta(x-x',y-y')$$

then

$$g(x,y) = \int_0^T f(x,y,t)dt$$

If the signal is also time invariant, i.e., $f(x,y,t) = f(x,y)$, then

$$g(x,y) = T\,f(x,y)$$

If planar motion exists in the x-y plane, $\{x_d(t), y_d(t)\}$, this yields

$$g(x,y) = \int_0^T f(x,y,t)dt = \int_0^T f(x-x_d(t), y-y_d(t))dt$$

Assuming 1D linear motion in the x direction only:

$$x_d(t) = vt, \qquad y_d(t) = 0$$

where v is the speed of the motion.
Letting $x' = vt$, we have $dt = dx'/v$ and the integral from 0 to T with respect to t becomes integral from 0 to $L \triangleq vT$ with respect to x', then

$$\begin{aligned} g(x,y) &= \int_0^T f(x,y,t)dt = \int_0^T f(x-vt,y)dt \\ &= \frac{1}{v}\int_0^L f(x-x',y)dx' = \int_{-\infty}^{\infty} f(x-x',y)h(x')dx' \\ &= f(x,y) * h(x) \end{aligned}$$

where the function

$$h(x) \triangleq \begin{cases} 1/v & \text{if } 0 \le x \le L \\ 0 & \text{else} \end{cases}$$

can be considered as the impulse response function, or the point spread function (PSF) of the imaging system.

Restoration by Inverse Filtering

Taking the Fourier transform of the above

$$G(f_x) = F(f_x)H(f_x)$$

where G, F, and H are the spectra of g, f and h, respectively. Specifically, we have

$$H(f_x) = \int_{-\infty}^{\infty} h(x)e^{-j2\pi x f_x} dx = \int_0^L e^{-j2\pi x f_x} dx = e^{-j\pi L f_x} \frac{sin(\pi L f_x)}{\pi f_x}$$

Note that $f(x)$ can be obtained by inverse transforming $F(f_x)$

$$F(f_x) = \frac{G(f_x)}{H(f_x)}$$

however, the points of $F(f_x)$ corresponding to $H(f_x) = 0$ at $f_x = k/L$, $(k = \pm 1, \pm 2, \cdots)$ can never be restored.
Moreover, this inverse filtering method is sensitive to noise that may exist in the imaging process.

Inverse Filtering example

- Direct application of
  $$\hat{F}(u,v) = \frac{G(u,v)}{H(u,v)}$$
  - Amplifies noise
  - Possible solution: Limit cutoff frequency

a b
c d

FIGURE 5.27 Restoring Fig. 5.25(b) with Eq. (5.7-1). (a) Result of using the full filter. (b) Result with H cut off outside a radius of 40; (c) outside a radius of 70; and (d) outside a radius of 85.

Power Spectra and The Wiener Filter

Cross-Correlation and Cross Spectrum

- Cross-correlation function for discrete space(stationary) random signals

$$C_{fg}(l,m) = E[f(x+l, y+m)g^*(x,y)]$$

- Cross-correlation function for discrete space deterministic signals

$$C_{fg}(l,m) = \frac{1}{LM}\sum_{x,y} f(x+l, y+m)g^*(x,y) = F(u,v)G^*(u,v)$$

- Cross power spectrum

$$\boxed{C_{fg}(l,m) \quad \underbrace{\Leftrightarrow}_{\text{Fourier Transf.}} \quad C_{fg}(u,v)}$$

Auto-Correlation and Power Spectral Density (PSD)

- Auto-correlation function for discrete space(stationary) random signals

$$C_{ff}(l,m) = E[f(x+l,y+m)f^*(x,y)]$$

- Auto-correlation function for discrete space deterministic signals

$$C_{ff}(l,m) = \frac{1}{LM}\sum_{x,y} f(x+l,y+m)f^*(x,y) = F^*(u,v)G(u,v) = |F|^2$$

- Auto power spectrum

$$\boxed{C_{ff}(l,m) \quad \underbrace{\leftrightarrow}_{\text{Fourier Transf.}} \quad C_{ff}(u,v)}$$

Power Spectra and Linear Systems: Basic Rules

$$f(x,y) \longrightarrow \boxed{h(x,y)} \longrightarrow g(x,y)$$

$$\underbrace{C_{gg}(u,v)}_{\text{PSD of Output}} = |H(u,v)|^2 \underbrace{C_{ff}(u,v)}_{\text{PSD of Input}}$$

$$\underbrace{C_{gf}(u,v)}_{\text{Cross PSD of Output with Input}} = H(u,v)C_{ff}(u,v)$$

$$\underbrace{C_{fg}(u,v)}_{\text{Cross PSD of Input with Output}} = H^*(u,v)C_{ff}(u,v)$$

Power Spectra and Linear Systems

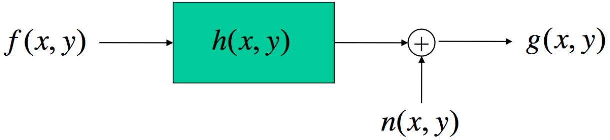

$$\boxed{\underbrace{C_{gg}(u,v)}_{\text{PSD of Output}} = |H(u,v)|^2 \underbrace{C_{ff}(u,v)}_{\text{PSD of Input}} + \underbrace{C_{nn}(u,v)}_{\text{PSD of noise}}}$$

Wiener Filter: Minimum Mean Squarred Error

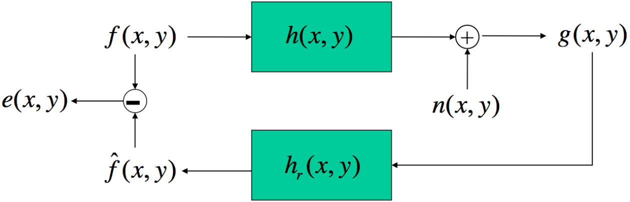

Design a restoration filter such that the reconstruction error is as small as possible

$$\min_{h_r(x,y)} E \parallel f(x,y) - \hat{f}(x,y) \parallel^2 \equiv \min_{h_r(x,y)} C_{ee}(u,v)$$

Wiener Filter: Minimum Mean squared Error

$$\min_{h_r(x,y)} E \parallel f(x,y) - \hat{f}(x,y) \parallel^2 \equiv \min_{h_r(x,y)} C_{ee}(u,v)$$

$$C_{ee}(u,v) = C_{\hat{f}\hat{f}}(u,v) + C_{ff}(u,v) - C_{f\hat{f}}(u,c) - C_{\hat{f}f}(u,v)$$

$$\begin{aligned} C_{\hat{f}\hat{f}}(u,v) &= |H_r(u,v)|^2 C_{gg}(u,v) \\ &= |H_r(u,v)|^2 \left[|H(u,v)|^2 C_{ff}(u,v) + C_{nn}(u,v) \right] \\ C_{f\hat{f}}(u,v) &= C_{gf}(u,v) H_r^*(u,v) \\ C_{\hat{f}f}(u,v) &= C_{fg}(u,v) H_r(u,v) \end{aligned}$$

Assumptions:

- The noise and the image are uncorrelated.
- The noise or the image has zero mean.
- The intensity levels in the estimate are a linear function of the levels in the degraded image.

Based on these conditions, the minimum of the error function is given in the frequency domain, minimizing the PSD $C_{ee}(u,v)$ for each frequency (u,v).

$$\frac{\partial C_{ee}}{\partial H_r} = 0 \Longrightarrow H_r(u,v) = \frac{C_{fg}(u,v)}{|H(u,v)|^2 C_{ff}(u,v) + C_{nn}(u,v)}$$

$$= \frac{C_{ff}(u,v)H^*(u,v)}{|H(u,v)|^2 C_{ff}(u,v) + C_{nn}(u,v)}$$

$$H_r(u,v) = \frac{H^*(u,v)}{|H(u,v)|^2 + \dfrac{C_{nn}(u,v)}{C_{ff}(u,v)}}$$

Wiener Filter Example

$$f(x,y) = \sqrt{2\pi}e^{-2\pi^2(x^2+y^2)}$$

$$h(x,y) = \sqrt{2\pi}\sigma e^{-}2\pi^2\sigma^2(x^2+y^2)$$

$$n(x,y)\ N(0,v^2)$$

$$H_r(u,v) = \frac{H^*(u,v)}{|H(u,v)|^2 + \frac{C_{nn}(u,v)}{C_{ff}(u,v)}}$$

$$H_r(u,v) = \frac{e^{-\frac{u^2+v^2}{2\sigma^2}}}{e^{-\frac{u^2+v^2}{\sigma^2}} + v^2 e^{u^2+v^2}}$$

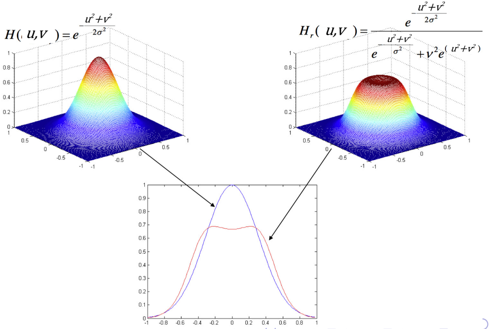

Using the fact: $H(u,v)H^*(u,v) = |H(u,v)|^2$

$$H_r(u,v) = \frac{H^*(u,v)}{|H(u,v)|^2 + \frac{C_{nn}(u,v)}{C_{ff}(u,v)}} = \left[\frac{1}{H(u,v)} \frac{|H(u,v)|^2}{|H(u,v)|^2 + \frac{C_{nn}(u,v)}{C_{ff}(u,v)}} \right]$$

The 2D Wiener filter:

$$\hat{F}(u,v) = \left[\frac{1}{H(u,v)} \frac{|H(u,v)|^2}{|H(u,v)|^2 + \frac{C_{nn}(u,v)}{C_{ff}(u,v)}} \right] G(u,v)$$

Often used approximation:

$$\hat{F}(u,v) = \left[\frac{1}{H(u,v)} \frac{|H(u,v)|^2}{|H(u,v)|^2 + K} \right] G(u,v)$$

where K is a specified constant that is added to all terms of $|H(u,v)|^2$.

Wiener Generalization

- Geometric Mean Filter generalization:

$$\hat{F}(u,v) = \left[\frac{H^*(u,v)}{|H(u,v)|^2}\right]^{\alpha} \left[\frac{H^*(u,v)}{|H(u,v)|^2 + \beta\left[\frac{C_{nn}(u,v)}{C_{ff}(u,v)}\right]}\right]^{1-\alpha} G(u,v)$$

- Special cases:
  - $\alpha = 1$: inverse filter
  - $\alpha = 0$: parametric of wiener filter (standard if $\beta = 1$)
  - $\alpha = 1/2; \beta = 1$: referred to as the Spectrum equalization filter
    - The product of two quantities raised to the same power
  - Trade-off between inverse filtering and Wiener filter controlled by α

Inverse and Wiener Filtering Comparison

a b c

FIGURE 5.28 Comparison of inverse- and Wiener filtering. (a) Result of full inverse filtering of Fig. 5.25(b). (b) Radially limited inverse filter result. (c) Wiener filter result.

- Wiener parameter K set experimentally
  - Wiener result much sharper that the band limited inverse filter
  - Full inverse results is useless

Degradation Function Estimation

- **Recall:** $G(u,v) = H(u,v)F(u,v) + N(u,v)$
- In the no (low) noise case
  - Estimation by observation:

$$H_s(u,v) = \frac{G_s(u,v)}{\hat{F}_s(u,v)}$$

  - Estimation by experimentation:
    - Input impulse of strength A

$$H(u,v) = \frac{G(u,v)}{A}$$

  - Estimation by modelling:
    - Example: atmospheric interference

$$H(u,v) = e^{-k(u^2+v^2)^{5/6}}$$

Estimation by Experiment Example

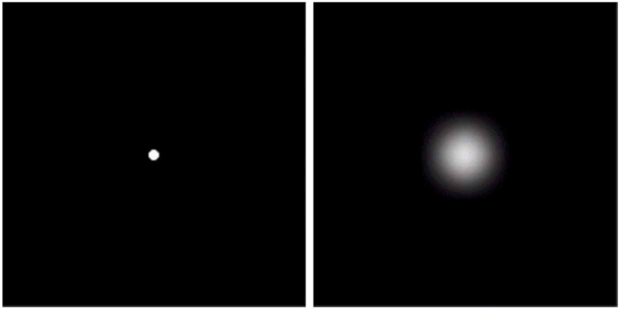

a b

FIGURE 5.24 Degradation estimation by impulse characterization. (a) An impulse of light (shown magnified). (b) Imaged (degraded) impulse.

Experimentally determinded point spread function (PSF)

Estimation by modelling: Atmospheric Turbulence Degradation

a b
c d

FIGURE 5.25 Illustration of the atmospheric turbulence model. (a) Negligible turbulence. (b) Severe turbulence, $k = 0.0025$. (c) Mild turbulence, $k = 0.001$. (d) Low turbulence, $k = 0.00025$.

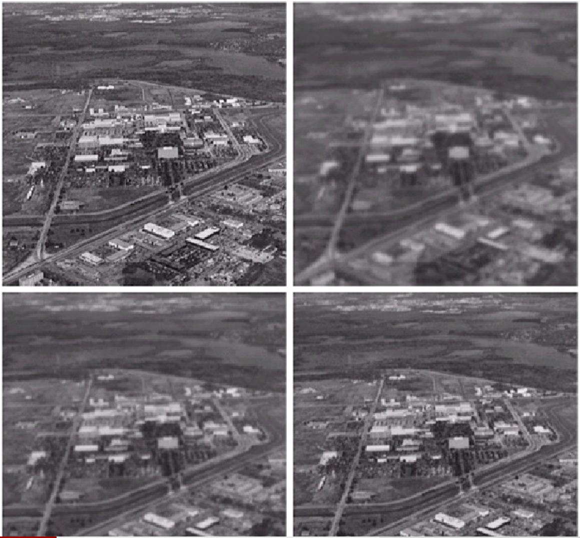

Modeling Image Motion

- Systems most often consider motion
  - Camera and/or subject motion
- If T is the exposure duration:

$$g(x,y) = \int_0^T f\left[x - x_0(t), y - y_0(t)\right] dt$$

  - Blurred image: $g(x,y)$
  - Planar motion trajectories: $x_0(t)$ and $y_0(t)$
- Fourier transform evaluation yields

$$\mathbf{G}(u,v) = \mathbf{H}(u,v)\mathbf{F}(u,v)$$

  where

$$\mathbf{H}(u,v) = \int_0^T e^{-j2\pi[ux_0(t)+vy_0(t)]} dt$$

Motion Example

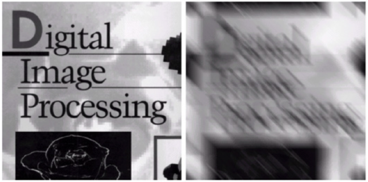

a b

FIGURE 5.26 (a) Original image. (b) Result of blurring using the function in Eq. (5.6-11) with $a = b = 0.1$ and $T = 1$.

- If $x_0(t) = at/T$ and $y_0(t) = 0$

$$\mathbf{H}(u,v) = \int_0^T e^{-j2\pi u x_0(t)} dt = \int_0^T e^{-j2\pi uat/T} dt = \frac{T}{\pi ua} \sin(\pi ua) e^{-j\pi ua}$$

- Two dimensional linear $(x_0(t) = at/T$ and $y_0(t) = bt/T)$ motion

$$\mathbf{H}(u,v) = \frac{T}{\pi(ua+vb)} \sin\left[\pi(ua+vb)\right] e^{-j\pi(ua+vb)}$$

Motion Blur with Additive Noise Example

- Left column:
  - Motion blur and additive noise corrupted observation
  - Top row has the largest noise corruption
- Center columnn:
  - Inverse filtering results
- Right column:
  - Wiener filtering result

Sparse Representation for Image Restoration*

Sparse representations for image denoising

$$\underbrace{\mathbf{y}}_{\text{measurements}} = \underbrace{\mathbf{x}_{orig}}_{\text{original image}} + \underbrace{\mathbf{w}}_{\text{white Gaussian noise}}$$

*Based on slides from Julien Mairal et. al

Energy minimization problem or MAP (Maximum a Posteriori) estimation:
Requires to obtain the lowest energy, and is known to be NP-hard

$$E(\mathbf{x}) = \underbrace{\| \mathbf{y} - \mathbf{x} \|_2^2}_{\text{relation to measurements}} + \underbrace{Pr(\mathbf{x})}_{\text{prior}}$$

Some classical priors

- Smoothness $\lambda \| \mathbf{L}\mathbf{x} \|_2^2$
- Total variation $\lambda \| \nabla \mathbf{x} \|_2^2$
- Wavelet sparsity $\lambda \| \mathbf{W}\mathbf{x} \|_1$
- ...

The MAP estimation problem can be formulated as an optimization problem:

$$\min_{\mathbf{x}} \| \mathbf{y} - \mathbf{x} \|_2^2 + \lambda \| \mathbf{x} \|_0,$$

where λ is the Lagrange multiplier.

Sparsity and redundancy

$Pr(\mathbf{x}) = \lambda \| \boldsymbol{\alpha} \|_0$ for $\mathbf{x} = \mathbf{D}\boldsymbol{\alpha}$

$$\underbrace{\begin{pmatrix} \mathbf{x} \end{pmatrix}}_{\mathbf{x}\in\mathbb{R}^N} = \underbrace{\begin{pmatrix} \mathbf{d}_1 & \mathbf{d}_2 & \cdots & \mathbf{d}_k \end{pmatrix}}_{\mathbf{D}\in\mathbb{R}^{N\times k}} \underbrace{\begin{pmatrix}\alpha_1\\ \alpha_2\\ \vdots \\ \alpha_k\end{pmatrix}}_{\boldsymbol{\alpha}\in\mathbb{R}^k,\text{sparse}}$$

Which dictionary to choose?

- Wavelets
- Curvelets
- Wedgelets
- Bandlets
- ...lets

Learned dictionaries of patches: Image restoration

Let $\mathbf{x}_0$ be a clean image and $\mathbf{y} = \mathbf{x_0} + \mathbf{w}$ its noisy version.
$\mathbf{w} \Rightarrow$ additive zero-mean white Gaussian noise, $\sigma = 0$.

Goal: Find a sparse approximation of every $\sqrt{n} \times \sqrt{n}$ overlapping patch of $\mathbf{y}$, where n is fixed a-priori.

$$\min_{\alpha_i, \mathbf{D} \in C} \sum_i \underbrace{\| \mathbf{x_i} - \mathbf{D}\alpha_\mathbf{i} \|_2^2}_{\text{reconstruction}} + \underbrace{\lambda \phi(\boldsymbol{\alpha}_i)}_{\text{sparsity}}$$

where $\boldsymbol{\alpha}_i$ is the sparsest representation of the i patch, $\mathbf{D}$ is an optimal dictionary and the index i mark the location of the patch in the image.

- $\phi(\alpha) = \| \boldsymbol{\alpha} \|_0$ ("ℓ_0 pseudo-norm")
- $\phi(\alpha) = \| \boldsymbol{\alpha} \|_1$ (ℓ_1 norm)

MOD (Method of Optimal Directions): [Engan et. al '99]

$$\{\mathbf{D}, \boldsymbol{\alpha}\} = \arg\min_{\mathbf{D}\in\mathcal{C},\boldsymbol{\alpha}} \sum_{i=1}^{P} ||\mathbf{x}_i - \mathbf{D}\boldsymbol{\alpha}_i||_2^2 + \mu_i||\boldsymbol{\alpha}_i||_0$$

Initialization of $\mathbf{D}$

ex: DCT

Sparse Coding

Fix $\mathbf{D}$ and $\forall i \in 1\cdots P$,

$$\{\boldsymbol{\alpha}_i\} \approx \arg\min_{\boldsymbol{\alpha}} ||\mathbf{x}_i - \mathbf{D}\boldsymbol{\alpha}||_2^2 + \mu_i||\boldsymbol{\alpha}||_0$$

using a Greedy approach

Dictionary Update

$$\{\mathbf{D}\} = \arg\min_{\mathbf{D}\in\mathcal{C}} \sum_i ||\mathbf{x}_i - \mathbf{D}\boldsymbol{\alpha}_i||_2^2$$

K-SVD (K mean - Singular Value Decomposition): [Elad & Aharon '06]

$$\{\mathbf{D}, \boldsymbol{\alpha}\} = \arg\min_{\mathbf{D}\in\mathcal{C},\boldsymbol{\alpha}} \sum_{i=1}^{P} ||\mathbf{x}_i - \mathbf{D}\boldsymbol{\alpha}_i||_2^2 + \mu_i||\boldsymbol{\alpha}_i||_0$$

Initialization of D → ex: DCT

Sparse Coding → Fix **D** and $\forall i \in 1 \cdots P$, $\{\boldsymbol{\alpha}_i\} \approx \arg\min_{\boldsymbol{\alpha}} ||\mathbf{x}_i - \mathbf{D}\boldsymbol{\alpha}||_2^2 + \mu_i||\boldsymbol{\alpha}||_0$ using a Greedy approach

Dictionary Update → Sequantially, $\forall j = 1 \cdots K$: Fix all $\mathbf{d}_{l\neq j}$, and minimize the reconstruction error respect to $\mathbf{d}_j$ and the non-zeros $\boldsymbol{\alpha}_i(j)$,

ℓ_1 : [Lee et al. '06]

$$\{\mathbf{D}, \boldsymbol{\alpha}\} = \arg\min_{\mathbf{D}\in\mathcal{C},\boldsymbol{\alpha}} \sum_{i=1}^{P} ||\mathbf{x}_i - \mathbf{D}\boldsymbol{\alpha}_i||_2^2 + \mu_i||\boldsymbol{\alpha}_i||_1$$

Initialization of D → ex: DCT

Sparse Coding → Fix $\mathbf{D}$ and $\forall i \in 1 \cdots P$,

$$\{\boldsymbol{\alpha}_i\} = \arg\min_{\boldsymbol{\alpha}} ||\mathbf{x}_i - \mathbf{D}\boldsymbol{\alpha}||_2^2 + \mu_i||\boldsymbol{\alpha}||_1$$

using LARS, coordinate descent,. . . ...

Dictionary Update →

$$\{\mathbf{D}\} = \arg\min_{\mathbf{D}\in\mathcal{C}} \sum_i ||\mathbf{x}_i - \mathbf{D}\boldsymbol{\alpha}_i||_2^2$$

Key ideas for denoising

- Consider each patch of size $n \times n (n = 8)$ in the image, including overlaps.
- Learn the dictionary on the corrupted image.
- The sparse Coding retrieve a sparse approximation of the noisy patches.
- Average the approximation of each patch to reconstruct the full image.

Dictionary trained on a noisy version of the boat image.

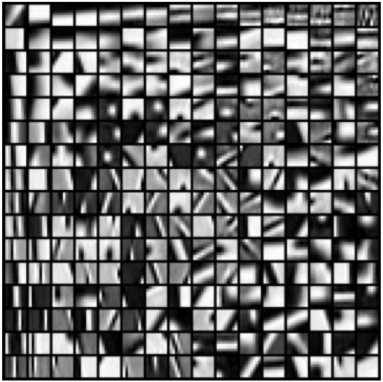

Example: Denoising of color images

- Original K-SVD (grayscale) on 3D: Dictionaries with 256 atoms and patches of size $8 \times 8 \times 3$
- K-SVD extension: Guarantee that the reconstructed patch will mantain the average color of the original one by changing the metric of the greedy algorithm (Sparse coding).

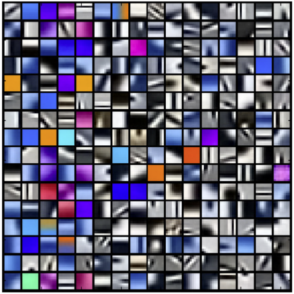

(a) $8 \times 8 \times 3$ patches

(b) Original

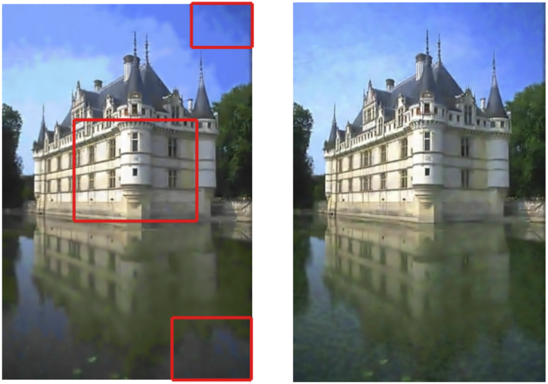

(c) Original algorithm, $\gamma = 0$, PSNR=28.78 dB

(d) Proposed algorithm, $\gamma = 5.25$, PSNR=30.04 dB

- Predominance of gray atoms: bias and color washing effect in (c)
- Correction of color artifacts with K-SVD extension in (d)

Inpainting for text removal

(Left) Original image, (center) Image with text, (left) Restored image.