

Digital Image Processing

Image Restoration Part 1

DR TANIA STATHAKI

READER (ASSOCIATE PROFESSOR) IN SIGNAL PROCESSING
IMPERIAL COLLEGE LONDON

What is Image Restoration?

Image Restoration refers to a class of methods that aim at reducing or removing various types of distortions of the image of interest. These can be:

- Distortion due to sensor noise.
- Out-of-focus camera.
- Motion blur.
- Weather conditions.
- Scratches, holes, cracks caused by aging of the image.
- Others.

Classification of restoration methods

Image restoration methods can be classified according to the type and amount of information related to the images involved in the problem and also the distortion.

- **Deterministic** or **stochastic** methods.
 - In deterministic methods, we work directly with the image values in either space or frequency domain.
 - In stochastic methods, we work with the statistical properties of the image of interest (autocorrelation function, covariance function, variance, mean etc.)
- **Non-blind** or **semi-blind** or **blind** methods.
 - In non-blind methods the degradation process is known. This is a typical so-called **Inverse Problem**.
 - In semi-blind methods the degradation process is partly-known.
 - In blind methods the degradation process is unknown.

Classification of implementation types of restoration methods

- **Direct** methods. The signals we are looking for (original undistorted" image and degradation model) are obtained through a single closed-form expression.
- **Iterative** methods. The signals we are looking for are obtained through a mathematical procedure that generates a sequence of improving approximate solutions to the problem.

Historical notes

- US and former Soviet Union space programs in 1950s and 1960s.
 - The 22 images produced during the Mariner IV flight to Mars in 1964 cost \$10M.
 - These were very valuable images which underwent various restoration techniques.
- Image Restoration is an active area of R&D.
 - Activity driven by mishaps (i.e., Humble Space Telescope - HSP).
 - New applications (i.e., SR of images and videos for Ultra High Definition - UHD displays), and new mathematical developments (i.e., sparsity).
- Notoriety in the media
 - “JFK” – Zapruder 8mm film underwent a number of restorations.
 - “No Way Out”, 1987, “Rising Sun”, 1993. These films’ whole plot largely relies on the successful restoration of some surveillance video.


A generic model of an image restoration system

- A generic common model for an image restoration system is given by the following mathematical equation:

$$g(x, y) = H[f(x, y)] + n(x, y)$$


where:

- (x, y) are the space coordinates
- $f(x, y)$ is the original (undistorted) image
- $H[\cdot]$ is a generic representation of the degradation function which is imposed onto the original image.
- $n(x, y)$ is a noise signal added to the distorted image.



A generic model of an image restoration system

- The objective of restoration is to design a system R which will operate on the observation $g(x, y)$ and give us an estimate \hat{f} of the original image.
- \hat{f} should be as close as possible to the original image f subject to an optimisation criterion.
- In the subsequent methods, we assume that the following information is available:
 - The degradation model H
 - The noise statistics
 - Some prior knowledge of f



Linear and space invariant (LSI) degradation model

In the degradation model:

$$g(x, y) = H[f(x, y)] + n(x, y)$$

we are interested in the definitions of linearity and space-invariance.

- The degradation model is linear if
$$H[k_1 f_1(x, y) + k_2 f_2(x, y)] = k_1 H[f_1(x, y)] + k_2 H[f_2(x, y)]$$
- The degradation model is space or position invariant if
$$H[f(x - x_0, y - y_0)] = g(x - x_0, y - y_0)$$
- In the above definitions we ignore the presence of external noise.
- In real life scenarios, various types of degradations can be approximated by linear, space-invariant operators.

Advantages and drawbacks of LSI assumptions

Advantages

- It is much easier to deal with linear and space-invariant models because mathematics are easier.
- The distorted image is the convolution of the original image and the distortion model.
- Software tools are available.

Drawbacks

For various realistic types of image degradations, assumptions for linearity and space-invariance are too strict and significantly deviate from the true degradation model.

Linear and space invariant (LSI) degradation model

- In the case of a LSI degradation model the output is the convolution between the input and the degradation model as follows:

$$\begin{aligned}g(x, y) &= \sum_{i=0}^{M-1} \sum_{j=0}^{N-1} f(i, j)h(x - i, y - j) + n(x, y) \\ &= f(x, y) ** h(x, y) + n(x, y)\end{aligned}$$

- Solving for $f(x, y)$, knowing the impulse response of the system $h(x, y)$ and the available data $g(x, y)$, is a deconvolution problem.
- The objective of restoration is to design a system R which will operate on the observation $g(x, y)$ and give us an estimate \hat{f} of the original image.
- \hat{f} should be as close as possible to the original image f subject to an optimisation criterion.
- In the subsequent methods, we assume that the following information is available:
 - The degradation model H
 - The noise statistics
 - Some prior knowledge of f

Motion blur: A typical type of degradation



Atmospheric turbulence: A typical type of degradation



Typical model for atmospheric turbulence

- Typical model for atmospheric turbulence: $h(x, y) = K \exp\left(-\frac{x^2 + y^2}{2\sigma^2}\right)$

negligible
distortion \Rightarrow



$\Leftarrow K = 0.0025$

$K = 0.001 \Rightarrow$



$\Leftarrow K = 0.00025$

Uniform out-of-focus blur: A typical type of degradation

- 2-D out-of-focus blur

$$h(x, y) = \begin{cases} \frac{1}{\pi R} & x^2 + y^2 \leq R^2 \\ 0 & \text{otherwise} \end{cases}$$

- Note that the model is defined within a circular disc.



Objective degradation/restoration metrics

- Blurred Signal-to-Noise Ratio (BSNR)

$$\text{BSNR} = 10 \log_{10} \left\{ \frac{\frac{1}{MN} \sum_x \sum_y [z(x, y) - \bar{z}(x, y)]^2}{\sigma_n^2}} \right\}$$

$$z(x, y) = g(x, y) - n(x, y)$$

$$\bar{z}(x, y) = E\{z(x, y)\}$$

- Improvement in Signal-to-Noise Ratio (ISNR)

$$\text{ISNR} = 10 \log_{10} \left\{ \frac{\sum_x \sum_y [f(x, y) - g(x, y)]^2}{\sum_x \sum_y [f(x, y) - \hat{f}(x, y)]^2} \right\}$$

Background: Linear convolution for one-dimensional signals

- In this section we need the concept of linear convolution in order to move forward to the concept of circular convolution.
- Linear convolution between a signal $x(n)$ of duration N samples and a signal $h[n]$ of duration L samples is a signal $y(n)$ of duration $n + L - 1$ samples defined as:

$$y(n) = \sum_k x(k)h(n - k)$$

$$x(n): [0, \dots, N - 1], h(n): [0, \dots, L - 1], y(n): [0, \dots, N + L - 2]$$

- The above relationship can be put in a vector-matrix form as:

$$\mathbf{y} = \mathbf{H}\mathbf{x} \Rightarrow \begin{bmatrix} y(0) \\ y(1) \\ \vdots \\ y(N + L - 2) \end{bmatrix} = \mathbf{H} \begin{bmatrix} x(0) \\ x(1) \\ \vdots \\ x(N - 1) \end{bmatrix}$$

\mathbf{y} is of dimension $(N + L - 1) \times 1$

\mathbf{H} is of dimension $(N + L - 1) \times N$

\mathbf{x} is of dimension $N \times 1$

Background:

Linear convolution for one-dimensional signals in matrix form Toeplitz matrices

- If we expand the linear convolution, we see that the matrix \mathbf{H} has the form shown below:

$$\mathbf{y} = \mathbf{H}\mathbf{x} \Rightarrow \mathbf{y} = \begin{bmatrix} h(0) & 0 & \dots & 0 & 0 \\ h(1) & h(0) & \dots & 0 & 0 \\ h(2) & h(1) & \dots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ h(L-2) & h(L-3) & \dots & 0 & 0 \\ h(L-1) & h(L-2) & \dots & 0 & 0 \\ 0 & h(L-1) & \dots & 0 & 0 \\ \vdots & \vdots & \dots & \vdots & \vdots \\ 0 & 0 & \dots & h(L-1) & h(L-2) \\ 0 & 0 & \dots & 0 & h(L-1) \end{bmatrix} \begin{bmatrix} x(0) \\ x(1) \\ \vdots \\ x(N-1) \end{bmatrix}$$

- Observe that the elements of the descending off-diagonals are identical.
- In linear algebra, a **Toeplitz matrix** or **diagonal-constant matrix**, is a matrix in which each descending diagonal from left to right has identical elements.

Background: Circular convolution for one-dimensional signals

- We form now the **extended** versions of $f(n)$ and $h(n)$, both of size $M \geq N + L - 1$ by zero padding. These can be denoted as $f_e(n)$ and $h_e(n)$.
- Furthermore, we assume that both $f(n)$ and $h(n)$ are periodic with period M .
- Since both $f_e(n)$ and $h_e(n)$ are periodic, their convolution is also periodic with period M .
- Is given by the expression

$$y_e(n) = \sum_{k=0}^{M-1} f_e(k)h_e(n-k) + n_e(n) = f_e(n) \circledast h_e(n)$$

Background: Circular convolution for one-dimensional signals

Using matrix notation we can write the following form

$$\mathbf{y} = \mathbf{H}\mathbf{f}$$

$$\mathbf{f} = \begin{bmatrix} f_e(0) \\ f_e(1) \\ \vdots \\ f_e(M-1) \end{bmatrix}, \mathbf{y} = \begin{bmatrix} f_e(0) \\ f_e(1) \\ \vdots \\ f_e(M-1) \end{bmatrix}$$

$$\mathbf{H}_{(M \times M)} = \begin{bmatrix} h_e(0) & h_e(-1) & \dots & h_e(-M+1) \\ h_e(1) & h_e(0) & \dots & h_e(-M+2) \\ \vdots & \vdots & \ddots & \vdots \\ h_e(M-1) & h_e(M-2) & \dots & h_e(0) \end{bmatrix}$$

- At the moment we have decided to ignore any external noise.
- Note that in most real life scenarios, the size of $h(n)$ is much smaller than the size of $f(n)$. Therefore, in the periodic extension of $h(n)$, most samples will be zero.
- Based on the above comments, \mathbf{H} will be a so called **sparse** matrix.

Circular convolution for one-dimensional signals

- Because $h_e(n)$ is periodic with period M we have that

$$\begin{aligned} \mathbf{H}_{(M \times M)} &= \begin{bmatrix} h_e(0) & h_e(-1) & \dots & h_e(-M+1) \\ h_e(1) & h_e(0) & \dots & h_e(-M+2) \\ \vdots & \vdots & \ddots & \vdots \\ h_e(M-1) & h_e(M-2) & \dots & h_e(0) \end{bmatrix} \\ &= \begin{bmatrix} h_e(0) & h_e(M-1) & \dots & h_e(1) \\ h_e(1) & h_e(0) & \dots & h_e(2) \\ \vdots & \vdots & \ddots & \vdots \\ h_e(M-1) & h_e(M-2) & \dots & h_e(0) \end{bmatrix} \end{aligned}$$

- Observe the special structure of \mathbf{H} : each row is formed by the row above if we push it one position to the right. The element that is thrown out of the matrix in the row above becomes the most left element in the row under consideration.
- This is a so called **circulant** matrix.

Eigenvalues of circulant matrices

We define $\lambda(k)$ to be

$$\lambda(k) = h_e(0) + h_e(M-1) \exp(j \frac{2\pi}{M} k) + \dots + h_e(1) \exp[j \frac{2\pi}{M} (M-1)k], \quad k = 0, 1, \dots, M-1$$

Since

- $\exp[j \frac{2\pi}{M} (M-i)k] = \exp(-j \frac{2\pi}{M} ik)$ and
- $\exp[j \frac{2\pi}{M} (-M+i)k] = \exp(j \frac{2\pi}{M} ik)$

we have that

$$\begin{aligned} \lambda(k) &= h_e(0) + h_e(M-1) \exp[-j \frac{2\pi}{M} (M-1)k] + \dots + h_e(1) \exp(j \frac{2\pi}{M} k) \\ &= M \left[\frac{1}{M} \sum_{n=0}^{M-1} h_e(n) \exp[-j \frac{2\pi}{M} nk] \right], \quad k = 0, 1, \dots, M-1 \end{aligned}$$

$$\lambda(k) = MH(k)$$

$H(k)$ is the discrete Fourier transform of $h_e(i)$.

Eigenvalues and eigenvectors of circulant matrices

I define $\mathbf{w}(k)$ to be

$$\mathbf{w}(k) = \begin{bmatrix} 1 \\ \exp(j \frac{2\pi}{M} k) \\ \vdots \\ \exp[j \frac{2\pi}{M} (M - 1)k] \end{bmatrix}$$

It can be seen that

$$\mathbf{H}\mathbf{w}(k) = \lambda(k)\mathbf{w}(k)$$

This implies that $\lambda(k)$ is an eigenvalue of the matrix \mathbf{H} and $\mathbf{w}(k)$ is its corresponding eigenvector.

We form a matrix \mathbf{W} whose columns are the eigenvectors of the matrix \mathbf{H} , that is to say:

$$\mathbf{W} = [\mathbf{w}(0) \quad \mathbf{w}(1) \quad \dots \quad \mathbf{w}(M - 1)]$$

Eigenvalues and eigenvectors of circulant matrices

$$\mathbf{W} = [\mathbf{w}(0) \quad \mathbf{w}(1) \quad \dots \quad \mathbf{w}(M-1)]$$

$$w(k, i) = \exp(j \frac{2\pi}{M} ki) \text{ and } w^{-1}(k, i) = \frac{1}{M} \exp(-j \frac{2\pi}{M} ki)$$

We can then diagonalize the matrix \mathbf{H} as follows

$$\mathbf{H} = \mathbf{W}\mathbf{D}\mathbf{W}^{-1} \Rightarrow \mathbf{D} = \mathbf{W}^{-1}\mathbf{H}\mathbf{W}$$

where

$$\mathbf{D} = \begin{bmatrix} \lambda(0) & 0 & \dots & 0 \\ 0 & \lambda(1) & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \lambda(M-1) \end{bmatrix}$$

Obviously \mathbf{D} is a diagonal matrix and

$$D(k, k) = \lambda(k) = MH(k)$$

Circular convolution and DFT

If we go back to the degradation model we can write

$$\mathbf{y} = \mathbf{H}\mathbf{f} \Rightarrow \mathbf{y} = \mathbf{W}\mathbf{D}\mathbf{W}^{-1}\mathbf{f} \Rightarrow \mathbf{W}^{-1}\mathbf{y} = \mathbf{D}\mathbf{W}^{-1}\mathbf{f} \Rightarrow$$
$$Y(k) = MH(k)F(k), k = 0, 1, \dots, M - 1$$

$Y(k), H(k), F(k), k = 0, 1, \dots, M - 1$ are the M –sample Discrete Fourier Transforms of $y(i), h(i), f(i)$, respectively.

So by choosing $\lambda(k)$ and $w(k)$ as above and assuming that $h_e(i)$ is periodic, we start with a matrix problem and end up with M scalar problems.

Circular convolution in two-dimensional signals (images)

- Suppose we have a two-dimensional discrete signal $f(i, j)$ of size $A \times B$ samples which is due to a degradation process.
- The degradation can now be modeled by a two dimensional discrete impulse response $h(i, j)$ of size $C \times D$ samples.
- We form the extended versions of $f(i, j)$ and $h(i, j)$, both of size $M \times N$, where $M \geq A + C - 1$ and $N \geq B + D - 1$, and periodic with period $M \times N$. These can be denoted as $f_e(i, j)$ and $h_e(i, j)$.
- From this point forward we will assumed that additive noise $n_e(i, j)$ is present in the model.
- For a linear, space-invariant degradation process we obtain

$$y_e(i, j) = \sum_{m=0}^{M-1} \sum_{n=0}^{N-1} f_e(m, n) h_e(i - m, j - n) + n_e(i, j)$$

Lexicographic ordering

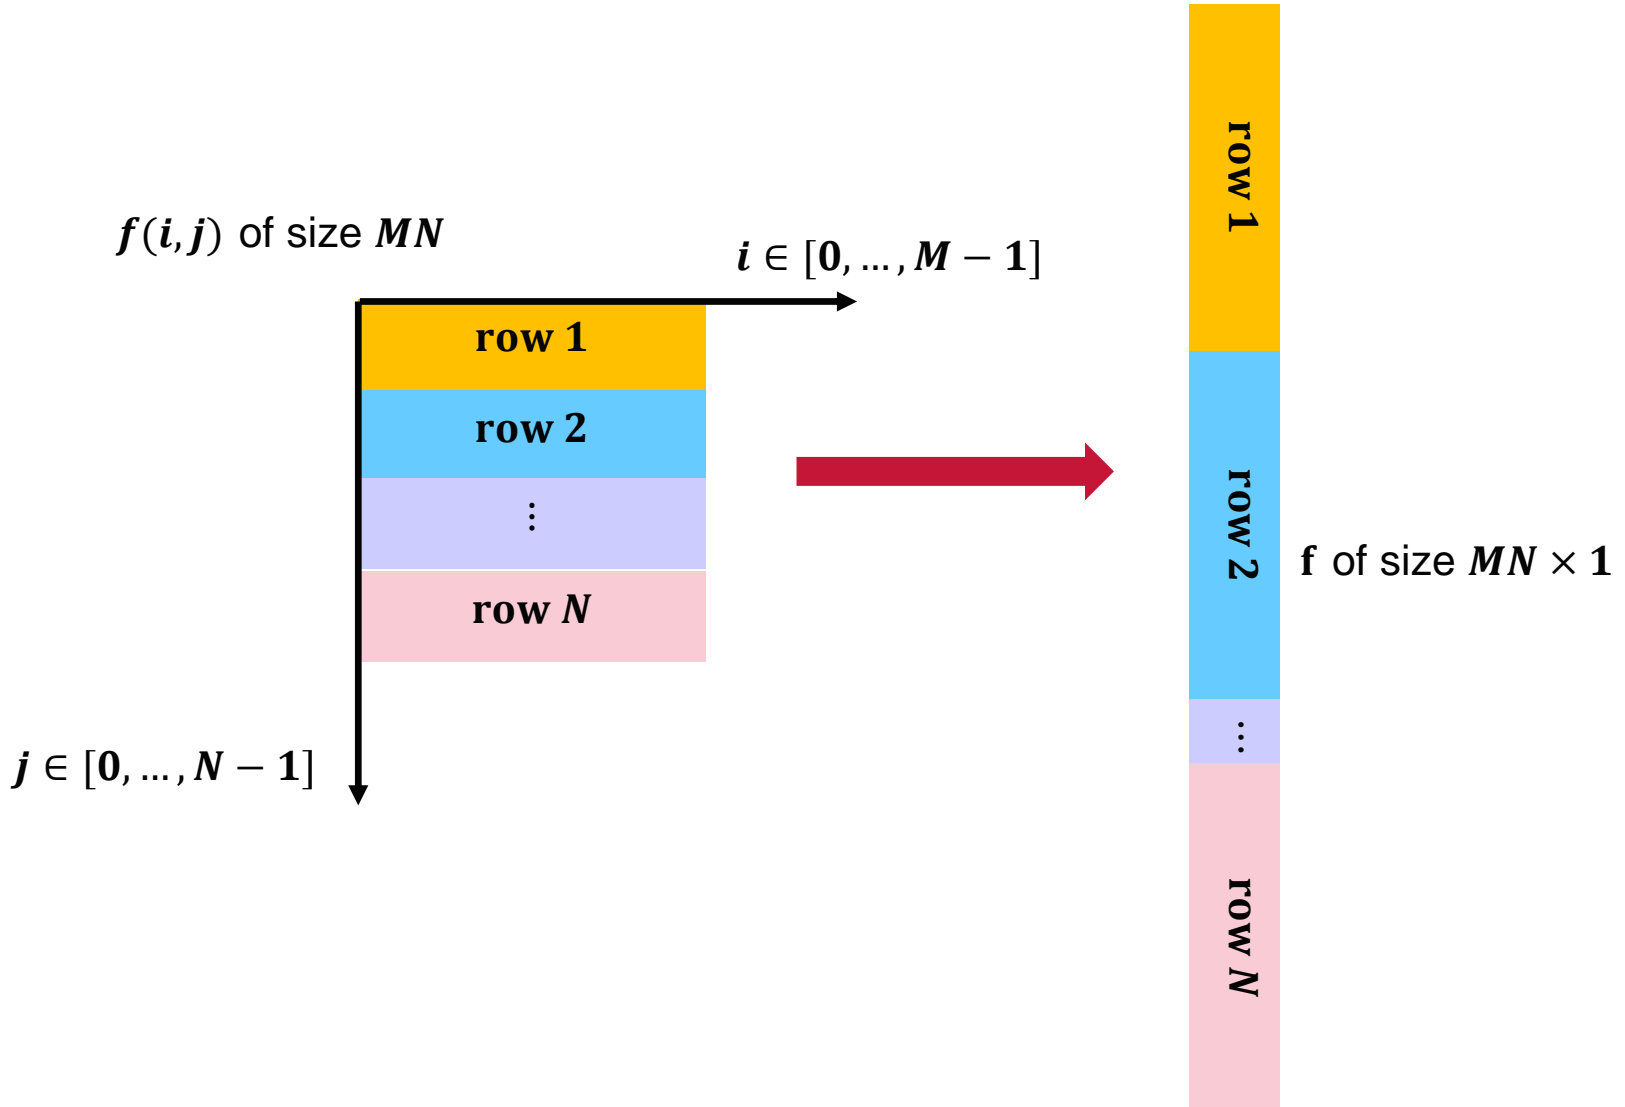
- In this section we use a mean for representing all the pixels in image as a vector.
- This can be done by stacking the rows of the image on top of each other.
 - **Please observe next slide.**
- This introduces an alternative way for representing the degradation equation in the restoration problem; it becomes a vector-matrix equation.
- The input-output relationship is written as:

$$\mathbf{y} = \mathbf{H}\mathbf{f} + \mathbf{n}$$

where f and y are MN –dimensional column vectors that represent the lexicographic ordering of images $f_e(i, j)$ and $h_e(i, j)$ respectively.

- If the degradation is LSI then this matrix has a specific form; it is a so-called **block-circulant** matrix.
- It is straightforward to describe this matrix in the spectral domain, by finding its eigenvalues and eigenvectors.

Lexicographic ordering of an image



Block-circulant matrices

- As mentioned, in two dimensions the input-output relationship is written as:

$$\mathbf{y} = \mathbf{H}\mathbf{f} + \mathbf{n}$$

where the matrix \mathbf{H} is a **block-circulant** matrix.

$$\mathbf{H} = \begin{bmatrix} \mathbf{H}_0 & \mathbf{H}_{M-1} & \dots & \mathbf{H}_1 \\ \mathbf{H}_1 & \mathbf{H}_0 & \dots & \mathbf{H}_2 \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{H}_{M-1} & \mathbf{H}_{M-2} & \dots & \mathbf{H}_0 \end{bmatrix}$$

$$\mathbf{H}_j = \begin{bmatrix} h_e(j, 0) & h_e(j, N-1) & \dots & h_e(j, 1) \\ h_e(j, 1) & h_e(j, 0) & \dots & h_e(j, 2) \\ \vdots & \vdots & \ddots & \vdots \\ h_e(j, N-1) & h_e(j, N-2) & \dots & h_e(j, 0) \end{bmatrix}$$

- The analysis of the diagonalisation of \mathbf{H} is a straightforward extension of the one-dimensional case.
- In that case we end up with the following set of $M \times N$ scalar problems.


$$Y(u, v) = MNH(u, v)F(u, v)(+N(u, v))$$

$$u = 0, 1, \dots, M-1, v = 0, 1, \dots, N-1$$


Periodic extension of images and degradation model

- In image restoration we often work with Discrete Fourier Transforms.
- DFT assumes periodicity of the signal in time or space.
- Therefore, periodic extension of signals is required.
- Distorted image is the convolution of the original image and the distortion model. We are able to assume this because of the linearity and space invariance assumptions.
- Convolution increases the size of signals.
- Periodic extension must take into consideration the presence of convolution: zero-padding is required.
- Every signal involved in an image restoration system must be extended by zero-padding and also treated as virtually periodic.

Wrong periodic extension of signals. Red and green signal are convolved




Correct periodic extension of signals Red and green signal are convolved



Correct periodic extension of images and degradation model

- The original image $f(x, y)$ is of size $A \times B$.
- The degradation model $h(x, y)$ is of size $C \times D$.
- We form the extended versions of $f(x, y)$ and $h(x, y)$ by zero padding, both of size $M \times N$.
 - $M \geq A + C - 1$
 - $N \geq B + D - 1$
- Example
 - Image 256×256
 - Degradation 3×3
 - With extension by zero padding both images have dimension at least $(256 + 3 - 1) \times (256 + 3 - 1) = 258 \times 256$.
 - They are also assume to be periodic.

Correct periodic extension of images and degradation model



Inverse filtering for image restoration

- Inverse filtering is a deterministic and direct method for image restoration.
- The images involved must be lexicographically ordered. That means that an image is converted to a column vector by stacking the rows one by one after converting them to columns.
- Therefore, an image of size $M \times N = 256 \times 256$ is converted to a column vector of size $(256 \times 256) \times 1 = 65536 \times 1$.
- The degradation model is written in a matrix form as
$$\mathbf{y} = \mathbf{Hf}$$
where the images are vectors and the degradation process is a huge but sparse matrix of size $MN \times MN$.
- The above relationship is ideal. The true degradation model is $\mathbf{y} = \mathbf{Hf} + \mathbf{n}$ where \mathbf{n} is a lexicographically ordered two dimensional noisy signal which corrupts the distorted image $y(x, y)$.

Inverse Filtering for image restoration

- We formulate an unconstrained optimisation problem as follows:

$$\text{minimise } J(\mathbf{f}) = \|\mathbf{n}(\mathbf{f})\|^2 = \|\mathbf{y} - \mathbf{H}\mathbf{f}\|^2$$

$$\begin{aligned} \|\mathbf{y} - \mathbf{H}\mathbf{f}\|^2 &= (\mathbf{y} - \mathbf{H}\mathbf{f})^T (\mathbf{y} - \mathbf{H}\mathbf{f}) = [\mathbf{y}^T - (\mathbf{H}\mathbf{f})^T] (\mathbf{y} - \mathbf{H}\mathbf{f}) \\ &= (\mathbf{y}^T - \mathbf{f}^T \mathbf{H}^T) (\mathbf{y} - \mathbf{H}\mathbf{f}) = \mathbf{y}^T \mathbf{y} - \mathbf{y}^T \mathbf{H}\mathbf{f} - \mathbf{f}^T \mathbf{H}^T \mathbf{y} + \mathbf{f}^T \mathbf{H}^T \mathbf{H}\mathbf{f} \end{aligned}$$

- We set the first derivative of $J(\mathbf{f})$ equal to 0.

$$\frac{\partial J(\mathbf{f})}{\partial \mathbf{f}} = \mathbf{0} \Rightarrow \frac{\partial \mathbf{y}^T \mathbf{y}}{\partial \mathbf{f}} - \frac{\partial \mathbf{y}^T \mathbf{H}\mathbf{f}}{\partial \mathbf{f}} - \frac{\partial \mathbf{f}^T \mathbf{H}^T \mathbf{y}}{\partial \mathbf{f}} + \frac{\partial \mathbf{f}^T \mathbf{H}^T \mathbf{H}\mathbf{f}}{\partial \mathbf{f}} = \mathbf{0}$$

Note that:

- $\frac{\partial(\cdot)}{\partial \mathbf{f}}$ indicates a vector of partial derivatives
- $\frac{\partial \mathbf{y}^T \mathbf{y}}{\partial \mathbf{f}} = \mathbf{0}$
- $\frac{\partial \mathbf{y}^T \mathbf{H}\mathbf{f}}{\partial \mathbf{f}} = \frac{\partial \mathbf{f}^T \mathbf{H}^T \mathbf{y}}{\partial \mathbf{f}}$

- Therefore,

$$\frac{\partial J(\mathbf{f})}{\partial \mathbf{f}} = \mathbf{0} \Rightarrow -2 \frac{\partial \mathbf{f}^T \mathbf{H}^T \mathbf{y}}{\partial \mathbf{f}} + \frac{\partial \mathbf{f}^T \mathbf{H}^T \mathbf{H}\mathbf{f}}{\partial \mathbf{f}} = \mathbf{0} \Rightarrow -2\mathbf{H}^T \mathbf{y} + 2\mathbf{H}^T \mathbf{H}\mathbf{f} = \mathbf{0} \Rightarrow$$

$$\mathbf{H}^T \mathbf{H}\mathbf{f} = \mathbf{H}^T \mathbf{y}$$

- If the matrix $\mathbf{H}^T \mathbf{H}$ is invertible then $\mathbf{f} = (\mathbf{H}^T \mathbf{H})^{-1} \mathbf{H}^T \mathbf{y}$
- If \mathbf{H} is square and invertible then $\mathbf{f} = \mathbf{H}^{-1} (\mathbf{H}^T)^{-1} \mathbf{H}^T \mathbf{y} = \mathbf{H}^{-1} \mathbf{y}$

Inverse Filtering for image restoration

- According to the previous analysis if \mathbf{H} (and therefore \mathbf{H}^{-1}) is block circulant the above problem can be solved as a set of $M \times N$ scalar problems as follows.

$$F(u, v) = \frac{H^*(u, v)Y(u, v)}{|H(u, v)|^2} \Rightarrow f(i, j) = \mathfrak{F}^{-1} \left[\frac{H^*(u, v)Y(u, v)}{|H(u, v)|^2} \right] = \frac{Y(u, v)}{H(u, v)}$$

- Diagonalisation of all matrices involved is necessary.

Computational issues concerning inverse filtering

Noise free case

- Suppose first that the additive noise $n(i, j)$ is negligible. A problem arises if $H(u, v)$ becomes very small or zero for some point (u, v) or for a whole region in the (u, v) plane. In that region inverse filtering cannot be applied.
- Note that in most real applications $H(u, v)$ drops off rapidly as a function of distance from the origin !

Solution: if these points are known they can be neglected in the computation of $F(u, v)$.

Computational issues concerning inverse filtering

Noisy case

- In the presence of external noise we have that

$$\hat{F}(u, v) = \frac{H^*(u, v)(Y(u, v) - N(u, v))}{|H(u, v)|^2} =$$
$$\frac{H^*(u, v)Y(u, v)}{|H(u, v)|^2} - \frac{H^*(u, v)N(u, v)}{|H(u, v)|^2} \Rightarrow$$
$$\hat{F}(u, v) = F(u, v) - \frac{N(u, v)}{H(u, v)}$$

- If $H(u, v)$ becomes very small, the term $N(u, v)$ dominates the result.

Pseudoinverse Filtering

- To cope with noise amplification we carry out the restoration process in a limited neighborhood about the origin where $H(u, v)$ is not very small.
- This procedure is called **pseudoinverse or generalized inverse filtering**.
- In that case we set

$$\hat{F}(u, v) = \begin{cases} \frac{H^*(u, v)Y(u, v)}{|H(u, v)|^2} & H(u, v) \neq 0 \\ 0 & H(u, v) = 0 \end{cases}$$

or

$$\hat{F}(u, v) = \begin{cases} \frac{H^*(u, v)Y(u, v)}{|H(u, v)|^2} = \frac{Y(u, v)}{H(u, v)} & |H(u, v)| \geq \varepsilon \\ 0 & \text{otherwise} \end{cases}$$

Pseudoinverse restoration examples



Figure 3: Degraded by a 7×7 pill-box blur, 20 dB BSNR.



Figure 5: Degraded by a 5×5 Gaussian blur ($\sigma^2 = 1$), 20 dB BSNR.




Figure 11: Result of Figure 3 restored by a generalized inverse filter with a threshold of 10^{-3} , ISNR = -32.9 dB




Figure 17: Result of Figure 5 restored by a generalized inverse filter with a threshold of 10^{-3} , ISNR = -36.6 dB

Pseudoinverse restoration examples



Figure 3: Degraded by a 7×7 pill-box blur, 20 dB BSNR



Figure 5: Degraded by a 5×5 Gaussian blur ($\sigma^2 = 1$), 20 dB BSNR



Figure 13: Result of Figure 3 restored by a generalized inverse filter with a threshold of 10^{-1} , ISNR = 0.61 dB



Figure 19: Result of Figure 5 restored by a generalized inverse filter with a threshold of 10^{-1} , ISNR = -1.8 dB

Constrained Least Squares (CLS) Restoration

- By introducing a so called **Lagrange multiplier** or **regularisation parameter** α , we transform the constrained optimisation problem to an unconstrained one as follows.

- The problem

$$\begin{aligned} & \underset{\mathbf{f}}{\text{minimise}} J(\mathbf{f}) = \|\mathbf{y} - \mathbf{Hf}\|^2 \\ & \text{subject to } \|\mathbf{Cf}\|^2 < \varepsilon \end{aligned}$$

is equivalent to

$$\underset{\mathbf{f}}{\text{minimise}} J(\mathbf{f}) = \|\mathbf{y} - \mathbf{Hf}\|^2 + \alpha\|\mathbf{Cf}\|^2$$

- The imposed constraint implies that the energy of the restored image at high frequencies is below a threshold.
- It is basically a smoothness constraint.
C a high pass filter operator
Cf a high pass filtered version of the image

Constrained Least Squares (CLS) Restoration

- We formulate an unconstrained optimisation problem as follows:

$$\text{minimise } J(\mathbf{f}) = \|\mathbf{y} - \mathbf{H}\mathbf{f}\|^2 + \alpha \|\mathbf{C}\mathbf{f}\|^2$$

$$\begin{aligned} \|\mathbf{y} - \mathbf{H}\mathbf{f}\|^2 + \alpha \|\mathbf{C}\mathbf{f}\|^2 &= (\mathbf{y} - \mathbf{H}\mathbf{f})^T (\mathbf{y} - \mathbf{H}\mathbf{f}) + \alpha (\mathbf{C}\mathbf{f})^T (\mathbf{C}\mathbf{f}) \\ &= \mathbf{y}^T \mathbf{y} - \mathbf{y}^T \mathbf{H}\mathbf{f} - \mathbf{f}^T \mathbf{H}^T \mathbf{y} + \mathbf{f}^T \mathbf{H}^T \mathbf{H}\mathbf{f} + \alpha \mathbf{f}^T \mathbf{C}^T \mathbf{C}\mathbf{f} \end{aligned}$$

- We set the first derivative of $J(\mathbf{f})$ equal to 0.

$$\frac{\partial J(\mathbf{f})}{\partial \mathbf{f}} = \mathbf{0} \Rightarrow -2\mathbf{H}^T \mathbf{y} + 2\mathbf{H}^T \mathbf{H}\mathbf{f} + 2\alpha \mathbf{C}^T \mathbf{C}\mathbf{f} = \mathbf{0}$$

- Therefore,

$$\frac{\partial J(\mathbf{f})}{\partial \mathbf{f}} = \mathbf{0} \Rightarrow -2 \frac{\partial \mathbf{f}^T \mathbf{H}^T \mathbf{y}}{\partial \mathbf{f}} + \frac{\partial \mathbf{f}^T \mathbf{H}^T \mathbf{H}\mathbf{f}}{\partial \mathbf{f}} = \mathbf{0} \Rightarrow -2\mathbf{H}^T \mathbf{y} + 2\mathbf{H}^T \mathbf{H}\mathbf{f} = \mathbf{0} \Rightarrow$$

$$(\mathbf{H}^T \mathbf{H} + \alpha \mathbf{C}^T \mathbf{C})\mathbf{f} = \mathbf{H}^T \mathbf{y}$$

- In frequency domain and under the presence of noise we have:

$$\hat{F}(u, v) = \frac{H^*(u, v)(Y(u, v) - N(u, v))}{|H(u, v)|^2 + \alpha |C(u, v)|^2}$$

Constrained Least Squares (CLS) Restoration: Observations

- In frequency domain and under the presence of noise we have:

$$\hat{F}(u, v) = \frac{H^*(u, v)(Y(u, v) - N(u, v))}{|H(u, v)|^2 + \alpha|C(u, v)|^2}$$

- The regularisation parameter α controls the contribution between the terms $\|\mathbf{y} - \mathbf{Hf}\|^2$ and $\|\mathbf{Cf}\|^2$.
- Small α implies that emphasis is given to the minimisation function $\|\mathbf{y} - \mathbf{Hf}\|^2$.
 - Note that in the extreme case where $\alpha = 0$, CLS becomes Inverse Filtering.
 - Note that with smaller values of α , the restored image tends to have more amplified noise effects.
- Large α implies that emphasis is given to the minimisation function $\|\mathbf{Cf}\|^2$. A large α should be chosen if the noise is high.
 - Note that with larger values of α , and thus more regularisation, the restored image tends to have more ringing.

Choice of α cont.

- The problem of the choice of α has been attempted in a large number of studies and different techniques have been proposed.

- One possible choice is based on a **set theoretic approach**.

- A restored image is approximated by an image which lies in the intersection of the two ellipsoids defined by

$$Q_{f|y} = \{f \mid \|y - Hf\|^2 \leq E^2\} \text{ and } Q_f = \{f \mid \|Cf\|^2 \leq \varepsilon^2\}$$

- The center of one of the ellipsoids which bounds the intersection of $Q_{f|y}$ and Q_f , is given by the equation

$$f = (H^T H + \alpha C^T C)^{-1} H^T y$$

with $\alpha = (E/\varepsilon)^2$.

- Finally, a choice for α is also:

$$\alpha = \frac{1}{\text{BSNR}}$$

Choice of α

- The variance and bias of the error image in frequency domain are

$$\text{Var}(\hat{f}(a)) = \sigma_n^2 \sum_{u=0}^{M-1} \sum_{v=0}^{N-1} \frac{|H(u, v)|^2}{(|H(u, v)|^2 + \alpha |C(u, v)|^2)^2}$$

$$\text{Bias}(\hat{f}(a)) = \sigma_n^2 \sum_{u=0}^{M-1} \sum_{v=0}^{N-1} \frac{|F(u, v)|^2 \alpha^2 |C(u, v)|^4}{(|H(u, v)|^2 + \alpha |C(u, v)|^2)^2}$$

- Note that the Mean Squared Error (MSE) $E(\alpha)$ in this problem is the expected value of the Euclidian norm of the difference between the true original image f and the estimated original image $\hat{f}(a)$, i.e., $E\{\|\hat{f}(a) - f\|^2\}$.
- It has been shown that the minimum Mean Squared Error (solid curve in next slide) is encountered close to the intersection of the above functions and is equal to

$$E\{\|\hat{f}(a) - f\|^2\} = \text{Bias}(\hat{f}(a)) + \text{Var}(\hat{f}(a))$$

- Observing the graphs for the variance and bias of the error in the next slide we can say that another good choice of α is one that gives the best compromise between the variance and bias of the error image.






Figure 3: Degraded by a 7×7 pill-box blur, 20 dB BSNR.



Figure 5: Degraded by a 5×5 Gaussian blur ($\sigma^2 = 1$), 20 dB BSNR.



Figure 26: CLS restoration of Figure 3 with $\alpha = 1$, ISNR = 2.5 dB



Figure 40: CLS restoration of Figure 5 with $\alpha = 1$, ISNR = 1.3 dB



Figure 3: Degraded by a 7×7 pill-box blur, 20 dB BSNR.



Figure 5: Degraded by a 5×5 Gaussian blur ($\sigma^2 = 1$), 20 dB BSNR.




Figure 30: CLS restoration of Figure 3 with $\alpha = 0.0001$, ISNR = -21.9 dB



e 44: CLS restoration of Figure 5 with $\alpha = 0.0001$, ISNR = -22.1 dB



Figure 3: Degraded by a 7×7 pill-box blur, 20 dB BSNR



Figure 27: Corresponding error image for Figure 26 ($|\text{original} - \text{restored}|$, scaled for display)




Figure 29: Corresponding error image for Figure 28 ($|\text{original} - \text{restored}|$, scaled for display)





Figure 31: Corresponding error image for Figure 30 ($|\text{original} - \text{restored}|$, scaled for display)

Wiener Filter Estimator (Stochastic Regularisation)

- The image restoration problem can be viewed as a system identification problem as follows:



- The objective is to minimize the expected value of the Euclidian norm of the error:

$$E\{(f - \hat{f})^T (f - \hat{f})\}$$

To do so the following conditions should hold:

- $E\{\hat{f}\} = E\{f\} \Rightarrow E\{f\} = WE\{y\}$
- The error must be orthogonal to the observation about the mean
 $E\{(\hat{f} - f)(y - E\{y\})^T\} = 0$

Wiener Filter Estimator (Stochastic Regularisation)

The following conditions should hold:

- i. $E\{\hat{f}\} = E\{f\} \Rightarrow E\{f\} = WE\{y\}$
- ii. The error must be orthogonal to the observation about the mean
 $E\{(\hat{f} - f)(y - E\{y\})^T\} = 0$

From i. and ii. we have that

$$\begin{aligned} E\{(Wy - f)(y - E\{y\})^T\} = 0 &\Rightarrow E\{(Wy + E\{f\} - WE\{y\} - f)(y - E\{y\})^T\} = 0 \\ &\Rightarrow E\{[W(y - E\{y\}) - (f - E\{f\})](y - E\{y\})^T\} = 0 \end{aligned}$$

If $\tilde{y} = y - E\{y\}$ and $\tilde{f} = f - E\{f\}$ then

$$E\{(W\tilde{y} - \tilde{f})\tilde{y}^T\} = 0 \Rightarrow E\{W\tilde{y}\tilde{y}^T\} = E\{\tilde{f}\tilde{y}^T\} \Rightarrow WE\{\tilde{y}\tilde{y}^T\} = E\{\tilde{f}\tilde{y}^T\} \Rightarrow WR_{\tilde{y}\tilde{y}}$$

Wiener Filter Estimator (Stochastic Regularisation)

- If the original and the degraded image are both zero mean then

$$R_{\tilde{y}\tilde{y}} = R_{yy} \text{ and } R_{\tilde{f}\tilde{y}} = R_{fy}$$

In that case we have that $WR_{yy} = R_{fy}$.

- If we go back to the degradation model and find the autocorrelation matrix of the degraded image then we get that

$$\begin{aligned} y &= Hf + n \Rightarrow y^T = f^T H^T + n^T \\ E\{yy^T\} &= HR_{ff}H^T + R_{nn} = R_{yy} \\ E\{fy^T\} &= R_{ff}H^T = R_{fy} \end{aligned}$$

- From the above we get the following result

$$W = R_{fy}R_{yy}^{-1} = R_{ff}H^T(HR_{ff}H^T + R_{nn})^{-1}$$

and the estimate for the original image is

$$\hat{f} = R_{ff}H^T(HR_{ff}H^T + R_{nn})^{-1}y$$

- Note that knowledge of R_{ff} and R_{nn} is assumed.

Wiener Filter Estimator (Stochastic Regularisation)

In frequency domain

$$W(u, v) = \frac{S_{ff}(u, v)H^*(u, v)}{S_{ff}(u, v)|H(u, v)|^2 + S_{nn}(u, v)}$$
$$\hat{F}(u, v) = \frac{S_{ff}(u, v)H^*(u, v)}{S_{ff}(u, v)|H(u, v)|^2 + S_{nn}(u, v)} Y(u, v)$$

- ❖ $S_{ff}(u, v) = |F(u, v)|^2$ is the Power Spectral Density of $f(i, j)$
- ❖ $S_{nn}(u, v) = |N(u, v)|^2$ is the Power Spectral Density of $n(i, j)$

Computational issues

- The noise variance has to be known, otherwise it is estimated from a flat region of the observed image.
- In practical cases where a single copy of the degraded image is available, it is quite common to use $S_{yy}(u, v)$ as an estimate of $S_{ff}(u, v)$.
This is very often a poor estimate.

Wiener Smoothing Filter

In the absence of any blur, $H(u, v) = 1$ and

$$W(u, v) = \frac{S_{ff}(u, v)}{S_{ff}(u, v) + S_{nn}(u, v)} = \frac{(SNR)}{(SNR) + 1}$$

- $(SNR) \gg 1 \Rightarrow W(u, v) \cong 1$
- $(SNR) \ll 1 \Rightarrow W(u, v) \cong (SNR)$

(SNR) is high in low spatial frequencies and low in high spatial frequencies so $W(u, v)$ can be implemented with a lowpass (smoothing) filter.

Relation with Inverse Filtering

If $S_{nn}(u, v) = 0 \Rightarrow W(u, v) = \frac{1}{H(u, v)}$ which is the inverse filter

If $S_{nn}(u, v) \rightarrow 0$

$$\lim_{S_{nn} \rightarrow 0} W(u, v) = \begin{cases} \frac{1}{H(u, v)} & H(u, v) \neq 0 \\ 0 & H(u, v) = 0 \end{cases}$$

which is the pseudoinverse filter.