DIGITAL IMAGE PROCESSING

Unit V – Image Restoration (Gonzalez & Woods)

Topic: A Model of the Image Degradation / Restoration Process

1. The Degradation Model

Image restoration is the process of recovering an original, uncorrupted image f(x,y) from an observed degraded image g(x,y). Unlike enhancement, which is largely subjective, restoration is an objective process that tries to reverse a known or estimated degradation using a mathematical model of how the degradation occurred.

The degradation process is modelled as a degradation function h(x,y) that operates on the input image f(x,y), together with an additive noise term η(x,y):

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

Taking the Fourier transform of both sides and using the convolution theorem, the model can be expressed equivalently and more conveniently in the frequency domain as:

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

where G, H, F and N are the Fourier transforms of g, h, f and η respectively. Given only g(x,y), and some knowledge of the degradation function H and the noise η, the objective of restoration is to obtain an estimate $\hat{f}$(x,y) that is as close as possible to the original image f(x,y).

```
                 f(x,y)  (original image)
                            |
                            v
                 +-------------------+
                 | Degradation H(u,v)|
                 +-------------------+
                            |
                            v
          (+) <----------- eta(x,y)   [additive noise]
                            |
                            v
            g(x,y)  (degraded / observed image)
                            |
                            v
                 +-------------------+
  | Restoration       |   <-- uses knowledge / estimate of H and eta
                 | filter(s)         |
                 +-------------------+
                            |
                            v
                 f^(x,y)  (restored image)
```

Fig. 1 General block diagram of the image degradation / restoration process

2. Noise Models

Noise $\eta(x,y)$ is generally assumed to be independent of spatial position and uncorrelated with the image itself. It is characterised statistically by its probability density function (PDF). The commonly used noise models in spatial-domain restoration are:

Noise type	PDF, p(z)	Mean / Variance
Gaussian	$p(z) = 1/(\sqrt{2\pi}\sigma) \cdot e^{[-(z-\mu)^2/2\sigma^2]}$	μ ; σ^2
Rayleigh	$p(z) = (2/b)(z-a)\, e^{[-(z-a)^2/b]}$, $z \geq a$	$a+\sqrt{\pi b/4}$; $b(4-\pi)/4$
Erlang (Gamma)	$p(z) = [a^b z^{(b-1)} / (b-1)!]\, e^{(-az)}$, $z \geq 0$	b/a ; b/a^2
Exponential	$p(z) = a\, e^{(-az)}$, $z \geq 0$	$1/a$; $1/a^2$
Uniform	$p(z) = 1/(b-a)$ for $a \leq z \leq b$; 0 otherwise	$(a+b)/2$; $(b-a)^2/12$
Impulse (salt-and-pepper)	p(z) = Pa for z=a (pepper); Pb for z=b (salt); 0 otherwise	depends on Pa, Pb

Gaussian noise arises mainly from sensor and circuit electronic noise; Rayleigh noise is useful in modelling range-imaging noise; exponential and gamma (Erlang) noise arise in laser-based imaging; uniform noise is used as a baseline test model since its parameters can be varied and its randomness quantified easily; and impulse (salt-and-pepper) noise appears due to faulty switching during image transmission or acquisition.

3. Estimating the Degradation Function H(u,v)

Three principal methods are used to estimate the degradation function, depending on what information is available:

1. Observation — a small, high-contrast area (e.g., a strong edge) is extracted from the degraded image itself; an unobserved-but-plausible reconstruction is compared to the degraded observation, and the ratio of their Fourier transforms gives an estimate of H(u,v) over that sub-image, which is then extrapolated for the entire image.
2. Experimentation — if equipment similar to the one that produced the degraded image is available, an impulse (bright dot) is imaged with the same system; since the Fourier transform of an impulse is a constant, the observed transform G(u,v) directly gives an estimate H(u,v) = G(u,v)/A.
3. Mathematical modelling — the physical process causing the degradation (e.g., atmospheric turbulence, uniform linear motion) is modelled mathematically. For example, the widely used atmospheric-turbulence model is:

$$H(u, v) = e^{[-k(u^2+v^2)^{(5/6)}]} \quad (k \text{ is a constant depending on the nature of the turbulence})$$

A second important model is that of uniform linear motion blur over a time T with velocity components a, b in the x and y directions:

$$H(u,v) = (T / \pi(ua+vb)) \cdot \sin[\pi(ua+vb)] \cdot e^{[-j\pi(ua+vb)]}$$

4. Restoration in the Presence of Noise Only — Spatial Filters

When the degradation is due to noise alone (H(u,v) = 1), restoration is done by spatial-domain denoising filters applied directly over an m × n neighbourhood Sxy centred at (x,y):

Property	Statement
Arithmetic mean filter	$\hat{f}(x,y) = (1/mn)\, \Sigma\, g(s,t), (s,t) \in S_{xy}$ — smooths noise but blurs edges
Geometric mean filter	$\hat{f}(x,y) = [\, \Pi\, g(s,t)\,]$ ^$(1/mn)$ — achieves smoothing comparable to the arithmetic mean but loses less image detail
Harmonic mean filter	$\hat{f}(x,y) = mn / \Sigma[1/g(s,t)]$ — works well for salt noise and Gaussian noise, but fails for pepper noise
Contraharmonic mean filter, order Q	$\hat{f}(x,y) = \Sigma\, g(s,t)$^$(Q+1) / \Sigma\, g(s,t)$^$Q$ — eliminates pepper noise for $Q > 0$, and salt noise for $Q < 0$ ($Q = 0$ reduces to arithmetic mean; $Q = -1$ reduces to harmonic mean)
Median filter	$\hat{f}(x,y) = \text{median}\{\, g(s,t)\, \}$ — excellent for bipolar and unipolar impulse (salt-and-pepper) noise, with less blurring than linear smoothing filters of the same size
Max / Min filters	$\hat{f}(x,y) = \max\{g(s,t)\}$ removes pepper noise (finds brightest point); $\hat{f}(x,y) = \min\{g(s,t)\}$ removes salt noise (finds darkest point)
Midpoint filter	$\hat{f}(x,y) = (1/2)[\max\{g(s,t)\} + \min\{g(s,t)\}]$ — works best for randomly distributed noise such as Gaussian or uniform noise
Alpha-trimmed mean filter	Delete the $d/2$ lowest and $d/2$ highest grey values in S_{xy}, then average the remaining $mn-d$ pixels — useful when the neighbourhood contains a mixture of noise types (e.g., salt-and-pepper together with Gaussian noise)

```
              Noisy image  g(x,y)
                          |
                          v
           +----------------------------+
           |  m x n moving window Sxy     |
           |  centred at each pixel (x,y) |
           +----------------------------+
        |   (compute mean / median / order-statistic
          |    of the mn samples inside the window)
                          v
              Restored estimate  f^(x,y)
```

Fig. 2 Spatial-domain neighbourhood (mean / order-statistics) filtering

5. Periodic Noise Reduction by Frequency-Domain Filtering

Periodic noise appears in the Fourier spectrum as discrete, symmetric impulse-like bursts and is best removed using selective frequency-domain filters:

- Band-reject filter — rejects a specific band of frequencies around a radius D_0 (with width W) about the origin, leaving other frequencies unchanged; used when the general location of the noise frequencies is known.
- Band-pass filter — the exact opposite of the band-reject filter ($H_{BP} = 1 - H_{BR}$); useful for isolating and visualising the effect of specific frequency bands, e.g., for noise-pattern extraction.

- Notch filter — rejects (or passes) frequencies in a small, predefined neighbourhood about a specific set of symmetric frequency-pair locations (u_0,v_0) and $(-u_0,-v_0)$; the most effective tool for removing periodic interference such as scan-line or sinusoidal noise, since it targets only the exact spikes seen in the spectrum.

6. Inverse Filtering

The simplest approach to restoration when H(u,v) is known is direct (naive) inverse filtering, which forms an estimate of F(u,v) by dividing the transform of the degraded image by the degradation function:

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

This shows the fundamental problem of inverse filtering: even if H(u,v) is known exactly, N(u,v) is not, and when H(u,v) has values close to zero (typically at high frequencies), the noise term N(u,v)/H(u,v) is amplified enormously, dominating the restored image. In practice, inverse filtering is therefore restricted to a limited radius around the origin (excluding the near-zero, noise-dominated region of H).

7. Minimum Mean Square Error (Wiener) Filtering

The Wiener filter incorporates both the degradation function and the statistical properties of the noise into the restoration process, minimising the mean square error between the estimate $\hat{f}(x,y)$ and the original f(x,y):

$$e^2 = E\{ (f - \hat{f})^2 \}$$

The resulting frequency-domain solution, known as the Wiener filter, is:

$$\hat{F}(u,v) = [\, 1/H(u,v) \cdot |H(u,v)|^2 / (|H(u,v)|^2 + Sn(u,v)/Sf(u,v)) \,] \cdot G(u,v)$$

where $|H(u,v)|^2 = H^*(u,v)H(u,v)$, $Sn(u,v) = |N(u,v)|^2$ is the noise power spectrum and $Sf(u,v) = |F(u,v)|^2$ is the power spectrum of the undegraded image. The ratio Sn/Sf is often replaced in practice by a single adjustable constant K, giving the commonly implemented approximate form:

$$\hat{F}(u,v) = [\, |H(u,v)|^2 / (H(u,v)(|H(u,v)|^2 + K)) \,] \cdot G(u,v)$$

Unlike inverse filtering, the Wiener filter does not produce a division by (near) zero, since the term $|H(u,v)|^2 + K$ stays bounded away from zero even where H(u,v) itself is small — giving a far more stable, noise-tolerant restoration.

8. Constrained Least Squares Filtering

Wiener filtering requires knowledge of the noise-to-signal power ratio, which is often unavailable except as an overall average value. Constrained least squares filtering needs only the mean and variance of the noise, and optimises restoration subject to a smoothness constraint using the Laplacian operator ∇^2:

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

where P(u,v) is the Fourier transform of the Laplacian operator and γ is a regularisation parameter adjusted iteratively so that the residual satisfies the known noise variance — giving a restoration that is optimal for each image rather than relying on ensemble-average statistics.

9. Summary — Choice of Restoration Approach

Situation	Recommended restoration approach
Noise only, no blur	Spatial mean / order-statistics filters (Sec. 4)
Periodic noise	Band-reject / notch filtering (Sec. 5)
Known H, negligible noise	Inverse filtering, radially limited (Sec. 6)
Known H and noise statistics	Wiener (minimum mean square error) filtering (Sec. 7)
Known H, only noise mean/variance known	Constrained least squares filtering (Sec. 8)