ABHYASMITRA.IN

COMPREHENSIVE UNIVERSITY STUDY NOTES

COMPUTER VISION

Unit III — Frequency Domain Processing and Edge Detection

Topics Covered in this Unit

Frequency Domain Concepts: 2D Fourier Transform, 2D Discrete Fourier Transform (2D DFT) & FFT
Frequency Spectrum Analysis: Magnitude Spectrum, Phase Spectrum, Power Spectrum & Centering
Spatial vs. Frequency Duality: The 2D Convolution Theorem & Filtering Workflow
Frequency Domain Filtering: Low-Pass (Ideal, Butterworth, Gaussian), High-Pass & Band-Pass/Reject
Inverse 2D Discrete Fourier Transform (2D IDFT) & Reconstruction Principles
Edge Detection Foundations: First-Order Gradient Operators (Sobel, Prewitt, Roberts Cross)
Second-Order Edge Detection: Laplacian, Laplacian of Gaussian (LoG / Mexican Hat) & Zero-Crossings
The Canny Edge Detector ('Gold Standard'): 5-Stage Pipeline (Smoothing, Gradient, NMS, Double Threshold, Hysteresis)
Comprehensive Edge Detector Comparison: Edge Localization, Noise Sensitivity & Computational Overhead
Practical Case Studies: Autonomous Vehicle Lane Detection (ROI + Canny + Hough) & Document Scanning Rectification
High-Yield University Exam Question Bank with Detailed Model Answers & Unit Summary

Compiled in a structured, easy-to-understand, textbook style format
Specially curated for B.Tech / B.E. / BCA / MCA / B.Sc Computer Science, AI & Data Science Students

COURSE SYLLABUS & MAPPING

[UNIT III SYLLABUS — PRESCRIBED UNIVERSITY CURRICULUM]

This section contains the official syllabus for Unit III. Verify with your university curriculum mapping.

Unit III Syllabus Breakdown:

Unit III Frequency Domain Processing and Edge Detection: Frequency Domain Concepts: Fourier Transform, DFT, FFT, Frequency spectrum (magnitude & phase). Frequency filtering: Low-pass, High-pass, Band-pass. Inverse transformation: Concept of Inverse Transformation, Inverse Discrete Fourier Transform (IDFT). Edge Detection: Sobel, Prewitt, Roberts Laplacian, LoG, Canny Edge Detector, the 'gold standard'- Gaussian smoothing, gradient, NMS, double thresholding, hysteresis, Comparison of edge detectors: performance vs. noise sensitivity. Practical application: Lane boundary detection, document scanning

- Course Code: ____________________ • Semester / Year: ____________________

TABLE OF CONTENTS

3.1 Frequency Domain Concepts & The 2D Fourier Transform 1
Frequency Spectrum, Magnitude, Phase and Power 1
The 2D Convolution Theorem (Spatial vs Frequency Duality) 1
3.2 Frequency Domain Filtering: Low-Pass, High-Pass & Band-Pass 1
1. Low-Pass Filters (LPF) — Smoothing & Noise Removal 1
2. High-Pass Filters (HPF) — Edge Sharpening 1
3. Band-Pass and Band-Reject Filters 1
3.3 Inverse Discrete Fourier Transform (2D IDFT) 1
3.4 Edge Detection: First-Order & Second-Order Operators 1
1. First-Order Gradient Operators 1
2. Second-Order Derivative: Laplacian of Gaussian (LoG / Mexican Hat) 1
3.5 The Canny Edge Detector: The 'Gold Standard' Algorithm 1
The 5-Stage Canny Edge Detection Pipeline 1
3.6 Practical Applications: Lane Detection & Document Scanning 1
1. Autonomous Vehicle Lane Boundary Detection 1
2. Document Scanner Perspective Rectification 1
3.7 High-Yield University Exam Question Bank & Model Answers 1
Part A: Short Answer Questions (2–3 Marks) 1
Part B: Long Answer Questions (8–10 Marks) 1
3.8 Unit III Summary & Quick Revision Sheet 1

3.1 Frequency Domain Concepts & The 2D Fourier Transform

Mathematical Formulation: 2D Discrete Fourier Transform (2D DFT)

The 2D Discrete Fourier Transform of an $M \times N$ digital image $f(x,y)$ decomposes the spatial spatial image into a linear combination of complex 2D sinusoidal basis functions:

$$F(u,v) = \sum_{x=0}^{M-1} \sum_{y=0}^{N-1} f(x,y) \times \exp(-j\, 2\pi\, (ux/M + vy/N))$$

where $u = 0,1,\ldots,M-1$ and $v = 0,1,\ldots,N-1$ (frequency variables).

Frequency Spectrum, Magnitude, Phase and Power

The Fourier Transform output $F(u,v)$ is a 2D matrix of complex numbers: $F(u,v) = R(u,v) + j\, I(u,v)$:

- **1. Fourier Magnitude Spectrum $|F(u,v)|$:** Represents the amplitude/energy of sinusoidal frequency components present in the image:

$$|F(u,v)| = \sqrt{R(u,v)^2 + I(u,v)^2}$$

- **2. Phase Spectrum $\varphi(u,v)$:** Encodes spatial location, structural orientation, and edge alignment information:

$$\varphi(u,v) = \arctan(I(u,v) / R(u,v))$$

Note: In image perception, the Phase Spectrum contains significantly more structural geometry information than the Magnitude Spectrum.

- **3. Power Spectrum $P(u,v)$:** $P(u,v) = |F(u,v)|^2 = R(u,v)^2 + I(u,v)^2$
- **4. Frequency Centering & Log Display:** In raw DFT, low frequencies are located at the four corners. Multiplying the input image by $(-1)^{(x+y)}$ prior to computing DFT shifts the zero-frequency DC component $F(0,0)$ to the center $(M/2, N/2)$. Because DC amplitude is orders of magnitude larger than high frequencies, dynamic range compression is applied for visual display:

$$D(u,v) = c \times \log(1 + |F(u,v)|)$$

The 2D Convolution Theorem (Spatial vs Frequency Duality)

The Convolution Theorem establishes that spatial domain convolution is mathematically equivalent to element-wise multiplication in the frequency domain, and vice versa:

Spatial Domain: $g(x,y) = f(x,y) * h(x,y)$ ↔ Frequency Domain: $G(u,v) = F(u,v) \times H(u,v)$

- • **Computational Advantage (FFT):** Spatial convolution for an $M \times N$ image with $K \times K$ mask requires $O(M N K^2)$ operations. Via Fast Fourier Transform (FFT, Cooley-Tukey algorithm), frequency filtering requires only $O(M N \log(M N))$ operations, making large-scale filtering dramatically faster.

3.2 Frequency Domain Filtering: Low-Pass, High-Pass & Band-Pass

Frequency domain filtering involves 3 foundational steps: (1) Compute 2D DFT $F(u,v)$, (2) Multiply by filter transfer function $H(u,v)$ to obtain $G(u,v) = F(u,v) \times H(u,v)$, (3) Compute Inverse DFT to obtain filtered image $g(x,y)$.

1. Low-Pass Filters (LPF) — Smoothing & Noise Removal

- • **Concept:** Attenuates high frequencies (edges, noise, fine textures) while passing low frequencies (smooth backgrounds, gradual intensity changes). Let $D(u,v) = \sqrt{(u - M/2)^2 + (v - N/2)^2}$ be the Euclidean distance from center:
- • **Ideal Low-Pass Filter (ILPF):** $H(u,v) = 1$ if $D(u,v) \leq D_0$, and 0 if $D(u,v) > D_0$. Produces severe ringing artifacts (concentric ripples) due to the Sinc function in the spatial domain.
- • **Butterworth Low-Pass Filter (BLPF) of order n:** $H(u,v) = 1 / [1 + (D(u,v) / D_0)^{2n}]$. Smooth transition eliminates harsh ringing artifacts; higher orders ($n \geq 3$) exhibit slight ringing.
- • **Gaussian Low-Pass Filter (GLPF):** $H(u,v) = \exp(- D(u,v)^2 / (2 D_0^2))$. Completely eliminates ringing artifacts because the inverse Fourier transform of a Gaussian is also a Gaussian.

2. High-Pass Filters (HPF) — Edge Sharpening

- • **Concept:** Attenuates low frequencies (DC component) while passing high frequencies, highlighting sharp edges, object boundaries, and fine structural contours:
- • **Mathematical Relation:** $H_{HPF}(u,v) = 1 - H_{LPF}(u,v)$
- • **Ideal HPF (IHPF):** $H(u,v) = 0$ if $D(u,v) \leq D_0$, and 1 if $D(u,v) > D_0$.
- • **Gaussian HPF (GHPF):** $H(u,v) = 1 - \exp(- D(u,v)^2 / (2 D_0^2))$. Produces crisp, artifact-free boundary sharpening.

3. Band-Pass and Band-Reject Filters

- • **Band-Reject Filter (BRF):** Attenuates a specific radial frequency band $[D_0 - W/2, D_0 + W/2]$, ideal for eliminating periodic sinusoidal interference patterns (e.g., electronic scanning lines).
- • **Band-Pass Filter (BPF):** $H_{BPF}(u,v) = 1 - H_{BRF}(u,v)$. Passes only a specific frequency band, isolating specific structural patterns.

3.3 Inverse Discrete Fourier Transform (2D IDFT)

Mathematical Formulation: 2D Inverse Discrete Fourier Transform (2D IDFT)

The 2D IDFT converts frequency domain coefficients $G(u,v)$ back into the spatial spatial domain image $g(x,y)$:

$$g(x,y) = \frac{1}{M \times N} \sum_{u=0}^{M-1} \sum_{v=0}^{N-1} G(u,v) \times \exp\left(+j\, 2\pi \left(\frac{ux}{M} + \frac{vy}{N} \right) \right)$$

COMPLETE FREQUENCY DOMAIN FILTERING PIPELINE

```
+------------------+         2D FFT        +------------------+
|  Spatial Image   | ---------------------> | 2D DFT Spectrum  |
|     f(x,y)       |                        |     F(u,v)       |
+------------------+                        +------------------+
                                                     |
                                                     v  [Element-wise Multipl.]
+------------------+         2D IFFT       +------------------+
| Enhanced Image   | <--------------------- | Filtered Output  |
|     g(x,y)       |                        | G(u,v)=F(u,v)*H  |
+------------------+                        +------------------+
```

Figure: Standard 4-step sequence for frequency domain image manipulation

3.4 Edge Detection: First-Order & Second-Order Operators

Definition: Edge in Digital Images

An edge is a local spatial discontinuity or significant gradient jump in image brightness intensity. Edges correspond to physical object boundaries, surface orientation changes, shadows, or material property variations.

1. First-Order Gradient Operators

Operator Name	Horizontal Mask (G_x)	Vertical Mask (G_y)	Distinctive Characteristics
Roberts Cross (2×2)	[[1, 0], [0, -1]]	[[0, 1], [-1, 0]]	Diagonal derivatives; computationally ultra-lightweight, but highly sensitive to noise.
Prewitt Operator (3×3)	[[-1, 0, 1], [-1, 0, 1], [-1, 0, 1]]	[[-1, -1, -1], [0, 0, 0], [1, 1, 1]]	Differential averaging; incorporates uniform smoothing perpendicular to gradient.

Operator Name	Horizontal Mask (G_x)	Vertical Mask (G_y)	Distinctive Characteristics
Sobel Operator (3×3)	[[-1, 0, 1], [-2, 0, 2], [-1, 0, 1]]	[[-1, -2, -1], [0, 0, 0], [1, 2, 1]]	Gaussian-weighted center coefficients (weight=2) provide superior noise suppression.

2. Second-Order Derivative: Laplacian of Gaussian (LoG / Mexican Hat)

- • **Problem with Raw Laplacian:** Raw Laplacian ($\nabla^2 f$) is extremely sensitive to noise and double-edges.
- • **LoG Mechanism:** First smooths the image with an isotropic Gaussian filter $G_\sigma(x,y)$, then applies the Laplacian:

$$LoG(x,y) = \nabla^2 [G_\sigma(x,y) * f(x,y)] = [\nabla^2 G_\sigma(x,y)] * f(x,y)$$

$$LoG(x,y) = - (1 / (\pi \sigma^4)) \times [1 - (x^2 + y^2) / (2 \sigma^2)] \times \exp(- (x^2 + y^2) / (2 \sigma^2))$$

- • **Zero-Crossing Detection:** Edges correspond to the zero-crossings (points where the second derivative changes sign from positive to negative) in the LoG filtered output.

3.5 The Canny Edge Detector: The 'Gold Standard' Algorithm

Canny's Three Optimal Edge Criteria (John Canny, 1986)

1. Low Error Rate: All real edges must be marked, and no spurious noise edges should be detected.

2. Good Localization: Distance between detected edge pixels and actual physical edges must be minimized.

3. Single Response: An edge must not produce multiple duplicate response pixels (thin 1-pixel edges).

The 5-Stage Canny Edge Detection Pipeline

- **Stage 1: Gaussian Smoothing:** Convolves image with a 5×5 Gaussian kernel with standard deviation σ to suppress sensor noise:

$$f_{smooth}(x,y) = G_\sigma(x,y) * f(x,y)$$

- **Stage 2: Gradient Magnitude and Orientation:** Computes horizontal (g_x) and vertical (g_y) derivatives via Sobel masks:

$$M(x,y) = \sqrt{g_x^2 + g_y^2}, \quad \theta(x,y) = \arctan(g_y / g_x)$$

The gradient angle θ is quantized into one of 4 discrete directions: 0° (horizontal), 45° (diagonal), 90° (vertical), or 135° (anti-diagonal).

- **Stage 3: Non-Maximum Suppression (NMS):** Thins thick gradient ridges into sharp 1-pixel edges. For every pixel (x,y), compares its gradient magnitude M(x,y) with two neighboring pixels along the gradient direction θ. If M(x,y) is NOT strictly greater than both neighbors, it is suppressed (set to 0).
- **Stage 4: Double Thresholding:** Applies two scalar thresholds—High Threshold T_high and Low Threshold T_low (typically T_high : T_low = 2:1 or 3:1):
  - • **Strong Edge Pixels:** M(x,y) ≥ T_high (unconditionally retained as true edges).
  - • **Weak Edge Pixels:** T_low ≤ M(x,y) < T_high (candidates for genuine edges).
  - • **Non-Edge Pixels:** M(x,y) < T_low (immediately discarded / set to 0).
- **Stage 5: Edge Tracking by Hysteresis:** Weak edge pixels are retained if and only if they are 8-connected to at least one strong edge pixel. This seamlessly traces continuous edge contours without breaking and without letting noise penetrate.

Edge Detection Algorithm	Mathematical Order	Noise Robustness	Edge Thickness & Localization	Computational Complexity
Roberts Cross	1st-order gradient (2×2)	Extremely poor; sensitive to noise.	Thick; poor sub-pixel localization.	Lowest (minimal arithmetic).
Prewitt Operator	1st-order gradient (3×3)	Moderate noise suppression.	Thick edges (requires thresholding).	Low.
Sobel Operator	1st-order gradient (3×3)	Good (Gaussian center weights).	Thick edges (2-3 pixels wide).	Low-Moderate.
Laplacian of Gaussian (LoG)	2nd-order derivative	High (pre-Gaussian smoothing).	Thin (zero-crossing), but closed loops.	Moderate.
Canny Edge Detector	Multi-stage optimal pipeline	Outstanding (adaptive hysteresis).	Optimal 1-pixel thin contours.	Higher (NMS + Hysteresis search).

3.6 Practical Applications: Lane Detection & Document Scanning

1. Autonomous Vehicle Lane Boundary Detection

- • **Pipeline Architecture:**
  - 1. Region of Interest (ROI) Masking: Extract trapezoidal polygon covering only the road surface ahead.
  - 2. Grayscale & Gaussian Denoising: Remove asphalt texture noise.
  - 3. Canny Edge Detection: Extract clean 1-pixel lane boundaries.
  - 4. Progressive Hough Transform (PHT): Detect collinear line segments in (r, θ) parameter space.

- 5. Slope Classification & Curve Fitting: Separate left lanes (negative slope) from right lanes (positive slope) and compute steering angles.

2. Document Scanner Perspective Rectification

- • **Pipeline Architecture:**
  - 1. Edge Detection & Contour Extraction: Identify the 4-sided document boundary.
  - 2. Corner Extraction: Locate 4 corner coordinates (top-left, top-right, bottom-right, bottom-left).
  - 3. Perspective Transformation (Homography Matrix H): Warp perspective distortion into a flat, deskewed rectangular document image.
  - 4. Adaptive Thresholding: Binarize text for clean OCR reading.

3.7 High-Yield University Exam Question Bank & Model Answers

Part A: Short Answer Questions (2–3 Marks)

- **Q1. What is the 2D Convolution Theorem? Why is it computationally advantageous?**

Answer: The 2D Convolution Theorem states that spatial convolution corresponds to element-wise multiplication in the frequency domain: f(x,y) * h(x,y) ↔ F(u,v) × H(u,v). Using FFT, filtering large images requires only O(N log N) operations compared to O(N × K^2) for spatial filtering with large kernels.

- **Q2. Why does an Ideal Low-Pass Filter produce Ringing Artifacts, while a Gaussian Low-Pass Filter does not?**

Answer: An Ideal LPF has an abrupt, discontinuous rectangular cutoff in the frequency domain, whose inverse Fourier transform is a Sinc function with infinite secondary ripple lobes, producing ringing artifacts. The inverse Fourier transform of a smooth Gaussian function is another smooth Gaussian without ripple lobes.

- **Q3. Differentiate between Magnitude Spectrum and Phase Spectrum.**

Answer: The Magnitude Spectrum |F(u,v)| represents the intensity and energy distribution of frequency components. The Phase Spectrum φ(u,v) encodes spatial position, geometric orientation, and structural edge relationships. In human image perception, phase preserves structural intelligibility.

- **Q4. What is Non-Maximum Suppression (NMS) in the Canny Edge Detector?**

Answer: NMS is an edge-thinning technique that scans gradient magnitude along the gradient direction θ and preserves pixel magnitude only if it is a local maximum compared to its two adjacent neighbors along θ, suppressing non-maxima to 0 to yield sharp 1-pixel wide edges.

Part B: Long Answer Questions (8–10 Marks)

- **Q5. Describe the Canny Edge Detection Algorithm in complete detail. Explain each of its 5 processing stages with mathematical formulations and diagrams.**

Model Answer Outline:

- 1. Introduction: Canny's 3 optimality criteria (Low error rate, Good localization, Single response).
- 2. Stage 1: Gaussian Smoothing (Noise reduction using 2D Gaussian kernel).
- 3. Stage 2: Gradient Magnitude M(x,y) and Direction θ(x,y) calculation via Sobel operators; 4-directional angle quantization (0°, 45°, 90°, 135°).
- 4. Stage 3: Non-Maximum Suppression (NMS) algorithm and ridge thinning mechanics.
- 5. Stage 4: Double Thresholding (Strong edges ≥ T_high, Weak edges between T_low and T_high, Non-edges < T_low).
- 6. Stage 5: Edge Tracking by Hysteresis (8-connectivity graph search connecting weak edges to strong anchors).
- 7. Comprehensive architectural block diagram tracing an image from raw input to clean binary edge map.

- **Q6. Compare Ideal, Butterworth, and Gaussian Low-Pass and High-Pass filters in the Frequency Domain. Derive their transfer functions.**

Model Answer Outline:

- 1. Concept of frequency domain filtering using Euclidean distance D(u,v) from frequency center.
- 2. Ideal Low-Pass Filter: H(u,v) formula, sharp rectangular cutoff, derivation of spatial Sinc impulse response, explanation of Ringing.
- 3. Butterworth Low-Pass Filter: H(u,v) = 1 / [1 + (D/D_0)^(2n)], role of order n, smooth rolloff trade-off.
- 4. Gaussian Low-Pass Filter: H(u,v) = exp(-D^2 / 2D_0^2), proof of zero ringing.
- 5. High-Pass duality: H_HPF(u,v) = 1 - H_LPF(u,v) for Ideal, Butterworth, and Gaussian variants.
- 6. Cross-sectional transfer function curve diagrams comparing cutoff sharpness and spatial domain behavior.

3.8 Unit III Summary & Quick Revision Sheet

Core Concept	Key Theoretical & Mathematical Takeaways for Exams
2D DFT & FFT	Decomposes 2D spatial image into sinusoidal frequencies. F(u,v) = R(u,v) + j I(u,v). FFT reduces complexity to O(N log N).

Core Concept	Key Theoretical & Mathematical Takeaways for Exams
Frequency Filters	Ideal (Ringing artifacts due to Sinc), Butterworth (Smooth rolloff, controlled ringing), Gaussian (Zero ringing, smooth Gaussian in both domains).
1st-Order Edge Detectors	Sobel, Prewitt, Roberts. Compute gradient magnitude $M = \sqrt{(g_x^2 + g_y^2)}$ and direction $\theta = \arctan(g_y/g_x)$.
2nd-Order Edge Detectors	Laplacian of Gaussian (LoG / Mexican Hat) finds edges at Zero-Crossings.
Canny Edge Detector	5 Stages: Gaussian Smoothing → Sobel Gradient → NMS → Double Thresholding → Hysteresis Tracking. Gold standard.