كلية العلوم

قســـــــم الانظمة الطبية الذكية

Lec 1 & 2

Introduction to Computer Vision

Subject: Computer Vision

Level: third

Lecturer: Asst. Lecturer Qusai AL-Durrah

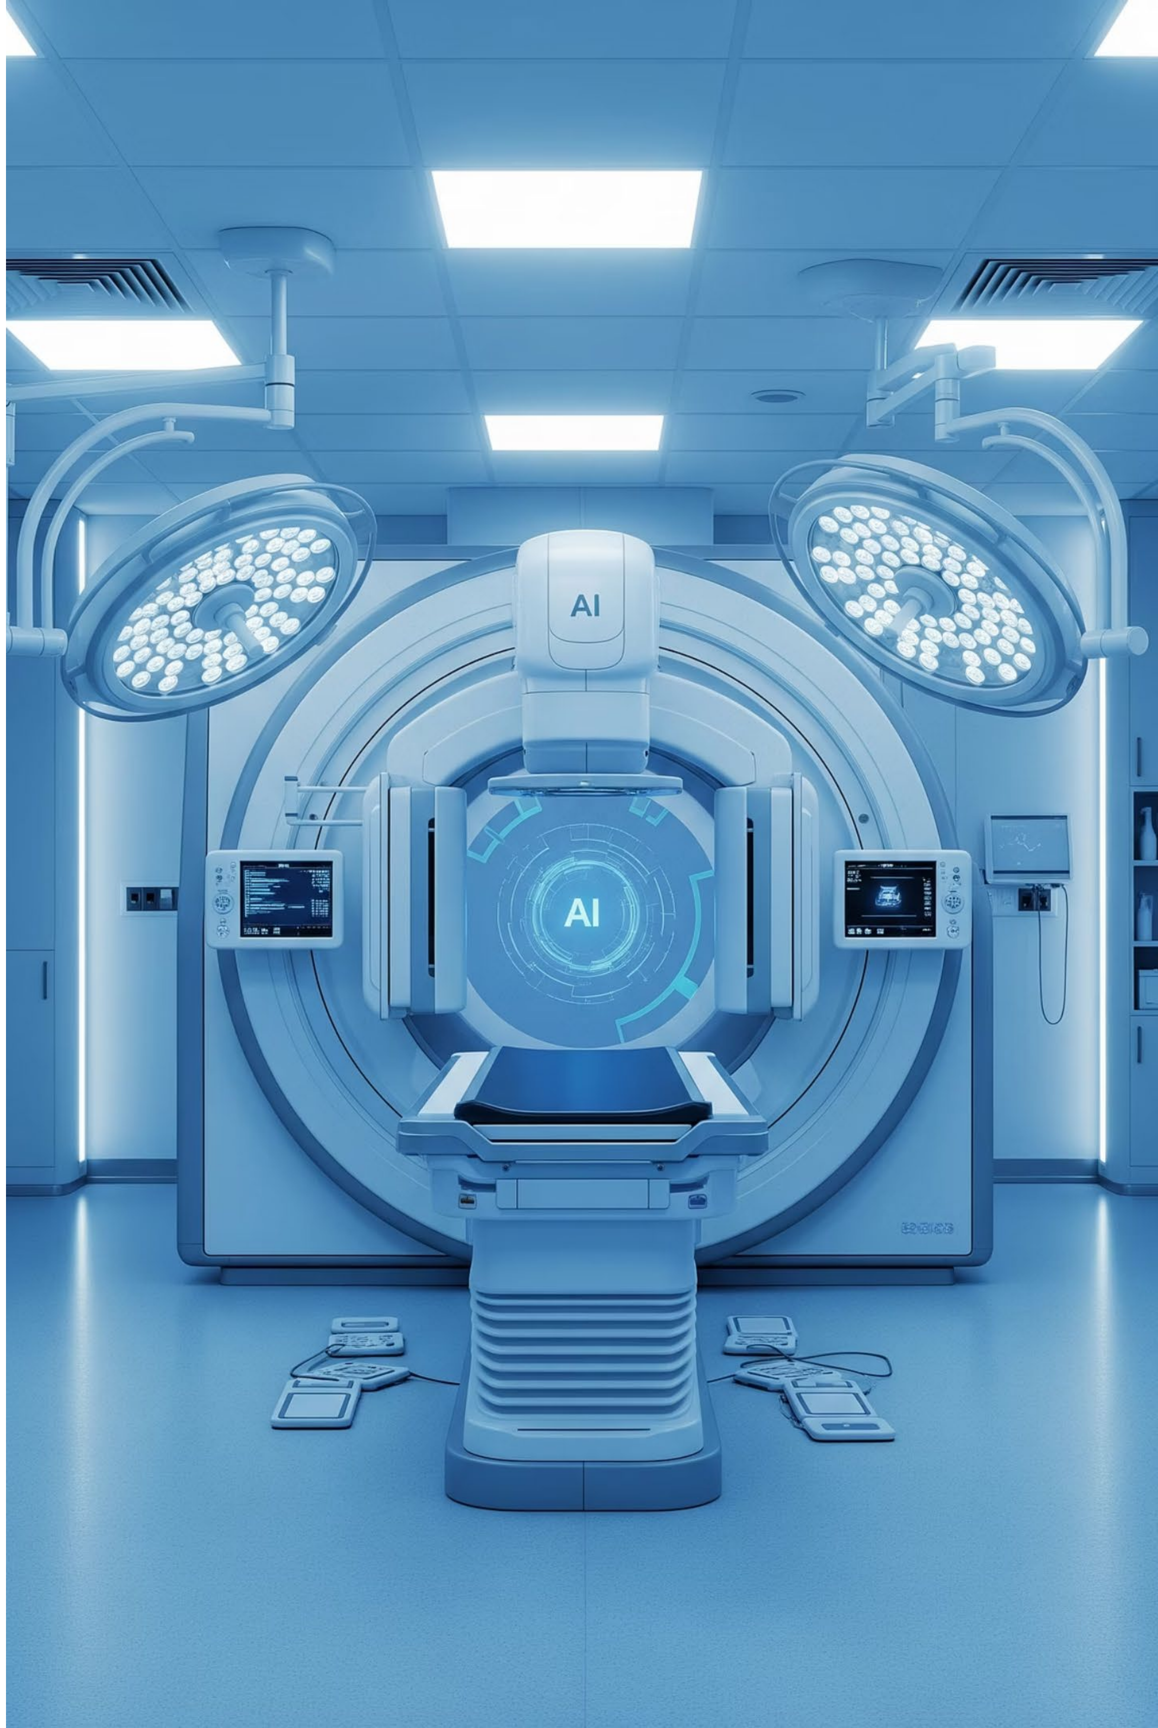

The Visual Data Challenge in Healthcare

Intelligent systems increasingly depend on visual data to make critical medical decisions. In modern healthcare environments, the most important diagnostic information often comes in visual form through multiple imaging modalities.

Medical imaging generates massive volumes of data from sources including X-ray radiography, computed tomography (CT), magnetic resonance imaging (MRI), ultrasound examinations, endoscopy video feeds, microscope slides for pathology, and specialized imaging in dermatology and ophthalmology.

The fundamental challenge facing healthcare systems today is not simply storing or displaying these images—it's extracting clinically meaningful information that can guide diagnosis, treatment planning, and patient care decisions.

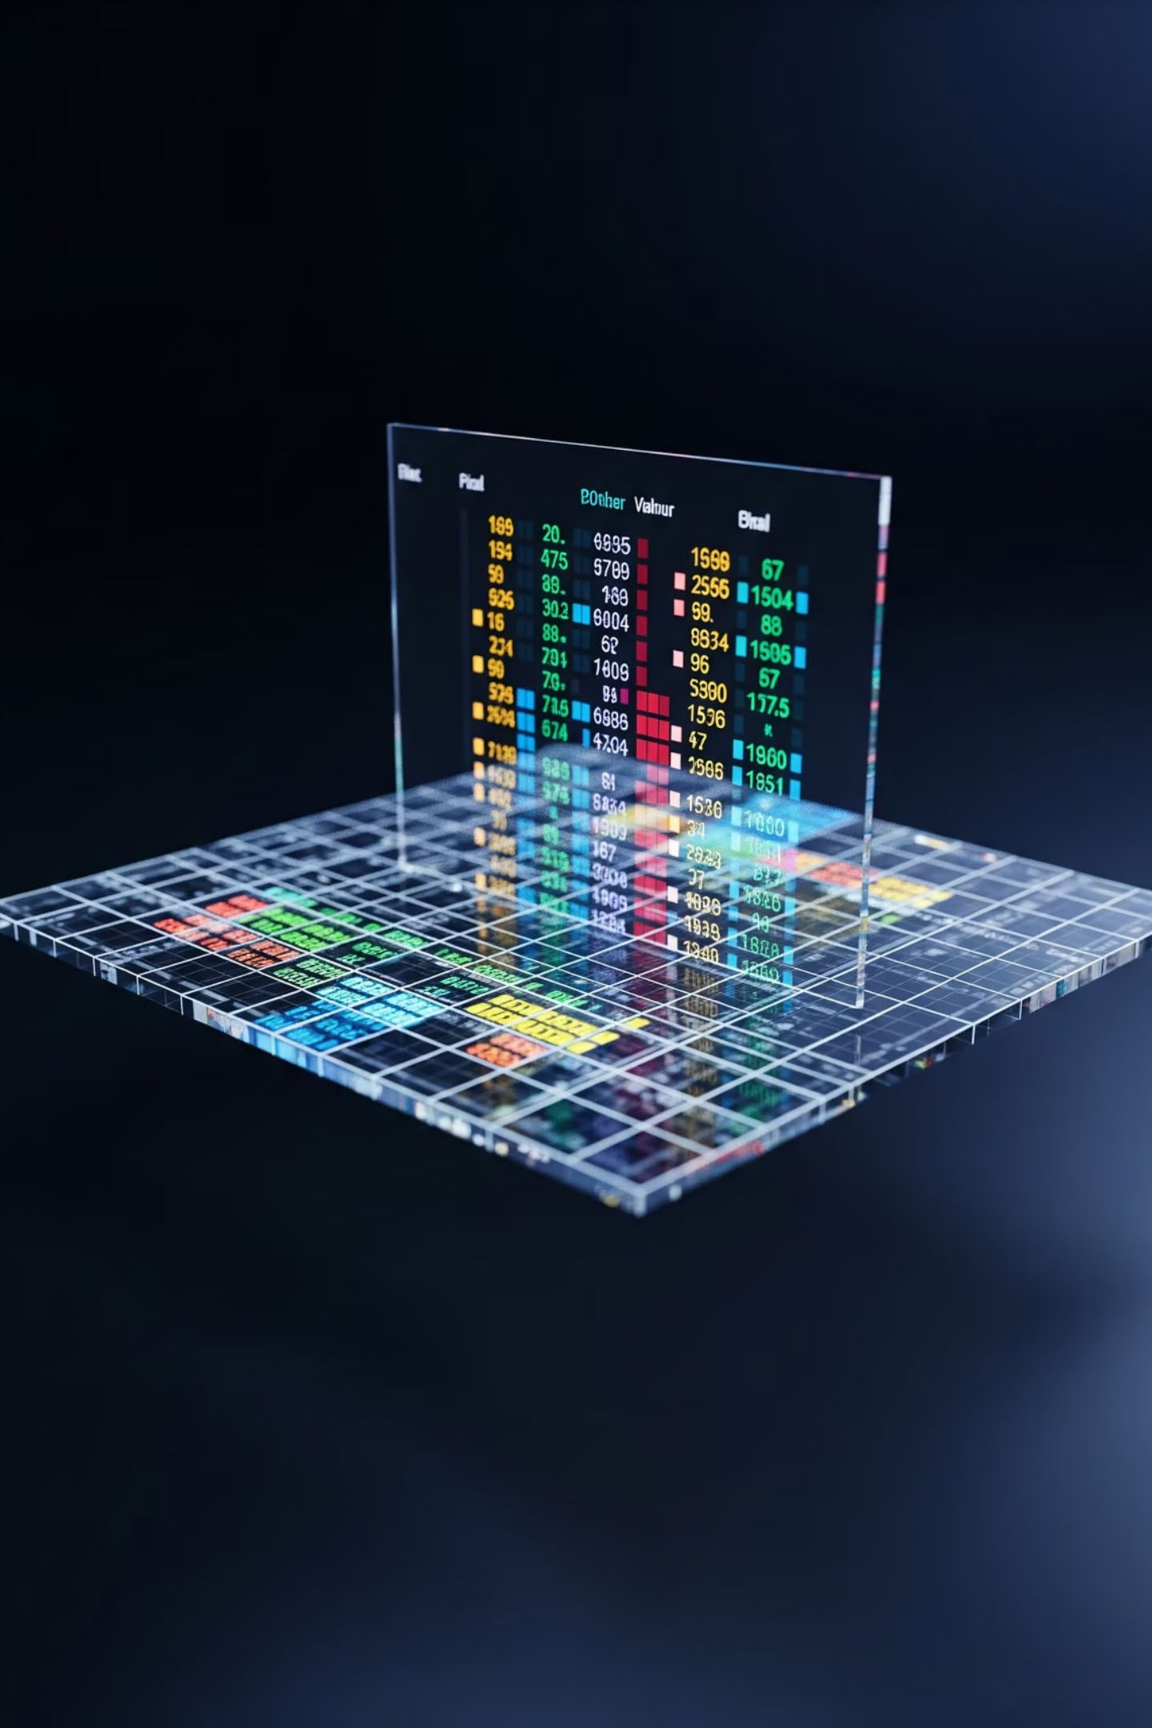

Understanding Digital Images

What Computers See

A digital image is not "a picture" to a computer. It is a matrix (or multiple matrices) of numerical values representing pixel intensities.

The Transformation Challenge

These numbers must be transformed into measurements, interpretations, and actionable decisions through sophisticated algorithms.

Clinical Applications

Systems must detect lesions, segment tumor boundaries, classify diseases, track movement, and estimate depth for robotic assistance.

Defining Computer Vision

> **Computer Vision is the discipline that enables a computer to extract meaningful information from images and video and use that information for inference and decision-making.**

Extract Meaningful Information

The system identifies location, boundaries, measurements, identity, or descriptive patterns (features) within visual data. This extraction process transforms raw pixels into structured, interpretable data.

Inference and Decision-Making

The system doesn't stop at modifying the image—it produces actionable outputs used for diagnosis, classification, detection, alerting, or robotic control in clinical settings.

Computer Vision vs Image Processing

Image Processing

Input

Image

Output

Image

Purpose

Improve visual quality or representation for viewing or preparation

Medical Examples:

- Denoising ultrasound images to reduce speckle artifacts
- Enhancing contrast in X-ray images for better visibility
- Correcting motion blur in endoscopy video frames
- Compressing medical scans for efficient storage

Computer Vision

Input

Image / Video

Output

Information + Decision

Purpose

Interpretation and understanding for clinical action

Intelligent Medical Systems Examples:

- "Tumor vs not tumor" classification decisions
- "Tumor boundary" segmentation mask generation
- "Estimate organ volume" precise measurements
- "Track patient movement" trajectory analysis
- "Detect anomaly" automated alert systems

The Relationship Between Processing and Vision

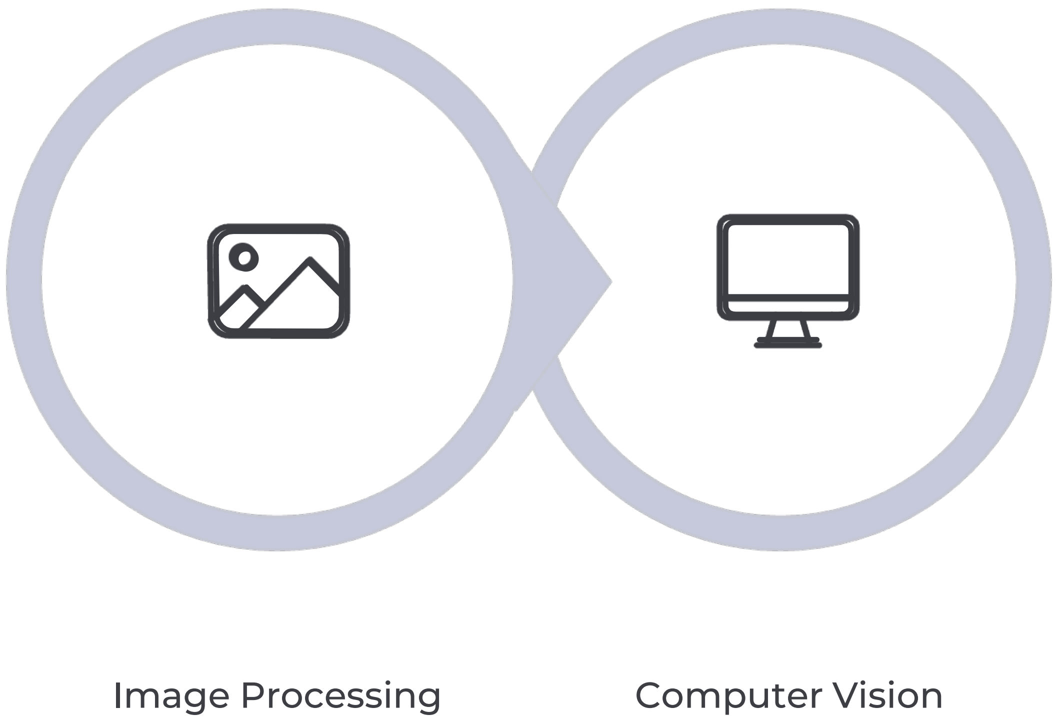

Image processing and computer vision work together in medical imaging workflows. Image processing serves as an essential early stage to improve image quality, reduce noise, and enhance relevant features. This preprocessed data then feeds into computer vision algorithms that produce the final interpretation, classification, or decision output used in clinical practice. Understanding this relationship is crucial for building effective intelligent medical systems.

Why Computer Vision is Difficult

The Fundamental Challenge

Computer vision is inherently challenging because pixels do not directly represent objects, diseases, or diagnoses. The gap between low-level pixel intensities and high-level semantic understanding requires sophisticated algorithms and extensive training data.

Medical imaging introduces additional layers of complexity that make computer vision particularly difficult in healthcare applications. These challenges must be understood and addressed to build reliable intelligent medical systems.

Medical Imaging Challenges

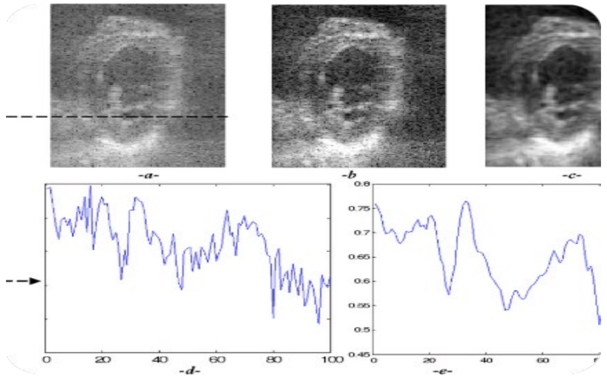

Noise and Artifacts

Medical images contain modality-specific artifacts: speckle in ultrasound, low-dose CT noise, motion artifacts, and scanner-dependent distortions that obscure clinical features.

Large Variability

Significant differences exist across patients, scanner manufacturers, imaging protocols, and acquisition settings, making generalization difficult.

Class Imbalance

Many diseases are rare in datasets, creating severe class imbalance that can bias machine learning models and evaluation metrics.

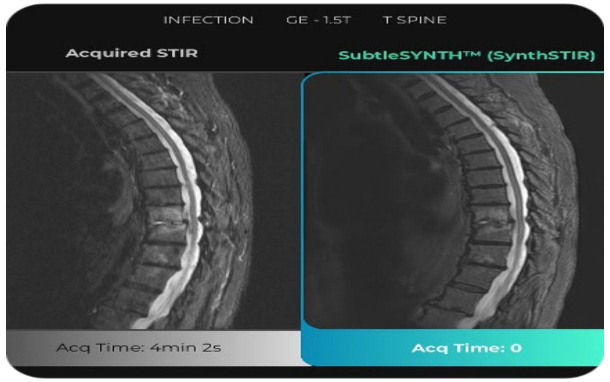

Subtle Patterns

Clinically important findings may be small, exhibit low contrast, or appear visually ambiguous even to trained radiologists.

Domain Shift

Models trained on data from one hospital or device often perform poorly when deployed on data from different sources, limiting clinical deployment.

The Computer Vision Pipeline

A vision system is best understood as a structured pipeline that transforms raw pixels into clinical decisions. Most real-world computer vision systems follow a systematic sequence of stages, each building upon the previous one. Understanding this pipeline is essential for designing, implementing, and debugging intelligent medical systems.

The following sections will explore each stage of this pipeline in detail, examining how raw sensor data progressively transforms into actionable medical intelligence.

Stage 1: Image Acquisition and Formation

Data Sources

Visual data originates from sensors: cameras for visible light imaging or specialized medical imaging devices (CT scanners, MRI machines, ultrasound transducers).

Physical Factors

Pixel intensities depend on illumination (in cameras), tissue properties (density, water content), scanner settings (voltage, timing), and sensor response characteristics.

Modality Differences

Different medical modalities produce distinct intensity distributions and noise characteristics, requiring modality-specific processing approaches.

Stage 2: Preprocessing

Signal Conditioning for Robust Analysis

Preprocessing reduces irrelevant variation and prepares data for subsequent analysis stages. This critical step strongly affects the quality of downstream segmentation and classification in medical workflows.

Denoising

Remove sensor noise and artifacts while preserving clinically relevant structures

Normalization

Standardize intensity ranges across different scanners and acquisition protocols

Enhancement

Improve contrast and visibility of diagnostically important features

Alignment

Resize, register, and align images to standard coordinate systems

Stage 3: Feature Extraction and Representation

Raw pixels are often too sensitive to noise and minor variations to serve as reliable inputs for decision-making. Instead, we compute representations that capture stable structural and textural properties of the image.

These feature representations form the foundation for robust segmentation, classification, and interpretation in medical imaging systems.

Key Feature Types:

- **Edges and Gradients:** Critical for detecting boundaries and characterizing shape (important for tumor margins, organ boundaries)
- **Texture Descriptors:** Capture local patterns essential for tissue characterization (distinguishing healthy vs diseased tissue)
- **Color Descriptors:** Important in endoscopy and dermatology imaging; less relevant in grayscale CT/MRI
- **Wavelets and Gabor Features:** Useful for analyzing frequency and orientation patterns at multiple scales

Why Features Matter

Good feature representations make the difference between systems that work reliably in clinical practice and those that fail when conditions change slightly.

Features should be:

- Discriminative (different for different classes)
- Robust (stable under noise and variation)
- Efficient (computable in reasonable time)

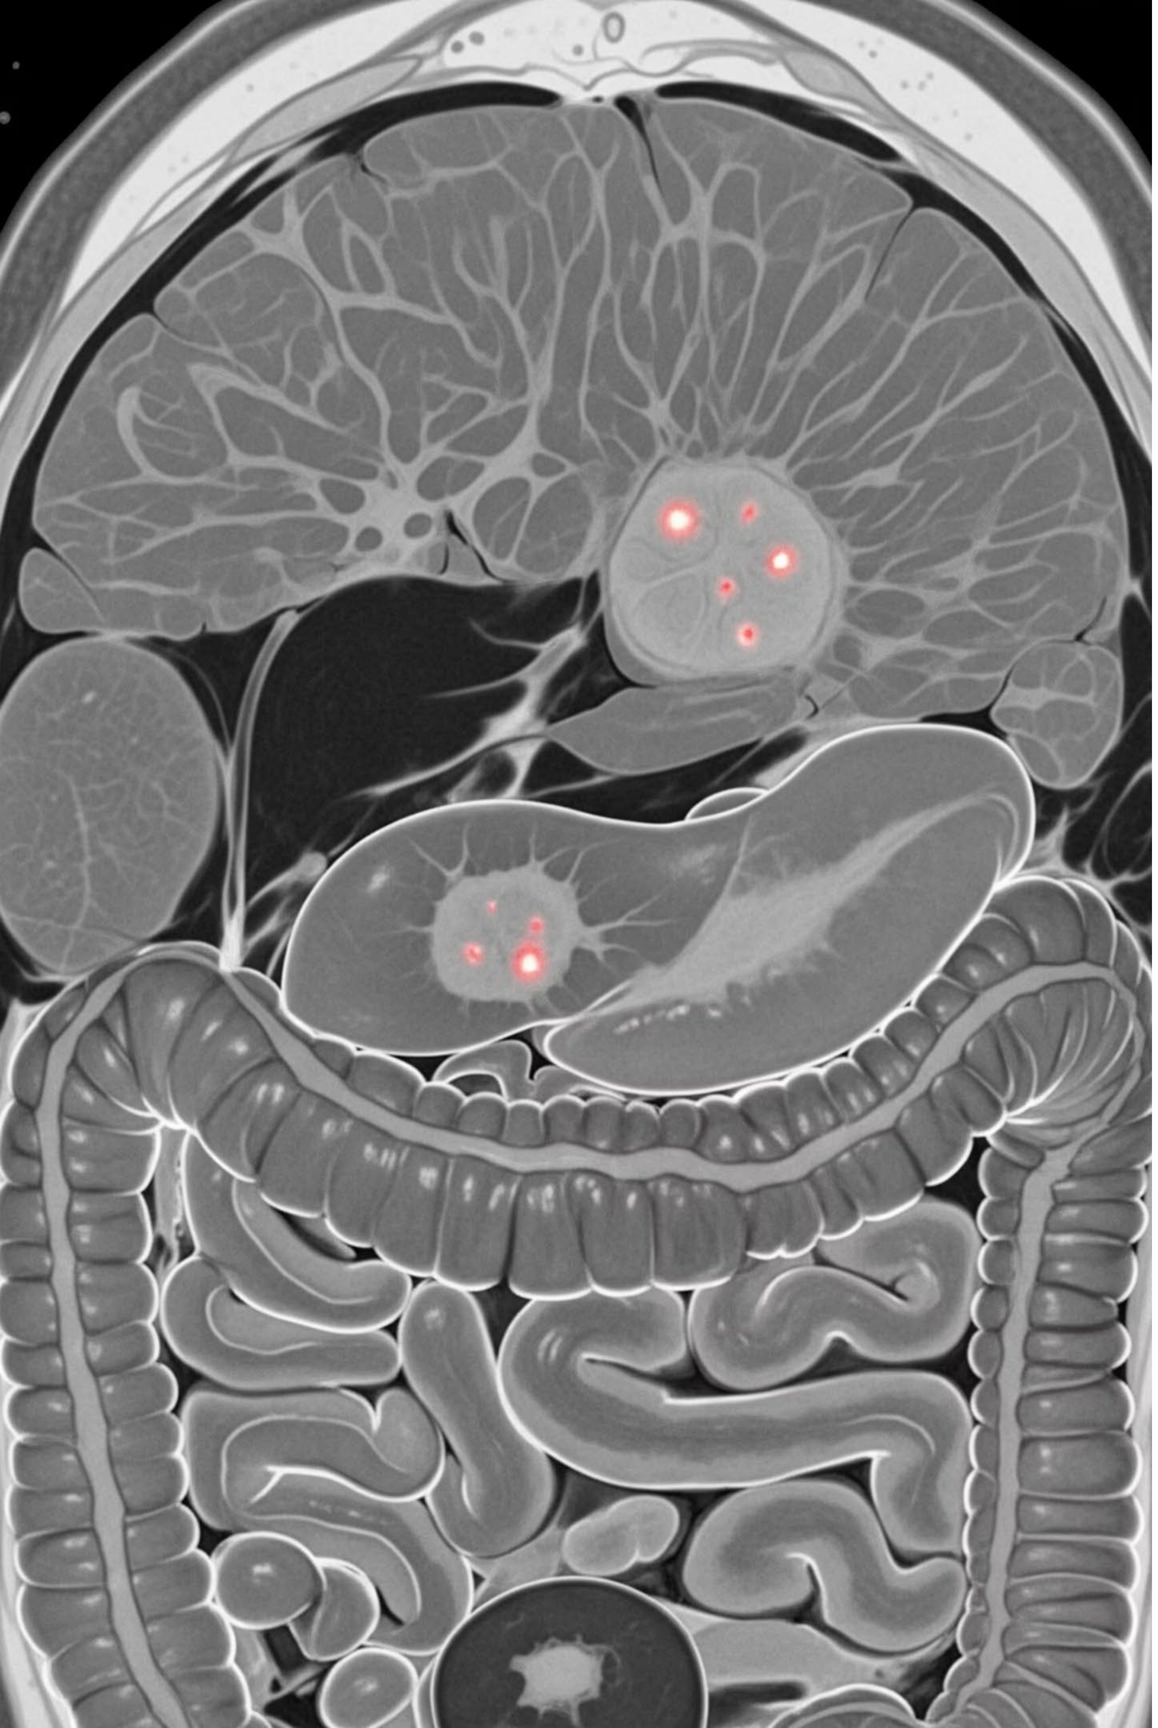

Stage 4: Segmentation

Structure Extraction

Segmentation divides an image into meaningful regions, separating structures of interest from background. This is one of the most important steps in medical image analysis, as it enables quantitative measurement, diagnosis support, and treatment planning.

Organ Segmentation

Delineating anatomical structures like liver, kidneys, heart, or brain regions for volume measurement and surgical planning

Tumor Segmentation

Precisely identifying tumor boundaries for treatment planning, monitoring response to therapy, and measuring growth

Vessel Segmentation

Extracting blood vessel networks for vascular disease assessment, surgical navigation, and perfusion analysis

Stage 5: Recognition and Classification

Where Inference Happens

This stage transforms extracted features and segmented structures into clinical decisions and diagnostic outputs. It represents the culmination of the vision pipeline, where visual data becomes actionable medical intelligence.

Classification

Assign categorical labels to images or regions: normal vs abnormal, disease type, severity grade, tissue class.

Detection

Find and localize objects of interest: lesion locations, anatomical landmarks, surgical instruments.

Recognition

Identify specific entities: patient identification through biometrics, face recognition for access control.

Interpretation

Extract higher-level meaning: severity scoring, disease progression assessment, treatment response evaluation.

Stage 6: Decision Output

Intelligent Medical Systems Perspective

The final stage delivers reliable, interpretable, and actionable outputs crucial for clinical workflows and patient care.

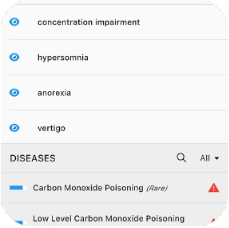

Diagnostic Decision Support

AI-assisted diagnosis recommendations, differential diagnoses, confidence scores, and supporting evidence.

Measurements and Reports

Quantitative measurements (volumes, distances, densities), structured reports, and trend analysis.

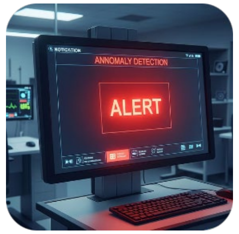

Alerts and Anomaly Detection

Automated critical finding detection, urgent case flagging, quality alerts, and real-time monitoring.

Tracking and Robotics

Patient movement monitoring, surgical instrument tracking, robotic assistance, and navigation systems.

PRACTICAL IMPLEMENTATION

Connection to Laboratory Work

Python Implementation

Practical work in this course will be implemented using Python, the dominant language for computer vision and machine learning in both research and industry.

Core Tools and Libraries:

- **NumPy:** Foundation for matrix operations and numerical computation, treating images as multi-dimensional arrays
- **OpenCV:** Comprehensive library for image and video processing, providing optimized implementations of classical computer vision algorithms
- **Machine Learning Libraries:** Later in the course, we'll use scikit-learn, TensorFlow, or PyTorch for classification and deep learning inference

Laboratory Exercises Overview

Throughout the course, students will implement key computer vision concepts in Python, building practical skills that translate directly to intelligent medical systems development.

01

Image Representation

Representing images as NumPy arrays, understanding pixel indexing, color channels, and data types

02

Filtering Operations

Applying filters through convolution and correlation, implementing smoothing, sharpening, and custom kernels

03

Edge Detection

Computing gradients and detecting edges using Sobel, Prewitt, and Canny edge detection algorithms

04

Segmentation Methods

Implementing thresholding techniques, clustering algorithms (K-means), and region-based segmentation approaches

05

Feature Extraction

Computing descriptors, texture filters, histogram features, and dimensionality reduction with PCA

06

Classification and Inference

Implementing K-Nearest Neighbors, Naive Bayes, and Artificial Neural Networks, progressing to modern machine learning and deep learning approaches

HOMEWORK ASSIGNMENT

Homework 1: Medical Vision Application Analysis

Assignment Requirements

Write 1–2 pages analyzing one Intelligent Medical Systems application of computer vision.

Submission Format: Word or PDF document

Choose one application area:

- Medical image diagnosis
- Biometric identification
- Security anomaly detection
- Patient tracking systems
- Robotic vision for surgery
- Optical character recognition (OCR)

Required Components:

1. **Problem Definition:** Clearly state what decision the system must make and why it matters clinically
2. **Input Data:** Specify the imaging modality or data type (X-ray, CT, MRI, ultrasound, video, photographs, etc.)
3. **Proposed Pipeline:** Describe 5–7 processing steps from acquisition through preprocessing, feature extraction, segmentation (if needed), inference/classification, to final output
4. **Key Challenges:** Identify at least 3 significant challenges (noise, artifacts, variability, domain shift, bias, class imbalance, etc.)
5. **Evaluation Metrics:** Specify appropriate metrics such as accuracy, precision/recall, F1-score, Dice coefficient, IoU for segmentation, or other relevant measures

Key Takeaways

Vision Transforms Data to Decisions

Computer vision extracts meaningful information from images and uses it for inference and decision-making, going beyond simple image enhancement to produce actionable intelligence.

Medical Imaging is Uniquely Challenging

Healthcare applications face noise, variability, class imbalance, subtle patterns, and domain shift—challenges that require sophisticated algorithms and careful validation.

Systematic Pipelines Enable Success

Understanding the complete pipeline from acquisition through preprocessing, feature extraction, segmentation, and classification to decision output is essential for building reliable systems.

Practical Skills Through Python

Laboratory work with NumPy and OpenCV will build hands-on experience implementing the concepts and algorithms that power intelligent medical systems.

Looking Ahead

This introduction has laid the foundation for understanding how computer vision transforms medical images into intelligent decisions. As we progress through the course, we'll dive deeper into each pipeline stage, implementing algorithms and building practical systems.

The journey from pixels to diagnosis requires mastering both theoretical concepts and practical implementation skills. Through lectures and laboratory work, you'll develop the expertise needed to contribute to the next generation of intelligent medical systems.

Computer vision is revolutionizing healthcare, enabling earlier disease detection, more precise treatment planning, and improved patient outcomes. Your work in this course prepares you to be part of this transformation.

Thank you

Google Classroom