AI in Computer Vision

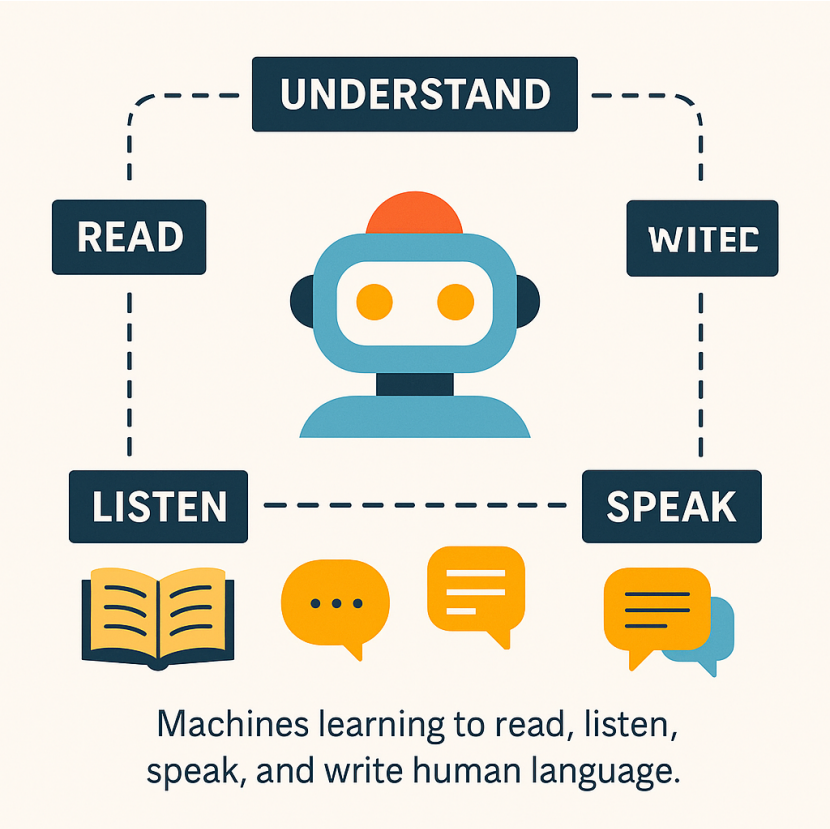

"Machines learning to see, understand, and act on what they perceive."

Learning Objectives

- Define Computer Vision and explain its role in Artificial Intelligence.
- Understand key AI models like CNNs, Transfer Learning, and Object

Detection frameworks.

- Recognize real-world applications in healthcare, surveillance, autonomous systems, and retail.
- Build a simple image classifier using Teachable Machine or similar tools.
- Evaluate the performance and limitations of AI-based vision models.
- Reflect on ethical concerns like privacy, bias, and trust in AI decisions.

1. What is Computer Vision?

Computer Vision (CV) is a field of Artificial Intelligence (AI) that enables computers to **see, interpret, and understand visual data** such as images and videos.

Analogy: Imagine a toddler learning to recognize cats and dogs. They look at many pictures, hear corrections, and over time learn what features define each. CV systems do something similar—just with math instead of neurons

Core Tasks:

- **Image Classification**: Assign a label to an image (e.g., "cat").
- **Object Detection**: Find and label multiple objects in a single image.
- **Facial Recognition**: Match faces to known identities.
- **Image Segmentation**: Identify object boundaries within an image.

Key Technique: Convolutional Neural Networks (CNNs)

CNNs use filters to detect patterns in images—starting from edges and progressing to complex shapes.

2. Applications in Healthcare

CV is revolutionizing medical analysis by acting like a super-powered doctor assistant

- **X-ray/MRI Analysis**: Automatically highlight abnormalities.

- **Tumor Detection**: Spot cancerous growths in early stages.
- **Surgical Assistance**: Real-time visual guidance during surgery.
- **Remote Monitoring**: Track patient health via cameras.

Story Example: Imagine a rural clinic with limited doctors. A CV-based app scans a chest X-ray and alerts the doctor if signs of pneumonia appear. Lives saved without needing an expert on site

3. Applications in Surveillance

Computer Vision adds eyes and brains to security cameras

- **Face Detection and Recognition**: Identify people in crowds.
- **Intrusion Detection**: Alert on unusual activity.
- **Person Tracking**: Follow individuals across multiple camera feeds.

4. Other Real-World Applications

- **Retail:**
  - Inventory checks using camera feeds.
  - Self-checkout using item recognition.
- **Manufacturing:**
  - Detect product defects on the assembly line.
- **Agriculture:**
  - Monitor crop health from drone images.
  - Identify weeds automatically.
- **Social Media:**
  - AR filters (like Instagram dog ears).
  - Tag friends in photos automatically.

5. Hands-on Session: Teachable Machine

- Used Google's **Teachable Machine** for image classification.
- Trained a model in real-time using the webcam.
- Exported the model for use in real projects (e.g., websites, apps).

Teachable Machine makes AI feel like magic—just show it examples and boom, it learns!

6. Model Evaluation and Limitations

Performance Measures:

- **Accuracy**: Overall correctness.
- **Precision and Recall**: For specific classes like "tumor".
- **Confusion Matrix**: Shows correct vs. incorrect predictions.

Limitations:

- Requires large labeled datasets.
- Models can be fooled by noise or slight changes.
- Poor generalization to new lighting or angles.

7. Ethical Concerns in AI Vision

Privacy: Cameras in public and private spaces raise consent issues. **Bias:** Face recognition may work better on certain skin tones. **Trust:** Decisions made by AI (e.g., medical scans) must be explainable and fair.

Reflective Story: Imagine a CV system refusing to recognize you because your face wasn't part of its training data. That's how algorithmic bias feels—and it must be addressed