

# Applications of Artificial Intelligence

NLP · Image Processing · Robotics · Vision · Navigation · Ambient Intelligence

Dr. Gopal Chandra Jana

Dept. of CSE, Sharda University

<https://www.gcjana.in/courses/shardauniversity/2501/CSE472/>

April 30, 2026

# Course Outline

- 1 AI: Present & Future
- 2 Natural Language Processing
- 3 Image Processing & Computer Vision
- 4 Robotics — Hardware
- 5 Robotics — Vision
- 6 Robotics — Navigation
- 7 Ambient Intelligence
- 8 Emerging Applications & Ethics
- 9 Summary & Resources

# What is Artificial Intelligence?

## Definition

AI is the simulation of human intelligence in machines — including **learning, reasoning, problem-solving, perception, and language understanding**.

### Core Capabilities:

- **Perception** — Vision, Speech, Text
- **Cognition** — Reasoning, Planning
- **Learning** — Adaptation, Generalization
- **Action** — Robotics, Automation
- **Communication** — NLP, Dialogue

### AI Subfields:

- Machine Learning (ML)
- Deep Learning (DL)
- Natural Language Processing
- Computer Vision
- Robotics & Autonomous Systems
- Expert Systems

# AI Timeline: Key Milestones

Year	Milestone
1950	Alan Turing proposes the Turing Test for machine intelligence
1956	Dartmouth Conference — term “Artificial Intelligence” coined
1997	IBM Deep Blue defeats world chess champion Kasparov
2011	IBM Watson wins Jeopardy!; Apple launches Siri
2012	AlexNet wins ImageNet — deep learning revolution begins
2016	AlphaGo defeats world Go champion Lee Sedol
2020	GPT-3 launched with 175 billion parameters
2022	ChatGPT reaches 100 million users in 60 days
2023	GPT-4, Gemini, Claude 3 — multimodal AI era

# AI Present & Future Outlook

## Current State

- **Narrow AI (ANI)** dominates today
- LLMs: GPT-4, Gemini, Claude
- Diffusion models for image/video
- RLHF for alignment
- Edge AI and TinyML

## Near Future (2025–2030)

- AGI research milestones
- Autonomous vehicles mainstream
- AI co-pilots in every profession
- Personalized AI-driven medicine

## AI Market 2024–2030

- 2024 Global AI market: **\$196 Billion**
- Projected 2030: **\$1.8 Trillion**
- 90% of enterprises plan AI adoption
- 97 million new AI jobs created

## Long-term Horizon (2030+)

- Artificial General Intelligence (AGI)
- Brain-computer interface AI
- Fully autonomous AI agents
- AI-managed smart cities
- Quantum AI integration

## What is NLP?

NLP enables computers to **understand, interpret, and generate** human language. It is the engine behind chatbots, translators, search engines, and voice assistants.

### Core NLP Tasks:

- Text Classification & Sentiment Analysis
- Named Entity Recognition (NER)
- Machine Translation
- Question Answering
- Text Summarization
- Conversational AI & Dialogue

### Key Technologies:

- Transformer Architecture (2017)
- BERT, GPT-4, T5, LLaMA
- Word Embeddings (Word2Vec, GloVe)
- Attention Mechanisms
- Transfer Learning & Fine-tuning
- Retrieval Augmented Generation (RAG)

# NLP Case Study 1: ChatGPT & Large Language Models

## OpenAI ChatGPT

- **Architecture:** GPT-4 Transformer-based LLM
- **Training:** RLHF on massive internet data
- **Parameters:** Estimated 1.7 Trillion

### Applications:

- Code generation and debugging (GitHub Copilot)
- Essay writing and content creation
- Customer service automation
- Legal document drafting
- Medical literature summarization

## Real-World Impact

- **100M users** in 60 days
- Used by **80%** of Fortune 500
- Legal contracts at \$0.001/word vs \$30/hr
- Powers Microsoft Copilot in Office 365

## Paradigm Shift

LLMs move AI from task-specific models to **general-purpose language engines** applicable across all domains without retraining.

# NLP Case Study 2: Machine Translation

## Google Neural Machine Translation (GNMT)

Uses deep Transformer encoder-decoder. Upgraded from LSTM in 2017, dramatically improving translation quality across all language pairs.

### Technical Pipeline:

- 1 Source sentence tokenization
- 2 Contextual embedding via Transformer encoder
- 3 Cross-lingual attention alignment
- 4 Target language generation
- 5 Beam search for best output

### Scale and Impact

- Supports **133 languages**
- **500 million** daily active users
- 100 billion words translated/day
- 55% error reduction vs phrase-based models
- Real-time spoken conversation translation

# NLP Case Study 3: Sentiment Analysis in Finance

## Bloomberg NLP & FinBERT

Financial institutions use NLP to analyze:

- Earnings call transcripts
- Real-time news article sentiment
- Social media signals (Reddit, Twitter/X)
- SEC filings and annual reports

### Workflow:

- 1 Scrape real-time news and social feeds
- 2 FinBERT classifies sentiment (Positive/Negative/Neutral)
- 3 Sentiment score feeds trading algorithm
- 4 Positions adjusted in milliseconds

## Business Outcomes

- Hedge funds using NLP: **3–8% alpha**
- J.P. Morgan COIN: reads **12,000 loan agreements** in seconds
- Bloomberg: 325,000+ news sources daily
- D.E. Shaw, Two Sigma: NLP quant strategies

## Key Insight

NLP converts **unstructured text** into **quantitative signals** at machine speed.

# NLP Case Study 4: Clinical NLP in Healthcare

## Electronic Health Record Analysis

**BioBERT** and **ClinicalBERT** process millions of clinical notes, discharge summaries, and medical literature for actionable medical insights.

### Clinical NLP Tasks:

- **ICD-10 Coding** — auto diagnosis coding
- **Drug interaction** detection
- **Clinical trial** patient eligibility matching
- **Radiology report** summarization
- **PHI De-identification** for HIPAA compliance

### Mayo Clinic AI

- Reads **6.6 million** patient records
- Identifies undiagnosed hypertrophic cardiomyopathy
- **95% accuracy** vs 79% human baseline
- Found **6,000 previously missed** patients

## WHO Endorsement

WHO supports AI clinical NLP for diagnosis assistance in low-resource settings where specialist physicians are scarce.

# NLP Case Study 5: Voice Assistants

## Major Platforms:

<b>Amazon Alexa</b>	100,000+ skills
<b>Google Assistant</b>	Deep web integration
<b>Apple Siri</b>	1.5B device users
<b>Microsoft Copilot</b>	Office 365 integration

## System Pipeline:

- 1 **ASR** — Automatic Speech Recognition
- 2 **NLU** — Natural Language Understanding
- 3 **DM** — Dialogue Management
- 4 **NLG** — Natural Language Generation
- 5 **TTS** — Text-to-Speech Synthesis

## Market Reach

- 570 million smart speakers worldwide (2024)
- Voice shopping: \$40 billion annually
- Healthcare voice: appointment booking, med reminders
- Vital accessibility tool for visually impaired users

## Emerging: Conversational AI

Modern systems use end-to-end neural pipelines with **contextual memory, multi-turn dialogue,** and emotional intelligence.

## Computer Vision

Enables machines to **interpret visual information** from images and videos using deep convolutional and transformer-based neural networks.

### Core Vision Tasks:

- Image Classification
- Object Detection and Localization
- Semantic and Instance Segmentation
- Face Recognition and Verification
- OCR and Document Analysis
- Image Generation (GANs, Diffusion)
- Depth Estimation and 3D Reconstruction

### Landmark Models:

- **AlexNet (2012)** — ignited deep learning
- **ResNet (2015)** — residual skip connections
- **YOLO (2016)** — real-time object detection
- **Vision Transformer/ViT (2020)** — attention-based
- **DALL-E/Stable Diffusion (2021+)** — generative AI
- **SAM (Meta, 2023)** — Segment Anything Model

# Vision Case Study 1: Medical Imaging AI

## DeepMind: Breast Cancer Detection

- Trained on **28,953 mammograms** (UK and USA)
- Modified ResNet with attention
- Outperforms **6 radiologists** head-to-head
- Reduces false positives: **5.7%**
- Reduces false negatives: **9.4%**

## Other FDA-Cleared Diagnostic AI:

- **IDx-DR** — diabetic retinopathy detection
- **PathAI** — pathology slide AI analysis
- **Viz.ai** — stroke detection in CT scans
- **Aidoc** — radiology workflow triage AI

## Healthcare AI Scale

- Radiologist shortage: 1 per 62,000 in India
- AI screens **10,000 images/hour**
- AI healthcare market: **\$150B by 2026**

## COVID-19 X-Ray AI

AI models achieved **90%+ accuracy** detecting COVID-19 pneumonia patterns in chest X-rays within weeks of the outbreak.

# Vision Case Study 2: Facial Recognition

## Technology

CNNs map faces to high-dimensional embedding space where the same person's faces cluster together, enabling fast identification and verification.

### Key Systems:

- **Apple Face ID** — 1-in-1M false accept rate; 30,000 IR dots
- **DeepFace (Meta)** — 97.35% accuracy on LFW benchmark
- **AWS Rekognition** — real-time video streams
- **Clearview AI** — 10B photo database for law enforcement

### Application Domains:

- Airport and border security (CBP, TSA)
- Smartphone biometric authentication
- Missing persons identification
- Banking KYC verification
- Retail customer analytics
- Attendance and access control

## Ethical Concern

Higher error rates for darker skin tones and women (MIT/NIST studies). Fairness, accountability, and transparency in facial AI is a critical research area.

# Vision Case Study 3: Precision Agriculture

## John Deere See & Spray

- **36 cameras** processing 10 fps in real-time
- YOLO-based AI distinguishes weeds from crops
- Sprays herbicide **only on weeds**
- Processing in milliseconds per plant
- Deployed on **2 million+** acres in USA

## Related Agri-Vision Systems:

- **DJI Agras** drones — NDVI crop health mapping
- **Taranis** — satellite-based disease detection
- **Plantix** — smartphone crop disease diagnosis

## Business Outcomes

- Herbicide use reduced by **77%**
- Saves **\$25 per acre** on chemicals
- Reduces environmental runoff
- ROI achieved in 2–3 growing seasons

## Global Impact

AI-powered agriculture will help feed **9.7 billion people** by 2050 while reducing food waste by 30%.

# Vision Case Study 4: Industrial Inspection

## Manufacturing Quality Control

- **Landing AI** (Andrew Ng) — visual inspection platform
- Detects micro-cracks, dents, and surface defects
- **99.9% accuracy** on PCB inspection lines
- BMW: AI cameras for car body surface defects
- Samsung: wafer defect detection at nanometer scale

## Infrastructure Inspection

- Drone + CV for bridge crack detection
- Power line and wind turbine inspection
- Pipeline corrosion monitoring
- Rail track anomaly detection
- Road surface mapping via dashcams

## Skydio Autonomous Drone Inspection

Skydio drones autonomously navigate structures, with AI identifying structural anomalies **10x faster** than manual inspection. Used by US Army for aircraft checks.

# Vision Case Study 5: Generative AI and Image Synthesis

## Generative AI — Creating Visual Content with AI

AI now creates, edits, and enhances photorealistic images and videos from text prompts using diffusion models and GANs.

### Key Generative Systems:

- **DALL-E 3** — OpenAI text-to-image
- **Stable Diffusion** — open-source generation
- **Midjourney** — artistic image synthesis
- **Adobe Firefly** — commercial creative AI
- **Sora (OpenAI)** — text-to-video generation
- **StyleGAN3** — photorealistic face synthesis

### Real-World Applications:

- Film VFX and virtual production
- Fashion and product design prototyping
- Game asset creation (80% cost reduction)
- Architectural visualization
- Medical training data augmentation
- Advertising and marketing content

## Deepfakes Concern

Generative AI raises misinformation risks. Active work on AI watermarking and detection tools (C2PA standard, 2024) is underway.

## What Makes a Robot AI-Powered?

Combines **perception, cognition, and actuation** — sense the environment, reason, and take purposeful physical actions.

### Sensor Hardware:

- **LiDAR** — 3D point cloud mapping
- **Depth Cameras** — Intel RealSense, Azure Kinect
- **IMU** — accelerometer + gyroscope
- **Force/Torque sensors** — precision manipulation
- **Ultrasonic** — proximity sensing
- **Thermal cameras** — heat detection

### AI Compute Hardware:

- **NVIDIA Jetson AGX Orin** — edge AI for robots
- **Qualcomm RB5** — mobile robotics SoC
- **Intel Movidius** — low-power neural inference
- **TPU/NPU** — neural network accelerators
- **FPGAs (Xilinx)** — low-latency real-time control
- **GPU arrays** — training and inference

# Robotics Case Study 1: Boston Dynamics

## Spot — Quadruped Robot

- **Hardware:** 12 joints, 17 sensors, stereo cameras, LiDAR
- **Payload:** 14 kg; Battery: 90 min runtime
- **AI:** Model Predictive Control + RL locomotion
- Climbs stairs, handles rough terrain
- Remote and fully autonomous modes

## Real Deployments:

- Chevron offshore oil rig inspections
- Construction site progress tracking
- Nuclear plant monitoring (radiation zones)
- COVID-19 patient monitoring in hospitals

## Atlas Humanoid Robot

- Height: 1.5m, Weight: 89 kg
- 28 actuated joints, electric actuation (2024)
- Performs backflips, parkour, tool use
- Learns new behaviors via RL

## Key Technology

**Model Predictive Control (MPC)** computes optimal joint trajectories in real-time for dynamic balance and locomotion.

# Robotics Case Study 2: Surgical Robots

## Intuitive Surgical: da Vinci Surgical System

Allows surgeons to perform minimally invasive procedures with **sub-millimeter accuracy** and AI-assisted safety features.

### Hardware Components:

- **Surgeon Console** — 3D viewer + hand controls
- **Patient Cart** — 4 robotic arms
- **Vision Cart** — 4K 3D endoscopic camera
- **EndoWrist instruments** — 540° articulation
- Haptic force-feedback for tissue sensing

### AI Features:

- Tremor cancellation (12Hz human tremor filtered)
- Motion scaling (10:1 ratio for micro-precision)
- Fluorescence imaging for blood vessel mapping

### Clinical Impact

- **10 million+** procedures globally
- Hospital stay: **50% reduction**
- Blood loss: **65% reduction**

# Robotics Case Study 3: Warehouse Automation

## Amazon Robotics Fleet

- **750,000+** robots in fulfillment centers
- **Proteus** — first fully autonomous mobile robot
- **Sparrow** — AI grasps items from totes
- **Cardinal** — AI-powered package sorting arm
- Fleet AI coordinates thousands of robots simultaneously

## Drive Unit Specs:

- 145 kg capacity, 1.5 m/s speed
- Magnetic floor navigation + AI path planning
- Li-ion battery, 2-hour runtime, auto-recharging

## Operational Gains

- Fulfillment: **15 min** → **3 min**
- Storage density: **+50%**
- Operating cost: **-20%** per unit
- **1.3 million** orders/day

## Also Notable

**Ocado (UK)** — world's most automated grocery warehouse; 3,000 robots coordinate via 4G network in the "hive."

# Robotics Case Study 4: Collaborative Robots (Cobots)

## Cobots — Working Alongside Humans

Cobots use AI to detect human presence and adapt behavior, enabling safe shared workspaces without expensive safety caging.

### Key Manufacturers:

- **Universal Robots UR5/UR10** — 6-DOF arms
- **ABB YuMi** — dual-arm human-safe robot
- **FANUC CR Series** — green cobots
- **Rethink Sawyer** — force-limited 7-DOF

### AI Safety Features:

- Force-torque collision detection (stops in <0.5s)
- Computer vision human tracking
- Speed and separation monitoring (ISO/TS

## BMW Plant Regensburg

- 7 YuMi cobots alongside human workers
- No safety cages required
- Tasks: screw fastening, component alignment
- **20% productivity gain**
- Reduced ergonomic strain for workers

## Market Outlook

Cobot market: **\$12.6 billion by 2030**, driven by SME manufacturing automation globally

## Robot Vision vs. Computer Vision

Robot vision is **active and embodied** — robots must interpret visuals to make real-time physical decisions under resource and latency constraints.

### Vision System Pipeline:

- 1 **Acquisition** — cameras, LiDAR, depth sensors
- 2 **Preprocessing** — calibration, undistortion
- 3 **Feature Extraction** — CNNs, ViTs, PointNets
- 4 **Scene Understanding** — segmentation, 3D maps
- 5 **Decision Making** — grasping, navigation

### Key Vision Capabilities:

- **6-DOF Pose Estimation** — object orientation
- **SLAM** — Simultaneous Localization and Mapping
- **Optical Flow** — motion detection
- **Stereo Vision** — passive depth perception
- **Instance Segmentation** — object separation
- **Gaze Estimation** — human-robot interaction

# Vision Case Study 1: Robotic Bin Picking

## Challenge: Unstructured Bin Picking

Industrial robots must identify randomly-oriented parts and grasp them — requiring 3D localization, grasp pose generation, and sub-second inference.

## AI Approach (GraspNet):

- 1 RGB-D camera captures bin scene
- 2 PointNet++ processes 3D point cloud
- 3 Neural network predicts grasp quality scores
- 4 Robot executes highest-scoring grasp
- 5 **Success rate: 93%** on novel objects

## Fanuc & NVIDIA Partnership

- **NVIDIA Isaac Sim** generates synthetic training data
- Domain randomization bridges sim-to-real gap
- No manual labeling needed
- Deployed at Toyota, GM, Volkswagen
- ROI achieved in **8 months**

## Innovation: Sim-to-Real

Synthetic data generation removes the need for expensive real-world training data labeling for robotic grasping tasks.

# Vision Case Study 2: Visual SLAM

## Simultaneous Localization and Mapping (SLAM)

Robots **build a map of an unknown environment** and **localize themselves** within it simultaneously, without any prior knowledge.

### Visual SLAM Pipeline:

- 1 Camera captures sequence of keyframes
- 2 Feature detection (ORB, SuperPoint)
- 3 Pose estimation via Bundle Adjustment
- 4 Loop closure detection to correct drift
- 5 Global map optimization

### Deep SLAM Approaches:

- CNN-based monocular depth estimation
- DROID-SLAM, MonoDepth2
- End-to-end ViT-based SLAM

### iRobot Roomba i7+

- iAdapt 3.0 uses full Visual SLAM
- Creates precise home floorplan
- Remembers room names via app
- Avoids obstacles using AI
- **40 million homes** worldwide

### NASA Perseverance Rover (Mars)

- Stereo cameras + AI terrain mapping

# Vision Case Study 3: Swarm Robot Coordination

## Distributed Vision in Robot Swarms

Each robot perceives locally; collective behavior emerges through AI coordination protocols, enabling tasks impossible for single robots.

### Vision-Based Coordination:

- ArUco marker-based relative localization
- Optical flow for collision avoidance
- Formation control with shared visual goals
- Distributed SLAM with map fusion

### AI Techniques:

- Multi-agent reinforcement learning (MARL)
- Shared neural policy networks
- Attention-based inter-agent communication

## Intel Drone Light Show (2022)

- **3,281 drones** simultaneously coordinated
- Each: RGB camera + GPS + radio
- Centralized AI assigns collision-free paths
- Sets Guinness World Record
- Used in Olympic ceremony productions

## Defense: DARPA OFFSET

Swarms of 250 drones and ground robots coordinate using shared visual AI for complex urban reconnaissance

## The Navigation Problem

Navigate robot from  $A$  to  $B$  while **avoiding obstacles**, adapting dynamically, and operating safely in real-time environments.

### Navigation Stack:

- 1 **Localization** — Where am I? (AMCL, EKF)
- 2 **Mapping** — What surrounds me? (OctoMap)
- 3 **Path Planning** — Route? ( $A^*$ , RRT, Dijkstra)
- 4 **Motion Control** — Execute? (PID, MPC)
- 5 **Obstacle Avoidance** — Avoid? (DWA, VFH)

### AI Navigation Approaches:

- **Deep RL** — end-to-end navigation policies
- **Imitation Learning** — learn from human demos
- **Neural SLAM** — learning-based mapping
- **Hierarchical Planning** — semantic + geometric
- **Social Force Models** — human-aware navigation
- **Transformer Planners** — waypoint prediction

# Navigation Case Study 1: Waymo Autonomous Vehicles

## Waymo One: Level 4 Autonomous Driving

### Sensor Suite:

- 29 cameras (360° full coverage)
- 5 LiDARs (short and long range)
- 6 radars for all-weather operation
- Redundant GPS + centimeter HD maps

### AI Navigation System:

- **Perception:** 3D object detection + prediction
- **Prediction:** Forecast all road agents' behavior
- **Planning:** Multi-scenario decision tree search
- **Behavior:** RL-trained driving policy

## Operational Stats 2024

- **20+ million miles** fully autonomous
- Active in Phoenix, San Francisco, LA
- **50,000 paid rides/week**
- Zero at-fault fatalities
- 85% fewer injury crashes vs humans

## Key Innovation

Uses **simulation** for 99% of training data — billions of virtual miles cover rare edge cases impossible to collect in real life.

# Navigation Case Study 2: Delivery Robots

## Last-Mile Autonomous Delivery

Ground robots solve expensive last-mile logistics using AI navigation on sidewalks and campuses.

### Key Players:

- **Starship Technologies** — 6-wheel sidewalk robot
- **Nuro R2** — road delivery (NHTSA approved)
- **FedEx SameDay Bot** — last-hop autonomous delivery
- **Kiwibot** — university campus delivery

### Navigation Challenges Solved:

- Pedestrian prediction and social navigation
- Traffic-controlled intersection crossing
- GPS-denied indoor transitions
- All-weather operation (rain, snow, uneven

### Starship Technologies

- **5 million deliveries** completed
- Operating in **25+ countries**
- Deployed on campuses worldwide
- 10x less energy than conventional delivery
- Cost: **\$1.50/delivery** vs \$8 by human

### Regulatory Progress

Many US states enacted **Personal Delivery Device (PDD) laws** enabling legal sidewalk robot operation

# Navigation Case Study 3: Aerial and Underwater Robots

## Autonomous UAV Navigation

### Skydio X2 Autonomy Engine:

- 6 × 4K cameras for 360° obstacle detection
- Real-time obstacle avoidance at 36 km/h
- AI subject tracking through complex environments
- Bridge, tower, and powerline inspection

## DARPA Subterranean Challenge

Autonomous robots navigate unknown underground tunnels creating 3D maps **without GPS**, using LiDAR-SLAM and AI.

## Autonomous Underwater Vehicles (AUVs)

### Kongsberg HUGIN AUV:

- Acoustic SLAM for 3D underwater mapping
- AI dynamic obstacle avoidance
- Behavior tree mission planning
- Pipeline inspection, seabed surveys, exploration

## Ocean Infinity

- **8 AUVs** operating simultaneously
- 10,000 km seabed mapped in 80 days
- Used in MH370 search operations

# Navigation Case Study 4: Autonomous Agricultural Robots

## Monarch Tractor: AI Autonomous Farming

- 6 cameras + LiDAR for full-field navigation
- GPS RTK for centimeter-level row positioning
- AI row detection and end-of-row turns
- Obstacle detection: humans, animals, equipment
- Fully autonomous or supervised operation modes

## Other Precision Agri-Robots:

- **Small Robot Company (Tom)** — per-plant farming
- **Burro** — autonomous vineyard transport
- **Tortuga AgTech** — AI strawberry harvesting
- **FFRobotics** — apple and citrus picking

## Business Impact

- Labor cost reduction: **40–60%**
- 24/7 operation during harvest season
- Reduced soil compaction vs heavy tractors
- Better yield via precise input application

## UN FAO

Agriculture employs **1 billion people**. AI robots will transform farming roles, augmenting human productivity rather than simply eliminating jobs.

## What is Ambient Intelligence (Aml)?

Vision of **intelligent environments** that are sensitive, responsive, and adaptive to people — environments that think, sense, and act proactively on your behalf.

### Core Pillars of Aml:

- **Ubiquitous Computing** — intelligence embedded everywhere
- **Context Awareness** — sensing user and environment state
- **Personalization** — adapting to individual preferences
- **Anticipation** — predicting needs proactively
- **Natural Interaction** — voice, gesture, gaze, touch

### Enabling Technologies:

- Internet of Things (IoT) sensor networks
- Edge AI and TinyML on microcontrollers
- 5G ultra-low latency connectivity
- Federated Learning (privacy-preserving AI)
- Wearable and embedded computing
- Smart actuators and transducers

# Aml Case Study 1: Smart Home Ecosystems

## Google Nest Ecosystem

- **Nest Thermostat** — learns schedule; saves 15% energy
- **Nest Hub** — proactive suggestions via sensor fusion
- **Nest Doorbell** — person/animal/vehicle AI detection
- AI builds behavioral model of household over time

## Amazon Alexa Smart Home:

- Context-triggered routines (time, location, occupancy)
- **Hunches**: AI proactively suggests automations
- **Guard**: AI break-in detection via microphones
- Sidewalk mesh network for neighborhood connectivity

## Apple HomeKit

- On-device AI for privacy protection
- Adaptive lighting by time and activity
- HomeKit Secure Video with ML
- **Matter** standard: cross-brand coordination

## Market Data

Smart home market: **\$174 billion by 2025**.  
Average smart home saves **\$1,200/year** in energy and utility costs.

# Aml Case Study 2: Smart Healthcare Environments

## AI-Powered Hospital and Home Care Monitoring

Aml creates **proactive care environments** monitoring patients continuously and intervening before conditions deteriorate.

### Hospital Smart Room (Caregility):

- Bed-exit prediction: pressure + depth camera AI
- Fall risk alert: **40% reduction in falls**
- Non-contact vitals monitoring via radar
- Delirium detection from speech AI analysis
- Smart IV: AI detects infiltration automatically

### Remote Patient Monitoring (RPM):

- **Current Health** — wearable AI monitoring patch
- Continuous SpO2, HR, respiratory rate
- ML predicts sepsis **6 hours early**
- **30% readmission reduction** in heart failure

## UK NHS Virtual Wards

AI-equipped home monitoring lets patients recover at home while hospital AI tracks 1,000+ patients simultaneously, escalating care automatically.

# Aml Case Study 3: Smart Cities

## Singapore Smart Nation Initiative

- **110,000 sensors** deployed city-wide
- AI traffic lights: **30% congestion reduction**
- Real-time crowd density AI monitoring
- Full **digital twin** of the city
- Predictive infrastructure maintenance

## DeepMind Green Light (70+ Cities)

- AI optimizes traffic signals in real-time
- Reduces stop-and-go events by **30%**
- Reduces fuel consumption by **10%**
- No new infrastructure required

## Barcelona Smart City

- IoT sensors in every street lamp
- AI parking: **21M litres fuel saved/year**
- Smart irrigation: **25% water savings**
- **\$97M annual** city-wide savings

## Emergency Response AI

Smart city AI predicts and reroutes emergency vehicles via dynamic signal preemption, reducing average response times by **20%** in pilot cities.

# Aml Case Study 4: Smart Education

## Adaptive Learning and Intelligent Tutoring

AI transforms education from one-size-fits-all to **personalized adaptive learning** at massive scale, serving millions simultaneously.

### Key Adaptive Systems:

- **Khanmigo** (Khan Academy) — GPT-4 based AI tutor
- **Duolingo Max** — AI explains errors, adapts difficulty
- **Carnegie Learning MATHia** — mathematics ITS
- **Coursera Coach** — personalized AI career mentor

### Smart Classroom Technologies:

- Emotion recognition for engagement monitoring

### Proven Outcomes

- Carnegie Learning: algebra learning **76% faster**
- Duolingo: AI learners progress **4x faster** vs classroom
- Squirrel AI (China): 1-on-1 AI tutoring at \$3/month serving 3M students
- Closes achievement gaps by up to **30%**

### Future Vision

# Aml Case Study 5: Cashierless Retail

## Amazon Go: Just Walk Out Technology

- Hundreds of cameras + weight sensors + depth cameras
- AI tracks every person and every item in store
- Pick item → added to virtual cart automatically
- Walk out → automatically charged via app
- No queues, no checkout, no cashier

### Technical Pipeline:

- 1 Entry via QR code at turnstile
- 2 Multi-person tracking with re-identification AI
- 3 Shelf item detection via weight + vision fusion
- 4 Transaction reconciliation, receipt to app

## Scale and Expansion

- **40+** **Amazon Go** stores in USA and UK
- Licensed to Whole Foods, Starbucks, Hudson News
- **35 airports** deployed cashierless checkout
- Checkout time: **4 min** → **26 sec**

## Also Notable

**Standard AI** retrofits existing stores using only their **existing camera infrastructure**, dramatically reducing deployment costs.

## AlphaFold 2 — DeepMind (2020)

- Predicted **200+ million protein structures**
- Solved a 50-year biology grand challenge
- Accuracy: **92.4 GDT** (like X-ray crystallography)
- Released freely to the global scientific community
- Described as a “leap for structural biology” by *Nature*

## AI Drug Discovery Pipeline:

- 1 Target identification via NLP literature mining
- 2 Molecular generation via diffusion models/GANs
- 3 ADMET property prediction (toxicity, bioavailability)
- 4 Clinical trial design and patient stratification

## Insilico Medicine Milestone

- AI designed drug for IPF lung disease
- Traditional: 4–5 years, \$100M+
- AI: **18 months, \$2.6M** only
- Entered Phase II clinical trials (2023)

## Projected Impact

AI could cut drug development from **12 years to 3–4 years** and reduce costs by **70%**.

# AI in Finance, Cybersecurity, and Climate

## AI in Finance

- **70%** of stock trades AI-executed
- Mastercard: analyzes **75B transactions/year** in 50ms
- **40% reduction** in fraud false positives
- Lemonade: AI settles claims in **2 seconds**
- Upstart AI: 27% more loans approved, same default rate

## AI in Cybersecurity

- Darktrace: ML baseline per device/user, detects zero-days
- **Autonomous Response** neutralizes threats in seconds

## AI for Climate Change

- DeepMind: **20% UK electricity demand** reduction
- Google Wind: **20%** improvement in wind energy value
- Global Fishing Watch: satellite AI detects illegal fishing
- Wildfire detection: **10 min** earlier warning
- Climate emulators: **1000x faster** than physics models

## Microsoft AI for Earth

\$50M program: 500+ projects tracking 1.2M species, real-time deforestation mapping, and carbon capture optimization

## Why Ethics is Non-Negotiable in AI

AI systems can **perpetuate, amplify, or create new bias and harm** if not carefully designed and governed with human values at the center.

### Key Ethical Challenges:

- **Bias & Fairness** — inequities in training data
- **Privacy** — pervasive surveillance potential
- **Transparency** — black-box decisions
- **Accountability** — who is liable for AI errors?
- **Safety** — edge case AI failures
- **Deepfakes** — synthetic media manipulation
- **Job Displacement** — economic disruption

### Documented AI Failures

- Amazon hiring AI **penalized women's resumes**
- COMPAS: **2x bias** against Black defendants in sentencing
- Facial recognition: **wrongful arrests** (Detroit, 2020)
- GPT models: confident **misinformation** generation
- Tesla Autopilot: under investigation for 17 fatalities

# EU AI Act & The Road to AGI

## EU AI Act (2024) Risk Classification

<b>Unacceptable</b>	Banned: social scoring, subliminal manipulation
<b>High Risk</b>	Medical, justice, employment AI — heavy regulation
<b>Limited Risk</b>	Chatbots — disclosure required
<b>Minimal Risk</b>	Spam filters, AI games

## Responsible AI Principles (IEEE):

- 1 Transparency and Explainability (XAI)
- 2 Fairness and Non-discrimination
- 3 Privacy and Data Protection
- 4 Safety and Reliability

## The Road to AGI

- GPT-4 passes bar exam (**90th percentile**)
- AlphaCode 2 outperforms **85%** of programmers
- Multimodal AI: vision + language + reasoning

## AGI Timeline Predictions

<b>Optimistic</b>	2027–2030 (Kurzweil, Altman)
<b>Moderate</b>	2040–2060 (most researchers)
<b>Skeptical</b>	May never arrive (LeCun, Marcus)

## What AGI Still Lacks

Common sense reasoning, causal understanding, open-ended curiosity, lifelong continuous learning,

# Comprehensive Summary

## Topics Covered:

<b>NLP</b>	Chatbots, translation, sentiment, healthcare
<b>Vision</b>	Medical, agriculture, inspection, generative AI
<b>Robotics HW</b>	Boston Dynamics, surgery, warehouse, cobots
<b>Robot Vision</b>	Bin picking, SLAM, swarm AI
<b>Navigation</b>	Waymo, delivery, AUVs, agri-robots
<b>Ambient AI</b>	Smart home, city, health, education, retail
<b>Finance/Cyber</b>	Fraud, trading, threat detection
<b>Emerging</b>	Drug discovery, climate, entertainment
<b>Ethics</b>	Bias, EU AI Act, AGI outlook

## Five Key Takeaways

- 1 AI is **already pervasive** — not a future technology
- 2 **Deep learning** powers most modern AI
- 3 AI creates **exponential value** with domain expertise
- 4 **Ethics and safety** are non-negotiable
- 5 The **Narrow-to-General AI** transition is the defining challenge of our era

## For AI Practitioners

Master Python, ML frameworks, domain knowledge, ethical reasoning, and maintain a **continuous learning mindset.**

# Recommended Resources

## Courses:

- Andrew Ng — Deep Learning Specialization (Coursera)
- Fast.ai — Practical Deep Learning for Coders
- MIT 6.034 — Artificial Intelligence (OCW)
- Stanford CS231n — Computer Vision
- Stanford CS224n — NLP with Deep Learning

## Books:

- *AI: A Modern Approach* — Russell & Norvig
- *Deep Learning* — Goodfellow, Bengio, Courville
- *Human Compatible* — Stuart Russell
- *The Alignment Problem* — Brian Christian

## Research & News:

- arXiv.org (cs.AI, cs.LG, cs.CV)
- Papers With Code (paperswithcode.com)
- MIT Technology Review
- DeepMind, OpenAI, Anthropic research blogs

## Practice Platforms

- **Kaggle** — competitions and real datasets
- **Hugging Face** — models, datasets, demos
- **Google Colab** — free GPU cloud compute
- **Weights & Biases** — experiment tracking

# Thank you for your attention!

---

Post your doubts:

<https://www.gcjana.in/courses/shardauniversity/2501/CSE472/#Post-doubts>

Dr. Gopal Chandra Jana · Dept. of CSE, Sharda University