Lecture 3

ARTIFICIAL INTELLIGENCE

Introduction to Artificial Intelligence (AI)

Artificial intelligence (AI) refers to the capability of machines or computer programs to perform tasks that require human intelligence, such as learning, problem-solving, and decision-making.

The field of AI has been around for several decades, but recent advances in technology have led to a rapid increase in its capabilities and applications.

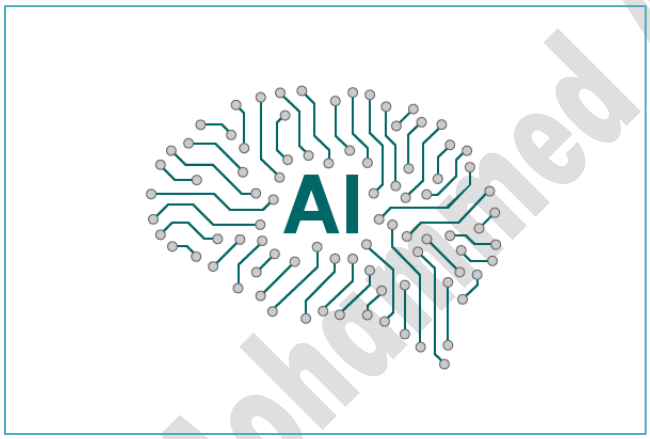

AI is categorized into three types:

1. Narrow AI (Weak AI): Specialized in a single task (e.g., Siri, spam filters).
2. General AI (Strong AI): Hypothetical systems performing any intellectual task humans can do.
3. Super intelligent AI: Future AI that surpasses human intelligence in all aspects (currently speculative).

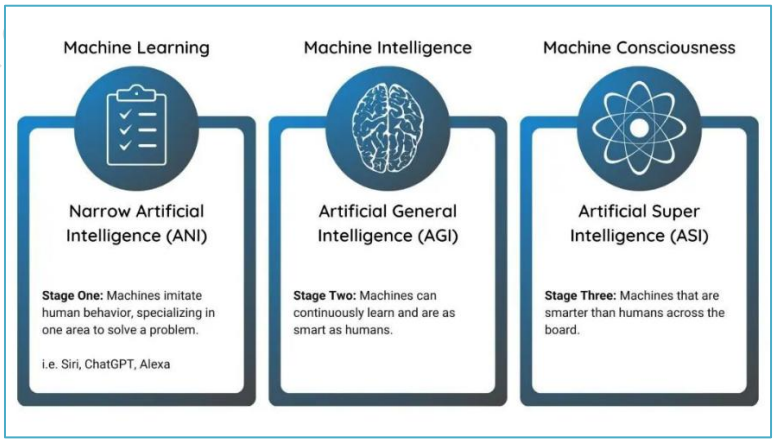

Core AI Capabilities:

- Learning: Machines improve performance over time through data.
- Reasoning: AI systems make decisions based on available information.
- Perception: Analyzing visual and sensory data for interpretation.

History to Artificial Intelligence (AI)

- AI originated in the **1950s** when researchers explored creating machines that could think like humans. ***Alan Turing*** played a key role, introducing the **Turing Test** to ***measure machine intelligence***.
- **In 1955**, Allen Newell and J.C. Shaw were developed "**The Logic Theorist**" that could ***Prove mathematical theorems***.
- **In 1957**, Newell and Herbert Simon were developed "**The General Problem Solver**" This program was able to ***solve a range of problems*** by searching through a set of rules and making logical deductions.
- **The 1960s** saw the rise of **expert systems** like **Dendral (1965)**, were developed by Edward Feigenbaum and Joshua Lederberg in 1965. This system was designed to ***identify the molecular structure of organic compounds***.
- **The 1970s and 1980s** saw further advances in AI, with researchers developing ***new algorithms and techniques for machine learning, natural language processing, and computer vision***.
  - development of the first ***neural network***, which was inspired by the structure of the human brain.
  - creation of ***rule-based expert systems***, which allowed computers to make decisions based on a set of predefined rules.
- **The 1990s–2000s** saw rapid progress in ***deep learning***, ***natural language processing***, and ***autonomous vehicles***.
- **Today,** AI is widely applied, from ***virtual assistants*** **and** ***self-driving cars*** to ***advanced robotics and medical diagnosis tools***.

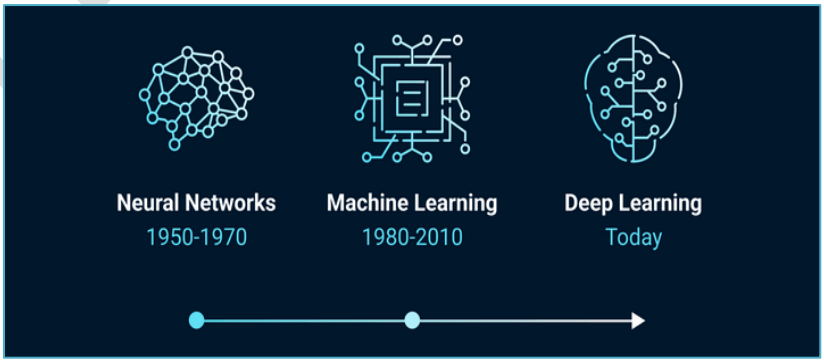

Al Techniques and Approaches

1. Machine Learning (ML)

ML is a subset of artificial intelligence (AI) that involves the use of algorithms and statistical models to enable machines to learn from data and improve their performance on specific tasks.

Types of ML:

1. Supervised Learning: Training a model on a labeled dataset, *where the correct output is known for each input*.
2. Unsupervised Learning: Training a model on an unlabeled dataset *and allowing it to identify patterns and relationships on its own*.
3. Reinforcement Learning: Training a model to make decisions *based on feedback from its environment*.

Applications of ML:

1. Healthcare: Analyze medical images and patient data to improve disease diagnosis and treatment.
2. Finance: Analyze market trends and make investment decisions.
3. Transportation: Optimize traffic flow and improve road safety.
4. Education: Personalize learning and improve student outcomes.

2. Deep Learning

Deep Learning is a subset of machine learning that involves the use of artificial neural networks to enable machines to learn from data and perform complex tasks.

Applications of Deep learning:

- Healthcare: Analyze medical images and identify potential diseases.
- Finance: Analyze market trends and make investment decisions.
- Transportation: Develop self-driving cars.

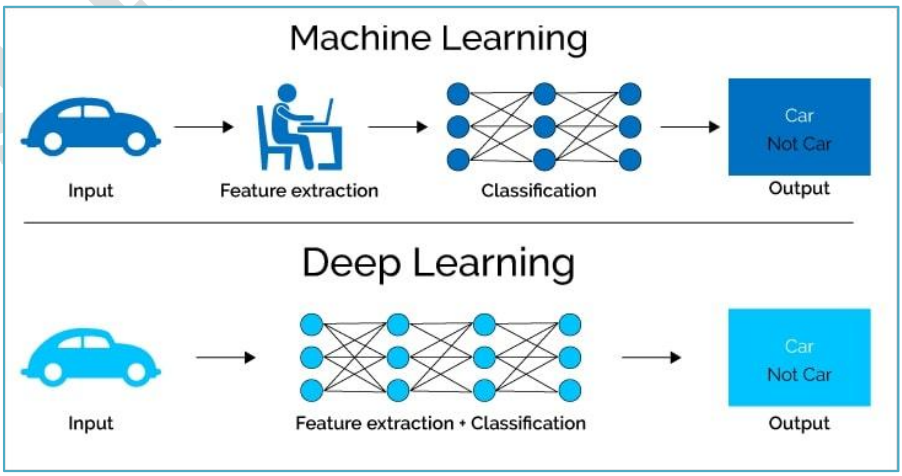

3. Natural Language Processing (NLP)

Natural language processing (NLP) is a subfield of artificial intelligence (AI) that focuses on the interaction between computers and humans using natural language.

Applications of NLP:

- Chatbots: that can understand and respond to natural language queries.
- Sentiment analysis: Analyze social media posts, customer reviews.
- Machine translation: Translate text from one language to another.
- Voice assistants: Voice assistants, such as Siri or Alexa.
- Text summarization: Generate summaries of longer texts, such as news articles.

4. Computer Vision

Computer vision is a field of study that focuses on enabling computers to *understand* and *interpret* visual data from the world around us.

Applications of Computer Vision:

- Self-driving cars: Detect and recognize objects such as pedestrians, other vehicles, and road signs, and make decisions based on that information.
- Facial recognition: Facial recognition systems, which are used for security and surveillance purposes.
- Medical imaging: Analyze and interpret medical images such as X-rays, CT scans, and MRI scans.
- Industrial automation: Industrial automation to monitor and control manufacturing processes.

5. Robotics

The integration of AI into machines that perform physical tasks autonomously.

Applications of Robotics:

- Industrial automation.
- Healthcare (e.g., surgical robots).
- Agriculture (e.g., harvesting robots).

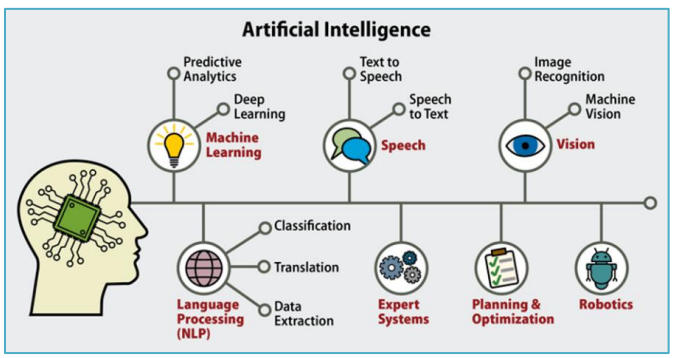

Challenges and Ethical Considerations

Challenges of AI

AI offers numerous benefits but comes with **several challenges:**

1. Data Dependency: AI requires large volumes of high-quality data.
2. Generalization: Ensuring AI models work effectively in diverse real-world scenarios.
3. Energy Consumption: Training large AI models is resource-intensive and raises sustainability concerns.

Ethical of AI

1. Bias and Fairness: Algorithms can reflect societal biases present in training data. Addressing fairness is crucial for trust in AI.
   Example: Discriminatory hiring algorithms or unfair loan approvals.
2. Transparency and Explainability: Explainable AI (XAI) is a growing field that focuses on making AI decision-making processes more transparent and understandable.
3. Privacy: AI relies on collecting and analyzing personal data, raising issues about how data is used and stored.
   Example: Social media platforms using AI to analyze and influence user behavior.
4. Job Displacement: Automation poses risks of job loss, especially in manufacturing, transportation, and customer service.
5. Autonomous Weapons: The use of AI in military applications poses ethical dilemmas, including autonomous weapons.

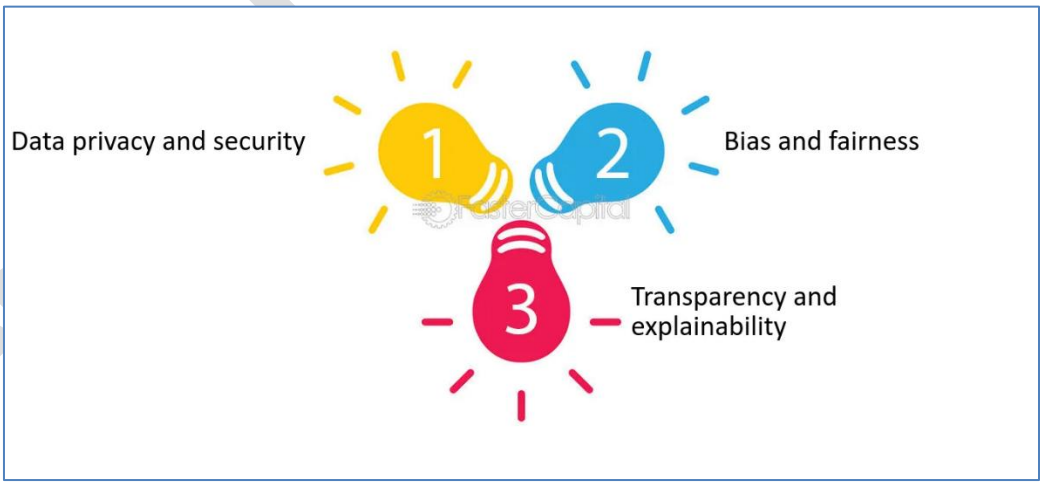

Al in Our Daily Lives

Artificial Intelligence (AI) is deeply embedded in our daily lives, often in ways we don't immediately recognize. From our smartphones to home automation systems, AI helps streamline tasks, improve efficiency, and enhance convenience.

AI in smartphones

- Facial Recognition: Unlock devices securely and quickly.
  Example: Apple's Face ID.
- Photography Enhancements: AI analyzes lighting and composition to improve image quality.
  Example: Google Pixel's computational photography.
- Voice Commands: Allow hands-free operation of devices.
  Example: Google's "OK Google" feature.
- Predictive Text and Smart Typing: AI suggests words and corrects grammar in real time.
  Example: SwiftKey and Gboard keyboards.
- Battery Optimization: AI learns user habits to improve battery life.
  Example: Adaptive Battery in Android devices

AI in virtual assistants

Virtual assistants are AI-driven programs designed to perform tasks, provide information, and interact with users using natural language processing (NLP).

Examples:

- Siri (Apple): A voice-activated assistant for iOS devices.
- Google Assistant: A highly integrated assistant for Android and Google devices.
- Alexa (Amazon): A smart assistant for home automation.

Applications of Al

1. Education

- Personalized Learning: AI adapts content to individual student needs (e.g., adaptive learning platforms like Coursera and Duolingo).
- Automated Grading: AI assists in grading multiple-choice and written assignments, saving educators time.
- Tutoring & Assistance: AI chatbots and virtual assistants provide 24/7 learning support.

2. Healthcare

- Medical Diagnosis: AI detects diseases through imaging (e.g., cancer detection in radiology).
- Predictive Analytics: AI forecasts disease outbreaks and patient risks.
- Robot-Assisted Surgery: AI enhances precision in complex surgeries.

3. Finance

- Fraud Detection: AI identifies suspicious transactions and prevents financial fraud.
- Algorithmic Trading: AI-driven trading bots analyze market trends and execute trades at high speeds.
- Customer Support: AI chatbots handle banking inquiries and financial planning.

4. Transportation

- Autonomous Vehicles: AI powers self-driving cars (e.g., Tesla, Waymo).
- Traffic Management: AI optimizes traffic flow and reduces congestion using smart traffic lights.
- Predictive Maintenance: AI forecasts vehicle maintenance needs, reducing breakdowns.

5. Marketing and Advertising

- Personalized Recommendations: AI tailors product recommendations (e.g., Amazon).
- Chatbots & Customer Engagement: AI chatbots provide instant customer support.
- Targeted Advertising: AI analyzes user behavior to optimize ad placement and audience targeting.

The Future of Al

Future trends in AI

- Democratization of AI: Cloud-based tools and open-source platforms make AI more accessible.
- AI for Sustainability: AI helps combat climate change by optimizing energy use and predicting disasters.
- AI in Everyday Life: From smart assistants to wearable tech, AI is seamlessly integrating into daily routines.
- Ethical & Responsible AI: Focus on fairness, transparency, and accountability in AI development.

Recent research in AI

- Explainable AI (XAI): Enhancing transparency and trust in AI decision-making.
- Federated Learning: AI models train on decentralized data while preserving privacy.
- Neuromorphic Computing: Brain-inspired AI hardware for energy efficiency.
- AI in Scientific Discovery: Accelerating research in medicine, materials, and astrophysics.

Emerging technologies

- Generative AI: AI creates content (e.g., text, images, music).
  - Large Language Models (LLMs) – AI like GPT-4 generates human-like text.
  - Diffusion Models – AI creates high-quality images from text prompts.
- Multimodal AI: Integrating text, images, and audio for enhanced AI capabilities.
- Edge AI: AI processing at the source for faster and more efficient applications.
- Quantum AI: Leveraging quantum computing for solving complex AI problems.