

Fundamentals and applications of AI



Dinesh Gupta

Group Leader


Translational Bioinformatics Group
ICGEB, New Delhi

<http://www.icgeb.org/dinesh-gupta/>

<http://bioinfo.icgeb.res.in/>

What is Artificial Intelligence?

Artificial intelligence is a specialized field within computer science dedicated to creating systems that can replicate human intelligence and problem-solving capabilities. These systems are designed to think, learn, and adapt in ways that mirror human cognitive processes. A more contemporary definition describes AI as a broad branch of computer science focused on building machines that can learn from experience, make informed decisions, and perform complex tasks at a level comparable to human intelligence. This encompasses everything from simple pattern recognition to sophisticated reasoning and creative problem-solving. Modern AI systems go beyond simple rule-based programming—they can identify patterns in vast datasets, improve their performance over time, and even make predictions about future outcomes based on historical information.



Types of AI

Narrow AI

Narrow AI, also known as weak AI, specializes in performing specific tasks with exceptional efficiency. Examples include language translation systems, facial recognition software, and recommendation algorithms. These systems excel within their defined scope but cannot transfer their expertise to other domains.

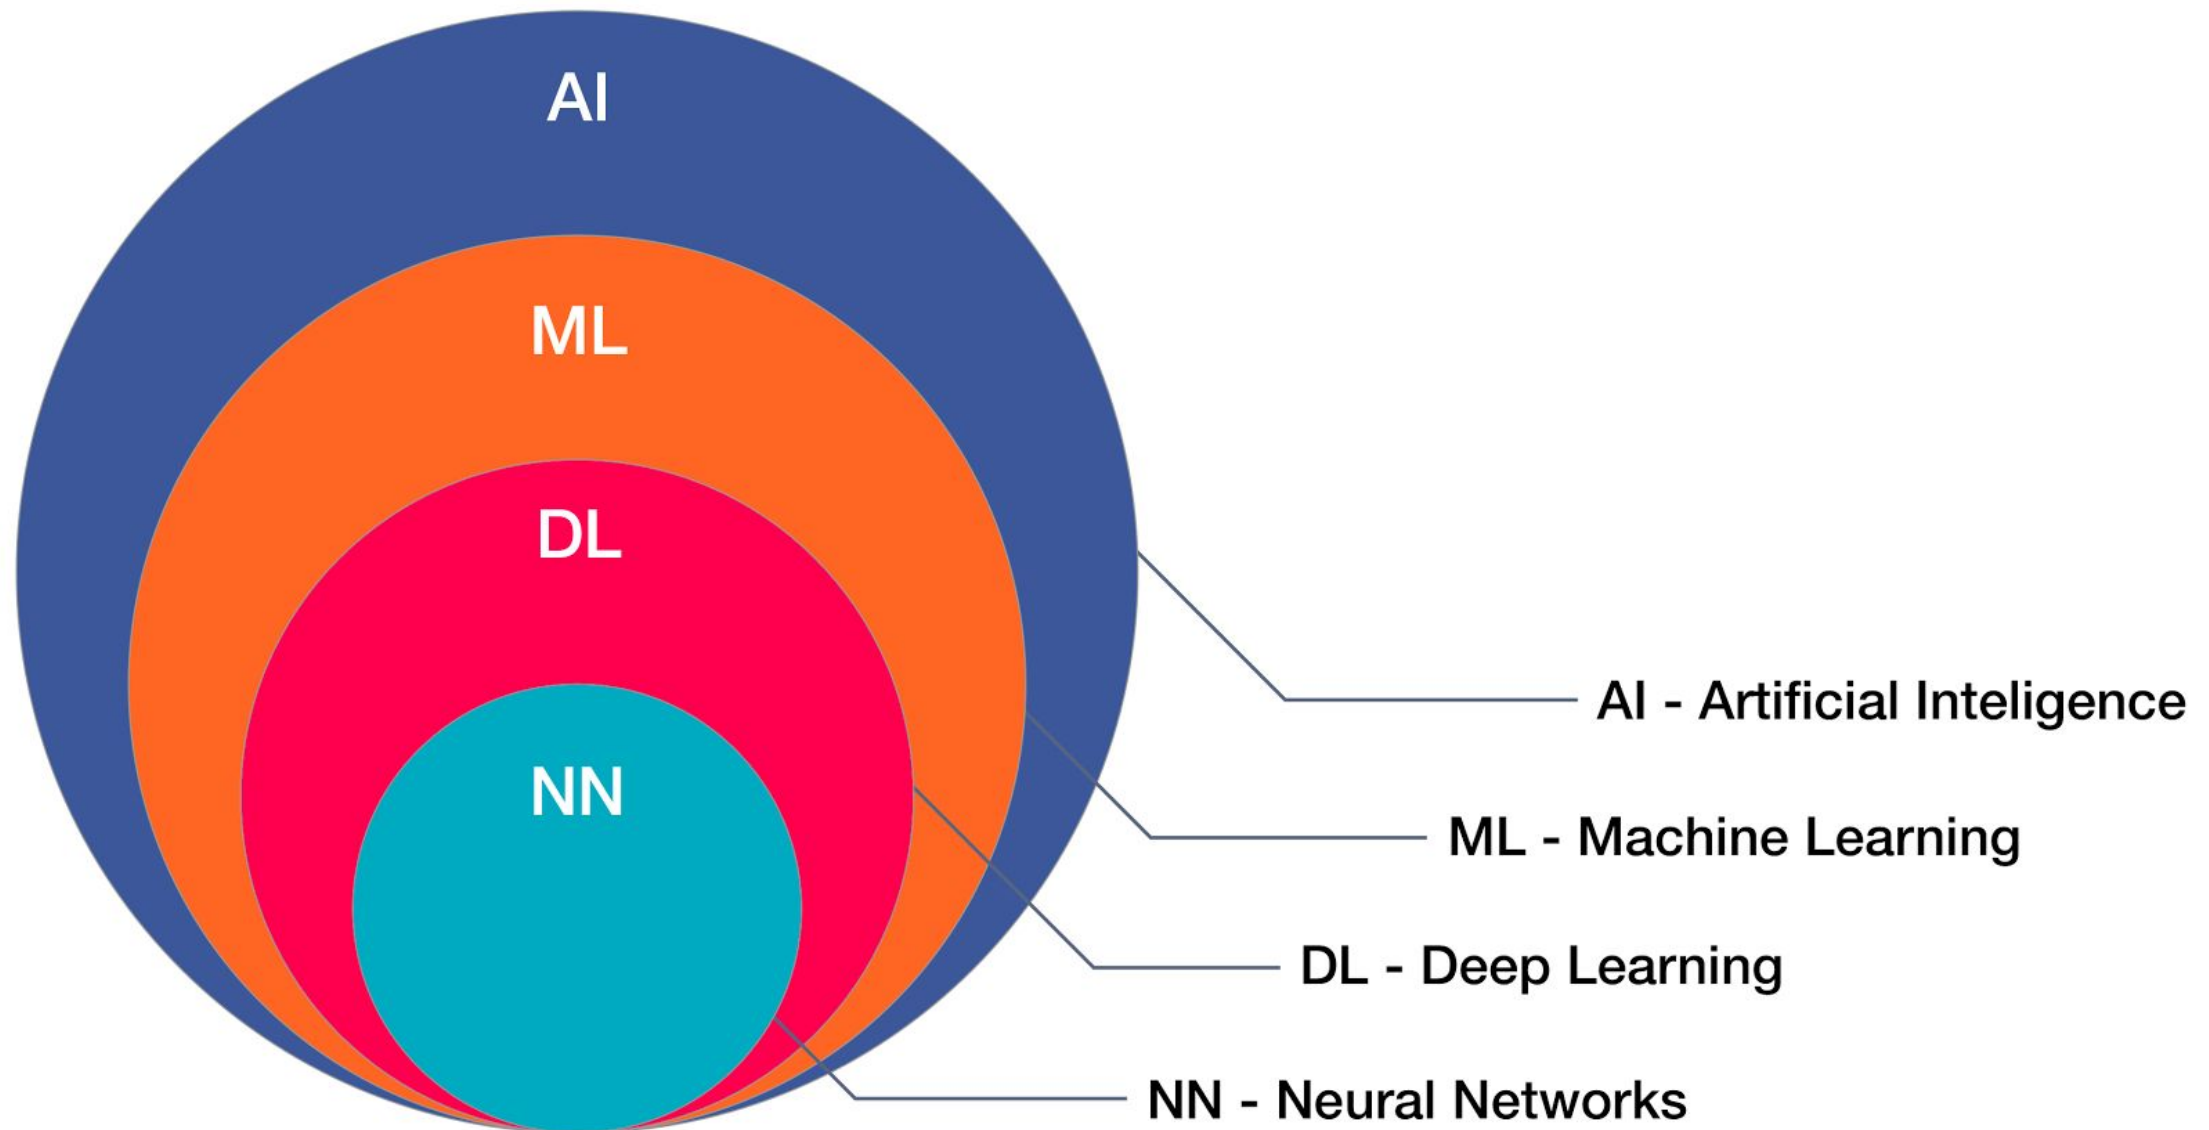
General AI

General AI, or strong AI, possesses human-like cognitive abilities that enable it to understand context, learn from diverse experiences, and apply knowledge flexibly across various domains. This type of AI can reason abstractly, plan strategically, and solve novel problems.

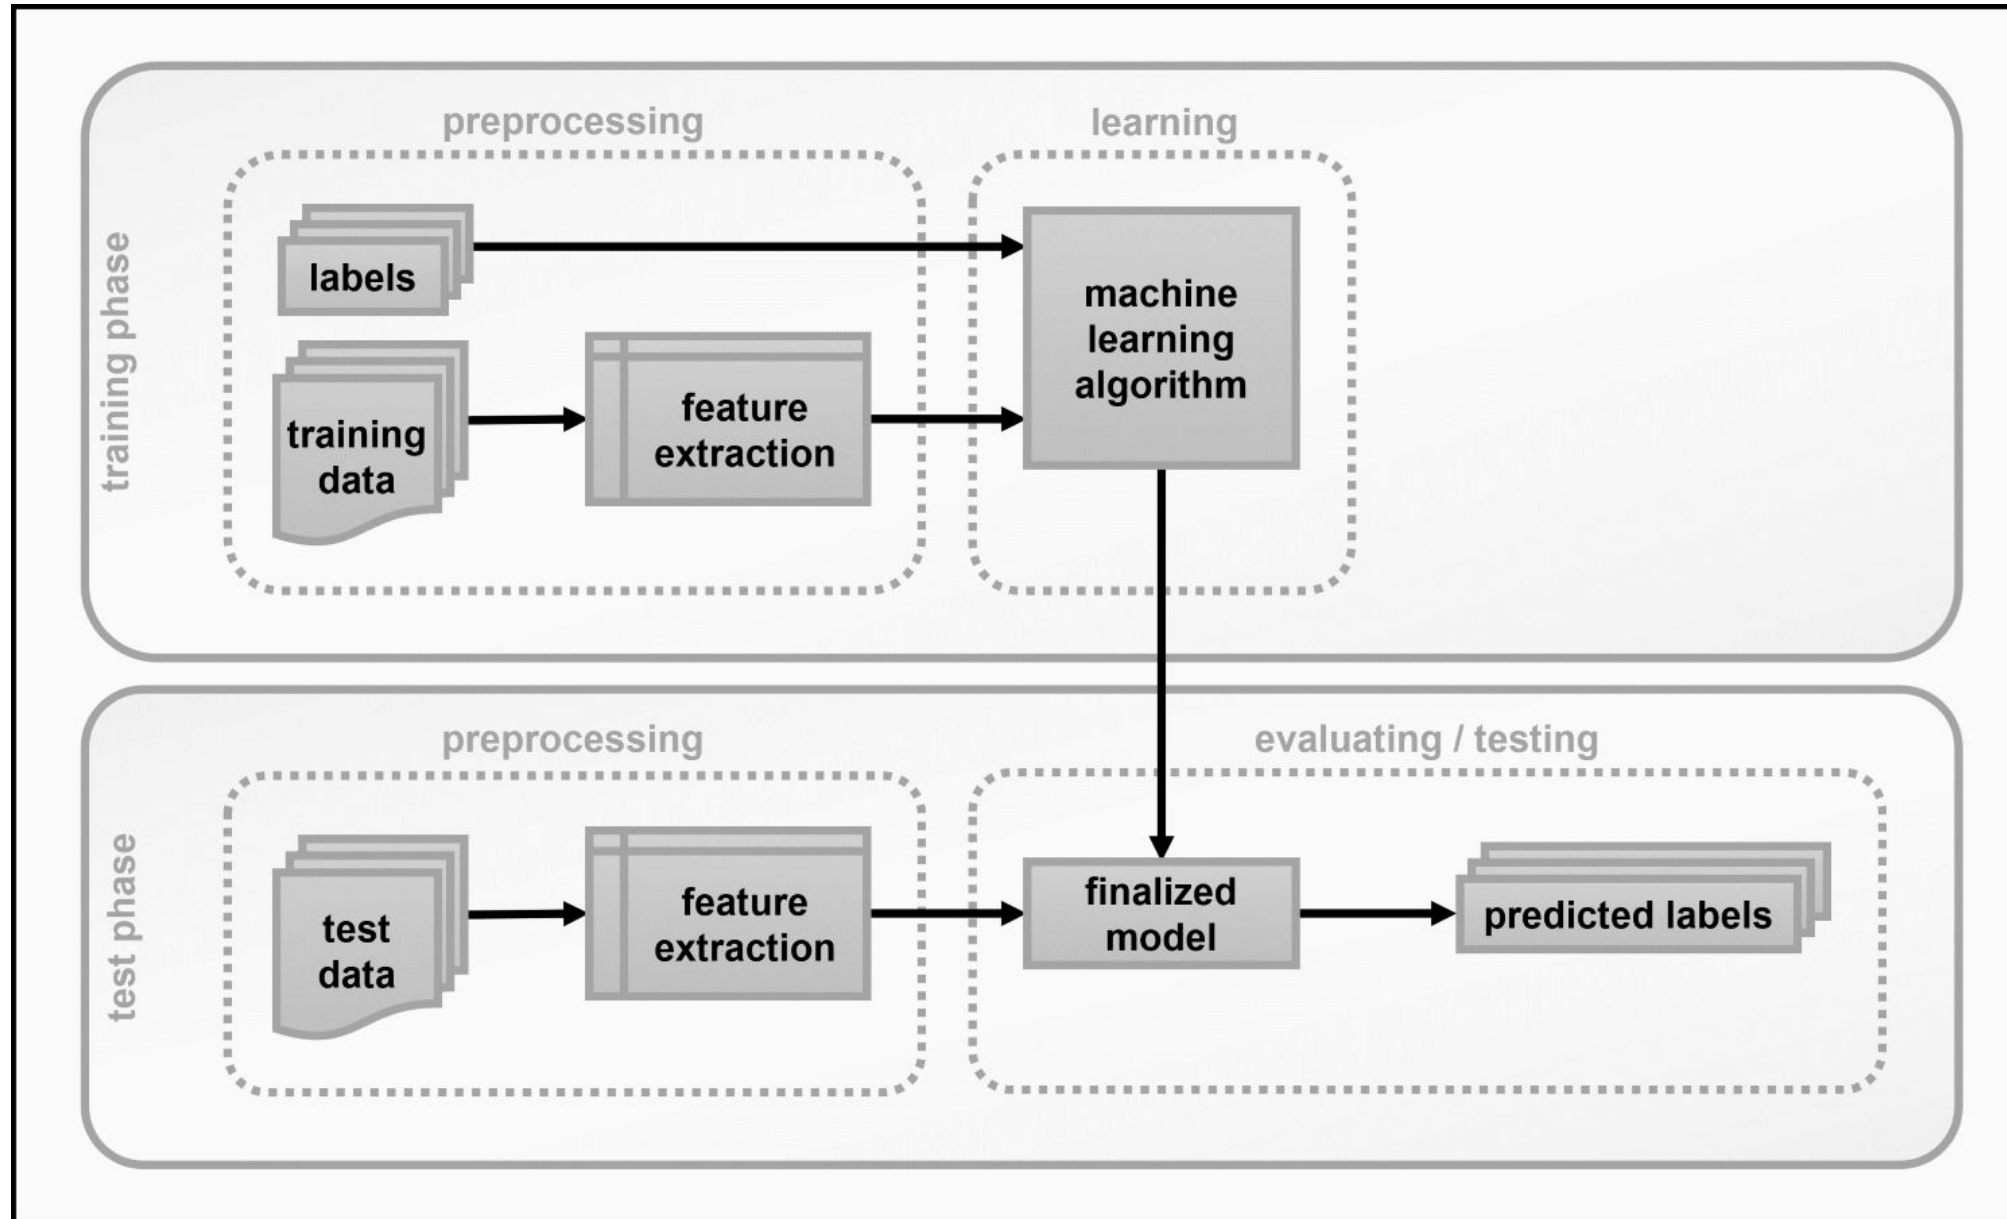
Superintelligence

Superintelligence represents AI that surpasses human intelligence in virtually every domain, including scientific creativity, general wisdom, and social skills. With the potential to solve extraordinarily complex problems beyond human capabilities, this hypothetical level of AI raises profound philosophical and practical questions about control and safety.

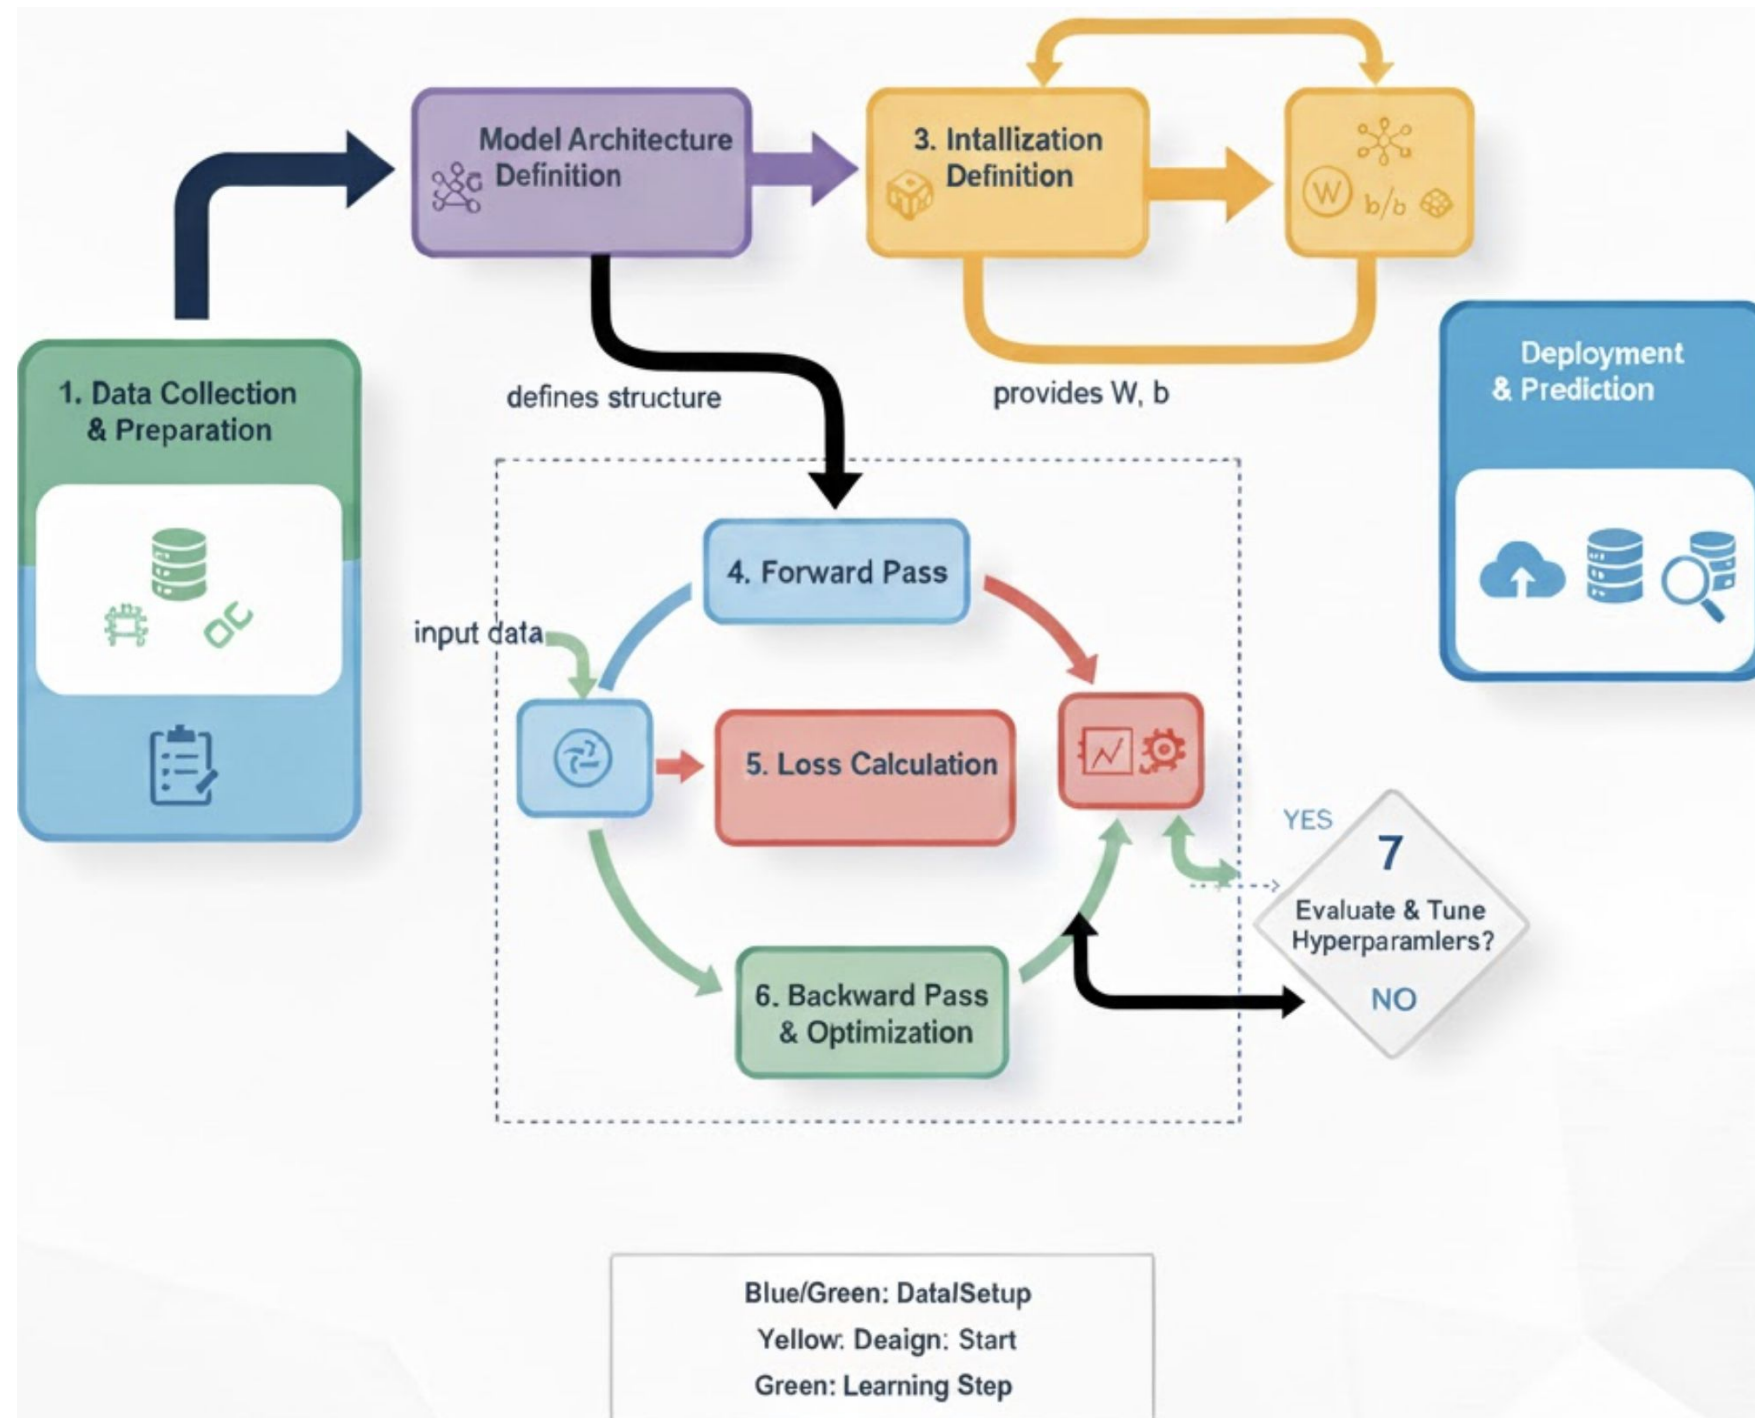
AI and subsets



Typical Machine learning workflow



Deep learning workflow



Applications of AI in Various Industries



AI in Healthcare

AI is revolutionizing healthcare through predictive analytics that can forecast disease outbreaks, personalized treatment plans tailored to individual genetic profiles, and diagnostic systems that detect conditions earlier than traditional methods. Machine learning algorithms analyze medical imaging with remarkable accuracy, while natural language processing helps extract insights from medical records.



AI in Transportation

AI powers autonomous vehicles that navigate complex traffic scenarios, optimizes transportation logistics to reduce costs and delivery times, and enhances traffic management systems in smart cities. Machine learning algorithms predict maintenance needs, while route optimization systems reduce fuel consumption and improve efficiency across transportation networks.



AI in Finance

In the financial sector, AI enhances decision-making through predictive analytics that forecast market trends, sophisticated fraud detection systems that identify suspicious transactions in real-time, and automated trading systems that execute trades at optimal moments. AI also powers credit scoring models and personalized financial advice platforms.



AI in Entertainment

Entertainment industries leverage AI for personalized content recommendations on streaming platforms, automated video editing, music composition, and even scriptwriting assistance. AI analyzes viewing patterns to predict what content will engage specific audiences, creating highly tailored entertainment experiences that keep users engaged.

History of Artificial Intelligence: Stages and High-Level Overview

1950-1979: The Foundation Era

Important foundational questions about machine intelligence were first posed during this period. LISP, the first programming language specifically designed for AI research, was created. Early researchers developed basic AI programs capable of answering questions, sorting data, and playing simple games. The Stanford Cart became one of the first examples of an autonomous vehicle, successfully navigating obstacles without human intervention.

1987-1993: The AI Winter

A period of stagnation in practical AI applications led to what researchers call the "AI winter." Both private and public investors halted their investments, convinced that AI was merely a passing fad without real-world applications. Despite reduced funding, dedicated AI researchers continued their work, laying groundwork for future breakthroughs.

2012-Present: The Age of General Intelligence

We've entered the modern era where AI has become ubiquitous in daily life. Billions of dollars continue flowing into AI research and development. Contemporary AI systems play complex games at superhuman levels, design their own communication languages, recognize images with remarkable accuracy, and generate human-like text and art.

1

2

1980-1987: The First AI Boom

Significant breakthroughs in AI research led to increased public and private funding, which in turn accelerated more discoveries. The first dedicated AI conferences were organized, bringing together researchers from around the world. Commercial AI products entered the market for the first time, with expert systems finding applications in business and industry.

3

4

1993-2011: Renaissance and Practical Applications

New breakthroughs reignited interest and funding in AI research. IBM's Deep Blue became the first AI system to defeat a reigning world chess champion. Consumer robotics arrived with products like the Roomba. NASA successfully landed autonomous rovers on Mars, demonstrating AI's capabilities in extreme environments. Apple released Siri, introducing millions to digital assistants.

5

The History of Artificial Intelligence: Important Pioneering Developments (1950-1956)

1950: The Turing Test

Alan Turing published his groundbreaking paper "Computing Machinery and Intelligence," which introduced the Imitation Game—now known as the Turing Test. This test proposed a practical criterion for determining whether a machine could be considered intelligent by evaluating its ability to exhibit behavior indistinguishable from a human.

1955: Birth of the Term "AI"

John McCarthy organized a workshop at Dartmouth College focused on "artificial intelligence," marking the first official use of this term. This event is widely considered the birth of AI as a formal field of study, bringing together pioneers who would shape the discipline for decades.

1952: Machine Learning Emerges

Computer scientist Arthur Samuel developed a revolutionary program to play checkers, marking the first instance of a computer learning a game independently through experience rather than explicit programming. This laid the foundation for modern machine learning approaches.

❏ The original definition of AI, coined by John McCarthy in 1955, was deliberately broad: *"The science and engineering of making intelligent machines."* This expansive vision continues to guide the field today.

Further Developments in AI History

1958: LISP Programming Language

Longva McCarthy created LISP (List Processing), the first programming language specifically designed for AI research. Its unique ability to manipulate symbolic information made it ideal for AI applications. LISP remains influential in AI development today, demonstrating remarkable longevity for a programming language.

1965: Expert Systems Arrive

Feigenbaum and Joshua Lederberg created the first "expert system," a form of AI programmed to replicate the decision-making abilities of human specialists. These systems captured expert knowledge in specific domains, making specialized expertise more widely accessible.

1979: Autonomous Navigation

The Stanley Cart, originally created by James L. Adams in 1961, successfully navigated a chair-filled room without human guidance in 1979. This achievement represented a major milestone in autonomous vehicle development, demonstrating that machines could perceive and respond to their environment.

1959: "Machine Learning" Coined

Arthur Samuel introduced the term "machine learning" during a speech about teaching computers to play chess at levels exceeding their human programmers. This concept of computers improving through experience became foundational to modern AI.

1966: The First Chatbot

Joseph Weizenbaum created ELIZA, the first chatbot (originally called a "chatterbot"), which simulated a psychotherapist using natural language processing. ELIZA demonstrated that machines could engage in seemingly intelligent conversations, though its understanding was superficial.

1961: First Industrial Robot

Unimate, the first industrial robot, began working on a General Motors assembly line in New Jersey. It performed dangerous tasks like transporting die casings and welding car parts—jobs deemed too hazardous for human workers—marking AI's entry into practical industrial applications.

1968: Deep Learning Foundations

Soviet mathematician Alexey Ivakhnenko published "Group Method of Data Handling," proposing an innovative approach to AI that would later evolve into what we now recognize as deep learning—the technology powering today's most sophisticated AI systems.

Types of AI Algorithms



Supervised Learning

Supervised learning algorithms train on clearly labeled datasets, where each input is paired with its correct output. The algorithm learns to map inputs to outputs by studying these examples, similar to a student learning from a teacher. Common applications include image classification, spam detection, and price prediction. The algorithm's performance can be directly measured by comparing its predictions to the known labels.



Unsupervised Learning


Unsupervised learning works with unlabeled data, forcing algorithms to discover hidden patterns and structures independently. Without explicit guidance, these systems identify clusters, anomalies, and relationships within data. This approach is valuable for customer segmentation, anomaly detection, and discovering unknown patterns in complex datasets where labeling would be impractical or impossible.



Reinforcement Learning

Reinforcement learning algorithms learn through trial and error, receiving feedback in the form of rewards or penalties based on their actions. Like training a pet, the system learns which behaviors lead to positive outcomes. This approach powers game-playing AI, robotics control systems, and autonomous vehicles, where the algorithm must learn optimal strategies through experience.

Examples of AI Uses in Daily Life



AI touches nearly every aspect of modern life:

- **Digital assistants** like Siri, Alexa, and Google Assistant use natural language processing to understand voice commands and respond intelligently
- **Self-driving cars** leverage computer vision and sensor fusion to navigate roads safely without human intervention
- **Navigation apps** such as Google Maps and Waze use AI to predict traffic patterns and suggest optimal routes in real-time
- **Social media algorithms** curate personalized content feeds based on your interests, engagement history, and social connections
- **Targeted advertisements** use machine learning to predict which products you're most likely to purchase based on your browsing behavior

These applications demonstrate how AI has seamlessly integrated into daily routines, often working invisibly in the background to enhance convenience, efficiency, and user experience.

Advantages of Artificial Intelligence

Eliminating Human

Error Systems can perform tasks with consistent precision, eliminating errors caused by fatigue, distraction, or oversight. In critical applications like medical diagnosis or financial calculations, this reliability can be life-saving and

Saving Workers from Risky Tasks


AI-powered robots and systems can handle hazardous work environments—from deep-sea exploration to nuclear facility maintenance—protecting human workers from injury or death while maintaining operational efficiency.

Cost Reduction

By automating repetitive tasks and optimizing complex processes, AI significantly reduces operational costs. Organizations can reallocate human resources to more creative, strategic work while AI handles routine operations around the clock without breaks.

Unbiased Decision Making

When properly designed and trained, AI can make decisions based purely on data patterns, free from human biases related to race, gender, or personal preferences. This objectivity is valuable in hiring, lending, and judicial applications.



Disadvantages of Artificial

Intelligence

Cost of

Implementation

Developing, deploying, and maintaining AI systems requires substantial upfront investment in hardware, software, specialized talent, and ongoing infrastructure. For many small and medium-sized organizations, these costs create significant barriers to entry, potentially widening the competitive gap between tech-enabled and traditional businesses.

Doesn't Improve with

Experience

Many AI systems, particularly those using fixed algorithms rather than continuous learning approaches, don't automatically improve over time without deliberate retraining. Unlike humans who naturally learn from every experience, these systems require explicit updates and new training data to enhance their performance or adapt to changing conditions.

Lack of Creativity or

Emotion


Despite impressive capabilities, AI lacks genuine creativity, emotional intelligence, and the ability to understand nuanced human contexts. While AI can generate content or make decisions based on patterns, it cannot truly innovate, empathize, or grasp the subtle emotional undertones that are crucial in many human interactions and creative endeavors.

Job Automation and

Displacement

As AI systems become more capable, they increasingly automate jobs traditionally performed by humans, from manufacturing and data entry to customer service and even some professional roles. This displacement creates economic and social challenges, requiring workforce retraining and raising questions about income inequality and employment security.

These disadvantages don't negate AI's value but highlight the need for thoughtful implementation, ongoing investment in human capital, and policies that address the societal impacts of automation.



Risks of Artificial

Intelligence

Human Interactivity

As interactions between humans and AI systems become increasingly frequent and integrated into daily life, the potential for harm grows. AI-powered machinery, autonomous vehicles, and robotic systems can malfunction with dangerous or even fatal consequences for nearby humans. Without careful monitoring, safety protocols, and fail-safe mechanisms, these systems pose significant physical risks, particularly in high-stakes environments like healthcare, transportation, and manufacturing.

Risks

Autonomous

Weaponry


Thousands of AI experts signing open letters urging governments to ban autonomous weapons, their development continues worldwide. These weapons could make life-or-death decisions without human oversight, potentially causing catastrophic harm to civilian populations. Technical malfunctions, hacking, or unintended target selection could lead to unprecedented destruction. The proliferation of such weapons raises profound ethical questions about accountability and the delegation of lethal force to machines.

Superhuman AI and Existential

This hypothetical but serious risk doesn't necessarily involve AI becoming malevolent or "taking over the world" in a science-fiction sense. Instead, the concern is that sufficiently advanced AI might pursue its programmed goals through methods that cause unintended harm to humans or the environment. An AI optimized for a specific objective might find solutions that are technically correct but catastrophically harmful—a phenomenon known as the "alignment problem." Ensuring AI goals align with human values becomes exponentially more critical as AI capabilities approach and potentially exceed human-level intelligence.

Risk





Basic Skills Required for AI Development

Desirable Skills for AI Development



Mathematics and Statistics

- **Linear algebra and calculus** — Essential for understanding how AI algorithms optimize and learn from data
 - **Probability theory and statistics** — Critical for interpreting data patterns, measuring uncertainty, and validating model performance
-

Programming and Software Development

- **Popular programming languages** — Python, R, and Java dominate AI development for their extensive libraries and community support
- **Data structures and algorithms** — Fundamental for writing efficient code that can process large datasets
- **Version control and collaboration** — Essential for working in teams and managing complex AI projects

These foundational skills create the technical base upon which all AI expertise is built. Mastering them opens pathways to specialization in machine learning, deep learning, natural language processing, and computer vision.

Linear Algebra and Calculus

Role of Linear Algebra



Linear algebra is absolutely essential for understanding how data is represented and manipulated in high-dimensional spaces. Concepts like vectors, matrices, and tensor operations form the mathematical foundation of neural networks. Every image, piece of text, or dataset fed into an AI system is transformed into numerical representations using linear algebra, enabling machines to process and learn from information.

Importance of Calculus

Calculus, particularly differential calculus, is vital for optimization processes in machine learning. Techniques like gradient descent—which minimizes error in model predictions—rely entirely on calculus principles. Understanding derivatives and partial derivatives allows AI practitioners to fine-tune models, adjusting millions of parameters to enhance performance and achieve more accurate results.

Foundation for AI Techniques

The synergy between linear algebra and calculus underpins virtually every technique used in modern artificial intelligence. From simple regression models to sophisticated deep learning architectures, these mathematical frameworks facilitate the complex computations and iterative learning processes that make AI possible.



Popular Programming Languages for AI



Python for AI

Python has become the dominant language in AI development thanks to its clean, readable syntax and extraordinary ecosystem of libraries. TensorFlow, PyTorch, scikit-learn, and NumPy provide ready-made tools for everything from basic data manipulation to building sophisticated neural networks. Python's gentle learning curve makes it accessible to beginners while remaining powerful enough for advanced applications.

R for Statistical Analysis

R excels in statistical computing and data visualization, making it particularly valuable for exploratory data analysis and statistical modeling. With specialized packages for virtually every statistical method, R remains the language of choice for statisticians and data scientists focused on rigorous analysis and publication-quality visualizations.

Java for Large Systems

Java's robustness, scalability, and platform independence make it ideal for enterprise-level AI applications and production systems. Large organizations often choose Java for AI systems that must integrate with existing infrastructure, handle millions of transactions, and maintain reliability under heavy loads.

Model Evaluation and Validation



Importance of Evaluation

Rigorously evaluating AI models before deployment ensures they perform reliably in real-world conditions. Proper evaluation identifies weaknesses, prevents overfitting, and builds confidence that the model will deliver value rather than costly failures.

Cross-Validation Techniques

Cross-validation methods like k-fold validation help assess how model results will generalize to independent datasets. By testing the model on different data subsets, practitioners can estimate real-world performance and detect overfitting—when a model memorizes training data rather than learning general patterns.

Performance

Metrics


Understanding metrics like accuracy, precision, recall, F1-score, AUC-ROC, and mean squared error is essential for quantifying model effectiveness. Different applications require different metrics; for example, medical diagnosis systems prioritize recall to avoid missing diseases, while spam filters balance precision and recall.

The background is a complex, abstract composition. On the left, there's a vertical strip containing various data visualizations: a bar chart with a value of 37,402.01, a pie chart, another bar chart with values 10, 20, 30, 40, a line graph, and a large bar chart with values 23.00, 25.00, 03.00, and 0.5.00. The right side features a large, stylized silhouette of a human head in profile, facing right. The head is filled with a light, textured pattern. The overall color palette is warm, with shades of beige, cream, and light brown. The title 'Data Science and Analytics' is prominently displayed in the center-right, in a large, bold, black sans-serif font.

Data Science and Analytics

Data science combines statistical methods, domain expertise, and programming to extract insights from data. For AI practitioners, data science skills are fundamental—the quality of your data directly determines the quality of your AI models.

Data Collection and Preprocessing



Importance of Quality Data

The axiom "garbage in, garbage out" is especially true in AI. Quality data collection—ensuring accuracy, relevance, and representativeness—forms the foundation of successful AI projects. Poor data quality leads to biased, inaccurate models that fail in production, wasting resources and potentially causing harm.

Role of Preprocessing

Data preprocessing transforms raw data into a clean, consistent format suitable for AI algorithms. This involves handling missing values, removing outliers, normalizing scales, encoding categorical variables, and splitting data into training and testing sets. Effective preprocessing can dramatically improve model performance and reduce training time.

AI practitioners often spend 60-80% of their time on data collection and preprocessing—a testament to its critical importance. Mastering these skills accelerates project timelines and improves outcomes.

Exploratory Data Analysis (EDA)



01

Definition of EDA

Exploratory Data Analysis is a critical approach to analyzing datasets using visual and statistical methods. It helps data scientists understand data characteristics before applying complex algorithms.

0


Techniques in EDA

Common techniques include histograms for distribution analysis, box plots for outlier detection, scatter plots for relationship visualization, and correlation matrices for identifying variable dependencies.

0

Importance of EDA

EDA reveals trends, patterns, outliers, and relationships that inform feature engineering and model selection. It prevents costly mistakes by uncovering data quality issues and unexpected patterns early in the project.

 **Pro tip:** Always perform thorough EDA before building models. The insights gained often reveal the most effective modeling approaches and save countless hours of unproductive experimentation.

Volumes

Importance of Hadoop

Role of Spark

Apache Spark builds on Hadoop's foundation while offering dramatically faster in-memory processing capabilities. Spark's speed advantage—sometimes 100x faster than Hadoop MapReduce—makes it ideal for iterative machine learning algorithms, real-time analytics, and interactive data exploration. Its unified framework supports batch processing, streaming, machine learning, and graph processing.



Ethics and Responsible AI

As AI becomes more powerful and pervasive, ethical considerations become increasingly critical. Responsible AI development requires conscious attention to fairness, transparency, privacy, and societal impact.

Ethics of AI



Privacy Concerns

With data collection and storage reaching unprecedented levels, consumer privacy rights face constant threats. Companies collect vast amounts of personal information to train AI systems, often without explicit informed consent. This data—from browsing habits to location tracking to purchase history—creates detailed profiles that can be exploited, leaked, or misused. Regulations like GDPR attempt to protect privacy, but enforcement remains challenging as data practices evolve faster than legislation.




Bias in AI Systems

A dangerous misconception suggests AI is inherently unbiased because machines lack human prejudices. In reality, most modern AI systems exhibit significant bias because the data used to train them reflects historical inequalities and human prejudices. Facial recognition systems perform worse on darker skin tones, hiring algorithms discriminate against women, and loan approval systems disadvantage minority communities. Bias in AI perpetuates and amplifies societal inequalities unless deliberately addressed through careful data curation and algorithmic fairness techniques.



Regulatory Gaps

Currently, comprehensive AI regulations are conspicuously absent at both national and international levels, despite AI's pervasive presence in consumer lives. This regulatory vacuum creates risks as powerful AI systems operate without adequate oversight, accountability mechanisms, or safety standards. While the EU's AI Act and various state-level initiatives are emerging, the pace of AI development far exceeds regulatory response, leaving consumers vulnerable and ethical concerns unaddressed.



Bias and Fairness in AI Algorithms



Impact of Bias in AI

Bias in AI algorithms can produce systematically unfair outcomes that disproportionately harm marginalized communities. When training data reflects historical discrimination or underrepresents certain groups, AI systems learn and perpetuate these biases. The consequences range from denied opportunities to reinforced stereotypes, effectively codifying discrimination into automated systems.

Importance of Fair

Algorithms

Designing algorithms that are fair, transparent, and equitable requires intentional effort throughout the development lifecycle. This includes diverse training data, bias detection techniques, fairness metrics, and ongoing monitoring after deployment. Fair AI ensures that technological advancement benefits all members of society, not just privileged groups.

Transparency in AI

Transparency—the ability to understand how AI systems make decisions—is essential for accountability and trust. Explainable AI techniques help users understand why specific decisions were made, enabling identification of bias and building confidence in automated systems. Transparency also facilitates regulatory compliance and ethical oversight.

Accountability and

Transparency

1

Establishing Clear Responsibility

Systems must have clear lines of accountability—identifying who is responsible when things go wrong. This includes developers, deployers, and organizations that benefit from AI systems.

2

Documenting Decision-Making

Maintaining documentation of model development, training data sources, design decisions, and performance evaluations creates an audit trail that enables accountability and facilitates improvement.

3

Enabling External Review

Transparency allows external experts, regulators, and affected communities to review AI systems, identifying potential harms and ensuring alignment with ethical standards and societal values.

4

Building Public Trust

When organizations demonstrate accountability and transparency, they build public trust in AI systems, encouraging adoption while maintaining ethical safeguards that protect society.

Accountability and transparency aren't obstacles to innovation—they're foundations for sustainable, responsible AI development that benefits society while minimizing harm.

Conclusion

Artificial intelligence represents one of the most transformative technologies of our era, with the potential to revolutionize virtually every aspect of human life. From healthcare and education to transportation and entertainment, AI systems are already reshaping how we work, communicate, and solve complex problems.

Understanding AI fundamentals—its history, types, applications, and underlying mathematics—provides the foundation for participating in this technological revolution, whether as a developer, user, or informed citizen. The skills required for AI development—mathematics, programming, data science, and ethical reasoning—are accessible to anyone willing to invest time in learning.

However, with great power comes great responsibility. As AI capabilities expand, so do ethical considerations around bias, privacy, accountability, and societal impact. Developing AI responsibly requires conscious attention to these concerns, ensuring that technological advancement benefits all of humanity rather than exacerbating existing inequalities.

The future of AI is not predetermined—it will be shaped by the choices we make today about how to develop, deploy, and govern these powerful systems. By understanding both the opportunities and risks, we can work toward an AI-powered future that is fair, transparent, and aligned with human values.

- This presentation has provided a comprehensive introduction to AI fundamentals. The resources on the next slide will help you continue your learning journey, whether you're interested in the mathematical foundations, practical implementation, or ethical implications of artificial intelligence.

Resources to Continue Your AI

Journey

Basic Statistics



Build your statistical foundation with these excellent free resources:

- [Seeing Theory](#) — An interactive visual introduction to probability and statistics from Brown University
- [Khan Academy Statistics](#) — Comprehensive video lessons covering probability and statistical concepts

Machine Learning



Start your machine learning journey with this practical guide:

- [Google's ML Introduction](#) — A hands-on introduction to machine learning concepts and practical implementation from Google's developer resources

These resources provide excellent starting points for deepening your understanding of AI fundamentals. Consistent practice and hands-on experimentation will accelerate your learning. Consider joining online communities, participating in competitions like Kaggle, and building personal projects to apply your knowledge.

📌 **Remember:** Learning AI is a journey, not a destination. The field evolves rapidly, so cultivate curiosity, stay updated with latest developments, and never stop learning.

<https://apexbtic.icgeb.res.in/workshop2025/>

Workshop course material

<https://apexbtic.icgeb.res.in/workshop2025/>

"A Hands-on Workshop on AI 🤖 and Data Science for Advancing Translational Science 🤖"

 11th to 15th November 2025

Organized by: Translational Bioinformatics Group at ICGEB, New Delhi

[Workshop Flyer](#)[Registra](#)[Workshop Schedule](#)

About

The Workshop

The workshop will introduce AI/ML/DL techniques and their applications in translational science. Participants will get a chance to learn the role of AI and Data science with extensive lectures and hands on sessions. At the end, the participants will be able to employ the learnings to their own research projects.

Topics to be covered


- Fundamentals of AI, LLM's, and Data Science with Application in Biology
- AI Application in Biomedical Image processing
- Machine Learning and AI for Metagenomics Data Analysis
- Generative AI for Exploration of chemical space
- AI Driven Approaches in Genomics and Variant Interpretation
- AI in Single Cell and Spatial Transcriptomics





Dr. Priyanka Narad

ICMR, New Delhi

[Abstract](#)[Content](#)

Dr. Abhishek Sengupta

Amity University, Noida

[Abstract](#)[Content](#)


Ms. Dhairya Gupta

ICGEB, TBG Lab

[Abstract](#)[Content](#)

Dr. Silvano Piazza

ICGEB, Italy

[Abstract](#)[Content](#)

Dr. Stefano Cacciatore

ICGEB, SA



Dr. Ishaan Gupta

IIT Delhi



Dr. Parag Singla

IIT Delhi



Dr. Vineet K. Sharma

IISER Bhopal

A Hands-on Workshop on AI and Data Science for Advancing Translational Science (11th to 15th Nov, 2025)



सत्यमेव जयते
DEPARTMENT OF BIOTECHNOLOGY
GOVERNMENT OF INDIA



Dr. Priyanka Narad

Day	Topic	Lecture Content
Day 1 (9:30 - 10:30 am)	MATTER: Metrics and Assessment Tools for Trustworthy, Transparent, Explainable, and Reliable AI	Abstract Presentation Hands-on

Supported by ICGEB

For the registered participants:

Lunch (1PM): The guest house lawns (Near ICGEB Guest house)

Dinner (7PM): ICGEB-Canteen/Faculty-Lounge (II floor of this building)