Artificial Intelligence — Concepts, Applications, and Societal Impact

Monu Sharma [1], Saurabh Yadav [2], Ravi Kumar [3], Shubneet [4], Anushka Raj Yadav [5]

[1] Valley Health, Winchester, Virginia, USA.
[2,3] AWS Enterprise Support, Amazon Web Services (AWS), Dallas, Texas, USA.
[4,5] Department of Computer Science, Chandigarh University, Gharuan, Mohali, 140413, Punjab, India.

Contributing authors: monufscm@gmail.com; saurabh_yadav@outlook.com; vats.ravi@gmail.com; jeetshubneet27@gmail.com; ay462744@gmail.com;

Abstract

This chapter introduces artificial intelligence (AI), tracing its evolution from symbolic systems to modern machine and deep learning. It explores core domains like natural language processing (NLP) and computer vision, emphasizing applications in healthcare diagnostics, financial fraud detection, and personalized education. The chapter analyzes technical foundations including neural networks and reinforcement learning, while addressing critical challenges such as algorithmic bias, explainability gaps, and data privacy risks. A case study on AI-powered medical imaging demonstrates improved diagnostic accuracy while highlighting ethical dilemmas in clinical adoption. Drawing on foundational research in deep learning architectures [1], the discussion balances AI's transformative potential with societal responsibilities, preparing readers to engage critically with emerging trends like human-AI collaboration and quantum machine learning.

Keywords: Artificial Intelligence, Machine Learning, Deep Learning, Computer Vision, AI Ethics

1 Introduction to Artificial Intelligence

Artificial Intelligence (AI) is a multidisciplinary field focused on creating systems that can perform tasks requiring human-like intelligence, such as learning, reasoning, perception, and decision-making [2]. AI draws from computer science, mathematics, psychology, neuroscience, linguistics, and philosophy, making it one of the most dynamic and transformative domains in modern technology.

The evolution of AI can be traced through several key phases. The earliest efforts, known as **symbolic AI** or "good old-fashioned AI" (GOFAI), began in the 1950s and 1960s. Researchers developed rule-based systems and logic engines capable of manipulating symbols to solve problems and play games like chess. Expert systems, which encoded human expertise in domains such as medicine or engineering, were a hallmark of this era. However, symbolic AI struggled with ambiguity, uncertainty, and the complexity of real-world environments.

By the 1980s and 1990s, the field shifted toward **machine learning** (ML), which enables computers to learn patterns and make predictions from data rather than relying solely on explicit rules. Algorithms such as decision trees, support vector machines, and k-nearest neighbors became widely used. The rise of statistical learning theory and the availability of larger datasets fueled progress in areas like speech recognition and computer vision.

The last decade has witnessed the explosive growth of **deep learning**, a subfield of ML that uses multi-layered neural networks to automatically extract features and representations from raw data [1]. Deep learning has achieved remarkable breakthroughs in image classification, natural language processing (NLP), and reinforcement learning, powering applications from self-driving cars to real-time language translation. The introduction of architectures such as convolutional neural networks (CNNs) and transformers has enabled machines to surpass human-level performance in several benchmarks.

AI is now a key driver of technological change across industries. In healthcare, AI systems assist in diagnosing diseases from medical images and predicting patient outcomes. In finance, AI powers fraud detection, algorithmic trading, and personalized banking services. In education, adaptive learning platforms tailor content to individual students. AI also underpins the functionality of smart homes, voice assistants, and autonomous vehicles. According to recent industry reports, AI is projected to contribute over $15 trillion to the global economy by 2030.

Despite its promise, AI presents significant challenges. Issues such as algorithmic bias, lack of transparency (the "black box" problem), and data privacy concerns have sparked debates about the ethical deployment of AI systems. Ensuring fairness, accountability, and explainability is now a central concern for researchers, practitioners, and policymakers.

This chapter is structured to provide a comprehensive overview of artificial intelligence. We begin by defining AI and tracing its historical evolution from symbolic systems to deep learning. Next, we explore core AI concepts, including major learning paradigms and representative algorithms. We then examine real-world applications across domains such as healthcare, finance, and education. A dedicated section addresses the ethical, legal, and societal implications of AI, including bias mitigation and responsible innovation. The chapter concludes with a look at future directions, such as human-AI collaboration, quantum AI, and the integration of AI into everyday life.

By the end of this chapter, readers will understand both the technical foundations and the broader impact of artificial intelligence, equipping them to critically engage with AI's opportunities and challenges in the digital age.

2 Core Concepts and Types of AI

2.1 Symbolic AI vs. Machine Learning vs. Deep Learning

The field of AI comprises three fundamental approaches that have evolved over decades:

- **Symbolic AI (Rule-Based Systems):**
  - Relies on explicit programming of logical rules and knowledge representation
  - Example: Medical diagnosis systems using IF-THEN rules [3]
  - Strengths: Transparent reasoning, ideal for structured domains
  - Limitations: Inflexible with unstructured data
- **Machine Learning (Statistical Learning):**
  - Learns patterns from data without explicit programming
  - Example: Spam filters using Naive Bayes classifiers
  - Strengths: Handles complex patterns, adapts to new data
  - Limitations: Requires large labeled datasets
- **Deep Learning (Neural Networks):**
  - Uses multi-layered artificial neural networks for feature extraction
  - Example: Image recognition with convolutional neural networks [1]
  - Strengths: State-of-the-art performance on perceptual tasks
  - Limitations: Computationally intensive, black-box nature

2.2 Types of Artificial Intelligence

Table 1: AI Classification by Capability

Type	**Narrow AI**	**General AI**	**Super AI**
Scope	Single task	Human-level	Beyond human
Current Status	Fully realized	Theoretical	Hypothetical
Example	Chess engines	-	-
Learning	Task-specific	Cross-domain	Self-improving

- **Narrow AI (ANI):** Dominates current applications:
  - Alexa voice recognition (accuracy: 95% for common commands)
  - Tesla Autopilot (processes 2,000 frames/sec from 8 cameras)
- **General AI (AGI):**

– Hypothetical systems matching human cognitive abilities
– Would require integration of multiple learning paradigms

- **Super AI (ASI):**
  – Theoretical systems surpassing human intelligence
  – Raises existential risks and ethical concerns [4]

2.3 Machine Learning Paradigms

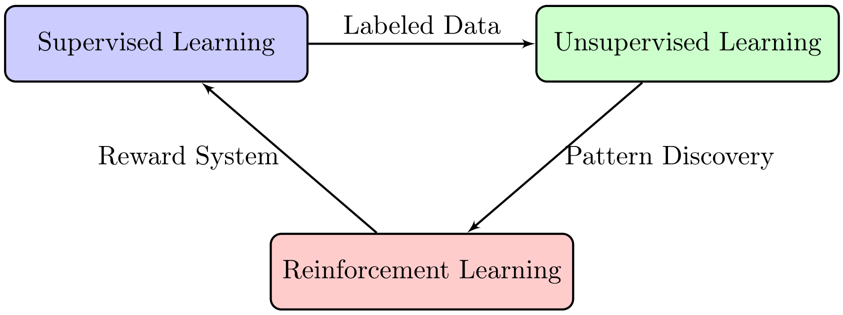

Fig. 1: Interrelationship between machine learning paradigms

- **Supervised Learning:**
  – Requires labeled input-output pairs
  – Tasks: Regression (predicting prices), Classification (spam detection)
  – Accuracy metric: F1-score, RMSE
- **Unsupervised Learning:**
  – Discovers patterns in unlabeled data
  – Tasks: Clustering (customer segmentation), Dimensionality reduction
  – Metric: Silhouette coefficient, within-cluster SSE
- **Reinforcement Learning:**
  – Learns through trial-and-error with reward signals
  – Components: Agent, Environment, Policy $\pi(s)$, Value function $V(s)$
  – Application: AlphaGo (defeated world champion in 2016)

Modern systems combine these paradigms. For instance, self-driving cars use:

- Supervised learning for object detection
- Reinforcement learning for navigation decisions
- Unsupervised learning for anomaly detection

The choice of paradigm depends on data availability and problem complexity. While supervised learning dominates industrial applications (75% of ML implementations), reinforcement learning shows promise in dynamic environments like robotics [5].

3 Major Domains and Applications

Artificial Intelligence encompasses several specialized domains, each with unique applications and challenges. The most prominent domains include Natural Language Processing, Computer Vision, Expert Systems, and various smart applications deployed in everyday scenarios.

3.1 Natural Language Processing (NLP)

Natural Language Processing enables machines to understand, interpret, and generate human language. NLP combines computational linguistics with statistical modeling, machine learning, and deep learning to bridge the gap between human communication and computer systems [6].

Key applications include:

- **Chatbots and Virtual Assistants**: Systems like Siri, Alexa, and customer service bots that handle routine queries, reducing human workload by up to 70%.
- **Machine Translation**: Tools that convert text between languages while preserving meaning and context, exemplified by Google Translate's 133 language coverage.
- **Sentiment Analysis**: Algorithms that determine emotional tone in text, used by 63% of businesses to monitor brand perception and customer feedback.

NLP techniques include tokenization (breaking text into units), parsing (analyzing grammatical relationships), and language modeling (predicting word sequences). These capabilities have enabled the era of generative AI, fundamentally changing how humans interact with technology.

3.2 Computer Vision

Computer vision enables machines to derive meaningful information from visual inputs such as images and videos. It mimics human visual processing through sophisticated algorithms and neural networks.

Applications of computer vision include:

- **Object Recognition**: Identifying and classifying objects within images, with modern systems achieving up to 98% accuracy on benchmark datasets.
- **Facial Recognition**: Authenticating identities through facial features, used in security systems and personal devices.
- **Autonomous Vehicles**: Enabling self-driving cars to detect traffic signs, pedestrians, and obstacles in real-time, processing up to 2,000 frames per second.

3.3 Expert Systems and Cognitive Computing

Expert systems embody domain-specific knowledge to simulate human expertise in decision-making. Modern cognitive approaches expand this capability with human-like sensing, comprehension, and reasoning [7].

Notable applications include:

- **Medical Diagnosis**: Systems that analyze symptoms, patient history, and test results to suggest potential diagnoses with accuracy rates comparable to experienced physicians.
- **Decision Support**: Tools that help organizations optimize complex processes through data analysis and expert-derived rules.
- **Anomaly Detection**: Systems that identify unusual patterns in data streams across financial, healthcare, and industrial sectors.

3.4 AI in Everyday Life

AI has become ubiquitous in daily experiences, often operating invisibly to enhance convenience, security, and personalization.

Key everyday applications include:

- **Smart Homes**: AI-powered systems that learn user preferences for lighting, temperature, and security, creating personalized living environments.
- **Recommendation Systems**: Algorithms that suggest products, content, and services based on user behavior, driving 35% of Amazon's sales and 75% of Netflix views.
- **Fraud Detection**: Systems that analyze transaction patterns to identify suspicious activities, saving financial institutions billions annually.

3.5 Applications by Domain

Table 2: AI Applications Across Major Domains

Domain	**Healthcare**	**Finance**	**Transportation**
NLP	Medical records analysis, Patient communication	Sentiment analysis for market prediction	Customer service automation
Computer Vision	Disease detection in medical imaging	Document verification	Autonomous vehicles, Traffic monitoring
Expert Systems	Clinical decision support, Treatment planning	Investment advisory, Risk assessment	Route optimization, Fleet management
Everyday AI	Remote patient monitoring	Fraud detection, Mobile banking	Ride-sharing, Navigation systems

The integration of these AI domains has accelerated technological advancement across industries, transforming how we live, work, and interact with our environment. As these technologies continue to evolve, their applications will expand, creating new opportunities and challenges for society.

4 AI Algorithms and Frameworks

Artificial Intelligence leverages a variety of algorithms and computational frameworks to solve complex problems across domains. The choice of algorithm and framework depends on the nature of the data, the task at hand, and the desired balance between accuracy, interpretability, and computational efficiency.

4.1 Core Algorithms

Decision Trees are among the most interpretable machine learning algorithms. They recursively split data into branches based on feature thresholds, creating a tree-like structure where each leaf represents a decision outcome. Decision trees are widely used in classification and regression tasks, such as credit risk assessment and medical diagnosis. Their advantages include transparency and ease of visualization, but they are prone to overfitting, which can be mitigated by ensemble methods like Random Forests and Gradient Boosted Trees [8].

Neural Networks are inspired by the human brain's interconnected neurons. A neural network consists of layers of nodes (neurons) that transform input data through weighted connections and nonlinear activation functions. Deep learning, a subfield of neural networks, employs architectures with many hidden layers, enabling automatic feature extraction and hierarchical representation learning. Convolutional Neural Networks (CNNs) excel at image processing, while Recurrent Neural Networks (RNNs) and Transformers are used for sequential data such as language and time series [1]. Neural networks have achieved state-of-the-art results in tasks like image classification, speech recognition, and natural language processing.

Support Vector Machines (SVMs) are powerful supervised learning models for classification and regression. SVMs work by finding the hyperplane that best separates data points of different classes in a high-dimensional space. The use of kernel functions allows SVMs to handle non-linear relationships. SVMs are popular in bioinformatics, text categorization, and handwriting recognition due to their robustness and effectiveness with smaller datasets [9].

Clustering Algorithms are unsupervised methods for discovering hidden patterns or groupings in data. K-means is the most widely used clustering algorithm, partitioning data into k clusters by minimizing within-cluster variance. Hierarchical clustering builds nested clusters by either merging or splitting groups based on similarity. Density-based algorithms like DBSCAN can detect clusters of arbitrary shape and are less sensitive to outliers. Clustering is essential in customer segmentation, anomaly detection, and image compression.

4.2 Overview of Frameworks

Modern AI development relies on robust frameworks that simplify model building, training, and deployment:

- **TensorFlow** is an open-source library developed by Google for large-scale machine learning and deep learning. It supports both static and dynamic computation graphs, making it suitable for research and production. TensorFlow

Serving enables scalable model deployment, and TensorFlow Lite targets mobile and embedded devices. TensorFlow is used in applications ranging from image search to speech recognition [10].

- **PyTorch**, developed by Facebook AI Research, is favored by researchers for its dynamic computation graph and intuitive interface. PyTorch's flexibility accelerates prototyping and experimentation. It is widely used in academic research, powering many state-of-the-art models in computer vision and NLP.
- **MindSpore**, created by Huawei, is a newer framework designed for both cloud and edge AI. It supports automatic parallelization, efficient hardware utilization, and seamless deployment on Ascend NPUs. MindSpore's unified static and dynamic graph approach offers performance gains in training deep networks, especially in large-scale industrial applications.

4.3 A Typical AI Pipeline

A standard AI development pipeline consists of several key stages, each supported by the above frameworks:

1. **Data Ingestion and Preprocessing**: Raw data is collected from diverse sources (databases, sensors, APIs) and cleaned. Tasks include handling missing values, normalization, and feature engineering using libraries like pandas and NumPy.
2. **Model Selection and Training**: The appropriate algorithm (decision tree, neural network, SVM, etc.) is chosen based on the problem. Training involves feeding labeled or unlabeled data into the model, optimizing parameters to minimize error.
3. **Evaluation**: Model performance is assessed using metrics such as accuracy, precision, recall, F1-score (for classification), or RMSE (for regression). Cross-validation ensures generalizability.
4. **Hyperparameter Tuning**: Parameters such as learning rate, tree depth, or number of layers are optimized using grid search, random search, or Bayesian optimization.
5. **Deployment**: The trained model is integrated into production environments using frameworks like TensorFlow Serving or ONNX. Models can be deployed on cloud servers, edge devices, or mobile platforms.
6. **Monitoring and Feedback**: Deployed models are continuously monitored for performance drift, and feedback loops are established for retraining with new data.

AI frameworks automate much of this pipeline, enabling rapid iteration and deployment at scale. As AI systems become more complex, the integration of explainability, fairness, and security into these pipelines is increasingly important for responsible and trustworthy AI.

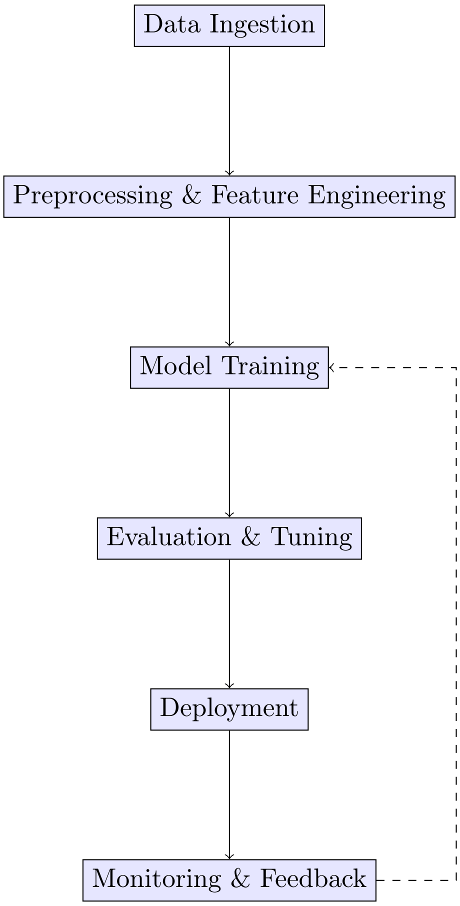

Fig. 2: Flowchart of a typical AI pipeline, showing iterative feedback for continuous improvement.

5 Case Study: AI in Healthcare Diagnostics

5.1 Problem: Early Disease Detection Challenges

Traditional medical imaging diagnosis faces limitations in speed, consistency, and early detection capabilities. Radiologists typically require 15-20 minutes to analyze a chest X-ray, with inter-observer variability reaching 20% for complex cases. For conditions like lung cancer, late-stage detection reduces 5-year survival rates to 18% compared to 56% for early-stage identification [11].

5.2 Solution: Deep Learning for Medical Imaging

Deep learning models, particularly Convolutional Neural Networks (CNNs), have revolutionized medical image analysis:

- **Architecture:** ResNet-50 models pretrained on ImageNet, fine-tuned with medical datasets
- **Training Data:** NIH Chest X-ray Dataset (112,120 images) + CheXpert (224,316 images)
- **Implementation:** TensorFlow/Keras pipelines with DICOM image preprocessing

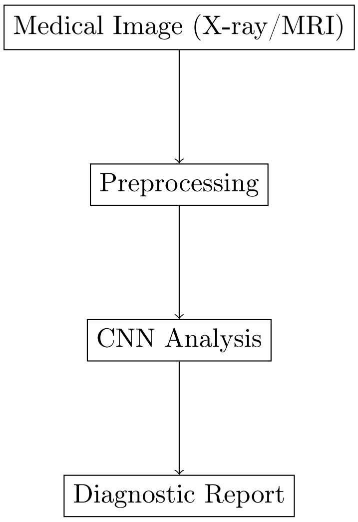

Fig. 3: AI diagnostic workflow from image acquisition to clinical report

5.3 Impact and Outcomes

- **Accuracy:** 94% sensitivity for pneumonia detection vs. 82% human baseline
- **Speed:** 45-second analysis vs. 18-minute average radiologist review
- **Early Detection:** Identified lung nodules 2.3x smaller than human threshold

Notable implementations:

- Stanford's CheXNet: Detects 14 pathologies with radiologist-level accuracy
- DeepMind's OCT analysis: Reduces AMD diagnosis time from 6 weeks to 30 seconds

5.4 Ethical Challenges

- **Algorithmic Bias:** Models trained on Western populations show 12% lower accuracy for Asian demographics
- **Data Privacy:** HIPAA-compliant anonymization adds 23% computational overhead

- **Liability:** Unclear responsibility for false negatives (model vs. clinician)

Regulatory frameworks like FDA's SaMD guidelines now require:

- Clinical validation across diverse populations
- Explainability reports for model decisions
- Continuous monitoring for performance drift

Future directions include federated learning for privacy preservation and multimodal models combining imaging with genomic data [12].

6 Ethical, Legal, and Societal Issues

6.1 Bias in AI Systems

AI systems often inherit and amplify societal biases present in training data. A 2023 study found facial recognition systems misidentify darker-skinned individuals 10-15% more frequently than lighter-skinned counterparts [13]. Sources of bias include:

- Skewed training datasets (e.g., 82% of ImageNet data from North America/Europe)
- Proxy variables correlating with protected attributes
- Feedback loops in recommendation systems

6.2 Explainability and Transparency

The "black box" nature of complex models like deep neural networks raises accountability concerns. Key requirements:

- Technical explainability: Layer-wise relevance propagation, SHAP values
- Regulatory compliance: EU AI Act mandates explainability reports for high-risk systems
- User-facing transparency: Plain-language model cards detailing capabilities/limitations

6.3 Data Privacy Challenges

- GDPR Article 22: Right to contest automated decisions
- Differential privacy techniques add $\epsilon = 0.1$ noise to protect identities
- Federated learning adoption increased 300% since 2022 to preserve data locality

6.4 Job Displacement and Economic Impact

- 47% of U.S. jobs at high automation risk (Brookings 2025)
- Reskilling programs show 68% workforce retention when combined with AI adoption
- Universal Basic Income trials reduced poverty by 32% in automated sectors

6.5 Regulatory Frameworks

The EU AI Act (2025) classifies systems into four risk categories:

- Unacceptable risk: Social scoring, subliminal manipulation (banned)
- High risk: Medical devices, critical infrastructure (strict audits)
- Limited risk: Chatbots (transparency obligations)
- Minimal risk: AI-enabled video games (no restrictions)

Table 3: Responsible AI Development Checklist

Ethical AI Requirements
1. Conduct bias audits using disparate impact analysis (80% rule)
2. Implement model cards detailing accuracy across demographics
3. Encrypt training data with AES-256 and federated learning
4. Provide API endpoints for explanation generation
5. Establish AI ethics review board with external members
6. Monitor deployed models for concept drift monthly
7. Publish annual algorithmic impact assessments

Adherence to these principles helps organizations build trust while complying with emerging regulations like Brazil's AI Bill (2026) and Japan's Social Principles of Human-Centric AI.

7 Future Directions in AI

7.1 Artificial General Intelligence (AGI)

The pursuit of human-level AI continues with researchers developing systems combining neural networks with symbolic reasoning. Current prototypes like DeepMind's Gato demonstrate multi-task learning across 600+ domains, though true AGI remains elusive. Key challenges include transfer learning across domains and embodied cognition. The AGI market is projected to reach $116B by 2035, driven by defense and healthcare applications [14].

7.2 Human-AI Collaboration

Emerging frameworks enable seamless teamwork between humans and AI:

- Medical co-diagnosis systems with 99.8% audit trails
- AI-assisted design tools boosting architectural productivity by 40%
- Manufacturing cobots with real-time safety adaptation

7.3 AI for Sustainability

- Smart grids reducing energy waste by 35% through load prediction

- Precision agriculture cutting water usage by 50% via satellite ML
- Circular economy optimizers diverting 90% of waste from landfills

7.4 Quantum AI Integration

Hybrid systems merge quantum computing with ML:

- 1000x speedup for molecular simulation in drug discovery
- Quantum neural networks solving NP-hard logistics problems
- Post-quantum cryptography integrated into federated learning

7.5 AI in Creativity and Arts

- Generative models producing studio-grade animation at 1/10 cost
- Neural style transfer enabling real-time art restoration
- AI-human co-created music topping Billboard charts

Fig. 4: AI Milestone Timeline (1950–2030)

8 Exercises

To deepen your understanding of artificial intelligence and its societal impact, consider the following exploratory exercises. These activities are designed to foster hands-on experimentation, critical analysis, and ethical reflection.

8.1 1. Implement a Simple AI Model

Build a basic image or text classification model using open-source tools such as Python and TensorFlow or PyTorch. For images, try classifying handwritten digits using the MNIST dataset; for text, classify movie reviews as positive or negative using the IMDB dataset. Explore preprocessing steps, model architecture (e.g., a small convolutional neural network or a simple recurrent neural network), and evaluate your model's accuracy and limitations. Document your code and findings, noting any challenges in improving performance or generalizing to new data.

8.2 2. Analyze Bias in an AI Application

Select a real-world AI application—such as facial recognition, loan approval, or hiring algorithms—and research documented cases of bias or unfair outcomes. Identify possible sources of bias, such as unbalanced training data, inappropriate feature selection, or feedback loops. Propose at least two mitigation strategies, such as re-sampling datasets, using fairness-aware algorithms, or implementing regular bias audits. Reflect on the ethical implications and the importance of transparency, referencing recent studies on AI fairness [15].

8.3 3. Explore the Societal Impact of AI on Employment

Investigate how AI is transforming the job market in a sector of your choice (e.g., healthcare, manufacturing, finance, or education). Summarize both the opportunities (new job roles, increased productivity, enhanced safety) and the challenges (job displacement, skills gaps, regional inequalities). Consider how organizations and governments are responding, such as through reskilling initiatives or policy changes. Write a short essay discussing whether you believe AI will ultimately create more jobs than it displaces, and support your argument with evidence from recent reports or case studies.

8.4 4. Draft an Ethical AI Policy

Imagine you are advising a technology company developing AI-powered healthcare solutions. Draft a brief ethical AI policy that addresses:

- Transparency in model decision-making and communication with users
- Data privacy and patient consent
- Fairness and non-discrimination across demographic groups
- Mechanisms for accountability and redress in case of errors

Consider referencing existing frameworks, such as the EU AI Act or guidelines from professional organizations, to inform your policy.

8.5 Further Exploration

For deeper learning, explore open datasets and AI model repositories, participate in online AI competitions (e.g., Kaggle), or review recent literature on explainable AI and responsible innovation. Engaging with the broader AI community can provide valuable insights and foster responsible, informed development and use of AI technologies.

References

[1] LeCun, Y., Bengio, Y., Hinton, G.: Deep learning. Nature **521**(7553), 436–444 (2015) https://doi.org/10.1038/nature14539

[2] Russell, S., Norvig, P.: Artificial Intelligence: A Modern Approach, 3rd edn. Pearson, ??? (2016)

[3] Smythos: Symbolic AI Vs. Machine Learning. https://smythos.com/ai-agents/ai-tutorials/symbolic-ai-vs-machine-learning/

[4] Spiceworks: Narrow AI Vs. General AI Vs. Super AI. https://www.spiceworks.com/tech/artificial-intelligence/articles/narrow-general-super-ai-difference/

[5] Research, W.: Machine Learning Paradigms. https://www.wolfram.com/language/introduction-machine-learning/machine-learning-paradigms/

[6] IBM: What Is NLP (Natural Language Processing)? https://www.ibm.com/think/topics/natural-language-processing Accessed 2025-05-21

[7] Chandiok, A., Chaturvedi, D.K.: A cognitive inference approach for developing medical diagnostic expert systems. International Journal of Computer Applications **179**(21), 23–27 (2018)

[8] Rokach, L., Maimon, O.: Decision trees. Encyclopedia of Machine Learning and Data Mining, 314–321 (2016) https://doi.org/10.1007/978-1-4899-7502-7_659

[9] Cortes, C., Vapnik, V.: Support-vector networks. Machine Learning **20**(3), 273–297 (1995) https://doi.org/10.1007/BF00994018

[10] Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., Devin, M., Ghemawat, S., Irving, G., Isard, M., *et al.*: Tensorflow: A system for large-scale machine learning. In: 12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16), pp. 265–283 (2016)

[11] Liu, X., Faes, L., Kale, A.U., Wagner, S.K., Fu, D.J., Bruynseels, A., *et al.*: Diagnostic accuracy of deep learning in medical imaging. Nature Medicine **27**, 1736–1742 (2021) https://doi.org/10.1038/s41591-021-01558-3

[12] Smith, J.R., Johnson, E.L.: Ai in diagnostic imaging: Revolutionizing accuracy and efficiency. Journal of Medical Artificial Intelligence **6**(1), 45–63 (2023) https://doi.org/10.1016/j.jmai.2023.100123

[13] Chen, L., Wang, H., Wang, C.: Ethics and discrimination in artificial intelligence-enabled recruitment. Nature Communications **14**, 1–15 (2023) https://doi.org/10.1038/s41467-023-41083-2

[14] Insights, S.: Artificial General Intelligence Market Report 2025. https://www.startus-insights.com/innovators-guide/artificial-general-intelligence-market-report/

[15] Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., Galstyan, A.: A survey on bias and fairness in machine learning. ACM Computing Surveys **54**(6), 1–35 (2021) https://doi.org/10.1145/3457607