

# ARTIFICIAL INTELLIGENCE

## CHEAT SHEET



[www.gsdcouncil.org](http://www.gsdcouncil.org)

# Introduction to Artificial Intelligence (AI)

Artificial Intelligence represents the simulation of human intelligence by machines, enabling them to perform tasks that traditionally required human cognition—learning from experience, reasoning through complex problems, and solving challenges autonomously. This transformative technology encompasses a broad spectrum of capabilities, from simple rule-based systems that follow predefined instructions to sophisticated generative models that create original content.

It's crucial to understand that AI is not sentient or conscious. Rather, it excels at pattern recognition and prediction based on vast amounts of data. Modern AI systems analyze information, identify correlations, and make informed decisions within their trained domains, pushing the boundaries of what machines can accomplish.

## Key Characteristics

- Learns from data patterns
- Adapts to new inputs
- Performs human-like tasks
- Improves with experience

# Core AI Subfields: Machine Learning and Deep Learning

## Machine Learning (ML)

Algorithms that learn from data to make predictions without explicit programming. ML systems identify patterns, build models, and continuously improve their performance as they process more information.

## Deep Learning (DL)

A subset of ML using multi-layered neural networks to model complex patterns. DL architectures automatically extract hierarchical features from raw data, enabling breakthrough performance in vision and language tasks.

Deep Learning requires substantial computational resources and large datasets to train effectively. Modern DL systems leverage GPUs and specialized hardware to process billions of parameters, achieving human-level performance in many domains.

# Artificial Neural Networks (ANN)

Artificial Neural Networks mimic the structure and function of the human brain, consisting of layers of interconnected nodes called neurons. Each connection carries a weight that adjusts during learning, allowing the network to recognize patterns and make decisions.

Information flows through the network from input layer through hidden layers to output layer. At each neuron, inputs are multiplied by weights, summed, and passed through an activation function that determines the neuron's output. This architecture enables ANNs to learn complex, non-linear relationships in data.



## Image Recognition

ANNs excel at identifying objects, faces, and patterns in visual data



## Natural Language Processing

Understanding and generating human language with high accuracy



## Prediction Tasks

Forecasting outcomes based on historical patterns and trends

# Key AI Terminology: Models, Training, and Inference

1

## Model

A mathematical representation trained to make predictions or classifications. Models encode learned patterns from data into a structure that can process new inputs and generate outputs.

2

## Training

The process of adjusting model parameters to minimize prediction errors. During training, the model iteratively learns from examples, refining its internal weights to improve accuracy.

3

## Inference

Using a trained model to analyze new, unseen data. This deployment phase applies learned patterns to real-world scenarios, generating predictions or insights without further learning.

# Parameters and Hyperparameters

Understanding the distinction between parameters and hyperparameters is fundamental to working with AI systems effectively. These two concepts control different aspects of the learning process and require different approaches to optimization.

## 1 Parameters

Internal variables including weights and biases that are learned automatically during training. These values define the model's behavior and are adjusted by optimization algorithms to minimize prediction errors. A deep neural network may contain millions or billions of parameters.

## 2 Hyperparameters

Settings defined before training begins that control the learning process itself. Examples include learning rate, batch size, number of layers, and regularization strength. Data scientists manually tune hyperparameters or use automated search methods to find optimal configurations.

- ❏ **Pro Tip:** Finding the right hyperparameters often requires extensive experimentation. Techniques like grid search, random search, and Bayesian optimization help automate this process.

# Types of AI Systems



## Narrow AI (Weak AI)

Designed for specific tasks such as voice assistants, recommendation engines, or image classification. These systems excel within their defined domain but cannot transfer knowledge to other areas. This form of AI operates under a constrained set of functions, mimicking human intelligence only for specific problems. Concrete examples include spam filters, facial recognition software, and translation tools. While highly effective in their specialized areas, they lack general cognitive abilities and common-sense reasoning.

- **Current state of AI technology:** This represents the vast majority of AI systems in use today, driving innovation across various industries.
- **Highly specialized performance:** Their effectiveness comes from being trained on vast amounts of data for a singular purpose, leading to high accuracy within that scope.
- **Task-specific optimization:** Their algorithms are fine-tuned to solve particular problems, making them invaluable for automation and data analysis.



## General AI (Strong AI)

Hypothetical AI capable of performing any intellectual task a human can. This level of AI would possess reasoning, learning, and adaptation across unlimited domains—a goal not yet realized. AGI would exhibit cognitive flexibility, problem-solving skills, and the ability to learn and adapt to new situations without explicit programming. Currently, AGI remains a subject of theoretical research and science fiction, as no systems have demonstrated human-level general intelligence. Achieving AGI would revolutionize industries but also raise profound ethical and philosophical questions about consciousness and control.

- **Theoretical future development:** AGI is considered the ultimate goal of many AI researchers, though its realization faces significant scientific and technical hurdles.
- **Human-level cognitive abilities:** This includes abstract thinking, creativity, common sense, and the capacity for self-improvement and self-awareness.
- **Cross-domain knowledge transfer:** Unlike Narrow AI, AGI would be able to apply learning from one domain to entirely different ones, similar to how humans learn.

# Common Machine Learning Algorithms

Machine learning encompasses diverse algorithmic approaches, each suited to different problem types and data characteristics. Understanding these methods helps practitioners select the right tool for their specific challenge.



## Supervised Learning

- Linear Regression
- Logistic Regression
- Support Vector Machines
- Decision Trees
- Random Forests



## Unsupervised Learning

- K-Means Clustering
- Hierarchical Clustering
- Principal Component Analysis
- Anomaly Detection



## Reinforcement Learning

- Q-Learning
- Deep Q-Networks
- Policy Gradients
- Actor-Critic Methods

# Neural Network Architectures



## Convolutional Neural Networks (CNNs)

Specialized for image and spatial data processing. CNNs use convolutional layers to automatically detect features like edges, textures, and objects through hierarchical pattern recognition.



## Recurrent Neural Networks (RNNs)

Designed for sequential data like time series and language. RNNs maintain internal memory to process sequences, making them ideal for tasks requiring temporal context and order.



## Transformers

State-of-the-art models using self-attention mechanisms for NLP and beyond. Transformers like BERT and GPT revolutionized language understanding by processing entire sequences in parallel.

# Natural Language Processing (NLP) Essentials

Natural Language Processing enables machines to understand, interpret, and generate human language. NLP combines linguistics, computer science, and AI to bridge the gap between human communication and machine comprehension, powering applications from search engines to virtual assistants.

Modern NLP leverages deep learning to capture semantic meaning, context, and linguistic nuances that traditional rule-based systems struggled with. The field has experienced revolutionary progress with transformer models that achieve near-human performance on many language tasks.



## Tokenization

Breaking text into words or subwords for processing



## Embeddings

Vector representations  
(Word2Vec, GloVe, BERT)



## Applications

Classification, sentiment analysis, translation

# Generative AI and Large Language Models (LLMs)

Generative AI represents a paradigm shift in artificial intelligence, moving beyond classification and prediction to content creation. These systems learn patterns from vast datasets and generate novel outputs—text, images, audio, code, and more—by remixing and recombining learned information in creative ways.

## Large Language Models

Models like GPT-4, Claude, and Gemini generate human-like text and support sophisticated conversational AI. They understand context, follow instructions, and produce coherent responses across countless topics and formats.

## Key Applications

- Intelligent chatbots and virtual assistants
- Content creation and copywriting
- Code generation and debugging
- Document summarization
- Language translation

📌 **Important:** While LLMs produce impressive outputs, they can generate plausible-sounding but incorrect information. Always verify critical information from authoritative sources.

# Data Fundamentals for AI

Data serves as the foundation for all AI systems. The quality, quantity, and structure of data directly determine model performance. Understanding data fundamentals helps practitioners build more effective AI solutions and avoid common pitfalls.



High-quality data is clean, representative, and relevant to the problem at hand. Poor data leads to biased, inaccurate models regardless of algorithmic sophistication.

## Dataset

Collection of data points used for training or evaluation. Datasets should be large enough to capture pattern diversity and split into training, validation, and test sets to assess generalization.

## Features

Individual measurable properties or characteristics used as input. Features represent the information the model uses to make predictions—selecting informative features is crucial for performance.

## Labels

Correct outputs used in supervised learning to guide training. Labels provide the "ground truth" that models learn to predict, requiring careful curation to ensure accuracy.

# Model Evaluation Metrics

Evaluating model performance requires appropriate metrics that align with business objectives and problem characteristics. Different tasks demand different evaluation approaches to properly assess model quality and reliability.

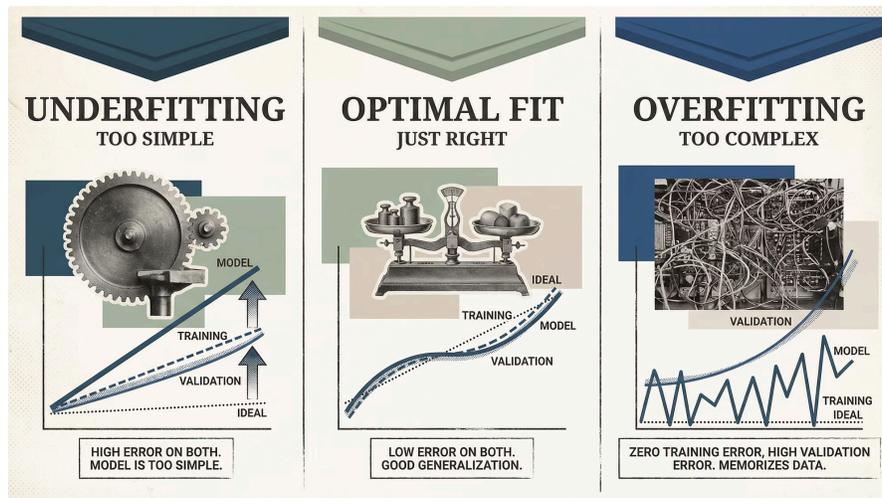
## Classification Metrics

- **Accuracy:** Overall correctness. This metric measures the proportion of correctly classified instances out of the total. It is calculated as  $(\text{True Positives} + \text{True Negatives}) / \text{Total observations}$ . Accuracy is intuitive and easy to understand but can be misleading for imbalanced datasets.
- **Precision:** Correctness of positive predictions. Precision quantifies the proportion of true positive predictions among all instances predicted as positive. It is calculated as  $\text{True Positives} / (\text{True Positives} + \text{False Positives})$ . This metric is crucial when the cost of false positives is high, such as in spam detection or medical diagnoses where incorrect positive results are problematic.
- **Recall:** Coverage of actual positives. Recall (also known as Sensitivity) measures the proportion of actual positive instances that were correctly identified by the model. It is calculated as  $\text{True Positives} / (\text{True Positives} + \text{False Negatives})$ . Recall is vital when the cost of false negatives is high, for example, in fraud detection or disease screening where missing an actual positive has severe consequences.
- **F1-Score:** Harmonic mean of precision and recall. The F1-Score provides a single metric that balances both precision and recall, offering a good measure for models with uneven class distribution. It is calculated as  $2 * (\text{Precision} * \text{Recall}) / (\text{Precision} + \text{Recall})$ . Use F1-Score when you need to consider both false positives and false negatives equally important.
- **ROC-AUC:** Trade-off between true and false positives. The Receiver Operating Characteristic (ROC) curve plots the True Positive Rate (Recall) against the False Positive Rate across various threshold settings. AUC (Area Under the Curve) quantifies the model's ability to distinguish between classes, with a higher AUC indicating better overall separability. It's particularly useful for evaluating binary classifiers across all possible classification thresholds, providing insight into the model's discriminative power independent of any specific threshold.

## Regression Metrics

- **MAE:** Mean Absolute Error. MAE measures the average magnitude of the errors in a set of predictions, without considering their direction. It is calculated as the average of the absolute differences between predicted and actual values. MAE is useful when all errors should be weighted linearly and is less sensitive to outliers compared to MSE.
- **MSE:** Mean Squared Error. MSE measures the average of the squares of the errors, giving more weight to larger errors. It is calculated as the average of the squared differences between predicted and actual values. While commonly used, MSE can be heavily influenced by outliers due to the squaring of error terms.
- **RMSE:** Root Mean Squared Error. RMSE is the square root of the MSE, bringing the error back to the same units as the target variable, making it more interpretable. It provides a good measure of accuracy where larger errors have a disproportionately larger effect. RMSE is widely used and often preferred over MAE when significant errors are particularly undesirable.
- **R-Squared:** Proportion of variance explained. R-Squared, also known as the Coefficient of Determination, measures the proportion of the variance in the dependent variable that can be predicted from the independent variable(s). Values range from 0 to 1, where 1 indicates that the model perfectly predicts the variation in the target. It helps assess how well the model's predictions approximate the real data points.

# Overfitting and Underfitting



## The Challenge

Balancing model complexity to capture true patterns without memorizing noise is a fundamental challenge in machine learning. Finding this sweet spot requires careful experimentation and validation.

## Understanding the Problems

**Overfitting** occurs when a model learns noise and specific details of training data, performing poorly on new data. The model becomes too specialized to its training examples.

**Underfitting** happens when a model is too simple to capture underlying patterns, resulting in poor performance on both training and new data.

## Solutions to Overfitting



### Regularization (L1, L2)

Penalizes large weights to constrain model complexity



### Dropout

Randomly disables neurons during training to prevent co-adaptation



### Early Stopping

Halts training when validation performance stops improving

# Feature Engineering and Data Preprocessing

Raw data rarely comes in a form ready for machine learning. Data preprocessing and feature engineering transform messy, real-world data into clean, structured inputs that models can effectively learn from. This critical phase often determines project success more than algorithm selection.

**1**

## Data Cleaning

Handling missing values through imputation or removal, detecting and addressing outliers, and removing duplicates to ensure data quality and consistency.

**2**

## Encoding

Transforming categorical data into numerical representations. One-Hot Encoding creates binary columns for each category; Label Encoding assigns integer values to categories.

**3**

## Scaling

Normalization and standardization ensure features have comparable ranges, preventing features with larger scales from dominating model learning and improving convergence.

# Optimization Techniques in AI

Optimization algorithms are the engines that drive machine learning, iteratively adjusting model parameters to minimize prediction errors. Understanding these techniques helps practitioners train models more efficiently and achieve better performance.

1

## Gradient Descent

Core algorithm that minimizes loss by updating parameters in the direction of steepest descent. Each iteration moves parameters toward lower error values.

2

## Stochastic Variants

Batch, Mini-batch, and Stochastic Gradient Descent trade off between computational efficiency and convergence stability by using different data subset sizes.

3

## Advanced Optimizers

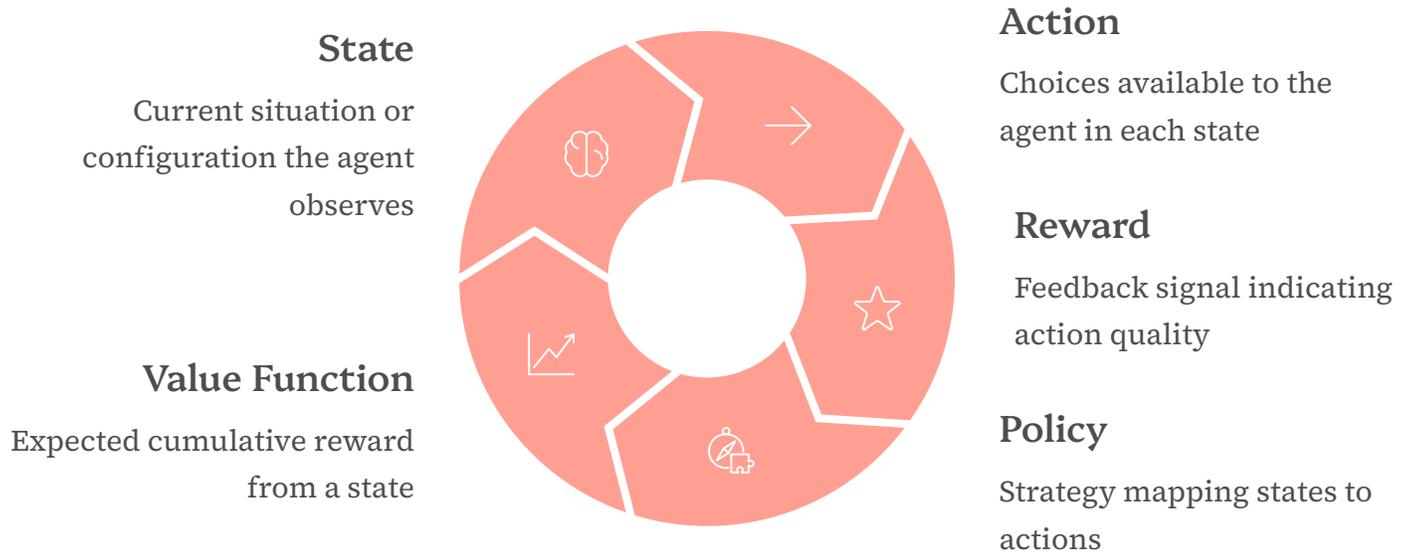
Adam, RMSprop, and AdaGrad incorporate adaptive learning rates and momentum to improve convergence speed and handle challenging optimization landscapes.

## Why It Matters

Optimizer selection significantly impacts training time, final performance, and model stability. Modern optimizers like Adam combine multiple techniques to reliably train deep networks with minimal hyperparameter tuning.

# Reinforcement Learning Concepts

Reinforcement Learning takes a fundamentally different approach from supervised learning. Instead of learning from labeled examples, an RL agent learns optimal behaviors through trial and error, receiving rewards or penalties based on its actions within an environment.



## Popular Algorithms

Q-Learning learns action values through experience; Deep Q-Networks combine Q-Learning with neural networks for complex environments; Policy Gradient methods directly optimize the policy for maximum expected reward.

# Ethical Considerations and Bias in AI

AI systems can inherit and amplify biases present in training data, leading to discriminatory outcomes that affect real people's lives. Historical biases in data—reflecting societal inequalities—become encoded in model predictions, perpetuating unfairness across hiring, lending, criminal justice, and other critical domains.

Transparency, fairness, and accountability are essential pillars of responsible AI development. Organizations must implement rigorous testing for bias, ensure diverse representation in training data and development teams, and establish clear accountability frameworks for AI decisions that impact individuals and communities.

## Bias Detection

Regular audits and fairness metrics to identify disparate impact across demographic groups

## Explainability

Making AI decisions interpretable and understandable to affected individuals and regulators

## Human Oversight

Maintaining meaningful human control over high-stakes AI decisions

Ongoing research focuses on developing techniques for bias mitigation, creating more interpretable models, and establishing governance frameworks that protect individual rights while enabling AI innovation.

# AI Tools and Frameworks Overview

## Popular Frameworks

- **TensorFlow:** Google's comprehensive ML platform. It's a powerful open-source library for numerical computation and large-scale machine learning, offering tools for building and deploying complex deep learning models across various platforms. Ideal for production environments and large-scale deployments.
- **PyTorch:** Research-friendly dynamic framework. Known for its flexibility and ease of use, PyTorch allows for dynamic computation graphs, making it popular in academic research and rapid prototyping of deep learning models. It excels in scenarios requiring custom model architectures and experimental development.
- **Scikit-learn:** Classical ML algorithms. A robust and simple library for traditional machine learning tasks like classification, regression, clustering, and dimensionality reduction. It's excellent for beginners and for quickly implementing standard ML algorithms on tabular data.
- **Keras:** High-level neural network API. Designed for fast experimentation with deep neural networks, Keras provides a user-friendly interface that runs on top of TensorFlow, PyTorch, or Theano. It's perfect for quickly building and testing deep learning models, especially for those new to the field.

## Cloud AI Services

- **AWS SageMaker:** End-to-end ML platform. AWS SageMaker provides a complete suite of tools for every step of the machine learning workflow, from data labeling and model training to deployment and monitoring. It's a highly scalable solution for enterprises looking to operationalize ML models within the AWS ecosystem.
- **Google AI Platform:** Integrated ML tools. Google Cloud's AI Platform offers a range of services for building, deploying, and managing ML models, leveraging Google's expertise in AI. It's particularly strong for integrating with other Google Cloud services and for applications requiring specialized AI services like vision and natural language processing.
- **Azure ML:** Enterprise ML solutions. Microsoft Azure Machine Learning is a cloud-based environment that helps you build, train, and deploy machine learning models quickly. It's well-suited for businesses already invested in Microsoft's ecosystem, providing strong integration with other Azure services and enterprise-grade security and governance features.

# The Future of AI: Trends and Opportunities

Artificial intelligence stands at an inflection point, with rapid advances opening unprecedented opportunities across industries and domains. The coming years will see AI become more capable, accessible, and integrated into our daily lives in ways we're only beginning to imagine.

## Multimodal AI



Systems that seamlessly combine text, image, audio, and video understanding, enabling richer human-computer interaction and more comprehensive reasoning across data types.

## Industry Transformation



Revolutionary growth in healthcare diagnostics, financial services, autonomous systems, drug discovery, and creative industries—AI augmenting human capabilities across professions.

## Human-AI Collaboration



Emphasis on AI democratization and accessibility, empowering broader populations to leverage AI tools. Focus on human-AI partnership rather than replacement.

"The goal is not to replace human intelligence, but to augment it—creating AI systems that enhance human creativity, productivity, and decision-making while addressing society's greatest challenges."

As AI capabilities grow, so does the importance of responsible development, ensuring these powerful technologies benefit humanity broadly and equitably.

# CERTIFIED ARTIFICIAL INTELLIGENCE FOUNDATION



## ABOUT GSDC CERTIFICATION



### EBOOK

Extensive and exclusive Ebook created by world's experts to help you with understanding core concepts.



### LEARNING MATERIALS

Get access to learning materials such as videos, ebooks, templates, and practice exams, which will help you clear the certification exam.



### CREATED BY EXPERTS

GSDC certifications are created and authored by world's leading experts in the field.

## LEARNING OBJECTIVE

- Gain insights into autonomous decision-making processes
- Apply knowledge using ready-to-implement templates
- Demonstrate ability to work with Agentic AI models
- Validate your skills wit

Enroll now with the code **LEARN20** To avail **20%** discount

**Enroll Now**