NEC Registration Examination

Computer Engineering (ACtE)

Chapter 9: Artificial Intelligence and Neural Networks (ACtE09)

Computer Engineering (ACtE)
Software Engineering
Information Technology

Complete MCQ-Focused Notes — Key Points • Diagrams • Easy Explanations • Practice MCQs

COMPLETE NOTES WITH SOLVED MCQ AND PYQS

EASY EXPLANATION



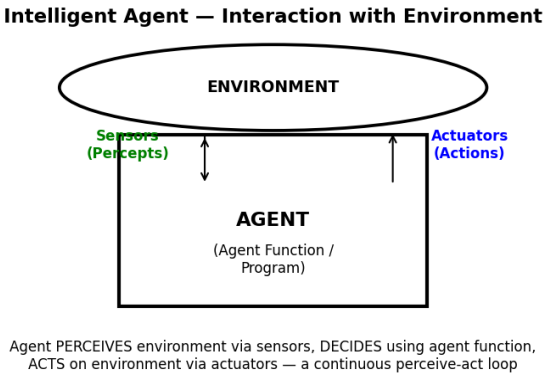
EASY EXPLANATION

9.1 Introduction to AI and Intelligent Agent

Concept & Foundations of AI

- **Artificial Intelligence:** the study of making machines **perceive, reason, learn, and act** in ways that would be considered intelligent if done by a human.
- **AI Perspectives:** `Thinking Humanly, Thinking Rationally, Acting Humanly (Turing Test), Acting Rationally` — the 4 classic ways to define AI.
- **Foundations of AI:** drawn from `Philosophy, Mathematics, Psychology, Neuroscience, Linguistics, Computer Engineering`.
- **Applications:** `expert systems, robotics, NLP, computer vision, game playing, autonomous vehicles, recommendation systems`.

Intelligent Agents



Agent PERCEIVES environment via sensors, DECIDES using agent function, ACTS on environment via actuators — a continuous perceive-act loop

- **Agent:** anything that **perceives** its environment via sensors and **acts** upon it via actuators.
- **PEAS description:** `Performance measure, Environment, Actuators, Sensors` — the standard framework to specify a task environment for an agent.

Types of Agents

Agent Type	Key Idea
Simple Reflex	Acts only on CURRENT percept using condition-action rules; no memory of the past
Model-Based Reflex	Maintains an INTERNAL MODEL/STATE of the world to handle partial observability
Goal-Based	Chooses actions that help achieve an explicit GOAL — considers the future
Utility-Based	Chooses actions that maximize a UTILITY function (best among multiple ways to reach a goal)

Environment Types

- **Deterministic vs Stochastic:** outcome fully determined by current state+action, vs involves randomness/uncertainty.
- **Static vs Dynamic:** environment doesn't change while agent deliberates, vs can change during deliberation.
- **Observable vs Semi/Partially-observable:** agent's sensors give **complete** state info, vs only **partial** info.
- **Single-agent vs Multi-agent:** only one agent acting, vs multiple agents (which may cooperate or compete).

Easy Explanation

A chess game (with a clock) is Deterministic + Static + (Semi-)Observable + Multi-agent; driving a car is Stochastic + Dynamic + Partially observable + Multi-agent — practicing classifying real scenarios this way is the most common PEAS/environment MCQ style.

Probable MCQs: AI Concepts & Agents

1. The PEAS framework stands for:

a) Process, Environment, Action, Sensor
b) Performance measure, Environment, Actuators, Sensors
c) Program, Engine, Agent, System
d) Perception, Efficiency, Automation, Speed
✓Answer: b) Performance measure, Environment, Actuators, Sensors
Why: PEAS is the standard way to describe a task environment for designing a rational agent.

2. An agent that chooses actions based on an explicit goal it wants to achieve is a:

a) Simple reflex agent
b) Model-based reflex agent
c) Goal-based agent
d) None of these
✓Answer: c) Goal-based agent
Why: Goal-based agents plan/act with a specific target outcome in mind, not just fixed rules.

3. An environment where the outcome of an action cannot be fully predicted is called:

a) Deterministic
b) Stochastic
c) Static
d) Observable
✓Answer: b) Stochastic
Why: Stochastic environments involve inherent randomness or uncertainty in outcomes.

4. A Utility-Based agent differs from a Goal-Based agent because it:

a) Has no goals at all
b) Selects the BEST action among several ways to reach a goal, using a utility measure
c) Never considers the future
d) Only works in single-agent environments
✓Answer: b) Selects the BEST action among several ways to reach a goal, using a utility measure
Why: Utility-based agents rank outcomes by desirability, not just checking if a goal is met.

5. A Simple Reflex agent makes decisions based on:

a) A model of the world's history
b) Only the current percept via condition-action rules
c) A long-term utility function
d) Communication with other agents
✓Answer: b) Only the current percept via condition-action rules
Why: Simple reflex agents react to the current percept alone, without maintaining internal state.

EASY EXPLANATION

9.2 Problem Solving and Searching Techniques

Problem Formulation

- A search problem is defined by: `Initial state, Actions, Transition model, Goal test, Path cost`.
- **Constraint Satisfaction Problem (CSP):** find values for variables satisfying a set of **constraints** (e.g. Sudoku, map coloring, N-Queens).

Uninformed (Blind) Search

Technique	Key Idea
BFS	Explores level by level using a QUEUE; finds shortest path (unweighted); high memory use
DFS	Explores as deep as possible using a STACK/recursion; low memory, but may not find shortest path
Depth-Limited Search (DLS)	DFS with a maximum depth limit — avoids infinite paths
Iterative Deepening Search (IDS)	Repeats DLS with increasing depth limits — combines BFS's completeness with DFS's low memory
Bidirectional Search	Searches simultaneously FORWARD from start and BACKWARD from goal, meeting in the middle

Informed (Heuristic) Search

- **Greedy Best-First Search:** expands the node that **appears closest to the goal** using heuristic h(n) alone — fast, but not guaranteed optimal.

A* Search: f(n) = g(n) + h(n)

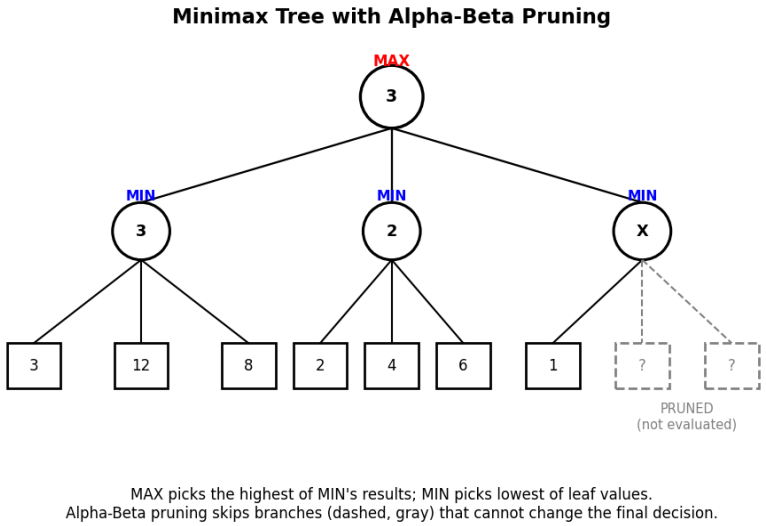
g(n) = actual cost from start to n; h(n) = estimated (heuristic) cost from n to goal

- A* is **optimal and complete** if the heuristic h(n) is **admissible** (never overestimates the true cost).
- **Hill Climbing:** greedily moves to the **neighboring state with the best value** — simple but can get stuck at **local maxima, plateaus, or ridges**.
- **Simulated Annealing:** like hill climbing, but **occasionally accepts worse moves** (probability decreasing over time, mimicking cooling metal) — helps escape local maxima.

Easy Explanation

A is essentially 'Greedy Best-First + Dijkstra' — it uses the ACTUAL cost so far (g) PLUS an ESTIMATE of what's left (h), giving both efficiency and optimality (when h never overestimates).*

Adversarial Search — Game Playing



Minimax: MAX player picks the move maximizing the outcome;
MIN player (opponent) picks the move minimizing it

- **Alpha-Beta Pruning:** an optimization of minimax that **skips (prunes) branches** that cannot possibly influence the final decision — **same result as minimax, less computation**.

? Probable MCQs: Search Techniques

1. Which uninformed search strategy is guaranteed to find the shortest path in an unweighted graph?

a) DFS
b) BFS
c) Depth-Limited Search
d) None of these
✓Answer: b) BFS
Why: BFS explores level by level, guaranteeing the shortest path is found first in an unweighted graph.

2. A* search evaluates nodes using the function:

a) f(n) = h(n) only
b) f(n) = g(n) only
c) f(n) = g(n) + h(n)
d) f(n) = g(n) − h(n)
✓Answer: c) f(n) = g(n) + h(n)
Why: A combines the actual cost so far g(n) with the heuristic estimate h(n).*

3. Hill Climbing search can get stuck at a:

a) Global maximum only
b) Local maximum, plateau, or ridge
c) Random state only
d) The goal state prematurely
✓Answer: b) Local maximum, plateau, or ridge
Why: Hill climbing's greedy nature can trap it at local maxima, flat plateaus, or ridges.

4. Simulated Annealing helps escape local maxima by:

a) Never accepting worse moves

EASY EXPLANATION

b) Occasionally accepting worse moves with decreasing probability over time

c) Always choosing randomly

d) Stopping the search immediately

✓Answer: b) Occasionally accepting worse moves with decreasing probability over time

Why: SA allows some worse moves early on (high 'temperature'), reducing this over time to converge.

5. Alpha-Beta pruning, compared to plain Minimax, produces:

a) A different (worse) final decision

b) The SAME final decision, with less computation

c) A random decision

d) Only works for single-player games

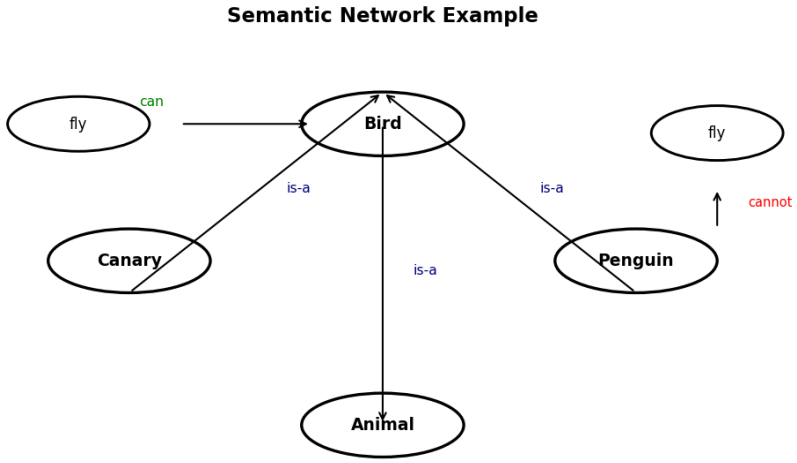
✓Answer: b) The SAME final decision, with less computation

Why: Alpha-beta pruning is a pure optimization — it yields the same result as minimax but explores fewer nodes.

EASY EXPLANATION

9.3 Knowledge Representation

Approaches to Knowledge Representation



- **Semantic Networks:** graph of **nodes (concepts) and labeled edges (relationships)** — e.g. 'is-a', 'has-property'; supports **inheritance** of properties down the hierarchy.
- **Frames:** structured representation with **slots** (attributes) and **fillers** (values) — similar to an object/record, groups related knowledge about an entity.

Propositional Logic (PL)

- **Connectives:** `¬ (NOT), ∧ (AND), ∨ (OR), → (IMPLIES), ↔ (IFF)`.
- **Tautology:** a formula that is **TRUE under every interpretation** (e.g. P ∨ ¬P).
- **Validity:** a formula is valid if it is a tautology; an argument is valid if the conclusion follows necessarily from the premises.
- **WFF (Well-Formed Formula):** a syntactically correct logical expression.
- **Resolution:** an inference rule — from `(A ∨ B)` and `(¬B ∨ C)`, derive `(A ∨ C)` — used for automated theorem proving (usually on formulas in **Conjunctive Normal Form, CNF**).

Predicate Logic (First-Order Predicate Logic, FOPL)

- Extends PL with **objects, predicates, and quantifiers**: `∀ (for all/universal), ∃ (there exists/existential)`.
- **Unification:** finding a **substitution** that makes two logical expressions identical — key step before applying resolution in FOPL.
- **Resolution refutation:** proves a goal by **assuming its negation**, adding it to the knowledge base, and deriving a contradiction (empty clause) via resolution.

Probabilistic Reasoning

Bayes' Rule: $P(A|B) = [P(B|A) \cdot P(A)] / P(B)$

- **Bayesian Network:** a **Directed Acyclic Graph (DAG)** where nodes = random variables, edges = direct dependencies — each node has a **conditional probability table (CPT)** given its parents.
- **Reasoning in Belief Networks:** computing the probability of some variables **given evidence** about others, by propagating probabilities through the network.

EASY EXPLANATION

Easy Explanation

Bayes' Rule in plain words: P(A|B) tells you how likely A is, GIVEN that you've observed B — it 'flips' a known conditional probability P(B|A) around using the prior probabilities P(A) and P(B).

Probable MCQs: Knowledge Representation

1. In a semantic network, 'is-a' links are primarily used to represent:

a) Random unrelated facts
b) Hierarchical/inheritance relationships between concepts
c) Only numerical data
d) Database queries
✓Answer: b) Hierarchical/inheritance relationships between concepts
Why: 'is-a' links establish a class hierarchy allowing property inheritance (e.g. Canary is-a Bird).

2. A tautology in propositional logic is a formula that is:

a) Always false
b) True under some interpretations only
c) True under EVERY interpretation
d) Undefined
✓Answer: c) True under EVERY interpretation
Why: A tautology holds true regardless of the truth values assigned to its variables.

3. In First-Order Predicate Logic, unification is the process of:

a) Converting logic to English
b) Finding a substitution that makes two expressions identical
c) Removing all quantifiers
d) Drawing a semantic network
✓Answer: b) Finding a substitution that makes two expressions identical
Why: Unification finds variable substitutions so that two logical expressions match, enabling resolution.

4. A Bayesian Network is structured as a:

a) Binary search tree
b) Directed Acyclic Graph (DAG) of random variables
c) Undirected cyclic graph
d) Simple linked list
✓Answer: b) Directed Acyclic Graph (DAG) of random variables
Why: Bayesian networks use a DAG structure to represent conditional dependencies between variables.

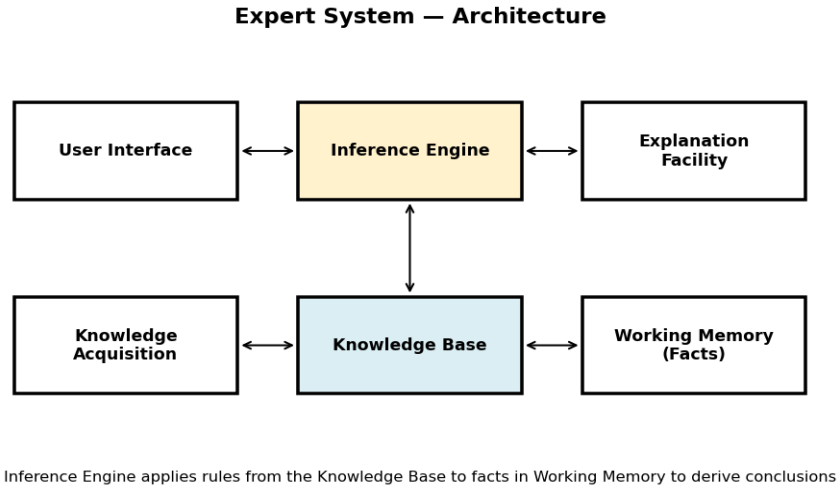
5. Bayes' Rule allows us to compute:

a) P(A) alone with no other information
b) P(A|B) from P(B|A) and the priors P(A), P(B)
c) Only deterministic outcomes
d) The syntax of a formula
✓Answer: b) P(A|B) from P(B|A) and the priors P(A), P(B)
Why: Bayes' Rule reverses a known conditional probability to compute the desired one.

EASY EXPLANATION

9.4 Expert System and Natural Language Processing

Expert Systems



Inference Engine applies rules from the Knowledge Base to facts in Working Memory to derive conclusions

- **Knowledge Base:** stores domain **facts and rules** (the expert knowledge).
- **Inference Engine:** applies **logical rules** to the knowledge base + working memory to **derive new conclusions** (forward or backward chaining).
- **Declarative knowledge:** '**knowing THAT**' — facts and relationships (e.g. 'a bird can fly').
- **Procedural knowledge:** '**knowing HOW**' — steps/rules to perform a task (e.g. how to diagnose a fault).
- **Knowledge acquisition:** process of **extracting expertise from human experts** (often the biggest bottleneck in expert system development).

Natural Language Processing (NLP)

- **NLU (Natural Language Understanding):** machine **interprets/comprehends** human language input.
- **NLG (Natural Language Generation):** machine **produces** human-readable language output.
- **Steps of NLP:** `Lexical/Morphological Analysis → Syntactic Analysis (Parsing) → Semantic Analysis → Discourse Integration → Pragmatic Analysis`.
- **Applications:** `machine translation, chatbots, sentiment analysis, speech recognition, text summarization`.
- **NLP Challenges:** `ambiguity (lexical, syntactic, semantic), context-dependence, sarcasm/idioms, multiple languages`.

Machine Vision & Robotics

- **Machine Vision stages:** `Image Acquisition → Preprocessing → Segmentation → Feature Extraction → Recognition/Interpretation`.
- **Robotics:** combines **sensing (perception), planning, and actuation (control)** to enable physical interaction with the real world — closely tied to AI's agent concept.

> **Easy Explanation**
>
> *Declarative vs Procedural: Declarative = 'the sky IS blue' (a fact); Procedural = 'HOW to change a tire' (a sequence of steps). Expert systems typically store rules combining both kinds of knowledge.*

Probable MCQs: Expert Systems & NLP

1. In an expert system, the component that applies rules to derive new conclusions is the:

a) Knowledge Base

b) Inference Engine

c) User Interface

d) Explanation Facility

✓Answer: b) Inference Engine

Why: The Inference Engine is responsible for reasoning over the Knowledge Base and Working Memory.

2. 'Knowing THAT a fact is true' describes which type of knowledge?

a) Procedural knowledge

b) Declarative knowledge

c) Heuristic knowledge only

d) None of these

✓Answer: b) Declarative knowledge

Why: Declarative knowledge represents facts and relationships, as opposed to step-by-step procedures.

3. Which NLP step involves determining the grammatical structure of a sentence?

a) Lexical analysis

b) Syntactic analysis (parsing)

c) Pragmatic analysis

d) Discourse integration

✓Answer: b) Syntactic analysis (parsing)

Why: Syntactic analysis (parsing) checks and determines the grammatical structure of the sentence.

4. NLG (Natural Language Generation) refers to a machine's ability to:

a) Understand spoken commands only

b) Produce human-readable language output

c) Recognize images

d) Perform arithmetic

✓Answer: b) Produce human-readable language output

Why: NLG is about generating coherent natural language text/speech as output.

5. In machine vision, the stage that identifies distinguishing characteristics in an image is called:

a) Image Acquisition

b) Segmentation

c) Feature Extraction

d) Preprocessing

✓Answer: c) Feature Extraction

Why: Feature Extraction pulls out key distinguishing characteristics used for recognition.

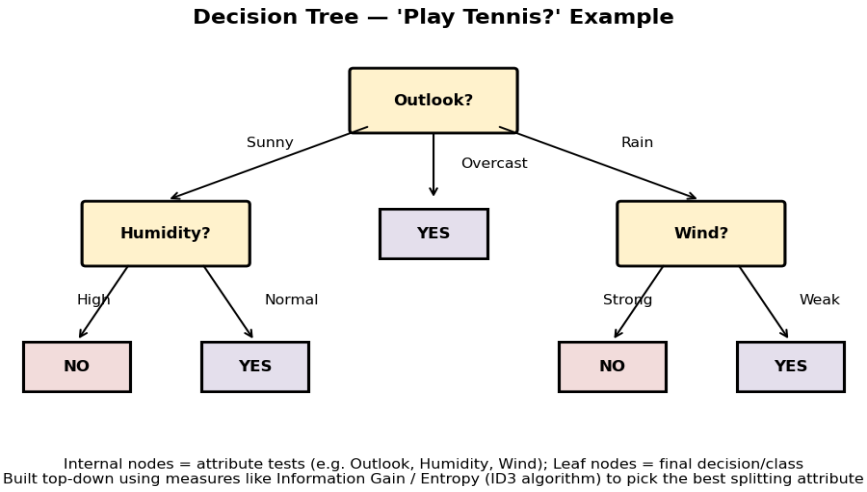
EASY EXPLANATION

9.5 Machine Learning

Types of Learning

Type	Key Idea
Supervised Learning	Learns from LABELED data (input-output pairs) — e.g. classification, regression
Unsupervised Learning	Learns patterns from UNLABELED data — e.g. clustering, association
Reinforcement Learning	Agent learns via REWARDS/PENALTIES from interacting with an environment (trial and error)

Inductive & Statistical Learning



Internal nodes = attribute tests (e.g. Outlook, Humidity, Wind); Leaf nodes = final decision/class
Built top-down using measures like Information Gain / Entropy (ID3 algorithm) to pick the best splitting attribute

- **Inductive Learning (Decision Tree):** builds a tree by **recursively splitting** data on the attribute that best separates classes — commonly measured via **Information Gain / Entropy** (ID3 algorithm).
- **Naive Bayes Model (Statistical Learning):** classifies using Bayes' theorem, **assuming all features are conditionally independent** given the class — simple, fast, surprisingly effective for text classification/spam filtering.

Fuzzy Learning

- **Fuzzy Logic:** allows **partial truth** (values between 0 and 1), unlike classical logic's strict True/False — models real-world vagueness (e.g. 'somewhat hot').
- **Fuzzy Inference System (FIS):** stages — `Fuzzification → Rule Evaluation (Inference) → Defuzzification` — converts crisp input to fuzzy, applies fuzzy rules, then converts back to a crisp output.

Genetic Algorithm (GA)

- Optimization technique inspired by **natural selection/evolution**.
- **GA Operators:** `Selection` (choose fittest individuals), `Crossover` (combine two parents' genes), `Mutation` (small random change) — applied over generations to evolve better solutions.
- **Encoding:** representing a candidate solution as a **chromosome** (often a bit string). **Fitness function:** evaluates **how good** a solution is — guides selection.

Easy Explanation

GA in one line: start with a random population of candidate solutions → repeatedly SELECT the fittest, CROSSOVER (mix) their traits, and randomly MUTATE → over generations, the population evolves toward better solutions.

? Probable MCQs: Machine Learning

1. A model trained on labeled input-output pairs uses which type of learning?

a) Unsupervised learning
b) Supervised learning
c) Reinforcement learning
d) Fuzzy learning
✓Answer: b) Supervised learning
Why: Supervised learning requires labeled data pairs (known correct outputs) during training.

2. The ID3 algorithm for building decision trees selects the splitting attribute based on:

a) Random selection
b) Information Gain / Entropy reduction
c) Alphabetical order
d) Fixed order given by the user
✓Answer: b) Information Gain / Entropy reduction
Why: ID3 picks the attribute that most reduces entropy (maximizes information gain) at each split.

3. The Naive Bayes classifier makes which key simplifying assumption?

a) All features are perfectly correlated
b) Features are conditionally independent given the class
c) There is only one feature
d) No probability is ever zero
✓Answer: b) Features are conditionally independent given the class
Why: Naive Bayes assumes conditional independence among features given the class label — a simplification that works well in practice.

4. In a Fuzzy Inference System, converting a crisp output back from fuzzy values is called:

a) Fuzzification
b) Rule evaluation
c) Defuzzification
d) Normalization
✓Answer: c) Defuzzification
Why: Defuzzification is the final step, converting fuzzy inference results into a crisp usable output.

5. In a Genetic Algorithm, the operator that combines genes from two parent solutions is called:

a) Mutation
b) Selection
c) Crossover
d) Fitness evaluation
✓Answer: c) Crossover
Why: Crossover mixes genetic material from two parents to produce offspring solutions.

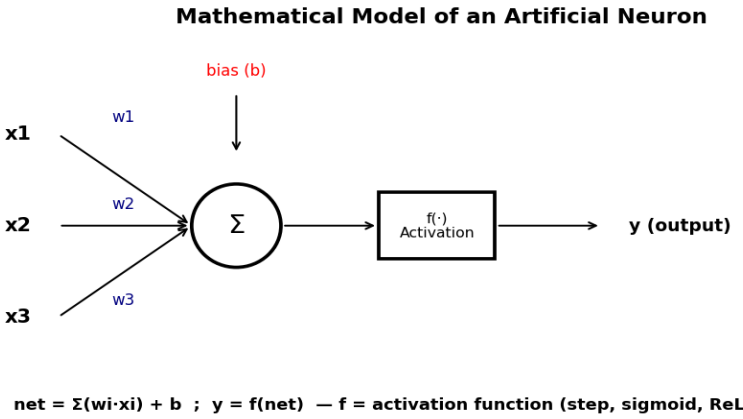
EASY EXPLANATION

9.6 Neural Networks

Biological vs Artificial Neural Networks

- Biological neuron: `dendrites (inputs) → soma/cell body (processing) → axon (output) → synapse (connection to next neuron)`.
- **McCulloch-Pitts Neuron:** the **earliest mathematical model** of a neuron (1943) — binary threshold unit: outputs 1 if weighted sum of inputs **exceeds a threshold**, else 0.

Mathematical Model of ANN



net = Σ(wi · xi) + b

output y = f(net)

Activation Function	Range / Shape
Step (threshold)	Binary output: 0 or 1
Sigmoid	Smooth S-curve, output in (0,1)
Tanh	Smooth S-curve, output in (−1,1)
ReLU	max(0, x) — output 0 for negative input, linear for positive

Perceptron & Learning Rules

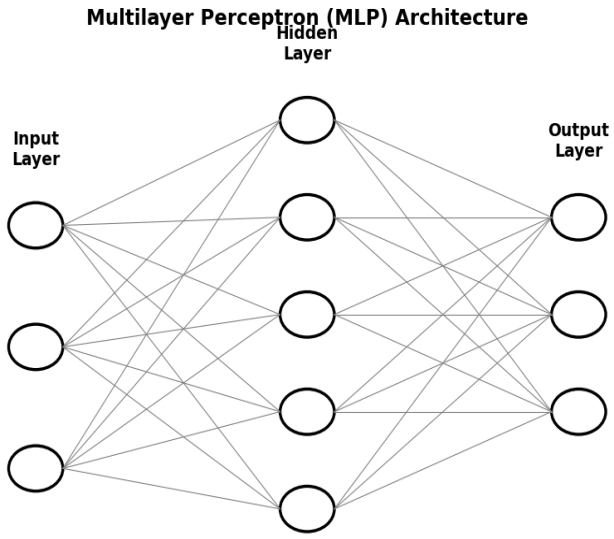
- **Perceptron:** single-layer network that can learn **linearly separable** patterns (e.g. AND, OR) — **cannot learn XOR** (a classic limitation highlighted by Minsky & Papert).

Delta Rule (Widrow-Hoff): $\Delta w = \eta \cdot (\text{target} - \text{output}) \cdot \text{input}$

η = Learning Rate (controls step size of weight updates)

- **Gradient Descent:** iteratively adjusts weights **in the direction that reduces error** (opposite to the gradient of the error function).
- **Hebbian Learning:** '**neurons that fire together, wire together**' — strengthens the connection between two neurons that are **simultaneously active**.
- **Adaline (Adaptive Linear Neuron):** uses a **linear activation** during training (with the delta rule) — predecessor to modern gradient-based learning.

Multilayer Networks & Backpropagation



Trained via BACKPROPAGATION: errors from output layer propagate backward, adjusting weights via gradient descent to minimize total error

- **Backpropagation Algorithm:** trains multilayer networks by **propagating the output error backward** through the layers, updating weights via gradient descent to minimize total error — solves the XOR problem that a single perceptron cannot.
- **Hopfield Network:** a **recurrent** (fully connected, symmetric-weight) neural network used as **associative/content-addressable memory** — can recall a complete stored pattern from a partial/noisy input.

Easy Explanation

Key limitation to remember: a SINGLE-layer Perceptron can only solve LINEARLY SEPARABLE problems (like AND/OR) and famously FAILS on XOR. Adding a HIDDEN layer (Multilayer Perceptron) + Backpropagation solves this — this XOR limitation is one of the most frequently tested facts in neural network MCQs.

Probable MCQs: Neural Networks

1. The McCulloch-Pitts neuron model produces which type of output?

a) Continuous real-valued output

b) Binary output (0 or 1) based on a threshold

c) Only negative values

d) Random output

✓Answer: b) Binary output (0 or 1) based on a threshold

Why: The McCulloch-Pitts model is a binary threshold unit — output is 0 or 1 based on whether the weighted sum exceeds a threshold.

2. A single-layer Perceptron is famously UNABLE to solve which classic problem?

a) AND

b) OR

c) XOR

d) NOT

✓Answer: c) XOR

Why: XOR is not linearly separable, which a single-layer perceptron cannot represent — this requires a multilayer network.

3. Hebbian Learning is summarized by the phrase:

a) 'Survival of the fittest'

b) 'Neurons that fire together, wire together'

c) 'Garbage in, garbage out'

d) 'Divide and conquer'

✓Answer: b) 'Neurons that fire together, wire together'

Why: Hebbian learning strengthens connections between neurons that activate simultaneously.

4. Backpropagation trains a neural network by:

a) Randomly setting all weights once

b) Propagating output error backward and updating weights via gradient descent

c) Ignoring the output error entirely

d) Only adjusting the input layer

✓Answer: b) Propagating output error backward and updating weights via gradient descent

Why: Backpropagation computes the error at the output and propagates it backward to update weights throughout the network.

5. A Hopfield Network is primarily used as a form of:

a) Feedforward classifier only

b) Associative/content-addressable memory

c) Decision tree

d) Genetic algorithm

✓Answer: b) Associative/content-addressable memory

Why: Hopfield networks store patterns and can recall a complete pattern from partial/noisy input, functioning as associative memory.

EASY EXPLANATION