CS24307 • MODULE 1

Module 1: Intelligent Agents & PEAS Framework

Topics 1 to 7 • Foundations, Evolution, Rationality & 5 Agent Types

Module 1 Table of Contents (Topics 1 to 7)

- **Topic 1:** Foundations & Definitions of AI
- **Topic 2:** Evolution, AI Winters & Milestones
- **Topic 3:** Intelligent Agents Architecture
- **Topic 4:** Concept of Rationality & Performance
- **Topic 5:** PEAS & Taxonomy of Environments
- **Topic 6:** 5 Classical Agent Architectures
- **Topic 7:** Real-World Domains & University Exam Bank

Topic 1: Foundational Definitions & The Four AI Perspectives

Artificial Intelligence (AI) is the science and engineering of synthesizing computational artifacts capable of performing cognitive tasks traditionally associated with human intellect: perception, inductive/deductive reasoning, learning from experiential data, goal-directed planning, natural language comprehension, and autonomous decision-making under uncertainty.

Dimension	Human-Centric Dimension (Empirical / Cognitive)	Rationality-Centric Dimension (Ideal / Normative)
Reasoning (Thought)	**Thinking Humanly:** Cognitive Science approach. Formulates computational models of human mental processes via fMRI neuroscience and introspection (e.g. Newell & Simon's General Problem Solver).	**Thinking Rationally:** Laws of Thought approach. Formalizes reasoning using strict deductive logic and syllogisms ($\forall x(P(x) \rightarrow Q(x))$) originating with Aristotle.
Behavior (Action)	**Acting Humanly:** The Turing Test approach (Alan Turing, 1950). An entity is intelligent if an interrogator cannot distinguish its conversational responses from a human.	**Acting Rationally (Modern Standard):** Rational Agent approach (Russell & Norvig). An agent operates to maximize its *expected performance measure* given its percept sequence and built-in knowledge.

🏛️ **Deep Dive: The 6 Sub-disciplines Required for the Total Turing Test**

1. **Natural Language Processing (NLP):** Comprehending, parsing, and generating human spoken/written communication.
2. **Knowledge Representation (KR):** Storing structured ontological assertions, episodic memories, and causal rules.
3. **Automated Reasoning:** Drawing logically sound inferences, answering queries, and resolving contradictions.
4. **Machine Learning (ML):** Adapting internal statistical weights to extrapolate patterns and survive novel scenarios.
5. **Computer Vision (Total Turing Test):** Perceiving and segmenting 3D physical objects from visual sensor streams.
6. **Robotics & Actuation (Total Turing Test):** Manipulating physical objects and navigating complex terrain.

Topic 2: Historical Evolution, Dartmouth 1956 & AI Winters



Figure 1.1: Chronological Evolution of Artificial Intelligence from Symbolic Logic to Deep Generative Models

Topic 3 & 4: Intelligent Agents & Formal Concept of Rationality

An **Agent** is anything that perceives its environment through physical or software sensors and acts upon that environment via actuators. Formally, an agent is mathematically characterized by an **Agent Function** f:

The Mathematical Agent Function:

$$\mathbf{f} : \mathcal{P}^* \longrightarrow \mathcal{A}$$

Where $\mathcal{P}^*$ represents the history of all percept sequences perceived up to the current discrete timestep t, and $\mathcal{A}$ is the set of executable actions. The concrete implementation of f running on physical compute hardware is the **Agent Program**.

⚖️ **Rationality vs. Omniscience**

• **Rationality:** Maximizes *expected performance* given the historical percept sequence $\mathcal{P}^*$ and prior built-in knowledge. It does NOT require perfection.
• **Omniscience:** Knows the *actual outcome* of its actions in advance (impossible in non-deterministic or partially observable real-world environments).

Topic 5: PEAS Specification & Environmental Taxonomy

To engineer an intelligent system, one must formally define its **PEAS Framework**:

System Domain	Performance Measure (P)	Environment (E)	Actuators (A)	Sensors (S)
Autonomous Taxi	Safety, speed, legal compliance, passenger comfort, fuel economy, profit.	Public roads, pedestrians, traffic lights, weather conditions.	Steering wheel, accelerator, brakes, horn, turn signals.	LiDAR, RADAR, cameras, GPS, odometer, engine sensors.
Medical Diagnosis	Patient health recovery, diagnosis accuracy, minimized costs and side effects.	Hospital ward, patient symptoms, lab test facilities.	Treatment plan, prescription dosage, surgery referral.	Blood pressure monitor, EHR records, lab results.
Chess AI (Stockfish)	Winning game ($+1$), draw (0), checkmating opponent, piece advantage.	8×8 Chess board, opponent player, clock timer.	Move piece command ($e2 \rightarrow e4$).	Board state matrix representation.
Automated Warehouse Robot	Items picked/hour, zero item drops, collision-free transit, battery life.	Warehouse aisles, shelving racks, human workers.	Differential drive wheels, robotic gripping arm.	Ultrasonic rangefinders, barcode scanners, depth camera.

Topic 6: The 5 Classical Agent Structural Architectures

Agent Architecture	Internal Decision Logic	Pros & Limitations
1. Simple Reflex Agent	Selects actions strictly based on the *current instantaneous percept* using condition-action rules (`IF car-in-front-braking THEN apply-brakes`). Ignores history.	Extremely fast; fails completely in partially observable environments (falls into infinite loops).
2. Model-Based Reflex Agent	Maintains an internal **State** tracking how the unobserved world evolves independently and how the agent's own actions affect the world.	Can handle partial observability; requires accurate world physics modeling.
3. Goal-Based Agent	Combines internal state with explicit **Goal Descriptions** to evaluate whether candidate action sequences achieve desired end states (uses Search & Planning).	Highly flexible (goals can change dynamically); slower than reflex lookup.
4. Utility-Based Agent	Maps states to real-valued scalar **Utility Values** $U(s) \in \mathbb{R}$ to make optimal trade-offs between competing goals and risk profiles under uncertainty.	Provides mathematically optimal decision making; computationally intensive.
5. Learning Agent	Divided into 4 modules: *Learning Element* (improves performance), *Critic* (evaluates against standard), *Performance Element* (acts), and *Problem Generator* (experiments).	Can operate in initially unknown environments and achieve superhuman competence over time.

Topic 7: Master University Examination Solved Question Bank (10 Solved Questions)

Q1. Prove why a Simple Reflex Agent cannot operate rationally in a partially observable environment like the Vacuum World. (8 Marks)

In a 2-room vacuum cleaner world where the agent can sense only its current location and cleanliness (not whether the other room is dirty), if both rooms are clean and the agent has no memory, its rule `IF clean THEN MoveRight` followed in Room B by `IF clean THEN MoveLeft` produces an **infinite oscillating loop**. Without internal state memory, it cannot deduce that both rooms have already been cleaned!

Q2. Formally classify the Environment of: (a) Poker, (b) Autonomous Driving, and (c) Crossword Puzzle across all 6 environmental dimensions. (10 Marks)

- **(a) Poker:** Partially Observable (hidden cards), Stochastic (card shuffling), Sequential (bets affect later rounds), Static (rules/pot do not change while thinking), Discrete (finite chips/cards), Multi-Agent (Competitive/Adversarial).
- **(b) Autonomous Driving:** Partially Observable (blind spots, occluded vehicles), Stochastic (erratic drivers, tire slip), Sequential, Dynamic (pedestrians and traffic move while vehicle decides), Continuous (steering angle, velocity), Multi-Agent (Cooperative & Competitive).
- **(c) Crossword Puzzle:** Fully Observable (entire grid visible), Deterministic (words do not change), Sequential, Static, Discrete (letters and grid cells), Single-Agent.

Q3. Detail the components of a Learning Agent with an architectural block diagram. (8 Marks)

A Learning Agent decomposes into four distinct structural units:
1. **Learning Element:** Responsible for making improvements by analyzing feedback from the critic.
2. **Critic:** Evaluates the agent's behavior against an external performance standard (provides reward/penalty signals).
3. **Performance Element:** The operational agent that selects external actions based on percepts (e.g. reflex, model, goal, or utility core).
4. **Problem Generator:** Suggests novel exploratory actions leading to new experiences rather than just exploiting known sub-optimal paths.

Q4. What is the Chinese Room Argument by John Searle? How does it challenge Strong AI? (8 Marks)

John Searle's Chinese Room Thought Experiment (1980): A monolingual English speaker sits in a closed room with a rulebook translating incoming Chinese characters into corresponding Chinese output responses. To an outside observer, the room passes the Turing Test in Chinese. However, the human inside does not understand a single word of Chinese; they are merely performing *syntactic symbol manipulation* without *semantic intentionality or true understanding*. Searle concludes that functional execution of a computer program cannot produce genuine human-like consciousness (Strong AI).

Q5. Explain the Subsumption Architecture proposed by Rodney Brooks. How does it contrast with Classical Symbolic AI? (8 Marks)

- **Classical Symbolic AI:** Relies on a centralized "Sense-Plan-Act" pipeline with explicit symbolic internal world models.
- **Subsumption Architecture (Behavior-Based AI):** Eliminates central representations ("The world is its own best model"). The system is built from parallel, asynchronous, reactive behavioral layers (e.g. Layer 0: Avoid Obstacles; Layer 1: Wander; Layer 2: Explore). Higher layers can *suppress* inputs or *subsume* outputs of lower layers, achieving robust physical reactivity in dynamic environments without computationally prohibitive planning bottlenecks!

Q6. Compare Goal-Based Agents and Utility-Based Agents with decision-theoretic examples. (8 Marks)

• **Goal-Based Agents:** Seek binary achievement of defined goals ($G(s) \in \{\text{True}, \text{False}\}$). They lack the ability to trade off conflicting criteria (e.g. speed vs safety) or quantify partial goal satisfaction.
• **Utility-Based Agents:** Map states to a real-valued scalar utility function $U : \mathcal{S} \to \mathbb{R}$. When multiple goals conflict (e.g., getting to the airport quickly vs spending minimal fuel vs avoiding toll roads), a utility agent computes the expected utility $\mathbb{E}[U] = \sum_s P(s \mid a)U(s)$ and selects the action that maximizes expected utility (Maximum Expected Utility Principle).

Q7. Differentiate between Epistemic Actions and Pragmatic Actions with robotic exploration examples. (6 Marks)

• **Pragmatic Actions:** Actions intended to change the physical world state to bring the agent closer to its goal (e.g., vacuum cleaner picking up dirt, autonomous car accelerating).
• **Epistemic Actions:** Actions whose primary purpose is to acquire information and reduce uncertainty in a partially observable world (e.g., exploring an unmapped room, running a medical diagnostic blood test, looking around a blind corner before turning).

Q8. What are the key limitations of Expert Systems that led to the Second AI Winter? (8 Marks)

1. **Knowledge Acquisition Bottleneck:** Hand-crafting thousands of IF-THEN rules from human experts was expensive and slow.
2. **Brittleness:** Expert systems could not handle edge cases outside their narrow rule sets.
3. **Lack of Commonsense Reasoning:** Inability to make simple real-world deductions.
4. **High Maintenance Costs:** Conflicting rules caused unpredictable system behavior as rule bases grew.

Q9. Explain the concept of Bounded Rationality proposed by Herbert Simon. (6 Marks)

Perfect rationality assumes unbounded computation to find optimal actions. In the real world, agents face strict computational, memory, and time constraints. **Bounded Rationality** dictates that an agent acts as rationally as possible given its finite computational resources (satisficing rather than global optimization).

Q10. Formulate the State-Space representation of the Water Jug Problem (4-Gallon and 3-Gallon Jugs to measure 2 Gallons). (8 Marks)

• **State Space:** Ordered pair (x, y) where $x \in \{0, 1, 2, 3, 4\}$ (water in 4-gal jug) and $y \in \{0, 1, 2, 3\}$ (water in 3-gal jug).
• **Initial State:** $(0, 0)$.
• **Goal State:** $(2, y)$ for any y.
• **Production Rules / Operators:**
1. Fill 4-gal jug: $(x, y) \to (4, y)$ if $x < 4$.
2. Fill 3-gal jug: $(x, y) \to (x, 3)$ if $y < 3$.
3. Empty 4-gal jug: $(x, y) \to (0, y)$ if $x > 0$.
4. Empty 3-gal jug: $(x, y) \to (x, 0)$ if $y > 0$.
5. Pour from 3 to 4: $(x, y) \to (\min(4, x + y), y - (\min(4, x + y) - x))$.
6. Pour from 4 to 3: $(x, y) \to (x - (\min(3, x + y) - y), \min(3, x + y))$.
• **Solution Trace:** $(0, 0) \to (0, 3) \to (3, 0) \to (3, 3) \to (4, 2) \to (0, 2) \to (2, 0)$ (Goal reached!).

Topic 7.5: Philosophical Foundations & Minds vs. Machines

The philosophical foundations of artificial intelligence examine the feasibility, limits, and metaphysical nature of synthetic cognition.

Philosophical Position	Core Claim	Key Objections & Philosophical Refutations
Strong AI (Functionalism)	An appropriately programmed computer with the right inputs and outputs literally *has a mind* and possesses genuine consciousness and understanding.	Searle's Chinese Room argument: Syntax alone cannot generate intentionality or semantics ($S \not\Longrightarrow \Sigma$).
Weak AI (Instrumentalism)	Computers can only simulate human cognition and act *as if* they were intelligent, providing powerful analytical tools without actual subjective experience.	Accepted by mainstream engineering; focuses on operational competence and mathematical performance measures.
Gödel's Incompleteness Argument (Lucas / Penrose)	By Gödel's First Incompleteness Theorem, any consistent formal axiomatic system contains true mathematical statements that cannot be proven within the system. Human mathematicians can perceive these truths; hence human minds are non-algorithmic.	Computers, like humans, are not infallible; Gödel's theorem applies only to consistent, closed formal systems, whereas humans learn heuristically from experience.
Dreyfus' Critique of Symbolic AI	Human expertise relies on tacit, bodily intuitive know-how and situational context (phenomenology) rather than explicit symbolic rules and propositional logic.	Led to modern embodied robotics, connectionism, and deep reinforcement learning from sensorimotor interactions.

🏛️ Step-by-Step Analysis: Alan Turing's 1950 Objections and Responses

In his landmark 1950 paper *"Computing Machinery and Intelligence"*, Alan Turing anticipated and systematically refuted nine major objections to machine intelligence:

1. **The Theological Objection:** "Thinking is a function of man's immortal soul; God gave an immortal soul to every man and woman, but not to any other animal or machine."
   Turing's Refutation: This restricts the omnipotence of God; why couldn't God grant a soul to a computational machine if He wished?
2. **The 'Heads in the Sand' Objection:** "The consequences of machines thinking would be too dreadful. Let us hope and believe that they cannot do so."
   Turing's Refutation: An emotional fear of losing human intellectual superiority rather than a rational scientific argument.
3. **The Mathematical Objection (Gödel):** Finite machines are subject to Gödelian limits on provability.
   Turing's Refutation: Humans frequently make errors, suffer from memory bounds, and fail to prove mathematical truths; machines need not be infallible gods to match human intellect.
4. **The Argument from Consciousness:** "Not until a machine can write a sonnet or compose a concerto because of thoughts and emotions felt, and not by the chance fall of symbols, could we agree that machine equals brain." (Jefferson).
   Turing's Refutation: The only way to know what a person feels is to *be* that person (Solipsism). In ordinary life, we accept behavioral communication as evidence of other minds.
5. **Arguments from Various Disabilities:** "A machine can do many things, but it will never be able to: be kind, have a sense of humor, fall in love, make mistakes, learn from experience, enjoy strawberries and cream."
   Turing's Refutation: These limitations are artifacts of current narrow engineering rather than fundamental physical laws of computation.
6. **Lady Lovelace's Objection (1842):** "The Analytical Engine has no pretensions to originate anything. It can do whatever we know how to order it to perform."
   Turing's Refutation: Machines can surprise humans by discovering unexpected proofs and emergence from complex non-linear interactions. Machine learning allows systems to acquire rules never explicitly programmed.
7. **Argument from Continuity in the Nervous System:** The biological brain is a continuous neural medium, whereas digital computers are discrete-state machines.
   Turing's Refutation: A discrete-state machine can approximate any continuous differential equation system to arbitrary mathematical precision.
8. **The Argument from Informality of Behavior:** "It is not possible to produce a set of rules purporting to describe what a man should do in every conceivable set of circumstances."
   Turing's Refutation: Humans follow probabilistic, statistical laws of behavior rather than rigid deterministic condition-action rule tables.
9. **The Argument from Extrasensory Perception (ESP):** Telepathy and telekinesis.
   Turing's Refutation: If telepathy exists, test conditions can be engineered in Faraday cages.

🏛️ **Complete Mathematical Formalism: The Agent-Environment Interaction Cycle**

Let discrete timesteps be $t \in \{0, 1, 2, \dots\}$. The interaction between agent and environment proceeds as:

$$\mathbf{e_t} \in \mathcal{E} \xrightarrow{\text{Sensor Matrix } \mathbf{S}} \mathbf{p_t} \in \mathcal{P} \xrightarrow{\text{History Sequence } \mathbf{h_t}=(\mathbf{p_0},\dots,\mathbf{p_t})} \mathbf{a_t} = \mathbf{f}(\mathbf{h_t}) \in \mathcal{A} \xrightarrow{\text{Transition Dynamics } \mathcal{T}} \mathbf{e_{t+1}}$$

The total performance score over horizon T is given by an evaluation functional:

$$\mathbf{V}(\mathbf{f}) = \mathbb{E}_{\mathcal{T}} \left[\sum_{\mathbf{t=0}}^{\mathbf{T}} \gamma^{\mathbf{t}} \mathbf{R}(\mathbf{e_t}, \mathbf{a_t}, \mathbf{e_{t+1}}) \right]$$

A globally rational agent function $\mathbf{f}^*$ maximizes expected cumulative discounted reward:

$$\mathbf{f}^* = \arg\max_{\mathbf{f}} \mathbf{V}(\mathbf{f}) = \arg\max_{\mathbf{f}} \mathbb{E} \left[\sum_{\mathbf{t=0}}^{\infty} \gamma^{\mathbf{t}} \mathbf{R_t} \right]$$

Q13. Compare the Symbolic Paradigm (Good Old-Fashioned AI / GOFAI) with the Connectionist Paradigm. (10 Marks)

Dimension	Symbolic Paradigm (GOFAI)	Connectionist Paradigm (Neural / Sub-symbolic)
Representation	Explicit, discrete symbols, predicates, logic rules, semantic ontologies.	Distributed continuous vector embeddings, synaptic weight matrices.
Inference Mechanism	Deductive logic, theorem proving, search algorithms, unification.	Non-linear activation propagation, matrix multiplication, backpropagation gradient descent.
Noise Tolerance	Brittle; single missing rule or syntactic typo causes total failure.	Graceful degradation; robust to noisy sensors and missing features.
Interpretability	Transparent white-box reasoning; generates step-by-step resolution proofs.	Black-box models; high dimensional weights difficult for humans to audit.
Strengths	Exact mathematical reasoning, planning, rule compliance, explanation.	Perception (Vision, Audio), continuous control, pattern recognition.

Q14. Explain the Concept of an Architecture in Russell & Norvig's formulation: $\text{Agent} = \text{Architecture} + \text{Program}$. (6 Marks)

The **Architecture** is the physical computational device (CPU, GPU, FPGA, micro-controller) along with hardware sensors (cameras, LiDAR, encoders) and actuators (motors, solenoids, speakers). The **Agent Program** is the concrete algorithm or software function that runs on the architecture, taking incoming sensor raw feeds and computing output actuator voltages. The architecture must provide sufficient compute throughput and low sensor latency to allow the agent program to make decisions within real-time deadlines!

Q15. Detail the 4 Classical Types of Environment Observability and give real-world examples. (8 Marks)

1. **Fully Observable:** The agent's sensors give it access to the complete state of the environment at each point in time (e.g., Chess, Go, 8-Puzzle).
2. **Partially Observable:** Sensor noise, occlusions, or hidden opponent state prevent complete knowledge (e.g., Autonomous driving with fog/blind spots, Poker with face-down cards).
3. **Unobservable (Sensorless / Conformant):** The agent has NO sensors at all. It must find a sequence of actions that achieves the goal regardless of the initial starting state (e.g., Automated vacuum cleaner running in a completely dark room finding clean states via blind sweeping).
4. **Semi-Observable:** Only high-level aggregated signals are available (e.g., Macro-economic central banking monitoring quarterly inflation indicators).

Topic 7.6: Detailed Mathematical Properties of State Space & Perception

Formal Agent Perception History & Action Space:

$$\mathcal{P} = \{\mathbf{p_1}, \mathbf{p_2}, \ldots, \mathbf{p_k}\} \implies \mathcal{P}^* = \bigcup_{\mathbf{t=0}}^{\infty} \mathcal{P}^{\mathbf{t}}$$

$$\text{Number of Possible Agent Functions over Horizon } \mathbf{T}: \ |\mathcal{A}|^{|\mathcal{P}|^{\mathbf{T}}}$$

Combinatorial Significance: For even modest $|\mathcal{P}| = 10$ and $|\mathcal{A}| = 5$ over $T = 10$ steps, there are $5^{10^{10}}$ candidate agent functions! This astronomical search space proves why tabular lookup reflex agents are physically impossible and why compact, generalizing representations (logic, heuristics, neural networks) are mandatory.

🏛️ Step-by-Step Solved Problem: Agent Performance Metric Engineering

An autonomous street cleaning robot is deployed in an urban city center. Compare two proposed performance measures:

- **Metric 1:** Amount of dirt collected in the robot's onboard disposal bin over 24 hours.
- **Metric 2:** Cleanliness of the entire street network averaged over 24 hours.

Critical Agent Analysis:

- Under **Metric 1**, a rational agent maximizes dirt in its bin. To maximize this metric, the agent might clean up a pile of dirt, dump it back onto the clean pavement, and sweep it up again repeatedly!
- Under **Metric 2**, the agent is rewarded strictly for the *desired state of the external environment* (clean streets), preventing perverse reward hacking!
- **Core Principle (Russell & Norvig):** Design performance measures according to what you want to achieve in the environment, NOT according to how you think the agent should behave!

Q16. Explain the Concept of Software Agents (Softbots) and Web Crawling Agents. (8 Marks)

A **Softbot** is an intelligent agent whose environment is entirely digital (operating systems, web networks, databases, cloud infrastructure) rather than physical hardware. Example (Search Engine Indexing Spider):
• **P:** Freshness of search index, crawl coverage, low server load, high page pagerank relevance.
• **E:** The World Wide Web (HTTP/HTTPS protocols, HTML/DOM trees, sitemaps, robots.txt constraints).
• **A:** HTTP GET requests, URL extraction, indexing pipeline dispatch, rate-limiting sleep calls.
• **S:** HTTP status codes, HTML content strings, header response times.

Topic 7.7: Ethical Foundations, Safety & Fairness in Autonomous Systems

As autonomous agents transition from laboratory benchmarks to real-world deployment, ethical constraints and safety bounds must be formalized as mathematical invariants within utility functions.

Ethical Dimension	Theoretical Formalization	Industrial Deployment Standard
Value Alignment (Nick Bostrom)	Ensuring an agent's objective utility function $\mathcal{U}$ strictly aligns with true human intent rather than literal narrow specifications (preventing the Paperclip Maximizer catastrophe).	Inverse Reinforcement Learning (IRL), Cooperative Inverse Reinforcement Learning (CIRL).
Algorithmic Fairness	**Demographic Parity:** $P(\hat{Y} = 1 \mid A = 0) = P(\hat{Y} = 1 \mid A = 1)$; **Equalized Odds:** True positive and false positive rates are equal across protected demographic attributes A.	AI-driven hiring tools, automated credit risk scoring, judicial recidivism prediction.
Explainability & Interpretability (XAI)	Generating post-hoc explanations for black-box neural decisions via Local Interpretable Model-agnostic Explanations (LIME) and SHAP (Shapley Additive exPlanations).	EU AI Act high-risk compliance, FDA medical device diagnostic validation.

🏛️ Step-by-Step Solved Problem: Formal Mathematical Fair Decision Metric

A credit underwriting agent evaluates loan applications from two demographic groups ($A = 0, A = 1$):

- Group $A = 0$: 1000 applicants, 200 approved ($\hat{Y} = 1$), 150 repaid ($Y = 1$).
- Group $A = 1$: 2000 applicants, 600 approved ($\hat{Y} = 1$), 450 repaid ($Y = 1$).

1. Demographic Parity Check:

$$P(\hat{Y} = 1 \mid A = 0) = \frac{200}{1000} = 0.20 \qquad P(\hat{Y} = 1 \mid A = 1) = \frac{600}{2000} = 0.30$$

$$\text{Disparate Impact Ratio: } \frac{\mathbf{0.20}}{\mathbf{0.30}} = \mathbf{0.667} < \mathbf{0.80} \implies \text{Violates } 80\% \text{ EEOC Rule!}$$

2. True Positive Rate (Equality of Opportunity):

$$\text{TPR}_{A=0} = \frac{\text{True Positives}}{\text{Actual Positives}} = \frac{150}{200} = 0.75 \qquad \text{TPR}_{A=1} = \frac{450}{600} = 0.75$$

$$\text{Equality of Opportunity is SATISFIED } (\text{TPR}_{\mathbf{0}} = \text{TPR}_{\mathbf{1}} = \mathbf{0.75})!$$