USN | | | | | | | | | | |

Internal Assessment Test 3 – August 2024

Sub:	**Artificial Intelligence**					Sub Code:	BAI 402	Branch:	AIML		
Date:	**1/8/2024**	Duration:	90 min	Max Marks:	50	Sem/Sec:	IV /A, B & C			OBE	
	Answer any FIVE FULL Questions								MARKS	CO	RBT
1	Describe the term Knowledge Acquisition (KA)? Describe the tasks and techniques used in KA?								10	CO2	L2
2	Describe uncertainty in AI? Explain the basic probability notation in detail.								10	CO3	L2
3	Describe Baye’s rule with an example. What are the uses of it?								10	CO3	L3
4.	Explain inference using full joint distribution in detail with suitable examples ?								10	CO2	L2
5	What is Wumpus World revisited? Explain how problem works in uncertainty?								10	CO3	L3
6	Elaborate expert systems with its architecture? How to represent them using domain knowledge?								10	CO3	L3

Faculty Signature CCI Signature HOD Signature

1. **Describe the term Knowledge Acquisition (KA)? Describe the tasks and techniques used in KA? (5+5)**

Knowledge Acquisition (KA) in AI refers to the process of gathering, organizing, and refining knowledge from various sources to build intelligent systems capable of reasoning, learning, and decision-making. In the context of AI, KA is essential for creating knowledge bases that enable expert systems, machine learning models, and other AI applications to perform effectively.

Tasks in Knowledge Acquisition for AI

1. **Identifying Knowledge Sources**: Determining where relevant knowledge can be found, such as subject matter experts, scientific literature, databases, and observation of real-world activities.
2. **Knowledge Elicitation**: Extracting knowledge from identified sources using various techniques. This often involves interactions with human experts, reviewing documentation, or analyzing data.
3. **Knowledge Representation**: Structuring the extracted knowledge in a form that can be used by AI systems, such as rules, ontologies, semantic networks, or knowledge graphs.
4. **Validation and Verification**: Ensuring that the acquired knowledge is accurate, consistent, and relevant. This involves testing the knowledge base against real-world scenarios and correcting any errors.
5. **Integration**: Combining the newly acquired knowledge with existing knowledge bases or integrating it into AI models and systems.
6. **Maintenance and Updating**: Regularly updating the knowledge base to incorporate new information, correct errors, and adapt to changing circumstances.

Techniques Used in Knowledge Acquisition for AI

1. **Interviews and Expert Consultations**: Engaging with domain experts to gather explicit knowledge and insights directly from them.
2. **Surveys and Questionnaires**: Collecting structured information from a broader group of experts or stakeholders.
3. **Document and Literature Analysis**: Reviewing books, research papers, manuals, and other written materials to extract relevant knowledge.
4. **Observation and Ethnography**: Watching how experts perform tasks in their natural environment to understand implicit knowledge and workflows.
5. **Protocol Analysis**: Recording and analyzing the thought processes of experts as they solve problems or perform tasks, often through think-aloud protocols.

6. **Task and Workflow Analysis**: Breaking down tasks and processes to understand the knowledge required at each step and how it is applied.
7. **Concept Mapping and Knowledge Graphs**: Creating visual representations of knowledge domains to show relationships between concepts and entities.
8. **Machine Learning and Data Mining**: Using algorithms to automatically extract patterns and insights from large datasets, often uncovering knowledge that may not be immediately apparent to human experts.
9. **Heuristic Elicitation**: Identifying and capturing rules of thumb, best practices, and guidelines that experts use to make decisions.
10. **Simulation and Scenario Analysis**: Using simulations to test and refine knowledge in a controlled environment, often using hypothetical scenarios to explore different outcomes.

Applications of Knowledge Acquisition in AI

- **Expert Systems**: Creating systems that mimic human expertise in specific domains, providing recommendations, diagnostics, or decision support.
- **Natural Language Processing (NLP)**: Building models that understand and generate human language, relying on extensive knowledge bases of linguistic information.
- **Robotics and Autonomous Systems**: Equipping robots with the knowledge required to perform tasks autonomously, such as navigating environments or interacting with objects.
- **Intelligent Tutoring Systems**: Developing educational software that adapts to the needs of individual learners, providing personalized instruction and feedback.
- **Healthcare AI**: Creating systems that assist in diagnosing diseases, recommending treatments, and managing patient care based on medical knowledge.
- **Business Intelligence**: Developing tools that help organizations make informed decisions by analyzing market trends, customer behavior, and operational data.

2.Describe uncertainty in AI? Explain the basic probability notation in detail. (5+5)

Uncertainty in AI refers to situations where the information available to an AI system is incomplete, ambiguous, or noisy, leading to uncertain outcomes and predictions. This uncertainty can arise from various sources such as sensor noise, incomplete data, or inherent randomness in the environment.

Types of Uncertainty in AI

1. **Aleatory Uncertainty**: This is due to inherent randomness or variability in the system or environment. For example, the outcome of a dice roll is inherently random.
2. **Epistemic Uncertainty**: This stems from incomplete knowledge about the system or environment. It can be reduced by acquiring more information. For example, uncertainty in a medical diagnosis due to limited patient data.
3. **Model Uncertainty**: Arises from the choice and limitations of the model used. Different models may produce different predictions based on the same data.

Basic Probability Notation

Probability theory provides a mathematical framework to handle uncertainty. Here are some key concepts and notations:

1. **Random Variable:** A variable that can take different values based on the outcome of a random phenomenon. It is usually denoted by a capital letter, e.g., X.

2. **Sample Space** (Ω): The set of all possible outcomes of a random experiment. For example, for a coin toss, the sample space is $\{H, T\}$, where H stands for heads and T for tails.

3. **Event**: A subset of the sample space. For instance, in a dice roll, the event of rolling an even number is $\{2, 4, 6\}$.

4. **Probability** (P): A measure of the likelihood of an event. The probability of event A is denoted by $P(A)$, where $0 \leq P(A) \leq 1$. If $P(A) = 0$, the event is impossible; if $P(A) = 1$, the event is certain.

5. **Conditional Probability:** The probability of an event A given that another event B has occurred, denoted by $P(A|B)$. It is defined as:

$$P(A|B) = \frac{P(A \cap B)}{P(B)}$$

provided $P(B) > 0$.

6. **Joint Probability**: The probability of two events A and B happening together, denoted by $P(A \cap B)$.

7. **Marginal Probability:** The probability of an event irrespective of the outcome of other events, denoted by $P(A)$ or $P(B)$.

8. **Bayes' Theorem:** A fundamental theorem for updating probabilities based on new evidence. It is given by:

$$P(A|B) = \frac{P(B|A)P(A)}{P(B)}$$

where $P(A)$ is the prior probability of A, $P(B|A)$ is the likelihood of B given A, and $P(B)$ is the marginal probability of B.

9. **Expectation (or Expected Value)**: The average value of a random variable X, denoted by $\mathbb{E}[X]$. For a discrete random variable, it is calculated as:

$$\mathbb{E}[X] = \sum_i x_i P(X = x_i)$$

For a continuous random variable, it is calculated as:

$$\mathbb{E}[X] = \int_{-\infty}^{\infty} x f(x)\, dx$$

where $f(x)$ is the probability density function of X.

10. **Variance:** A measure of the spread of a random variable around its mean, denoted by $\mathrm{Var}(X)$. It is calculated as:

$$\mathrm{Var}(X) = \mathbb{E}[(X - \mathbb{E}[X])^2]$$

Handling Uncertainty in AI
AI systems use various methods to manage and reason about uncertainty, including:

1. **Probabilistic Models**: Models like Bayesian networks and Hidden Markov Models (HMMs) that explicitly represent and manipulate probabilities.
2. **Fuzzy Logic**: A form of logic that deals with reasoning that is approximate rather than fixed and exact. It is used in systems where the information is imprecise or uncertain.
3. **Monte Carlo Methods**: Techniques that rely on repeated random sampling to obtain numerical results, often used in simulations and optimizations.
4. **Ensemble Methods**: Combining multiple models to improve prediction accuracy and reduce uncertainty.
5. **Bayesian Inference**: Using Bayes' theorem to update the probability estimate for a hypothesis as additional evidence is acquired

3.**Describe Baye's rule with an example. What are the uses of it? (5+5)**

Bayes' Rule

Bayes' rule, also known as Bayes' theorem, provides a way to update the probability estimate for a hypothesis based on new evidence. It is a fundamental concept in probability theory and statistics, forming the basis for Bayesian inference. The rule is expressed mathematically as:

$$P(A|B) = \frac{P(B|A) \cdot P(A)}{P(B)}$$

where:

- $P(A|B)$ is the posterior probability: the probability of hypothesis A given evidence B.
- $P(B|A)$ is the likelihood: the probability of evidence B given that hypothesis A is true.
- $P(A)$ is the prior probability: the initial probability of hypothesis A before seeing the evidence B.
- $P(B)$ is the marginal likelihood: the total probability of evidence B under all possible hypotheses.

Example of Bayes' Rule

Medical Diagnosis Example:

Suppose a doctor is trying to diagnose whether a patient has a certain disease D based on a positive test result T^+.

- Let $P(D) = 0.01$ be the prior probability that a randomly chosen patient has the disease (1% prevalence).
- Let $P(T^+|D) = 0.99$ be the likelihood that a patient with the disease tests positive (99% sensitivity).
- Let $P(T^+|\neg D) = 0.05$ be the probability that a patient without the disease tests positive (5% false positive rate).
- Let $P(\neg D) = 0.99$ be the prior probability that a randomly chosen patient does not have the disease.

We want to find $P(D|T^+)$, the probability that the patient has the disease given a positive test result.

First, we calculate $P(T^+)$, the total probability of a positive test result:

$$P(T^+) = P(T^+|D) \cdot P(D) + P(T^+|\neg D) \cdot P(\neg D)$$
$$P(T^+) = (0.99 \cdot 0.01) + (0.05 \cdot 0.99)$$
$$P(T^+) = 0.0099 + 0.0495$$
$$P(T^+) = 0.0594$$

Now, applying Bayes' rule:

$$P(D|T^+) = \frac{P(T^+|D) \cdot P(D)}{P(T^+)}$$
$$P(D|T^+) = \frac{0.99 \cdot 0.01}{0.0594}$$
$$P(D|T^+) \approx \frac{0.0099}{0.0594}$$
$$P(D|T^+) \approx 0.1667$$

So, given a positive test result, the probability that the patient has the disease is approximately 16.67%.

Uses of Bayes' Rule in AI

1. **Spam Filtering**:
   - Bayes' rule is used in email spam filters to classify emails as spam or not spam based on the occurrence of certain words and features.
2. **Medical Diagnosis**:
   - AI systems use Bayes' rule to update the probability of diseases based on new symptoms and test results.

3. **Machine Learning**:
   - In Bayesian machine learning, Bayes' rule is used to update the model parameters as new data is observed, helping in techniques like Bayesian networks, Gaussian processes, and Bayesian neural networks.
4. **Robotics**:
   - Robots use Bayes' rule for localization and mapping, updating their position based on sensor readings and motion models.
5. **Natural Language Processing (NLP)**:
   - Bayes' rule is employed in NLP tasks such as text classification, sentiment analysis, and machine translation to update probabilities based on new linguistic evidence.
6. **Recommendation Systems**:
   - Bayesian methods are used to update user preferences and improve recommendations based on user interactions and feedback.
7. **Computer Vision**:
   - In computer vision, Bayes' rule helps in object recognition and image segmentation by updating the probability of object presence based on pixel data.
8. **Anomaly Detection**:
   - AI systems use Bayes' rule to detect anomalies by updating the probability of normal vs. abnormal behavior based on new observations.

4. Explain inference using full joint distribution in detail with suitable examples?(5+5)

Inference Using Full Joint Distribution in AI

Inference using the full joint distribution involves computing probabilities for various events by considering the complete probabilistic model that captures all possible combinations of variables in a given domain. The full joint distribution is a comprehensive representation that defines the probability of every possible state (combination of values) of the random variables involved.

Full Joint Distribution

For n random variables $X_1, X_2, \ldots, X_n$, the full joint distribution $P(X_1, X_2, \ldots, X_n)$ assigns a probability to each possible combination of values that these variables can take. The sum of all these probabilities is 1.

Example Scenario

Consider a simplified medical diagnosis scenario with three binary variables:

- F: Patient has a fever (True/False)
- H: Patient has a headache (True/False)
- D: Patient has a disease (True/False)

Example Scenario

Consider a simplified medical diagnosis scenario with three binary variables:

- F: Patient has a fever (True/False)
- H: Patient has a headache (True/False)
- D: Patient has a disease (True/False)

Full Joint Distribution Table

The full joint distribution $P(F, H, D)$ could be represented in a table as follows:

F	H	D	$P(F, H, D)$
T	T	T	0.02
T	T	F	0.08
T	F	T	0.01
T	F	F	0.09
F	T	T	0.04
F	T	F	0.16
F	F	T	0.02
F	F	F	0.58

Marginal Probability

The marginal probability of a variable is obtained by summing over the probabilities of all possible states of the other variables. For example, the marginal probability of the patient having a disease ($D = T$) is:

$$P(D = T) = \sum_{F,H} P(F, H, D = T)$$

Using the table:

$$P(D = T) = 0.02 + 0.01 + 0.04 + 0.02 = 0.09$$

Conditional Probability

The conditional probability of one variable given another is obtained by dividing the joint probability by the marginal probability of the given variable. For example, the probability of the patient having a disease given they have a fever ($P(D = T|F = T)$) is:

$$P(D = T|F = T) = \frac{P(F = T, D = T)}{P(F = T)}$$

First, find $P(F = T)$:

$$P(F = T) = \sum_{H,D} P(F = T, H, D) = 0.02 + 0.08 + 0.01 + 0.09 = 0.20$$

Now, find $P(F = T, D = T)$:

$$P(F = T, D = T) = \sum_{H} P(F = T, H, D = T) = 0.02 + 0.01 = 0.03$$

Thus:

$$P(D = T|F = T) = \frac{0.03}{0.20} = 0.15$$

Inference Example

Suppose we want to know the probability that a patient has a headache given they have a disease ($P(H = T|D = T)$):

$$P(H = T|D = T) = \frac{P(H = T, D = T)}{P(D = T)}$$

First, find $P(H = T, D = T)$:

$$P(H = T, D = T) = \sum_{F} P(F, H = T, D = T) = 0.02 + 0.04 = 0.06$$

We already have $P(D = T) = 0.09$. Thus:

$$P(H = T|D = T) = \frac{0.06}{0.09} = 0.67$$

Uses in AI

1. **Diagnosis Systems**: AI systems use full joint distributions for probabilistic reasoning in medical diagnosis, fault detection, and troubleshooting.
2. **Decision Making**: Enables rational decision-making under uncertainty by evaluating the probabilities of different outcomes.
3. **Predictive Modeling**: Full joint distributions are used to make predictions about future events based on observed data.
4. **Risk Assessment**: Helps in assessing and mitigating risks by modeling the probabilistic dependencies between different risk factors.

5 . What is Wumpus World revisited? Explain how problem works in uncertainty? (5+5)

Wumpus World Revisited

Wumpus World is a classic example used in artificial intelligence to illustrate reasoning and decision-making under uncertainty. It is a grid-based environment where an agent must navigate through rooms to find gold while avoiding pits and a deadly creature called the Wumpus.

Description of the Problem

- **Grid Structure**: The environment is typically represented as a 4x4 grid.
- **Starting Point**: The agent starts at a designated position in the grid (usually the bottom-left corner).
- **Goals**: The agent's primary goal is to find the gold and exit the world safely.
- **Hazards**: The environment contains several hazards:
  - **Wumpus**: A creature that can kill the agent if they enter its room.
  - **Pits**: Deadly traps that the agent must avoid.
- **Indicators**:
  - **Breeze**: Felt in rooms adjacent to pits.
  - **Stench**: Smelled in rooms adjacent to the Wumpus.
  - **Glitter**: Seen in the room containing the gold.

Uncertainty in Wumpus World

Uncertainty arises because the agent cannot directly see the pits or the Wumpus but must infer their locations based on the indicators (breeze and stench). The agent must use logical reasoning and probabilistic inference to make decisions.

How the Problem Works in Uncertainty

1. **Perception**: The agent perceives the current state by sensing breeze, stench, and glitter in the room it occupies.
2. **Inference**: Based on the percepts, the agent infers the possible locations of pits and the Wumpus. It uses rules like:
   - If there is a breeze, at least one adjacent room contains a pit.
   - If there is a stench, at least one adjacent room contains the Wumpus.
3. **Decision Making**: The agent decides its next action based on its current beliefs about the environment. Actions can include moving to an adjacent room, shooting an arrow to kill the Wumpus, or grabbing the gold.
4. **Updating Beliefs**: After taking an action, the agent updates its beliefs based on new percepts received. This involves:
   - **Bayesian Updating**: Using Bayes' rule to update the probability of pits and the Wumpus in each room.
   - **Logical Inference**: Using logical rules to deduce safe and unsafe rooms.

Example of Reasoning Under Uncertainty

Suppose the agent perceives a breeze in its current room. It concludes that one or more adjacent rooms may contain pits but doesn't know which ones. The agent uses probabilistic reasoning to estimate the likelihood of each adjacent room having a pit and chooses to move to the room with the lowest perceived risk.

Tools for Handling Uncertainty

- **Probabilistic Models**: The agent uses models like Bayesian networks to represent the probabilistic relationships between different events (e.g., the presence of a breeze and the likelihood of adjacent pits).
- **Markov Decision Processes (MDPs)**: These are used to model the decision-making process, where the agent chooses actions to maximize its expected utility while accounting for uncertainty.
- **Logic-Based Systems**: The agent uses logical inference to derive certain conclusions from its percepts (e.g., if no stench is perceived, adjacent rooms cannot contain the Wumpus).

Key Concepts in Handling Uncertainty

- **Partial Observability**: The agent cannot observe the entire environment directly but must infer hidden information based on partial observations.
- **Probabilistic Inference**: The agent uses probabilities to quantify uncertainty and update its beliefs as it gathers more information.
- **Decision Theory**: The agent uses principles from decision theory to choose actions that maximize its expected utility, balancing risk and reward.

6.Elaborate expert systems with its architecture? How to represent them using domain knowledge? (5+5)

Expert Systems
An expert system is a type of artificial intelligence program that uses a knowledge base of human expertise to aid in problem-solving. These systems are designed to emulate the decision-making abilities of a human expert in specific domains, such as medical diagnosis, engineering, finance, and more.

Architecture of Expert Systems

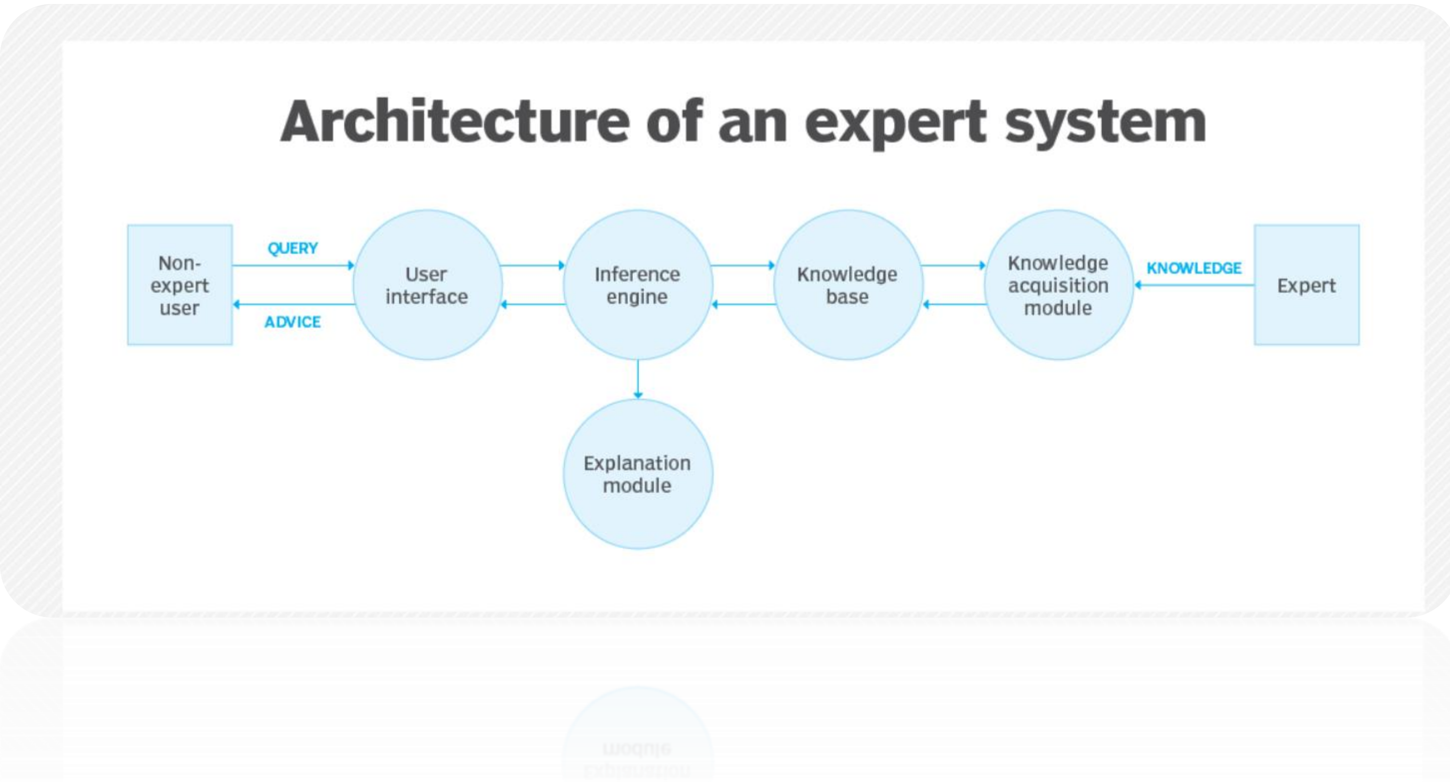

The architecture of an expert system typically consists of the following components:

1. **Knowledge Base**:
   - Contains domain-specific knowledge in the form of facts and rules.
   - **Facts**: Statements that describe objects, situations, or events in the domain.
   - **Rules**: If-Then statements that represent the relationships and logic used by experts to make decisions.
2. **Inference Engine**:
   - The core component that applies logical rules to the knowledge base to derive new information or make decisions.
   - Uses techniques such as forward chaining (data-driven) and backward chaining (goal-driven) to reason through the knowledge base.
3. **User Interface**:
   - Allows users to interact with the system, inputting data, and receiving advice or conclusions.
   - Can include graphical interfaces, natural language processing, or command-line interfaces.
4. **Explanation Facility**:
   - Provides explanations of the reasoning process to users, explaining how the system arrived at a particular conclusion.
   - Enhances the transparency and trustworthiness of the system.

5. **Knowledge Acquisition Module**:
    - Assists in the creation and updating of the knowledge base by acquiring new knowledge from human experts or other sources.
    - Tools and interfaces that help experts input their knowledge into the system.

Representation of Domain Knowledge in Expert Systems

Domain knowledge in expert systems is typically represented in several ways, including:

1. **Rules**:
    - Represented as If-Then statements.
    - Example: IF a patient has a fever AND a rash THEN the patient might have measles.
2. **Frames**:
    - Structures that represent stereotypical situations, objects, or events.
    - Consist of slots (attributes) and values.
    - Example: A frame for a medical patient might include slots for name, age, symptoms, diagnosis, and treatment.
3. **Semantic Networks**:
    - Graph structures that represent relationships between concepts.
    - Nodes represent concepts or objects, and edges represent relationships.
    - Example: A network representing relationships between diseases and symptoms.
4. **Ontologies**:
    - Formal representations of a set of concepts within a domain and the relationships between those concepts.
    - Provide a shared and common understanding of a domain that can be communicated across people and applications.
5. **Production Rules**:
    - Sets of rules that represent the logic of the domain.
    - Example: IF condition THEN action.

Example: Medical Diagnosis Expert System

1. **Knowledge Base**:
    - Facts: Patient A has a fever.
    - Rules: IF a patient has a fever AND a cough THEN the patient might have the flu.
2. **Inference Engine**:
    - Uses forward chaining to infer new facts from existing ones.
    - Example: Given Patient A has a fever and Patient A has a cough, the system infers that Patient A might have the flu.
3. **User Interface**:
    - Allows a doctor to input symptoms and receive a possible diagnosis.
    - Example: A form where the doctor enters symptoms like fever and cough.
4. **Explanation Facility**:
    - Explains that the diagnosis was made because the patient has both a fever and a cough, which matches the rule for flu.
5. **Knowledge Acquisition Module**:
    - A tool for medical experts to input new rules or update existing ones as new medical knowledge becomes available.

Detailed Example: Car Troubleshooting Expert System

Knowledge Base

- **Facts**:
    - The car does not start.
    - The battery is charged.
- **Rules**:
    - IF the car does not start AND the battery is charged THEN check the spark plugs.
    - IF the spark plugs are faulty THEN replace the spark plugs.

Inference Engine

- Uses backward chaining to determine why the car does not start.

- Begins with the goal (why the car does not start) and works backward to find supporting facts and rules.

User Interface

- Allows the mechanic to input the car's symptoms and receive troubleshooting advice.
- Example: A question-answer interface where the system asks, "Is the battery charged?" and the mechanic responds.

Explanation Facility

- Explains that the advice to check the spark plugs is based on the facts that the car does not start and the battery is charged.

Knowledge Acquisition Module

- A tool for mechanics to input new troubleshooting rules as they encounter new car issues and solutions.