
# Introduction to Artificial Intelligence

## Subject 2: Formal Logic for Knowledge Representation and Natural Language Understanding

Dr. Mohammed Shambour

# Learning Objectives

Providing an overview of how formal logic serves as a foundation for representing knowledge and understanding natural language in AI systems.

By the end of this lecture, students will be able to:

- Explain how logic is used in AI knowledge representation.
- Differentiate between propositional and predicate logic.
- Express facts and relationships using logical expressions.
- Translate natural language into formal logic.
- Understand challenges in mapping linguistic meaning to logic.

# Why Logic in AI?

AI needs structured representations of knowledge to reason and make decisions.

**Logic enables:**

- Transfer knowledge accurately and clearly.
- Deduction of new facts (inference).
- Verification of truth and consistency.

**Example: If all babies are smart, and Noor is a baby, then Noor is smart.**

- **Premise 1 (General Rule):** All babies are smart. (If X is a baby, then X is smart)
- **Premise 2 (Specific Fact):** Noor is a baby.
- **Conclusion (Logical Inference):** Therefore, Noor is smart.

# Knowledge Representation (KR)

## ❑ Definition:

Knowledge Representation (KR) is the process of representing information about the world in a form that a computer or AI system can understand and reason about.

## ❑ Purposes of Knowledge Representation

- **Store domain knowledge**
  - It provides a structured way to **store information** about a particular area (like medicine, law, or geography).
  - Example: A medical AI system stores knowledge about diseases, symptoms, and treatments.

- **Enable reasoning and inference**
  - Once knowledge is stored, the AI can **draw conclusions**, **predict outcomes**, or **solve problems** based on that knowledge.
  - Example: If "All planets in our solar system orbit the Sun" and "Mars is a planet in our solar system, Then " the AI can infer "Mars orbits the Sun."

- **Support communication with humans**
  - KR helps AI **explain its reasoning** or **understand human language and queries** better.
  - Example: Chatbots use structured knowledge to answer user questions accurately.

# Components of a Knowledge Representation System

- Knowledge Base (KB): Contains facts and rules.
- Inference Engine: Derives new information using reasoning.
- Query Processor: Interprets and answers user queries.
- Learning Module: Updates and improves the knowledge base.

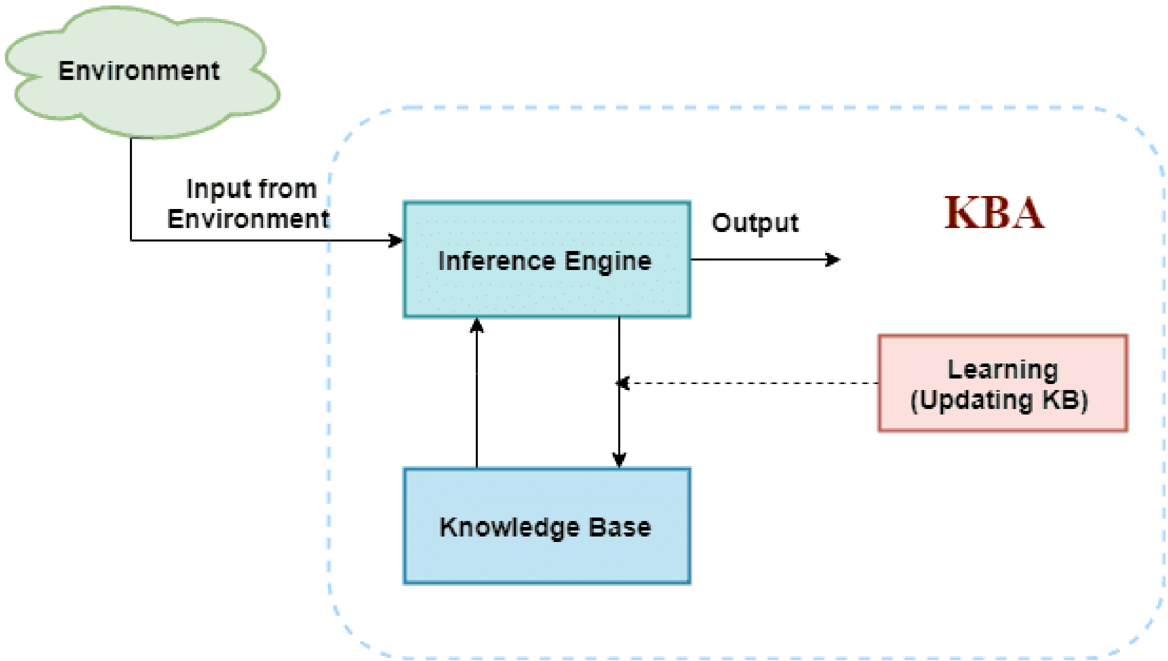


Architecture of knowledge-based agent

https://tutorialforbeginner.com/knowledge-based-agent-in-ai

# Overview of Logic in AI

Logic provides a mathematical language for reasoning.

**Main types:**

**1. Propositional Logic** – Works with complete statements (propositions) as indivisible units and uses logical (e.g. AND (∧), OR (∨), NOT (¬)), connectors between them.

**2. Predicate Logic** – introduces variables (e.g. x, y, z) and quantifiers (e.g. ∀: for all, ∃: there exists) for more meaningful statements.

# Syntax vs Semantics

**Syntax** defines how we can combine symbols into valid sentences.

**Semantics** assigns meaning to those sentences.

**Example:**

- Syntax: $P \rightarrow Q$
- Semantics: If it rains (P), then the ground is wet (Q).

<u>Syntax = structure; Semantics = meaning.</u>

# Propositional Logic

Uses simple propositions that can be True or False.

Example statements:

- p: It is raining.
- q: The ground is wet.

Compound examples:

- $\neg p$: Not raining.
- $p \land q$: It is raining AND the ground is wet.
- $p \rightarrow q$: If it rains, the ground is wet.

# Propositional Logic Syntax

Basic elements:

- Symbols: p, q, r, etc.
- Logical connectives: AND (∧), OR (∨), NOT (¬), IMPLIES (→), IFF (↔)

Example formula:

- $\neg(p \land q) \rightarrow r$
- Meaning: If not both p and q are true, then r is true.

# Truth Tables and Semantics

- Truth tables define how the truth value of a compound statement depends on its components.

- Interpretation: The statement 'If p then q' is false only when p is true and q is false.

| p | q | $p \to q$ |
|---|---|---|
| T | T | T |
| T | F | F |
| F | T | T |
| F | F | T |

# Logical Equivalences

Equivalences simplify complex logical expressions.

For propositional expressions P, Q and R:

$\neg (\neg \mathbf{P}) \equiv \mathbf{P}$

$(\mathbf{P} \vee \mathbf{Q}) \equiv (\neg \mathbf{P} \rightarrow \mathbf{Q})$

the contrapositive law: $(\mathbf{P} \rightarrow \mathbf{Q}) \equiv (\neg \mathbf{Q} \rightarrow \neg \mathbf{P})$

de Morgan's law: $\neg (\mathbf{P} \vee \mathbf{Q}) \equiv (\neg \mathbf{P} \wedge \neg \mathbf{Q})$ and $\neg (\mathbf{P} \wedge \mathbf{Q}) \equiv (\neg \mathbf{P} \vee \neg \mathbf{Q})$

the commutative laws: $(\mathbf{P} \wedge \mathbf{Q}) \equiv (\mathbf{Q} \wedge \mathbf{P})$ and $(\mathbf{P} \vee \mathbf{Q}) \equiv (\mathbf{Q} \vee \mathbf{P})$

the associative law: $((\mathbf{P} \wedge \mathbf{Q}) \wedge \mathbf{R}) \equiv (\mathbf{P} \wedge (\mathbf{Q} \wedge \mathbf{R}))$

the associative law: $((\mathbf{P} \vee \mathbf{Q}) \vee \mathbf{R}) \equiv (\mathbf{P} \vee (\mathbf{Q} \vee \mathbf{R}))$

the distributive law: $\mathbf{P} \vee (\mathbf{Q} \wedge \mathbf{R}) \equiv (\mathbf{P} \vee \mathbf{Q}) \wedge (\mathbf{P} \vee \mathbf{R})$

the distributive law: $\mathbf{P} \wedge (\mathbf{Q} \vee \mathbf{R}) \equiv (\mathbf{P} \wedge \mathbf{Q}) \vee (\mathbf{P} \wedge \mathbf{R})$

Purpose: Simplify logical expressions and proofs.

# Inference in Propositional Logic

**Inference is the process of deriving new conclusions from existing knowledge.**

Example rule: Modus Ponens

- If **p implies q**
- and **p is true**,
- then **q must be true**.

**Modus Ponens** is one of the most fundamental rules of inference in logic — it's how we *derive conclusions* from conditional statements.

Example:

- If it rains, the ground gets wet.
- It rains.
- Therefore, the ground gets wet.

# Limitations of Propositional Logic

- Cannot describe relations between objects.
- Cannot express quantities such as 'all' or 'some'.
- Lacks variable binding and object structure.

Example: 'All birds can fly.' cannot be expressed in propositional logic, You'd need a separate proposition for each bird — "Bird1 can fly," "Bird2 can fly," etc.

It can't generalize for all birds because it requires variables and quantifiers, which only predicate logic supports.

# Predicate Logic (First-Order Logic)

Extends propositional logic by adding variables, functions, and quantifiers.

Predicates describe properties or relationships,
e.g., For every cat x, there exists a person y such that y owns x.

- Quantifiers:
  - Universal (∀): for all
  - Existential (∃): there exists

$\forall x\ (Cat(x) \rightarrow \exists y\ (Owns(y, x)))$

# Components of Predicate Logic

| **Component** | **Example** |
|---|---|
| **Variables** | x, y |
| **Predicate** | Cat(x) → Predicate meaning “x is a cat.”<br>Owns(y, x) → Predicate meaning “y owns x.” |
| **Quantifier** | ∀x → Universal quantifier: “For all x”<br>∃y → Existential quantifier: “There exists some y” |

# Examples of Predicate Logic

1. **Universal:** $\forall x\,(\mathrm{Bird}(x) \rightarrow \mathrm{Fly}(x))$
   "All birds can fly."

2. **Existential:** $\exists x\,(\mathrm{Student}(x) \wedge \mathrm{Smart}(x))$
   "There exists a student who is smart."

3. **Nested:** $\forall x\,\exists y\,(\mathrm{Parent}(y, x))$
   "Everyone has a parent."

# Translating Natural Language to Logic (I)

| **Natural Language** | **Predicate Logic** |
|---|---|
| All humans are mortal. | $\forall x\ (Human(x) \rightarrow Mortal(x))$ |
| Some students are intelligent. | $\exists x\ (Student(x) \land Intelligent(x))$ |
| If it rains, the streets get wet. | $Rain \rightarrow Wet(Streets)$ |

# Translating Natural Language to Logic (II)

**More Complex Examples:**

- "Every teacher teaches some student."

→ $\forall x\ (Teacher(x) \rightarrow \exists y\ (Student(y) \land Teaches(x, y)))$

- "No cat likes water."

→ $\neg\exists x\ (Cat(x) \land Likes(x, Water))$

- "Only doctors can prescribe medicine."

→ $\forall x\ (Prescribe(x, Medicine) \rightarrow Doctor(x))$

# Natural Language Ambiguity

**Problem:** Natural language is ambiguous; logic is not.

Examples:

- “Visiting relatives can be annoying.” (Who visits whom?)
- “Every student read a book.” (Same book or different ones?)

**AI Approach:** Use logic + context + probability models (semantic parsing).

# Knowledge Representation in NLP

Logic supports natural language understanding by:

- Semantic parsing (structure to meaning)
- Question answering
- Dialogue systems
- Ontologies and knowledge graphs

Example: Convert sentence → logical form → inference → answer.

# The Two AI Paradigms

- **Symbolic AI (Logic-Based or Classical AI)**
  - **What:**
    Represents knowledge using symbols, rules, and logic (e.g., “if-then” statements, ontologies).
  - **How:**
    Uses explicit reasoning — algorithms manipulate symbols based on predefined logical rules.
  - **Examples:**
    - Expert systems
    - Knowledge graphs
    - Rule-based reasoning systems (e.g., Prolog, early AI)
  - **Strengths:**
    - Transparent and explainable reasoning
    - Works well in structured, rule-based domains
    - Easy to verify and debug
  - **Weaknesses:**
    - Rigid (fails when data or rules change)
    - Cannot learn from new data automatically
    - Requires manual knowledge engineering

- **Statistical AI (Machine Learning-Based)**
  - **What:**
    Learns **patterns from data** rather than relying on explicit rules.
  - **How:**
    Uses **statistical methods, probability, and neural networks** to make predictions or decisions.
  - **Examples:**
    - Deep learning
    - Decision trees
    - Support vector machines
    - Bayesian models
  - **Strengths:**
    - **Learns automatically** from data
    - Handles **uncertainty, noise, and complexity**
    - Scales well to large and unstructured data (images, text, etc.)
  - **Weaknesses:**
    - Often a **“black box”** — hard to interpret
    - Can produce **illogical or biased** results
    - Needs **large, high-quality datasets**

# Exercise

Translate the following from English into first order logic

1. Every AI system uses data.
2. Some algorithms are efficient.
3. If a model overfits, it performs poorly on new data.
4. Every hardworking student who attends his exams will pass
5. Ali revises his lectures and does his homework
6. Ali attends all his exams

That’s all for Today