


Knowledge-Based Agents (Logical Agents)



Material from Dr. Marie desJardins. Some material adopted from notes by Andreas Geyer-Schulz and Chuck Dyer

A Knowledge-Based Agent


- A knowledge-based agent needs (at least):
 - A **knowledge base**
 - An **inference system**
- A knowledge base (KB) is a set of representations of facts about the world.
 - Each individual representation is a **sentence** or **assertion**
 - Expressed in a **knowledge representation language**
 - Usually starts with some background knowledge
 - Can be general (world knowledge) or specific (domain knowledge)
- Many existing ideas apply – is it closed-world, etc.

2

A Knowledge-Based Agent

- Operates as follows:

1. TELLS the knowledge base what it perceives.
2. ASKS the knowledge base what action to perform.
3. Performs the chosen action.



3

Architecture of a Knowledge-Based Agent

- **Knowledge Level**
 - The most abstract level
 - Describe agent by saying what it knows
 - Example: A taxi agent might know that the Golden Gate Bridge connects San Francisco with the Marin County.
- **Logical Level**
 - Level at which **knowledge** is encoded into **sentences**.
 - Example: `Links(GoldenGateBridge, SanFrancisco, MarinCounty)`
- **Implementation Level**
 - The physical representation of the sentences in the logical level.
 - Example: `'(links goldengatebridge sanfrancisco marincounty)'`

4

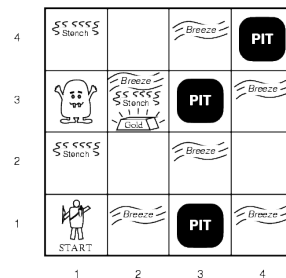
The Wumpus World Environment

- The Wumpus computer game
 - Agent explores a cave consisting of rooms connected by passageways.
 - Lurking somewhere in the cave is the Wumpus, a beast that eats any agent that enters its room.
 - Some rooms contain bottomless pits that trap any agent that wanders into the room.
 - Occasionally, there is a heap of gold in a room.
 - The goal is to collect the gold and exit the world without being eaten (or trapped).

5

A Typical Wumpus World

- The agent always starts in the field [1,1].
- The task of the agent is to find the gold, return to the field [1,1] and climb out of the cave.



6

Agent in a Wumpus World: Percepts

- Agent perceives
 - **Stench** in the square containing the wumpus and in adjacent squares (not diagonally)
 - **Breeze** in the squares adjacent to a pit
 - **Glitter** in the square where the gold is
 - **Bump**, if it walks into a wall
 - **Woeful** scream everywhere in the cave, if the wumpus is killed
- The percepts are given as a five-symbol list.
- If there is a stench and a breeze, but no glitter, no bump, and no scream, the percept is:
[Stench, Breeze, None, None, None]
- The agent cannot perceive its own location

7

Wumpus Agent Actions

- **go forward**
- **turn right 90 degrees**
- **turn left 90 degrees**
- **grab**: Pick up an object that is in the same square as the agent
- **shoot**: Fire an arrow in a straight line in the direction the agent is facing.
 - The arrow continues until it either hits and kills the wumpus or hits the outer wall.
 - The agent has only one arrow, so only the first Shoot action has any effect
- **climb**: leave the cave. This action is only effective in the start square
- **die**: This action automatically happens if the agent enters a square with a pit or a live wumpus

8

Wumpus Goal

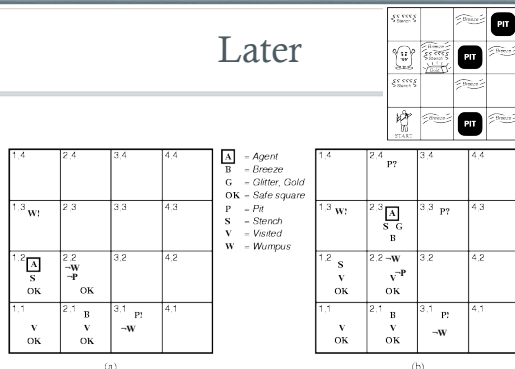
- Agent's goal is to:
 - Find the gold
 - Bring it back to the start square as quickly as possible
 - Don't get killed!
- Scoring
 - 1000 points reward for climbing out with the gold
 - 1 point deducted for every action taken
 - 10000 points penalty for getting killed

9

Wumpus Agent's First Step



Later



11

Wumpuses Online

- <http://www.cs.berkeley.edu/~russell/code/doc/overview-AGENTS.html>
 - Lisp version from Russell & Norvig
- <http://www.dreamcodex.com/wumpus.php> – Java-based version you can play online
- <http://codenautics.com/wumpus/> – Downloadable Mac version


12

Representation, Reasoning, and Logic

- Point of knowledge representation is to express knowledge in a **computer usable** form
- Needed for agents to act on it (to do well, anyway)
- A knowledge representation language is defined by:
 - Syntax:** all possible sequences of symbols that form sentences
 - Example: noun referents can be a single word or an adjective-then-noun
 - Semantics:** facts in the world to which the sentences refer
 - What does it *mean*?
- Each sentence makes a claim about the world
- An agent is said to “believe” a sentence about the world

13

The Connection Between Sentences and Facts



Semantics maps sentences in logic to facts in the world. The property of one fact following from another is mirrored by the property of one sentence **being entailed** by another.

“Dr M is sick with the flu” \models “Dr M is sick”


14

Entailment and Derivation

- Entailment:** $KB \models Q$ x \models y: x semantically entails y
 - Q is entailed by KB (a set of premises or assumptions) if and only if there is no logically possible world in which Q is false while all the premises in KB are true.
 - Or, stated positively, Q is entailed by KB if and only if the conclusion is true in every logically possible world in which all the premises in KB are true.
- Derivation:** $KB \vdash Q$ x \vdash y: y is provable from x
 - We can derive Q from KB if there is a proof consisting of a sequence of valid inference steps starting from the premises in KB and resulting in Q

15

Logic as a KR Language



16

Ontology and Epistemology


- Ontology** is the study of what there is—an inventory of what exists. An ontological commitment is a commitment to an existence claim.
- Epistemology** is a major branch of philosophy that concerns the forms, nature, and preconditions of knowledge.

Language	Ontological Commitment (What exists in the world)	Epistemological Commitment (What an agent believes about facts)
Propositional logic	facts	true/false/unknown
First-order logic	facts, objects, relations	true/false/unknown
Temporal logic	facts, objects, relations, times	true/false/unknown
Probability theory	facts	degree of belief 0...1
Fuzzy logic	degree of truth	degree of belief 0...1

17

No Independent World Access

- The reasoning agent often gets its knowledge about the facts of the world as a *sequence of logical sentences*.
- Must draw conclusions from them without (other) access to the world.
- Thus it is very important that the agent's reasoning is sound!



18

KB Agents - Summary

- Intelligent agents need **knowledge about the world** for making good decisions.
- The knowledge of an agent is stored in a knowledge base in the form of **sentences** in a **knowledge representation language**.
- A knowledge-based agent needs a **knowledge base** and an **inference mechanism**. It operates by storing sentences in its knowledge base, inferring new sentences with the inference mechanism, and using them to deduce which actions to take.
- A **representation language** is defined by its syntax and semantics, which specify structure of sentences and how they relate to world facts.
- The **interpretation** of a sentence is the fact to which it refers. If this fact is part of the actual world, then the sentence is true.