CPE/CSC 480 ARTIFICIAL INTELLIGENCE

FINAL EXAM

FALL 2003

PROF. FRANZ J. KURFESS

CAL POLY, COMPUTER SCIENCE DEPARTMENT

This is the Fall 2003 final exam for the CPE/CSC 480 class. It is a take-home exam, and you may use textbooks, course notes, or other material, but you must formulate the text for your answers yourself. You are not allowed to discuss the questions and answers of the exam with other students or anybody else. The use of calculators or computers is allowed for numerical calculations, but not for the execution of algorithms or programs to compute solutions for exam questions.

If you need clarifications about questions, you can contact me via email or during my office hours on Thursday Dec. 11, from 2 pm to 5 pm. The deadline for the exam is Thursday, December 11, 2002, at 5:00 pm. You must submit a printed and signed copy of the exam, which you can either leave in the drop box in front of the CSC department office (it is emptied at 4:00 pm), or give it to me on Thursday before 5:00 pm in my office.

Student Name:

Signature: **Date:**

PART 1: MULTIPLE CHOICE QUESTIONS

Mark the answer you think is correct. Unless otherwise noted, there is only one correct answer. Each question is worth 3 points.

a) What is the *horizon problem* in the context of game-playing programs?

□ a degree of uncertainty, introduced by the presence of an opponent or by chance elements
□ the outcome of a move may not be visible due to search limitations
□ the need for arbitration (e.g. by a referee) in some types of games
□ the elimination of branches that will never be explored

b) Which statement describes the *syntax* of a formal language for knowledge representation?

□ It describes how a particular sentence relates to the facts in the world.
□ It allows the generation of new sentences that follow from a set of given sentences.
□ It specifies the admissible configurations of sentences in that language.
□ It makes sure that only truth-preserving sentences are admitted in the language.

c) What is a *logical proof* for a sentence?

□ A collection of sound inference rules.
□ A world in which that sentence is true under a particular interpretation.
□ A set of sentences from which the sentence to be proved can be entailed.
□ A sequence of applications of inference rules, leading from sentences in the knowledge base, to the sentence to be proved.

d) What is the role of a *universal quantifier* □ in a predicate logic sentence?

□ It allows statements about some objects in a collection.
□ It allows general statements about every object in a collection.
□ It is used in the specification of the semantics for terms.
□ It can be used to make statements about quantitative aspects of objects, such as length, weight, temperature, etc.

e) Which statement is the best characterization of *data* (as opposed to knowledge)?

□ Its representation relies on separate, possibly dynamic relations between individual concepts or objects.
□ Relations between individual items are fixed according to a scheme applied to a large collection of items.
□ It is a generic term with a relatively broad meaning, except for the field of information theory, where it is precisely defined.
□ Since computers are mostly concerned with syntactic manipulation, knowledge is irrelevant for Artificial Intelligence.

f) Which statement is the best characterization of *inference*?

- □ It describes methods and procedures for drawing conclusions on the basis of existing knowledge.
- □ It relies on the storage of relations between individual items according to a fixed scheme applied to a large collection of elementary items.
- □ It requires a set of formal inference methods and clearly specified syntax and semantics.
- □ It is concerned with methods and techniques for the storage of knowledge and information in a format that is suitable for treatment by computers.

g) What is the purpose of developing a *ontology* for a knowledge-based system?

- □ It is used to represent elementary, undisputed facts and rules of a domain.
- □ It is a method used to formalize the semantics of sentences in a domain.
- □ It helps with a systematic description of objects and their relations in a domain.
- □ It greatly improves the efficiency of inference procedures by steering the selection of applicable inference rules.

h) Which of the following statements is the best description of *composite objects* in the context of knowledge-based systems?

- □ Their internal structure is of importance for the task or domain under consideration.
- □ They describe values for properties of objects, often expressed in a quantitative way.
- □ They are objects whose temporal and spatial extent is especially important.
- □ They are similar to objects in Software Engineering, and can be represented easily in conventional programming languages like Java or C++.

i) Which of the following statements is easy to express in *propositional logic*?

- □ "Squares in the wumpus world are either empty, or contain the wumpus or a pit, but not both."
- □ "The scream of the dying wumpus can be heard throughout the cave."
- □ "If the wumpus is dead, it still smells."
- □ "If there is a breeze in a square, one of the adjacent squares must be a pit."

j) Which of the following statements is difficult to express in (first order) *predicate logic*?

- □ "If the wumpus is dead, it still smells."
- □ "Squares in the wumpus world are either empty, or contain the wumpus or a pit, but not both."
- □ "The agent has only one arrow to shoot at the wumpus."
- □ "For most configurations of the environment, the agent has a reasonable chance of finding the gold, and climbing out of the cave."

PART 2: SHORT QUESTIONS

In this part of the exam, you should answer the questions in the space provided.

1. The evaluation of a logical sentence □ with respect to a set of other sentences in a knowledge based can be done in two ways: through the *entailment* relation KB |= □ or through *derivation* KB |-$_{i}$ □ with respect to an inference procedure i. Explain the difference between *entailment* and *derivation*.

[10 points]

Aspect	Entailment	Derivation
Expla- nation	• • • • • •	• • • • • •
Differences	• • • • •	• • • • •

2. Explain the general concept of *probably approximately correct (PAC)* learning, and discuss its possible use for an agent in the Wumpus world. The main emphasis for the agent is to learn from its experiences in different configurations of the Wumpus world environment.

10 points

- PAC Learning Explanation

- PAC Learning for a Wumpus World Agent

PART 3: HEXAGONAL WUMPUS WORLD

In this part of the exam, you need to explore and evaluate the capabilities of an agent in a Wumpus World with a hexagonal instead of a square-based topology. Due to this topology, the navigation task for the agent is also somewhat different: it moves in a different way, and the assumptions for reasoning about properties of nodes may have to be changed.

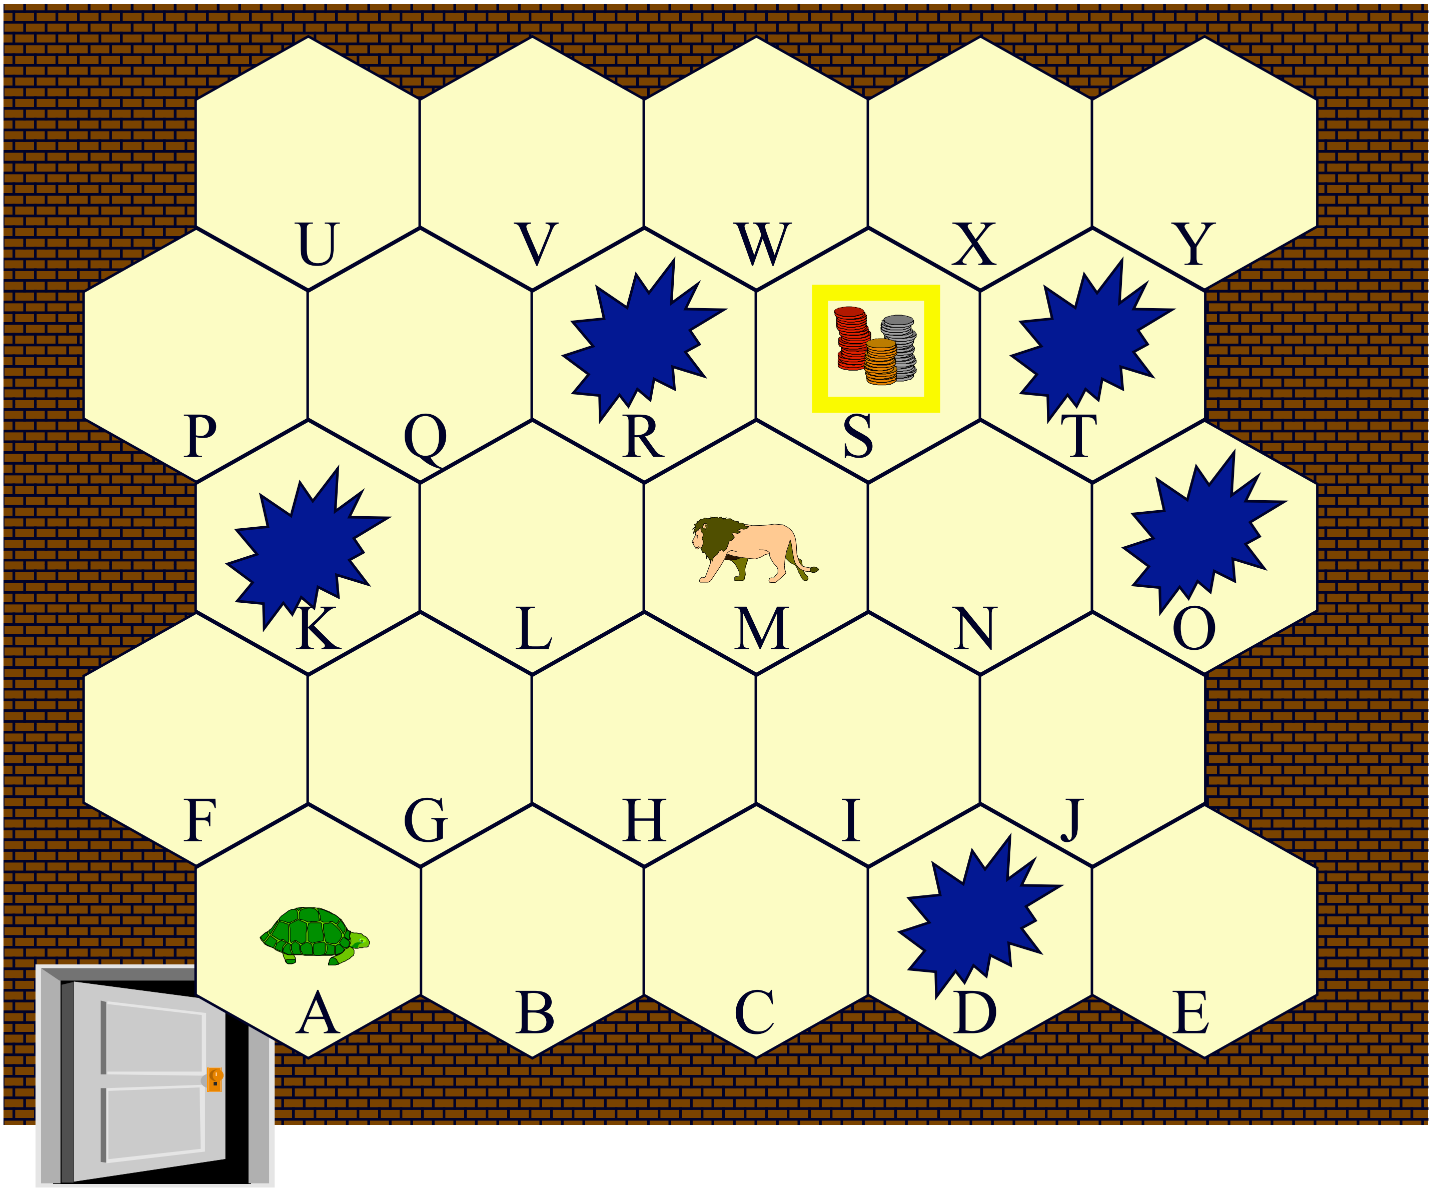

The basic setup is the same: The agent has to find the gold, pick it up, and return to the starting point. It has as percepts {`Breeze, Smell, Glitter, Scream`} ("bump" is not necessary here), and as actions {`Move-Next, Backtrack, Shoot(n), Grab, Climb`}. The `Move-Next` action transports the agent from the current node to another node. The agent can `Backtrack` if there is no unvisited node adjacent to the current node, but there still are nodes in the respective queue (safe or risky). The `Shoot(n)` action requires as parameter one of the adjacent nodes, indicating the direction in which the arrow will fly.

In this part, you are provided with an agent that has the following characteristics:

- The agent works with two queues for nodes to be explored: A *safe queue* for nodes that are safe to enter, and a *risky queue* containing nodes that may be dangerous for the agent (e.g. a pit, or the wumpus may be on the node).
- The agent can move to the next node in its safe queue that is adjacent to the current node (it can not "jump" or "teleport"), or it backtracks if there are safe nodes, but not adjacent to the current node. If the safe queue is empty, it will select the next adjacent node from the risky queue, or backtrack. It explores adjacent nodes in a clockwise fashion, starting with the node to the right of the "noon" position.
- The agent keeps a knowledge base with useful statements for later use (e.g. the possibility of a pit at a node).

And here are a few assumptions to clarify ambiguous situations:

- Once the wumpus is dead, its smell still can be perceived.
- The wumpus can not be on the same node as a pit.

If necessary, you can make additional assumptions, but please state them below.
The detailed tasks you need to solve are described in the following sections, with the respective number of points indicated. The total number of points for this task is 40.

a) In the table below, describe the following important aspects of the agent outlined above. If you are making additional assumptions, please state them under "Additional Assumptions" below.

12 points

- *Behavior*
  What are the observable activities of the agent for an outside observer? Could that observer judge the "intelligence" of the agent without knowing anything about how the agent works internally?
- *Internal storage and retrieval of information (memory)*
  How does the agent keep track of previous percepts and possibly other important information? Does it construct an internal representation of the environment?
- *Reasoning and decision making*
  How does the agent make decisions about which action to take? Is it possible for the agent to plan ahead for several steps?
- *Learning*
  Is the agent capable of learning from observations or experiences? If not, would it be possible to incorporate a learning component into the agent? What advantages and problems might result from that?

Additional Assumptions:

-
-
-
-
-
-

Agent Aspects

Behavior	Observable: • • • • Appearance of “intelligent” behavior: • • • • Performance: • • • •
Memory	Necessary memory capabilities: • • • • Useful additional memory capabilities: • • • •
Reasoning	Necessary reasoning capabilities: • • • • Useful additional reasoning capabilities: • • • • Planning capabilities: • • • •

Learning	Necessary learning capabilities: • • • • Useful additional learning capabilities: • • • •

b) What type of agent is described above? Explain your answer!

4 points

□ simple reflex agent
□ model-based agent
□ goal-based agent
□ utility-based agent
□ knowledge-based agent
□ learning agent

Explanation:

c) What is the simplest type of agent that can successfully solve the task of finding the gold and returning to the starting point? The characteristics of this type may be different from the ones listed above. Explain your answer!

4 points

□ simple reflex agent
■ model-based agent
□ goal-based agent
□ utility-based agent
□ knowledge-based agent

Explanation:

d) Indicate in the diagram a possible path for the agent from the starting point to the goal, and describe the important steps taken and critical decisions made by the agent. Use the diagram and the table below. The agent initially is at node A and the gold at node S. The wumpus is at node M, and nodes D, K, O, R, and T are pits.

12 points

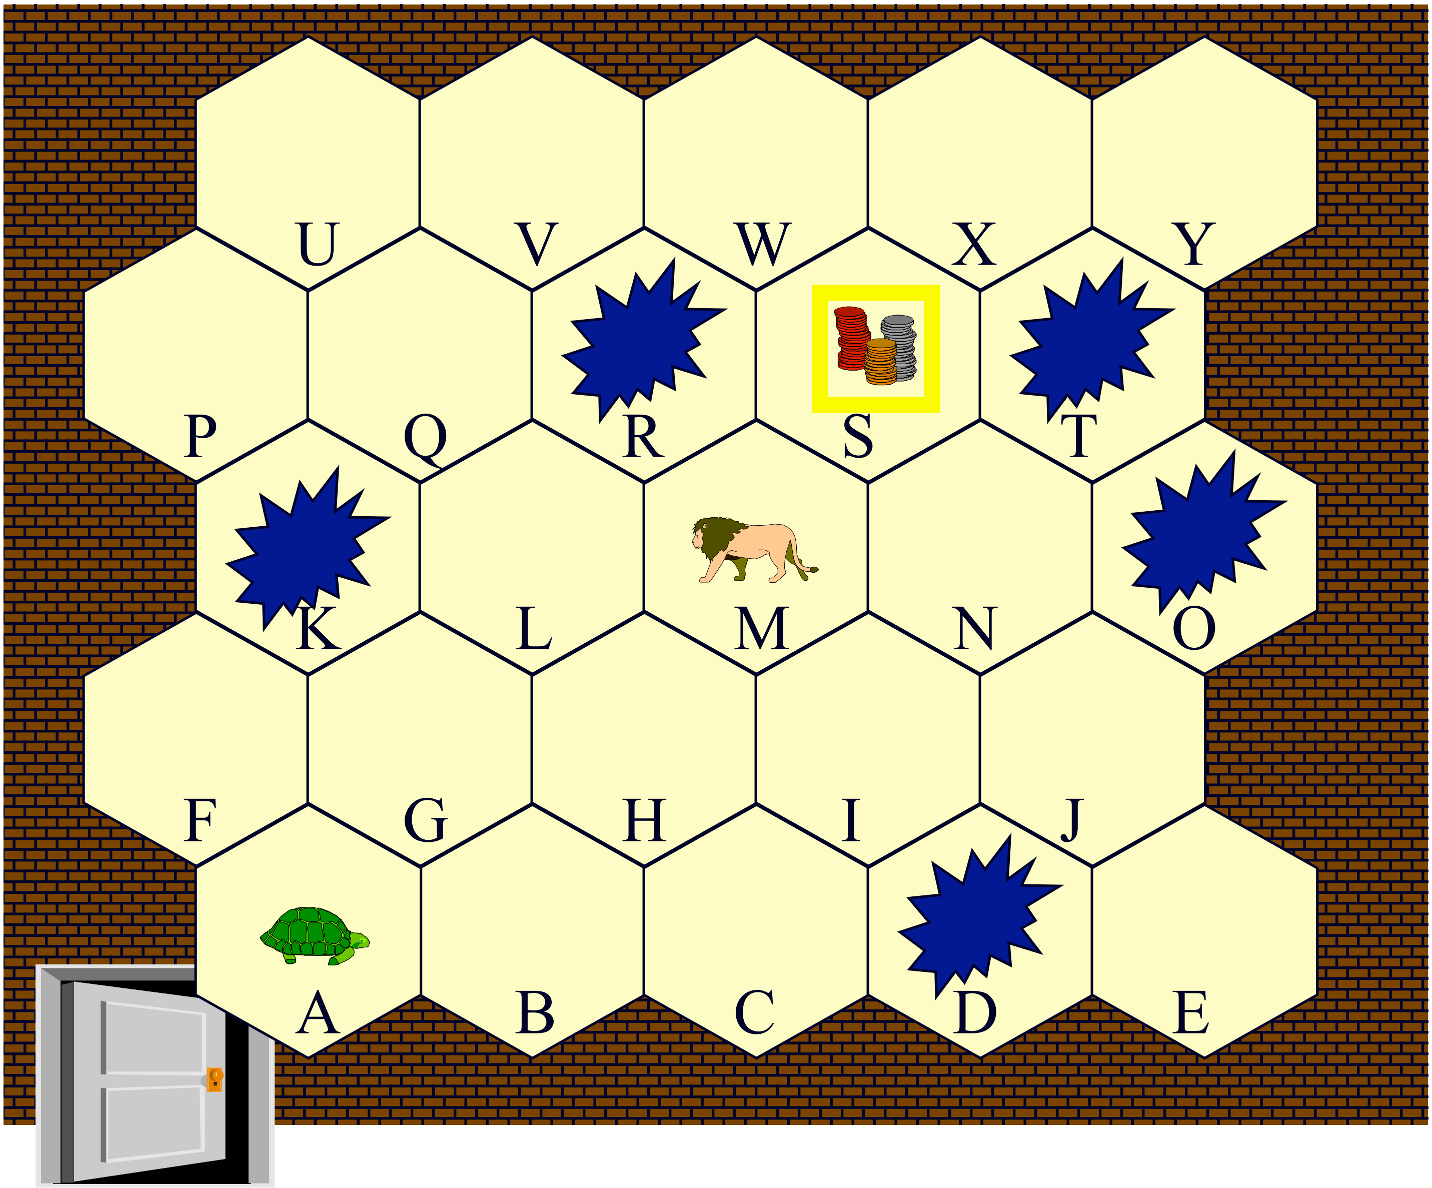

Please use the following notation:

- mark the nodes where a breeze or a smell can be perceived with appropriate colors (on paper, use highlighters; in the electronic version, make sure that the printout shows the difference between breezy and smelly squares)
- mark the nodes known to be safe to the agent with "**ok**"
- mark the nodes visited by the agent with a checkmark
- mark the nodes that the agent thinks could be pits with "**?Pit**" (or "**?P**")

- mark the nodes where the agent thinks the wumpus might be with "**?Wumpus**" (or "**?W**")
- if the agent ascertains the presence of a pit or the wumpus, cross out the question mark, and write down an exclamation mark (e.g. "**! ~~?~~Pit**")
- if the agent ascertains the absence of a pit or the wumpus, cross out the question mark, and write down a negation sign (e.g. "**¬ ~~?~~Pit**")

If you want to, you can replicate the diagram below to indicate the progress of the agent.

e) Use the same notation to record the activities of the agent in the table, listing the steps (e.g. movements, important reasoning steps) of the agent. The table lists the following data structures: A list of nodes visited, a list of nodes with the agent's return path, a two-part fringe consisting of "safe" and "risky" unexplored nodes, and a set of statements that reflect the knowledge acquired by the agent. You can assume that knowledge from previous steps is accessible to the agent, and that "newer" knowledge overrides "older" knowledge. As an alternative, you can replicate the complete knowledge base of the agent in each step; this is possibly less confusing, but requires more space.

12 points

Step	*Node*	*Percept*	*Visited*	*Return Path*	*Safe Queue*	*Risky Queue*	*Knowledge*	*Action*
0	**A**	None			**[G,B,F]**	**[]**	¬Pit(A,B,G,F) ¬Wumpus(A,B,F,G)	Move-Next
1								
2								
3								
4								
5								
6								
7								
8								
9								
10								
11								
12								
13								
14								
15								

16
17
18
19
20

f) Is there a safe path for the agent in this environment? Explain your answer.

3 points

g) Is the agent described above guaranteed to find a safe path, provided there exists one? Explain your answer.

3 points

Total Points: