VICTORIA UNIVERSITY OF WELLINGTON
Te Whare Wananga o te Upoko o te Ika a Maui

EXAMINATIONS — 2003

COMP 423

INTELLIGENT AGENTS

Time Allowed: 3 Hours

Instructions: Do all of questions 1 to 6
Do any **three** of question 7 to 11 The exam will be marked out of 180.

Non-electronic foreign language dictionaries are permitted.

Clean copies of the papers will be distributed for the exam.

Questions

Part 1 (Do all questions)

1.	Intelligent Agents	[10]
2.	Precision and Recall	[5]
3.	Naive Bayes	[10]
4.	Graph Plan	[10]
5.	Your Project	[15]
6.	Wrapper Induction for Information Extraction	[40]

Part 2 (Do any **three** questions)

7.	Games	[30]
8.	Mobile Robots	[30]
9.	Help Agents	[30]
10.	Spam Filtering	[30]
11.	Programming by Demonstration	[30]

Part I

Do all of questions 1 to 6.

Question 1. Intelligent Agents [10 marks]

The term “intelligent agent” is not precisely defined, but is usually specified in terms of a set of desirable properties that a system should have a reasonable number of in order to be considered an intelligent agent.

- List a set of such properties.
- Briefly discuss two of the systems that we looked at during the course in terms of these properties.

Question 2. Precision and Recall [5 marks]

Suppose a system for retrieving names from a web page is run on a page that contains 1000 names, and returns a set of 600 tokens, of which just 450 are actually names.

- What is the recall of the system on this page?
- What is the precision of the system on this page?

Question 3. Naïve Bayes inference [10 marks]

Suppose an agent needs to classify an email message as “spam” or “not spam” on the basis of the presence or absence of words in the message using the Naïve Bayes formula:

$p(\text{“spam”} \mid \text{set of words in message}) =$

$$\frac{\prod_i \left[\, p(\text{word}_i \mid \text{“spam”})\; p(\text{“spam”}) \,\right]}{\prod_i \left[\, p(\text{word}_i \mid \text{“spam”}) \,\right]\; p(\text{“spam”}) \;+\; \prod_i \left[\, p(\text{word}_i \mid \text{“non-spam”}) \,\right]\; p(\text{“non-spam”})}$$

(a) [3 marks] What probabilities could the system measure from a training set of classified email messages?

(b) [2 marks] What assumption do we make about the probabilities to use the Naive Bayes formula.

(c) [5 marks] Suppose we have a set of 1000 messages, of which 500 are spam, and we have counted the number of spam and non-spam messages that contain each of the two words “dollars” and “click” (in the table below), compute the estimated probability that a message containing the words “dollars” but not “click” is a spam message?

Word	Number of Spam with word	Number of non-Spam with word
dollars	50	5
click	200	1

If you do not have a calculator, you may leave your answer as an expression containing fractions.

Question 4. Graph Plan [10 marks]

Graph Plan is an iterative planning algorithm. It represents states as collection of nodes representing propositions, and actions as nodes linked from the proposition nodes representing preconditions and linked to nodes representing effects.

In each iteration,

- It extends a graph by one level - a layer of action nodes and a layer of proposition nodes,
- It identifies any mutually exclusive actions and propositions
- If all goal propositions are in the final layer, it searches for a successful plan in the current graph.

(a) Draw the graph for the planning problem below when it has been extended one level.
  - Show the nodes for the actions and the propositions. (Don't forget "maintenance" actions.)
  - Show any mutex relations between pairs of actions and between pairs of propositions.

(b) Why won't it find a plan for the goal, even though both of the propositions are in the final layer?

Initial State: Prop1 and Prop2

Goal State: Prop3 and Prop4

Possible actions:

Action1:	**pre**:	Prop1
	add:	Prop3
	del:	Prop2
Action2:	**pre**:	Prop2
	add:	Prop4
	del:	-

Note: Actions are mutex if

- one deletes an effect of the other
- one deletes a precondition of the other
- their preconditions are inconsistent.

Question 5. Your Project [20 marks]

(a) [12 marks] The course considered a number of papers related to information extraction agents. Identify three technologies or algorithms in these papers that could be be used to improve the agent you created for your project.

- Specify the paper (title or authors), the technology or algorithm,
- Very briefly explain how it would improve your agent.

(b) [8 marks] How would you measure the performance of your agent?

- If precision and recall is appropriate, describe how you would measure these parameters.
- If not, state what parameters would be appropriate and justify your choice.

??Note: please answer this question, even if you didn't complete your project.

Question 6. Wrapper Induction for Information Extraction [40 marks]

(a) [5 marks] What is a wrapper, and how is it used for Information Extraction?

Consider the following example page:

```
<HTML><TITLE>Some Country Codes</TITLE>
<BODY><B>Some Country Codes</B><p>
<B>Congo</B> <I>242</I><BR>
<B>Egypt</B> <I>20</I><BR>
<B>Belize></B><I>501</I><BR>
<HR><B>End</B></BODY></HTML>
```

(b) [7 marks] What wrapper would be learned for the example above by the BuildHLRT algorithm (described in paper titled 'Wrapper Induction for Information Extraction")?

(c) [7 marks] What wrapper would be learned for the example above by the STALKER algorithm (described in paper titled 'A Hierarchical Approach to wrapper Induction")?

(d) [15 marks] In the paper titled "A scalable Comparison-Shopping Agent for the World-Wide Web", the agent called Shopbot is able to identify product description formats (wrappers). What format it would learn from this example? Briefly describe how it learns the formats.

(e) [10 marks]Compare the three systems and list the advantages and disadvantages of each system.

Part II

Do **three** of questions 7 to 11.

Question 7. Games [30 marks]

Describe *two* applications of intelligent agents to computer games. For each application,

- Describe the problem or need that the agent is intended to meet
- Outline the technology/techniques it uses
- Evaluate the effectiveness of the approach for improving that category of computer game.
- Either suggest improvements that could be made to the system, or outline alternative uses of the approach that you think would be more effective.

Question 8. Mobile Robots [30 marks]

The design of an intelligent mobile robot to operate in places with lots of people (eg, a museum guide, footpath cleaning robot, cocktail party waiter, etc) must deal with a wide range of problems at didfferent levels of abstraction.

- Briefly (one or two sentences each) describe at least 6 problems.
- For two of the problems, outline an approach to solving the problem in more detail.

Question 9. Help Agents [30 marks]

Microsoft Office has a help agent (the paper clip “Clippit” and friends) which seems to be generally disliked by users, in spite of the great need of most users for intelligent help. Clippit is based on research that developed an intelligent user help system for Excel.

- Describe the goals of the help system.
- Describe the key ideas of the technology underlying “Clippit”.
- Explain some of the problems and difficulties in creating a helpful help system (rather than an annoying one).

Question 10. Spam Filtering [30 marks]

Dealing with the massive quantity of spam email requires intelligent agents that will distinguish spam messages from real messages. The paper by Sahami *et. al.* described an agent that used a Naive Bayes classifier.

- Explain the key differences between this approach, and the more common (and simpler) approach using hand-built rules.
- Outline how the system works, and identify what data has to be provided to the system.
- Discuss the requirements on a good spam filter and how it could be evaluated.
- Describe the advantages and limitations of the Naive Bayes classifier approach.

Question 11. Programming by Demonstration [30 marks]

The two papers by Paynter/Witten and Lau *et. al.* both described user interface agents that would help users to automate repetitive tasks.

- Describe the key differences between the approaches in the two papers.
- Describe the kinds of tasks that each system is able to do.
- Discuss the limitations of these systems — how useful could they be in practice?

****