Part I - Multiple Choice. 10 points total.

Circle *the one* answer that best answers the question.

1. Which of the following is an effect axiom in situation calculus?

a. Holding(g,s) → Poss(Release(g),s)
b. Poss(Grab(g),s) → Holding(g,Result(Grab(g),s))
c. At(Agent,[1,2],Result(Go([1,1],[1,2),S0))
d. Effect: ~Holding(g)

2. Inference by enumeration

a. is based on conditional probabilities between atomic events
b. is based on a list of random variables along
c. is based on the full joint distribution of atomic events
d. is based on a list of atomic events alone

3. An ontology

a. is a method for representing uncertainty
b. is a subset of first-order logic
c. is an inference mechanism
d. provides a vocabulary for expressing knowledge

4. An inference procedure

a. is a knowledge representation
b. provides rules for deriving new facts from existing facts
c. is a proof
d. is a type of Bayesian network

5. An inference method is *sound* if it

a. can derive any sentence that is entailed
b. only derives entailed sentences
c. is efficient in both time and space
d. is not NP-complete

6. An inference method is *complete* if it

a. can derive any sentence that is entailed
b. only derives entailed sentences
c. is efficient in both time and space
d. is not NP-complete.

7. The most general unifier of *Q(y,G(A,B)), Q(G(x,x),y) is*

a. *{x/A, y/A}*
b. *{x/A, y/B}*
c. *{x/A, y/(G(A,B)}*
d. no unifier

8. Which of the following sentences is not always TRUE?

a. *A* => A
b. *A* => *B*
c. *A V B V ~B*
d. none of the above

9. Which of the following illustrates the *product rule*?

a. *P(A,B)= P(B|A)P(A)*
b. *P(A,B)=P(A|B)P(B|A)*
c. *P(A,B)=P(A)P(B)*
d. none of the above

10. Conditional probabilities can be defined in terms of unconditional probabilities using which equation?

a. *P(A,B)=P(A)P(B)*
b. *P(B|A) = (P(A|B)P(B)) / P(A)*
c. *P(A|B) = P(A,B)/P(B)*
d. none of the above

Part II – Problem Solving. 75 points total.

1. [15 points] Representation: AI as Search

Describe how to formulate STRIPS, POP, and Resolution Theorem Proving as heuristic search.

a) STRIPS

- What does each state represent?

- What is the initial state?

- Describe the goal state/test.

- Explain how the successor function would operate.

- Give an example heuristic function.

b) POP

- What does each state represent?

- What is the initial state?

- Describe the goal state/test.

- Explain how the successor function would operate.

- Give an example heuristic function.

C) Resolution Theorem Proving

- What does each state represent?

- What is the initial state?

- Describe the goal state/test.

- Explain how the successor function would operate.

- Give an example heuristic function.

2. [15 points] Resolution Theorem-Proving

Everyone at the Halloween party wore a costume. While there, everyone also watched the "1571 Video". Only people at the party had seen that video. Nordenburg never wore a costume. Prove that Nordenburg didn't see the "1571 Video".

A. [5 points] Represent the above scenario in *First-Order* Logic, using the following predicates:

Partied(x): x went to the Halloween party
Costumed(x): x wore a costume
Watched(x): x watched the "1571 Video"

B. [5 points] Put the resulting sentences into Conjunctive Normal Form.

C. [5 points] Using proof by refutation and resolution as the single inference rule, show the resolution proof that proves or disproves the goal (show any unifications required).

3. [15 points] Planning

Luke has to travel from Pittsburgh to Harrisburg to deliver a letter to Ed at his office. Traveling between the two cities will make Luke hungry and tired. Plus, Luke always wakes up hungry. Ed's office offers both food and bed. Luke's goal is to deliver the letter, be well fed and well rested. The actions available to Luke include: *Go* from one city to another, *Deliver* something to someone, *Eat*, and *Sleep*.

A. [5 points] Representation

Represent the initial state, the goal state, and the actions in the above scenario, using STRIPS notation. Use lower case for variables and upper case for constants. Assume a simplified world such that these actions are all that you will need to solve this planning problem.

B. [5 points] Planning with State-Space Search

- Show the first expansion in the search space, using a *progression* planner.

- Show the first expansion in the search space, now using a *regression* planner.

C. [5 points] Partial-Order Planning

- Show a first expansion in the search space, using regression .

- Show an example of a partially ordered solution.

4. [15 points] Probabilities

A. [10 points] Refer to the following *full joint distribution table* over three binary random variables: whether a class size is large or small, whether that class's teacher is old or young, and whether that class has a reputation of being hard or easy.

	Size=large		Size=small	
	Teacher=old	Teacher=young	Teacher=old	Teacher=young
Reputation=hard	.108	.012	.072	.008
Reputation=easy	.016	.064	.144	.576

a) What is P(Size=large, Reputation=hard, Teacher=young)?

b) What is P(Reputation=hard V Size=large)?

c) What is P(Reputation=hard | Size=large)?

d) What is **P**(Size)?

e) How many more probabilities would you need to specify if you added a fourth binary random variable, that exhibited absolute independence from the existing three?

B. [5 points] B.After your yearly physical, your doctor has both good and bad news. The bad news is that you tested positive for a serious disease and that the test is 99% accurate (i.e., the probability of testing positive when you do have the disease is .99, as is the probability of testing negative when you don't have the disease). The good news is that this is a rare disease, striking only 1 in 10,000 people of college student age.

a) Formalize the 3 prior and conditional probabilities in the English description above.

b) You want to compute the probability that you actually have the disease, given your physical's results. Formalize this in terms of a conditional probability . (You don't have to actually compute it!)

5. [15 points] Bayesian Networks

P(C=true) = 0.1

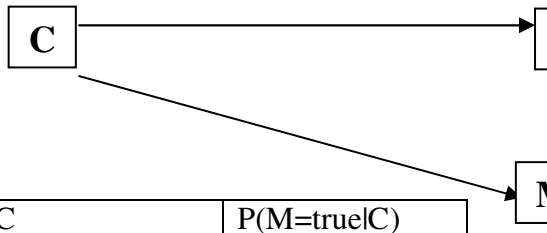

C	P(L=true|C)
True	.2
False	.7

C	P(M=true|C)
True	.4
False	.5

L	M	P(B=true|L,M)
True	True	.2
True	False	.9
False	True	.3
False	False	.8

1. [8 points] Calculate the value of P(C,L,¬M,¬B) for the Bayesian net shown above. You do not have to reduce the answer to a single number, but you should show enough work such that someone with a calculator could easily figure out the answer.

2. [4 points] Explain concretely and quantitatively why using a Bayesian network to compute the value in part 1 is a better approach than using inference by enumeration.

3. [3 points] Suppose I had constructed the network in part 1 by first adding node B (rather than first adding node C as shown above). Explain in one or two sentences whether this would have yielded a "better" network?

Part III - *Short* Answers. 15 points total.

1. [2 points] Give one reason why you might want/need to use *situation calculus* instead of the *STRIPS* approach to planning, and one reason why you might want/need to choose STRIPS instead of the situation calculus.

2. [3 points] Represent "Dr. Litman is a professor and Pitt is a university in Maryland" in *propositional logic*, and use your example to explain the difference between *syntax* and *semantics*.

3. [2 points] Explain the terms *causal link* and *promotion.*

4., [2 points] *When* can you and *why* would you want to use a Horn clause version of a logic?

5. [2 points] What does it mean for a partial plan to be *complete* and *consistent*?

6. [2 points] What is a *frame axiom* and what representation/inference method uses it?

7. [2 points] What is *Bayes' Rule* (i.e., write it down), then explain why Bayes' Rule is useful for computing P(Flu | Fever)?