Name:

Examination # **2**

Date: **11/13/97 26**

Please indicate your section, or if you do not remember, put a tick mark on the last line.

Time	Instructor	
10-11	Davis	
11-12	Davis	
10-11	Szolovits	
11-12	Szolovits	
1-2	Shrobe	
1-2	Berwick	
2-3	Berwick	
2-3	Lozano-Perez	
3-4	Lozano-Perez	

Please indicate the TA you are assigned to.

Calista Tait	
Dennis Ruhl	
Irene Wilson	
Jonathan Loflin	
Ozlem Uzuner	
Serkan Arslanalp	

Problem 1	
Problem 2	
Problem 3	
Problem 4	
Total	

Problem 1 True of False (20 Points)

Circle T or F for true or false. Each correct answer will score 2 points.

Each wrong answer will score -2 points, so you should not guess.

1. T F Many rule-based systems can answer some questions about their own behavior, such as certain questions focused on "how" and "why," but rule-based systems generally cannot explain the origins of the rules.

2. T F The main purpose for introducing the rete algorithm was to connect rule based systems to information in relational databases.

3. T F If you use a beam search, without backtracking, to a specified depth, you may miss the shortest path from start to goal.

4. T F If you use alpha-beta search, to a specified depth, you may miss a better move that would be found by mini-max alone, looking ahead to the same specified depth.

5. T F If you face an opponent that uses mini-max to a specified depth to play a game, you can reduce the ability of that opponent to win by playing what you believe to be your second-best move, rather than your best move.

6. T F The labeled arcs on a semantic transition tree grammar often provide domain specific contraints on sentence interpretation.

7. T F In semantic transition tree grammars, recursions handle grammatical constructions that would otherwise require loops.

8. T F There is no need for backup when using semantic transition tree grammars to analyze English sentences

9. T F Relational databases provide all the functions of frame systems, but frame systems are more commonly used in AI applications because there is rarely a need to handle huge quantities of data.

10. T F The performance of a nearest-neighbor recognition system will, in general, go down if you add measurements in dimensions that are unrelated to identity.

Problem 2 Nearest Neighbors (25 points)

A local machining company, Self-adjusting Widgets (SAW), has contracted with you to help them analyze some problems in machine learning. SAW produces widgets on a completely automated assembly-line. At some point widgets must be sorted into two grades: "GREAT" and "GOOD". The customers of SAW are picky about their widgets, if they pay extra for a GREAT widget they want to get a GREAT widget (i.e. **it is important to never classify a GOOD widget as a GREAT widget**). SAW loses the additional profit they might make by charging more for a GREAT widget, but they are not too worried about it (i.e. **you can classify some of the GREAT widgets as GOOD widgets, but fewer is better**).

In order to classify these widgets SAW intends to take a two measurement of each widget. SAW has two problems. First they need to pick suitable sensors to make these measurements. Once they have settled on sensors, they need to select an algorithm to perform the automatic classification. They want you to help them with both of these problems.

SAW sends you several files of data that have been generated on an assembly-line. Each file contains measurements of widgets taken with two sensors which are being evaluated. Each file of data is displayed as a set of points on a two dimensional graph: GREAT widgets are displayed as a 'diamond'; GOOD widgets are displayed as a 'plus'. For each file of data you are asked to determine if Nearest Neighbor classification can be used to automatically classify widgets on the assembly line. If you determine that nearest neighbor is a possible approach you must also specify which type of nearest neighbor algorithm is most appropriate.

The data that SAW sends you is to be considered *training data*. It is meant to show the natural types of variation in both the GREAT and GOOD widgets. You are being asked to comment on whether a nearest neighbor system trained with this data (has this data in memory) can be used to correctly classify widgets which are *not* in the *training data*. You can assume that this new data, which is often called the *test* data, has the same types of variations as the training data, but is not precisely the same.

Note: In the questions below you will be asked about Nearest Neighbor using a Euclidean distance between points. The Euclidean distance between two dimensional points X and Y is:

$$d(X,Y) = \sqrt{(X_1 - Y_1)^2 + (X_2 - Y_2)^2}$$

where X_1 is the value of dimension 1 and X_2 is the value of dimension 2. You will also be asked about Nearest Neighbor using weighted Euclidean distance. The weighted Euclidean distance between two points X and Y is:

$$d(X,Y) = \sqrt{w_1(X_1 - Y_1)^2 + w_2(X_2 - Y_2)^2}$$

where w_1 and w_2 are used to differentially weight the two components of the distance.

The graphs representing the five data files are on the following pages.

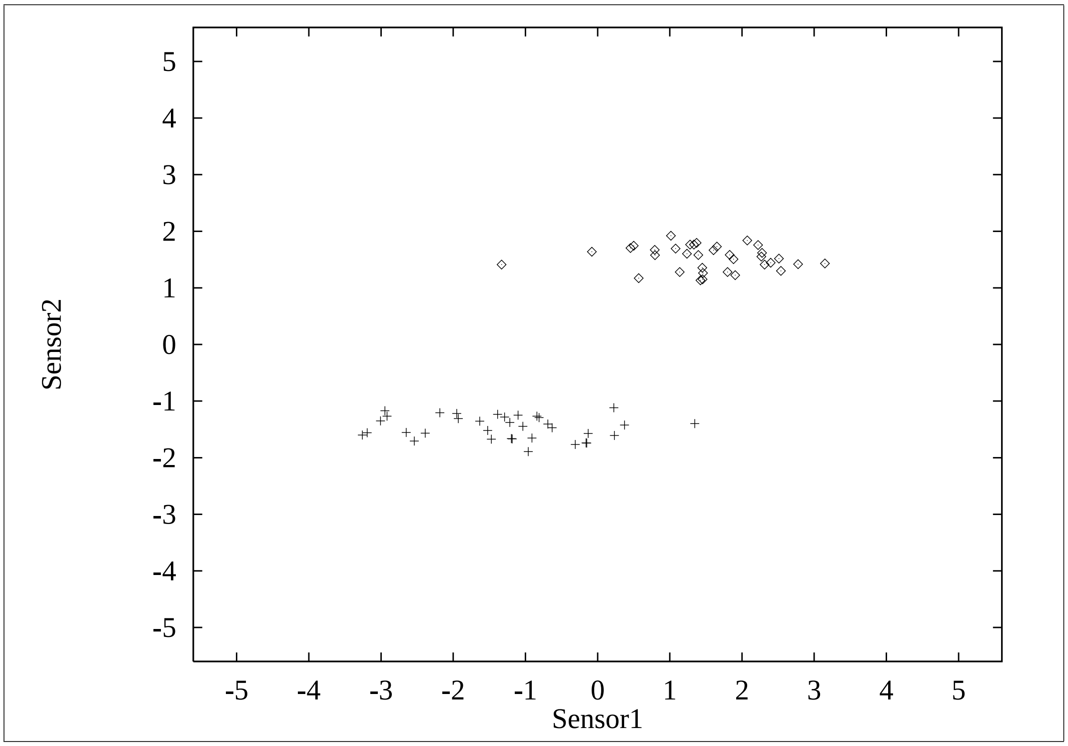

Data from File A.

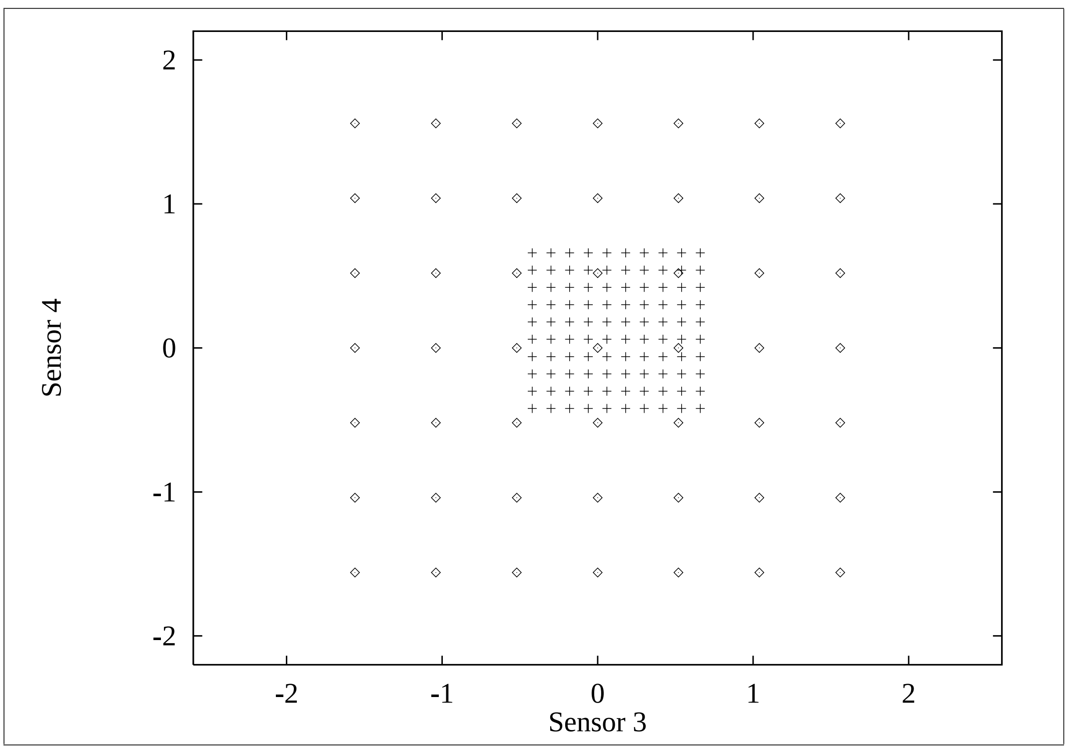

Data from File B.

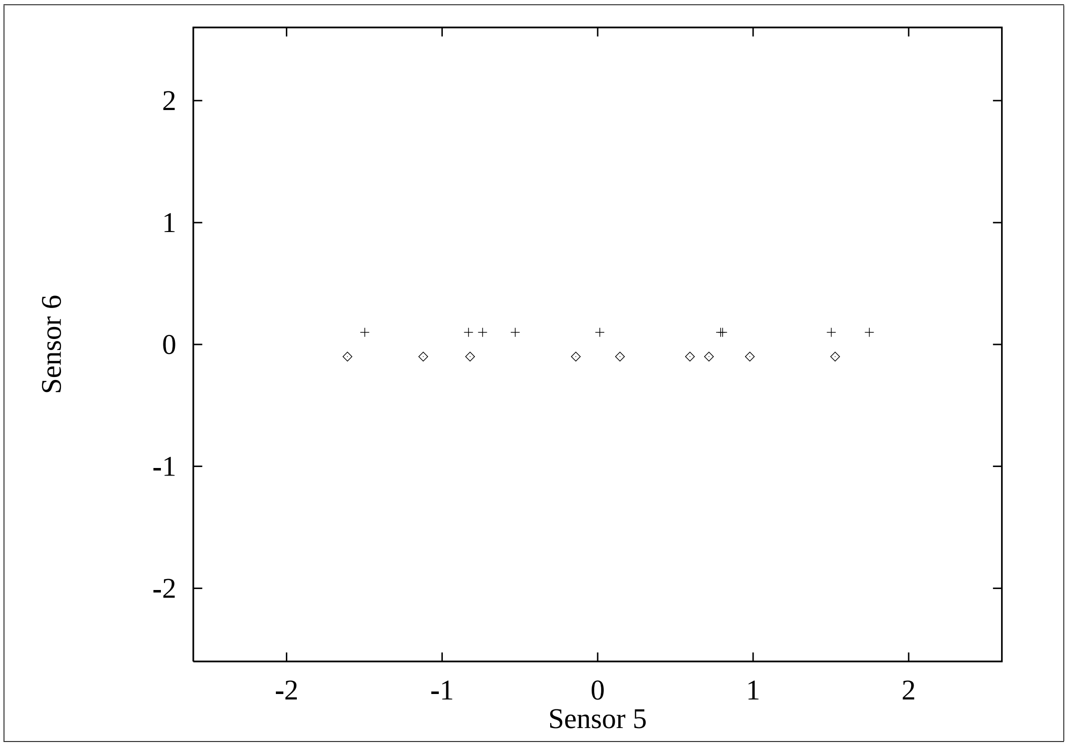

Data from File C.

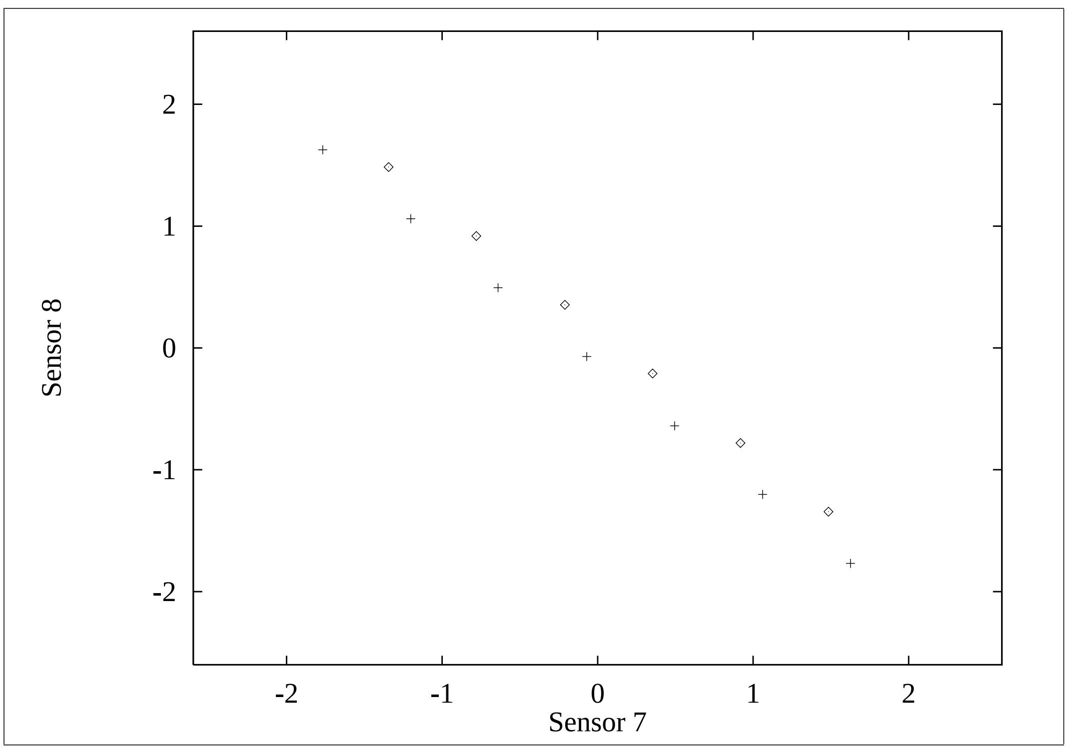

Data from File D.

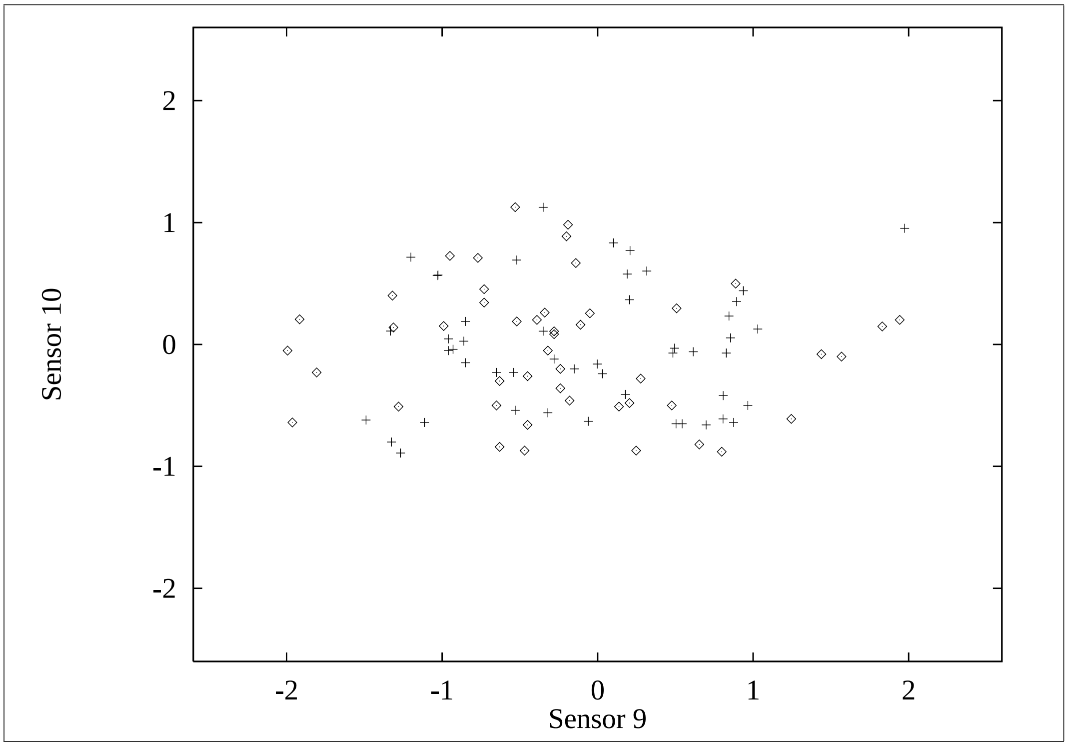

Data from File E.

Questions

For each of the five data files you must first determine if nearest neighbor classification is appropriate. You have one of two options:

- **Yes**: Nearest neighbor classification should work well. The probability of error will be low.
- **No**: This type of problem is too hard. Using nearest neighbor classification will result in too many errors.

If your answer to the above question was **Yes**, SAW is curious about the "correct" nearest neighbor approach for this file. They ask you to recommend the *simplest* possible nearest neighbor classifier that will work for this data. The different possible approaches are listed below. The simplest technique is first and the most complex last. You need only give the corresponding number.

1. One Nearest Neighbor using Euclidean distance between points.
2. One Nearest Neighbor using a weighted Euclidean distance (the difference between each sensor value is multiplied by a sensor specific weight).
3. Five Nearest Neighbors (i.e. find the 5 nearest neighbor and take the majority vote).
4. Five Nearest Neighbors *plus* weighted Euclidean distance.

If you chose one of the weighted Eucldean distance techniques, give a set of weights for each sensor that would work (note: w_i is the weight for sensor i).

Write your answers in the table below:

Data File	Appropriate?	Which Approach?	Weights (If necessary)	
File A			$w1 =$	$w2 =$
File B			$w3 =$	$w4 =$
File C			$w5 =$	$w6 =$
File D			$w7 =$	$w8 =$
File E			$w9 =$	$w10 =$

Problem 3 Classy Objects d'Art (25 Points)

Arriving home after a particularly gruelling 6.034 recitation, you find a letter from a lawyer telling you that you have inherited some valuable works of art. These artworks, worth millions of dollars, have mistakenly been residing in art galleries around the world, but now they are yours.

It is only a matter of time before journalists start phoning you, asking all sorts of dumb questions about the objects you now own, and wanting background information about the artists who created them

Naturally, you decide to use the frame system in `frames.scm` (part of Problem Set 6) to represent knowledge about your new property. This way, journalists can quiz the frame system leaving you to appreciate the finer things in life.

A 6.034 professor drops by and tells you that the primary object classes in your frames systems should be `artwork` and `artist`; that artworks are either 2-dimensional or 3-dimensional; and that 2-d artworks are either paintings, drawings, or prints. As a consultancy fee, the professor takes away a Cezanne painting.

The lawyer's letter tells you that most of the art you've inherited is by Spanish artists. The famous 2-d artworks you've inherited include Picasso's *"Guernica"*, Dali's *"Autumnal Cannibalism"*, Miro's *"Birth of the World"*, all of which are paintings, and *"The Knight, Death, and the Devil"*, a print by Dürer. The only 3-d artworks mentioned by name are Michaelangelo's *"David"* and Rodin's *"The Kiss"*.

All you remember from high school art classes is that Michaelangelo was Italian, and Dürer was German.

As a reminder, the comments at the start of `frames.scm` are:

```
; Some examples of using this system follow:
;
; To assert that all insects are arthropods:
;    (define-class 'arthropod 'insect)
;
; To assert that all arthropods are invertebrates:
;    (define-class 'invertebrate 'arthropod)
;
; To assert that all arachnids are arthropods:
;    (define-class 'arthropod 'arachnid)
;
; To assert that all insects have six legs:
;    (define-class-default 'insect 'num-legs 'six)
;
; To assert that all arthropods have jointed legs:
;    (define-class-default 'arthropod 'leg-type 'jointed)
;
; To assert that all bees are insects:
;    (define-class 'insect 'bee)
;
; To assert the details of bees coloring:
;    (define-class-default 'bee 'color 'striped)
;
; To assert that Eric is a bee:
;    (define-instance 'eric 'bee)
;
; To assert that Buzzy is a bee:
;    (define-instance 'buzzy 'bee)
;
; To assert that Eric has lost some legs:
;    (assert-property 'eric 'num-legs 'three)
;
;
; After these assertions, we could represent the knowledge in the system
; as a tree-diagram with "a-kind-of" (ako) links joining a subclass to its
; superclass above, and "is-a" (isa) links joining an instance to the class
; it is a member of. And using :: to denote default key values:
```

```
;
;    invertebrate
;      |
;     (ako)
;      |
;    arthropod
;    ::leg-type=jointed
;     |         |
;    (ako)    (ako)
;     |         |
;  arachnid    insect
;              ::num-legs=six
;                   |
;                 (ako)
;                   |
;                  bee
;                  ::color=striped
;                  |          |
;                 (isa)      (isa)
;                  |          |
;                  buzzy      eric
;                             ::num-legs=three
;
; To find whether Buzzy is a bee:
;    (class-member? 'buzzy 'bee)
;
; To find what class Eric is:
;    (get-class 'eric)
;
; To find how many legs Buzzy has:
;    (get-property 'buzzy 'num-legs)
;
; To find whether bees are invertebrates:
;    (subclass? 'bee 'invertebrate)
;
; Note that the answers to the last two examples illustrate inheritance,
; because the relevant values are specified as class-defaults and are
; inherited by all subclasses that do not explicitly over-ride the defaults
```

Part A (6 points)

The same 6.034 professor offers to help you by drawing a tree-diagram of the classes and objects that should be in your frames system. After a few minutes the professor takes a call on his cellular phone, and then rushes off, muttering something about a better-paid consultancy job. The partially complete diagram is on the next page: the professor has used `::by=` to indicate the creator of an artwork and `::nat=` to indicate the nationality of an artist. Fill in the rest of the diagram, showing all the information that was given to you by the professor and the lawyer, and the two facts you recall from your high-school art lessons. You should use `kdd`, `botw`, and `ac` as abbreviations for the artworks by Dürer, Miro, and Dali, respectively, and simply refer to the Rodin as `kiss`. You should assume that all artists in the collection are Spanish unless told otherwise, and don't include any information about the Cezanne that the professor took.

To avoid page-flipping, you can rip the first page of this question out of the booklet.

Part B (7 points)

Using `'nat` for *nationality* and `'by` for *was created by*, and the same class and object labels as in the diagram, show the `frames.scm` expressions you would write to:

B1 (1 point)

Represent the fact that Picasso was an artist.

B2 (1 point)

Represent the fact that Picasso painted *Guernica*.

B3 (1 point)

Represent the fact that Autumnal Cannibalism is a painting.

B4 (1 point)

Represent the fact that Dürer was German.

B5 (1 points)

Represent the default assumption that artists in the collection are Spanish.

B6 (2 points)

Represent the fact that artworks are either 2d or 3d.

Part C (11 points)

The phone starts ringing. It's a journalist who wants answers to some questions. For each question, indicate the Scheme expression that you'd advise the journalist to type to your frames system, and also what response you would expect from the frames system. If the questions requires a yes/no answer, it's ok to return `#t` (true) for "yes" and `#f` (false) for "no".

C1 (1 point)

Who painted *Guernica*?

C2 (1 point)

What is *David*?

C3 (1 point)

Are drawings artworks?

C4 (1 point)

What nationality was Rodin?

C5 (2 points)

Is *Autumnal Cannibalism* an artwork?

C6 (2 points)

Was Dürer German?

C7 (3 points)

Was the sculptor of *The Kiss* Spanish?

Part D (1 point)

Another journalist calls. She asks you the following questions:

- How many 3d artworks are there in your collection?
- What are the names of the Spanish artists?

A 6.034 TA points out that, to answer questions like these, you have to write Scheme code to search the knowledge in the frames system and count or list the number of objects that match the question. The TA hints that there might be a more appropriate technology. Check the one you think is most appropriate for questions like the ones this journalist asked:

1. Bayes nets
2. Forward-chaining rule-based systems.
3. Relational databases
4. A* search
5. Steam engines

Problem 4 Icy Road Networks at Bay (30 points)

NOTE: this problem has four parts, A–D. EACH part may be worked completely INDEPENDENTLY of the others, so you may plan your time accordingly.

Inspector Smith is awaiting the arrival of Mr Holmes and Dr Watson. They are late, and Smith has another appointment (lunch). Looking out the window she wonders whether the British road network has become icy. Both Holmes and Watson are notoriously bad drivers, so if the roads are icy they may crash.

An assistant enters and tells her that Dr Watson has had a car accident. "Watson? Well, could be worse... icy roads! Then Holmes has most probably crashed too. I think I'll take that lunch appointment now."

"Icy roads?", the assistant replies, "It is not that cold, and further all the roads are salted." Inspector Smith is relieved. "Bad luck for Watson. Let us give Holmes ten minutes more."

As a top-notch member of Que Sera Computer Consulting ("Whatever will be, will be, (TM)") you've decided to build a Bayes net to model this situation and that can reason like Inspector Smith. You proceed as follows:

- All the events can be represented as variables with two states, *true* (t) or *false* (f);
- There are three variables:
  1. *Icy roads*, or I
  2. *Holmes crashes*, or H
  3. *Watson crashes*, or W
- With each event you can associate a *probability* value between 0.0 and 1.0
- Watson and Holmes drive different cars
- All the roads are either icy or entirely ice-free
- The assistant's statements are always accurate

I has the effect of increasing the probability of both W and H — icy roads cause Watson or Holmes (or both) to crash. To finish constructing the Bayes network we must set some initial probability assessments at each node, as shown in the figure below. The probability for icy roads is 0.7. Both Holmes and Watson are bad drivers, so we set the probability of a crash if the roads *are* icy at 0.8 and the probability that they will crash if the roads are *not* icy at 0.1 — they really are very bad drivers!

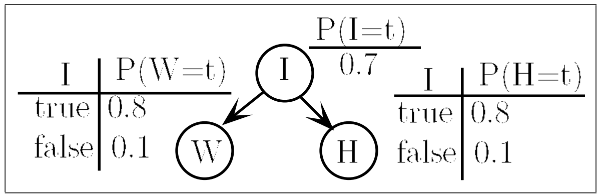

REMEMBER that each part A–D may be done INDEPENDENTLY of the others.

Part A Warm-up!(4 pts)

Why is there *no* directed arc between H and W in the network?

Why No Arc:

1.5in

Part B The Inspector Calls (10 pts)

(i) (3 pts)

Now we can use the Bayes network to do some probability calculations.

What is the numerical value of $P(W = t, I = t)$?

1.75in

(ii) (3 pts)

Calculate the probability of icy roads *given* the knowledge that Watson has crashed, that is, $P(I = t|W = t)$. Write down the calculation symbolically, a formula in terms of probabilities such as $P(I = t)$ that are specified in the probability tables in the figure above. Then write down the formula for $P(W = t)$ also in terms of the entries in the tables.

Symbolic formula for $P(I = t|W = t) =$

1.5in

Symbolic formula for $P(W = t) =$

1.5in

(iii) (4 pts)

In terms of probabilities of fully specified states, write down the *simplest* symbolic formula needed to calculate the probability that Holmes has crashed, conditional on W, that is, incorporating the fact that Watson has crashed, but that we do not yet know whether or not the roads are icy. For the purposes of this problem, the probability of a *fully specified state* is simply a probability in a form like $P(H = x, I = y, W = z)$, that is, where the lowercase variables x, y, z range over possible values like t, f for the uppercase random variables like H, I, etc.

2.0in

Part C Connecting nets (6 pts)

(i) (5 pts)

Suppose that besides being a bad driver and crashing on icy roads, Watson also sometimes drives a Saab with bad tires, denoted by the binary-valued variable BS. We set the value of $P(BS = t)$ to be 0.8. We assume for this problem that BS and I are independent random variables. We may also assume that when Watson drives a Saab with bad tires on *non*-icy roads, there is a 0.85 probability that he will crash (and so a 0.15 probability that he will not crash when driving a Saab with bad tires on non-icy roads). (Note that this means that BS and I are independently causative factors for Watson crashing.)

Show how to update the existing Bayes network to include this new BS factor, using the space below. Add/delete links as appropriate. Naturally, there will have to be some probability table(s) associated with the new node as well as some potential changes to the existing tables; please draw in those tables as well. Fill in those entries for which you can pick values from the original network or from the text above. If a value is not known, leave it blank.

Modify this picture:

New table:

Original table:

P(I=t)
0.7

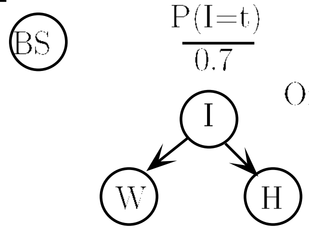

Original table:

I	P(H=t)
true	0.8
false	0.1

Original table:

Original table:

I	P(W=t)
true	0.8
false	0.1

Revised table:

(ii) (1 pt)

The Inspector's assistant told her that the roads are, in fact, *not* icy. That is, after the assistant speaks, the values for W and I are now both known: $W = t$ (Watson has crashed) and $I = f$ (the roads are not icy). As we saw in the earthquake/alarm question in problem set 7, this means that our initial estimate of the likelihood that Watson drives a Saab with bad tires, $P(BS = t)$, might have to be re-estimated, given our new knowledge.

Given that the roads are now known to be *not* icy ($I = f$), and that Watson *has* crashed ($W = t$) which of the following statements best characterizes the probabilistic relationship between $P(BS = t)$ and $P(BS = t|I = f, W = t)$? Circle *one* answer below.

- The two will have the **same value**
- $P(BS = t|I = f, W = t)$ will be **smaller** than $P(BS = t)$
- $P(BS = t|I = f, W = t)$ will be **greater** than $P(BS = t)$
- There is no way to tell the relation between the two, given the information provided

Part D Noisy networks (10 pts)

Of course, we can imagine many possible causes for Watson crashing in addition to icy roads and Watson driving a Saab with bad tires. For instance, we could add a new possible cause, such as E, indicating whether Watson is Exhausted or not. In such cases, as we have seen, it often becomes useful to posit a *noisy-or* model to reduce the number of probability parameters we need to use for calculations. In the following problem, assume that just E and I cause crashes, neglecting BS.

(i) (3 pts)

If the "causal probability" from I to W is 0.7, and the causal probability from E to W is 0.8, then what is the numerical value of the probability $P(W = t|I = t, E = t)$ under the noisy-or assumptions, assuming no leak term is present?

$P(W = t|I = t, E = t)$=

1.0in

(ii) (3 pts)

Now suppose we also add a *Leak* variable, L, to cover the possibility where the "symptom" W appears (Watson crashes), but neither I nor E is true. That is, the $P_c(L)$ is like a probability in that it is always positive.

Please write out symbolically the formula for the noisy-or calculation using the *Leak* binary-valued variable, L. Use $P_c(X|Y)$ to denote the "causal probability" estimates that Y causes X for each conditional probability.

Symbolic formula for $P(W = t|I = t, E = t) =$

1.75in

(iii) (4 pts)

Using the same causal probabilities for I and E as in Part (i) above, suppose the probability $P(W = t|I = t, E = t)$ has been calculated and found to be equal to 0.90 (when Watson is both exhausted and the roads are icy, there is a higher probability of Watson crashing).

Is this probability of 0.90 consistent with the noisy-or assumptions, given a leak variable? Explain why or why not.

Why or why not consistent with noisy-or:

2.0in