6.034 Quiz 4

4 December 2019

Name	**Overlord Hunter**
Email	***SOLUTIONS***

For 1 extra credit point: Circle the TA whose recitations you attend so that we can more easily enter your score in our records and return your quiz to you promptly.

Sydney Gibson **Rui Li** **Allison Tam**

Udgam Goyal **Jennifer Madiedo** **Héctor Vazquez**

Jenna Hong **Jack Murphy** **Eric Wong**

Damon Jones **Mira Partha**

Problem number	Maximum	Score	Grader
1 - Bayes	50	**50**	
2 - Boosting	50	**50**	
Total	**100**	**100**	

SRN	7	**7**	

There are **14** pages in this quiz, including this one, but not including tear-off sheets. Tear-off sheets with duplicate drawings and data are located after the final page of the quiz. **We do not collect tear-off sheets, so please show your work on the quiz pages, not the tear-off sheets.**

As always, the quiz is open book, open notes, open just about everything, including a calculator, but no computers or cell phones.

This page is intentionally blank (except for this sentence).

Quiz 4, Problem 1: Bayesian Inference (50 points)

World-renowned detective Sherlock Holmes has discovered that his partner, Dr. John Watson, has gone missing! Holmes suspects foul play and decides to investigate Watson's disappearance using a BayesNet.

The following (Boolean) variables are at play:

- **A**: Irene **A**dler is acting suspiciously
- **M**: Professor **M**oriarty is acting suspiciously
- **W**: **W**atson has been kidnapped
- **B**: Professor Moriarty is robbing a **b**ank
- **N**: Watson left a **n**ote for Holmes on their dining table
- **E**: The note on the dining table is **e**ncrypted

Holmes devises the following BayesNet:

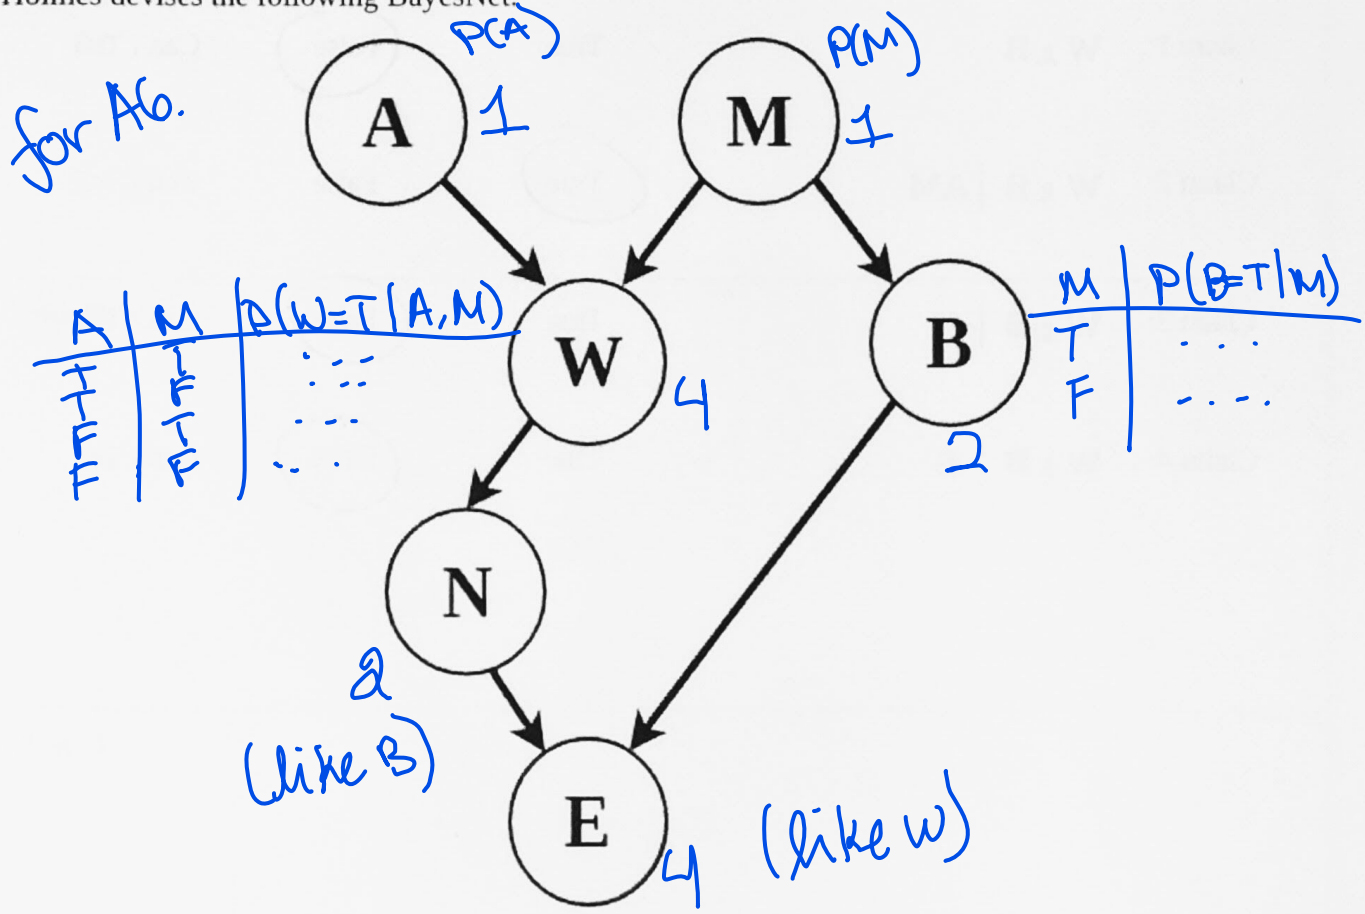

Part A: Where's Watson? (30 points)

A1 (2 points) Using the BayesNet assumption, simplify the following expression:

$$P(N \mid AMWB)$$

Or, if it cannot be simplified, write 'Not possible' in the box provided.

$P(N \mid W)$

By Bayes Net Assumption A, M, & B are non-desc. of N so cond. indep given W.

A2 (20 points): Holmes wants to use d-separation to investigate the following claims.

Claim 1:	$\mathbf{W} \perp\!\!\!\perp \mathbf{B}$	True	(False)	Can't Tell
Claim 2:	$\mathbf{W} \perp\!\!\!\perp \mathbf{B} \mid \mathbf{AM}$	(True)	False	Can't Tell
Claim 3:	$\mathbf{W} \perp\!\!\!\perp \mathbf{B} \mid \mathbf{N}$	True	(False)	Can't Tell
Claim 4:	$\mathbf{W} \perp\!\!\!\perp \mathbf{B} \mid \mathbf{E}$	True	(False)	Can't Tell

Show your work in the boxes on the following page for partial credit.

Claim 1:

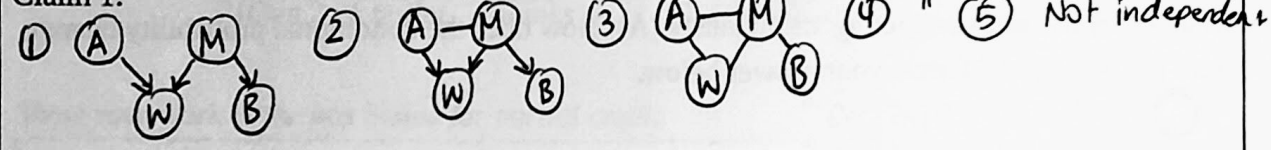

Claim 2:

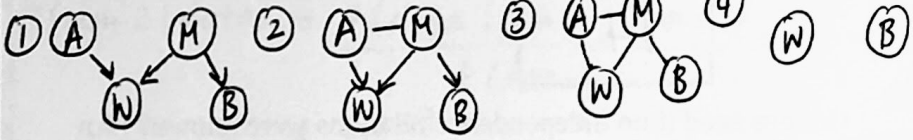

Claim 3:

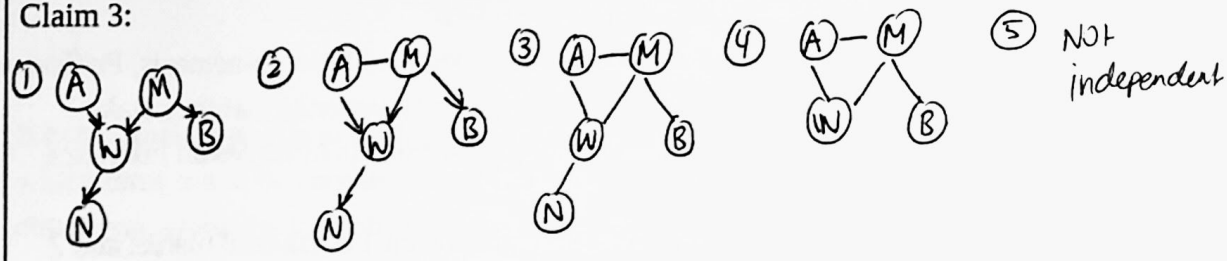

Claim 4:

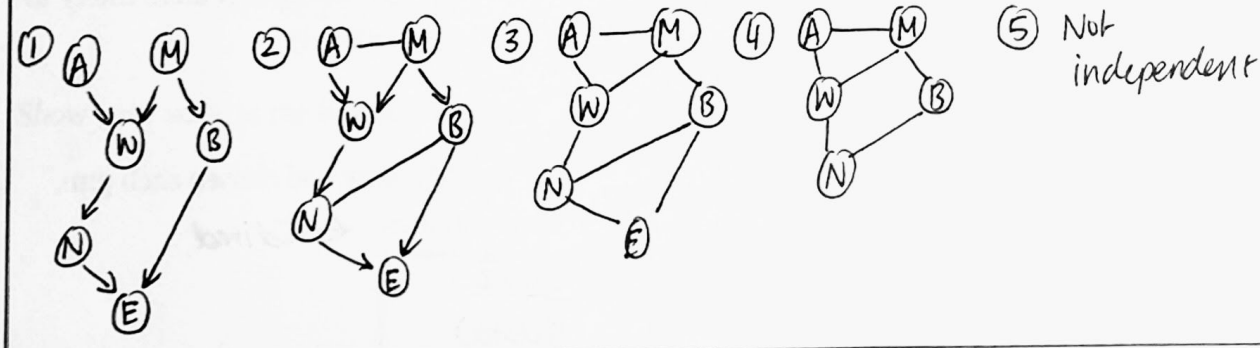

A3 (2 points) Trying to deduce the culprit, Holmes considers the probability that Moriarty is acting suspiciously, given that Watson has been kidnapped ($P(\mathbf{M}=\text{true} \mid \mathbf{W}=\text{true})$). He then discovers that Irene Adler has **not** been acting suspiciously (**A**=false). How does the conditional probability change with the new given? Circle your answer below.

(a.) $P(\mathbf{M}=\text{true} \mid \mathbf{A}=\text{false}, \mathbf{W}=\text{true}) > P(\mathbf{M}=\text{true} \mid \mathbf{W}=\text{true})$

b. $P(\mathbf{M}=\text{true} \mid \mathbf{A}=\text{false}, \mathbf{W}=\text{true}) < P(\mathbf{M}=\text{true} \mid \mathbf{W}=\text{true})$

c. $P(\mathbf{M}=\text{true} \mid \mathbf{A}=\text{false}, \mathbf{W}=\text{true}) = P(\mathbf{M}=\text{true} \mid \mathbf{W}=\text{true})$

d. Can't Tell

By explaining away if A=false, then the P(M=T) ↑ b/c we know W=T so its even more likely that M=T vs. when we knew nothing about A

A4 (6 points) Given the independence relations shown in the BayesNet, how many parameters does Holmes need to fully describe this system?

14

How many parameters would Holmes need if no independence relations were known? (An unsimplified answer is fine.)

$2^6 - 1 = 63$

↑ -1 b/c of exhaustion

Part B: Bayes to the Rescue! (20 points)

Holmes discovers that his arch-nemesis, Professor Moriarty, is holding Watson hostage. Moriarty challenges Holmes to the following coin game for Watson's freedom:

Moriarty has two coins, a fair coin and a biased coin. He will select a coin and flip it 3 times. Holmes's objective is to figure out, based on the observed sequence of flips, which coin Moriarty chose. Holmes is given the following information:

Coin 1 (fair)	**Coin 2 (biased)**
$P(\text{H}) = 0.5$ $\quad P(\text{T}) = 0.5$	$P(\text{H}) = 0.8$ $\quad P(\text{T}) = 0.2$

P (Moriarty selects Coin 1) = 0.6
P (Moriarty selects Coin 2) = 0.4

Moriarty secretly chooses one of the coins and flips it three times. He and Holmes observe the following sequence of flips:

HHT

B1 (8 points): Calculate the following probabilities, leaving the marginal probability of observing the given sequence, P(HHT), uncalculated in your answer:

P (Coin 1 | HHT) = $\boxed{\frac{0.075}{P(HHT)}}$

P (Coin 2 | HHT) = $\boxed{\frac{0.0512}{P(HHT)}}$

Show your work in the box below for partial credit.

$$P(\text{Coin 1} \mid HHT) = \frac{P(HHT \mid \text{Coin 1}) \cdot P(\text{Coin 1})}{P(HHT)} = \frac{(0.5 \cdot 0.5 \cdot 0.5) \cdot (0.6)}{P(HHT)} = \frac{0.075}{P(HHT)}$$

$$P(\text{Coin 2} \mid HHT) = \frac{P(HHT \mid \text{Coin 2}) \cdot P(\text{Coin 2})}{P(HHT)} = \frac{(0.8 \cdot 0.8 \cdot 0.2) \cdot (0.4)}{P(HHT)} = \frac{0.0512}{P(HHT)}$$

B2 (2 points) Given the observed data, which coin did Moriarty most likely select?

Coin 1 Coin 2

B3 (8 points) Wanting to further test Holmes's powers of deduction, Moriarty flips the coin a fourth time. The full flip sequence is now:

HHTH

Calculate the following updated probabilities, leaving the marginal probability of observing the given sequence, *P*(HHTH), uncalculated in your answer:

P(Coin 1 | HHTH) = $\frac{0.0375}{P(HHTH)}$

P(Coin 2 | HHTH) = $\frac{0.04096}{P(HHTH)}$

Show your work in the box below for partial credit.

$P(\text{Coin 1} \mid HHTH) = \frac{P(HHTH \mid \text{Coin 1}) \cdot P(\text{Coin 1})}{P(HHTH)} = \frac{(0.5^4)(0.6)}{P(HHTH)} = \frac{0.0375}{P(HHTH)}$

$P(\text{Coin 2} \mid HHTH) = \frac{P(HHTH \mid \text{Coin 2}) \cdot P(\text{Coin 2})}{P(HHTH)} = \frac{(0.8^3 \cdot 0.2)(0.4)}{P(HHTH)} = \frac{0.04096}{P(HHTH)}$

B4 (2 points) Given the observed data, which coin did Moriarty most likely select?

Coin 1 Coin 2

Problem 2: Friends Who Boost Together (50 points)

Part A: The One with the Thanksgiving (35 points)

For Thanksgiving dinner this year, Monica decides to go all out. Her friends place bets on which parts of the dinner will go wrong: Turkey (**T**), Potatoes (**P**), Gravy (**G**), Cranberry sauce (**C**), Stuffing (**S**), or Dessert (**D**). Treating each friend as a **weak binary classifier**, the table below shows where their predictions went wrong. Correct predictions are +; incorrect predictions are -.

	T	**P**	**G**	**C**	**S**	**D**
h1 (Ross)	-	+	-	-	-	+
h2 (Rachel)	+	+	-	+	+	-
h3 (Chandler)	-	+	+	-	+	-
h4 (Phoebe)	-	+	+	-	+	+
h5 (Joey)	+	+	-	-	-	+
Actual	+	+	+	-	+	-

A1 (2 points) Fill out the table below to show each friend's misclassifications.

	Misclassified Training Points
h1	T, G, S, D
h2	**G, C**
h3	**T**
h4	**T, D**
h5	**G, S, D**

A2 (30 points) Fill out the table below to complete the first 3 rounds of Adaboost, choosing the classifier with the error rate **furthest from ½**. Break ties according to which classifier comes first in this list: **h1, h2, h3, h4, h5.**

For partial credit, show your work on the next page.

	Round 1	**Round 2**	**Round 3**
Weight of T (w_T)	**1/6**	**1/2**	**5/16**
Weight of P (w_P)	**1/6**	**1/10**	**1/4**
Weight of G (w_G)	**1/6**	**1/10**	**1/16**
Weight of C (w_C)	**1/6**	**1/10**	**1/4**
Weight of S (w_S)	**1/6**	**1/10**	**1/16**
Weight of D (w_D)	**1/6**	**1/10**	**1/16**
Error rate of h1 (ε_{h1})	**4/6**	**8/10**	**1/2**
Error rate of h2 (ε_{h2})	**2/6**	**2/10**	**5/16**
Error rate of h3 (ε_{h3})	**1/6**	**5/10**	**5/16**
Error rate of h4 (ε_{h4})	**2/6**	**6/10**	**6/16**
Error rate of h5 (ε_{h5})	**3/6**	**3/10**	**3/16**
Chosen classifier (h)	**Chandler (h_3)**	**Ross (h_1)**	**Joey (h_5)**
Error rate of classifier (ε)	**1/6**	**4/5**	**3/16**
Voting power (α)	**½ ln(5)**	**½ ln(1/4)**	**-½ ln(3/13)**

For partial credit on A2, you can show your work in the box below.

Round 1

All weights = 1/6

$\varepsilon_{h_i} = \sum_{\text{misclassified}} w_i$

$\varepsilon h_1 \to 4/6$

$\varepsilon h_2 \to 2/6$

$\varepsilon h_3 \to 1/6$

$\varepsilon h_4 \to 2/6$

$\varepsilon h_4 \to 3/6$

Error = $\min(|\varepsilon - 1/2|)$

$\alpha = \frac{1}{2}\ln\left(\frac{1-1/6}{1/6}\right) = \frac{1}{2}\ln 5$

$H(\vec{x}) = \text{Sign}(\frac{1}{2}\ln 5 \cdot h_3)$

Round 2

got wrong in round 1

got right in round 1

wrong $w_i = \frac{w_{old}}{2\varepsilon}$, right $w_i = \frac{w_{old}}{2(1-\varepsilon)}$

$\sum_{wrong} w_i = \sum_{right} w_i = 1/2$

$w_T = 1/2$

$w_P = 1/10$

$w_G = 1/10$

$w_C = 1/10$

$w_S = 1/10$

$w_D = 1/10$

$\varepsilon h_1 = 8/10$

$\varepsilon h_2 = 3/10$

$\varepsilon h_3 = 5/10$

$\varepsilon h_4 = 6/10$

$\varepsilon h_5 = 3/10$

chose h_1 before h_2 doing $\min(|\varepsilon - 1/2|)$ and h_1 before h_2

$\alpha = -\frac{1}{2}\ln\left(\frac{8/10}{1-8/10}\right) = -\frac{1}{2}\ln 4$

$H(\vec{x}) = \text{Sign}(\frac{1}{2}\ln 5 \cdot h_3 - \frac{1}{2}\ln 4/5 \cdot h_1)$

Round 3

applies again

$w_T = 5/16$

$w_P = 1/4$

$w_G = 1/16$

$w_C = 1/4$

$w_S = 1/16$

$w_D = 1/16$

$\varepsilon h_1 = 1/2$

$\varepsilon h_2 = 5/16$

$\varepsilon h_3 = 5/16$

$\varepsilon h_4 = 6/16$

$\varepsilon h_5 = 3/16$

$\frac{1}{2}\ln\left(\frac{1-3/16}{3/16}\right) = \frac{1}{2}\ln\left(\frac{13}{3}\right)$

A3 (3 points) What ensemble classifier **H(x)** would you generate after 3 full rounds of Adaboost?

H(x) = SIGN(½ ln(5)*h₃ + ½ ln(1/4)*h₁ - ½ ln(3/13)*h₅)

also $H(x) = \text{Sign}(\frac{1}{2}\ln 5 \cdot h_3 - \frac{1}{2}\ln 4 \cdot h_1 + \frac{1}{2}\ln 13/3 \cdot h_5$

Part B: Central Perk (5 points)

Joey is attempting to settle the hotly debated argument: *Did Ross cheat on Rachel, or were they "on a break"?* Joey identifies seven key events **(A-G)** and chooses 3 random people **(P_1, P_2, P_3)** to check **the** correctness of each event. He treats each person as a **weak classifier** and records their misclassifications in the table below.

Classifier	Misclassifications
Person 1 (P_1)	A, D
Person 2 (P_2)	B, E, G
Person 3 (P_3)	C, F

$\alpha_{P_3} > \alpha_{P_2}$ so P_3 has higher voting power
=> everything P_3 misclassifies will be misclassified by $H(x)$

B1 (3 points) Joey constructs this following classifier:

$$H(x) = 3*P_2(x) + 4*P_3(x)$$

Is this a perfect classifier (i.e., correctly classifies all training points)? If NO, indicate which training points are misclassified using this classifier.

YES | **NO, H(x) misclassifies:** C, F

B2 (2 points) Assume Person 1 misclassified <u>one</u> additional training point (B, C, E, F, or G). If this is the case, is it possible to create a perfect classifier by running Adaboost?

YES | **NO** | **CAN'T TELL**

There is no combo of α's that will allow perfect classification b/c will now have overlapping errors

Part C: Smelly Cat 2.0 (10 points)

Phoebe is writing a song about Adaboost, but she got her facts all mixed up. For each claim below, either write **TRUE,** or write **FALSE and explain why it's false in <u>10 words or less</u>.**

1. Adaboost is fast and prone to overfitting.

FALSE: Adaboost is robust to overfitting - generalizes well with each iteration/round

2. When choosing classifiers with the smallest error rate less than ½, the weights of all training points in a particular round can be negative.

FALSE: Weights of training points cannot be negative.

3. If more than two weak classifiers misclassify the same training point, the created classifier cannot be a perfect classifier.

FALSE: Could be a perfect classifier if other weak classifiers (which correctly classify the training point) have larger voting power.

4. If you let Adaboost run on any dataset, it will not terminate until instructed to do so.

FALSE: Adaboost can terminate if all points are correctly classified or if no more weak classifiers are "good enough".

5. The following is a possible set of weights for an unknown round of Adaboost: 1/10, 1/10, 2/10, 6/10.

FALSE: These weights are not possible - no combination of weights adds up to ½. Also accepted that there cannot be weights greater than ½.

Problem 3: Spiritual and Right Now (6 points + 1 bonus)

For each question, **write in the box** provided the letter corresponding to the **one** best answer **and circle** the answer. There is **no penalty for wrong answers,** so it pays to guess in the absence of knowledge.

B

1. Davis explained that the problem with the general problem solver (GPS) was:
   A. It couldn't generalize.
   B. It didn't have domain-specific knowledge for problems.
   C. There was no way to quantify the difference between two states.
   D. It had no sense of universal subgoaling.

B

2. According to Pratt, which of the following statements is **TRUE?**
   A. Only fully autonomous cars can save lives.
   B. The hardest part of autonomous vehicle research is prediction.
   C. Billions of miles of simulated driving data solves the need for training data.
   D. The main motivators behind autonomous driving research are safety and traffic mitigation.

D

3. What did Katz mean when he said the human vision system is biased?
   A. Humans recognize other humans faster than they recognize animals.
   B. Humans recognize objects faster than they recognize images of objects.
   C. Humans recognize ImageNet images more accurately than machines recognize ImageNet images.
   D. Humans recognize objects that make sense in a scene.

D

4. According to Berwick, what is the critical step that occurs uniquely in human children aged 3-4?
   A. They learn to emulate the cadence of their native language.
   B. They understand words in their native language.
   C. They learn to string together words in their native language to ask for things they want.
   D. They understand and can create hierarchical phrases in their native language.

A

5. Which of the following was **NOT** an example Holmes discussed as an application of concept patterns in the Genesis story understanding system?
   A. Concept patterns enable alternate endings to be generated from a story.
   B. Concept patterns enable comparison between stories.
   C. Concept patterns enable different viewpoints of the same story.
   D. Concept patterns enable summarization of stories.

C

6. Holmes illustrated the Genesis system's capabilities with a story about:
   A. Henry VIII
   B. Graduate advising
   C. International conflict
   D. A criminal and a policeman

A

7. Davis mentioned which companies that developed technology and let consumers figure out applications?
   A. Apple and Boston Dynamics
   B. iRobot and SharkNinja
   C. Android and Netflix
   D. Bank of America and GE

Tear-off sheet

We do not collect tear-off sheets, so please show your work on the quiz pages, not the tear-off sheet.

Problem 1 (Bayes)

The following (Boolean) variables are in the BayesNet:

- **A**: Irene **A**dler is acting suspiciously.
- **M**: Professor **M**oriarty is acting suspiciously.
- **W**: **W**atson has been kidnapped.
- **B**: Professor Moriarty is robbing a **b**ank.
- **N**: Watson left a **n**ote for Holmes on their dining table.
- **E**: The note on the dining table is **e**ncrypted.

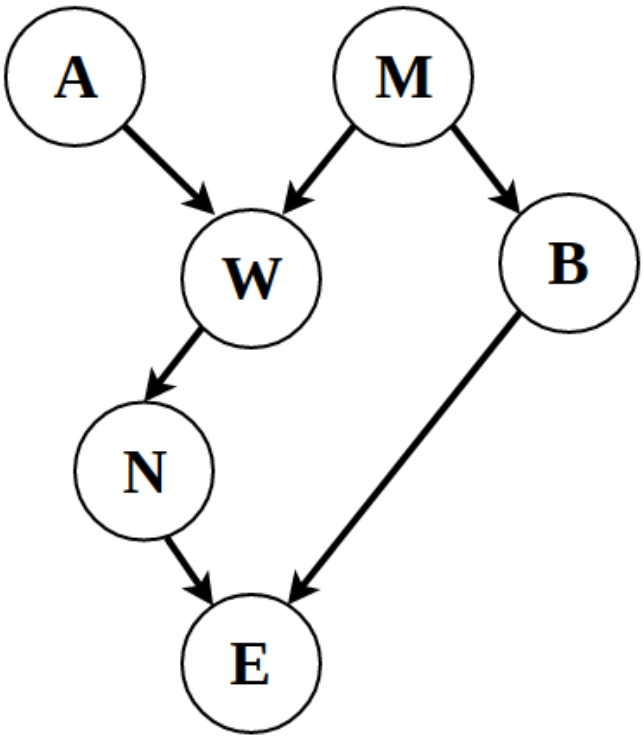

Tear-off sheet

We do not collect tear-off sheets, so please show your work on the quiz pages, not the tear-off sheet.

Problem 2 (Boosting)

	T	P	G	C	S	D
h1 (Ross)	-	+	-	-	-	+
h2 (Rachel)	+	+	-	+	+	-
h3 (Chandler)	-	+	+	-	+	-
h4 (Phoebe)	-	+	+	-	+	+
h5 (Joey)	+	+	-	-	-	+
Actual	+	+	+	-	+	-