Name:

Birth Date:

Matriculation Number:

Exam
Artificial Intelligence 2

July 20, 2021

	To be used for grading, do not write here									
prob.	1.1	1.2	2.1	2.2	3.1	3.2	4.1	4.2	Sum	grade
total	10	15	15	10	15	15	5	10	95	
reached										

Exam Grade: Bonus Points: Final Grade:

Organizational Information

Please read the following directions carefully and acknowledge them with your signature.

1. Please place your student ID card and a photo ID on the table for checking
2. The grading information on the cover sheet holds with the proviso of further checking.
3. no resources or tools are allowed except for a pen.
4. You have 90 min(sharp) for the test
5. You can reach 95 points if you fully solve all problems. You will only need 90 points for a perfect score, i.e. 5 points are bonus points.
6. Write the solutions directly on the sheets.
7. If you have to abort the exam for health reasons, your inability to sit the exam must be certified by an examination at the University Hospital. Please notify the exam proctors and have them give you the respective form.
8. Please make sure that your copy of the exam is complete (17 pages including cover sheet and organizational information pages) and has a clear print. **Do not forget to add your personal information on the cover sheet and to sign this declaration (next page).**

Declaration: With my signature I certify having received the full exam document and having read the organizational information above.

Erlangen, July 20, 2021 .....

(signature)

Organisatorisches

Bitte lesen die folgenden Anweisungen genau und bestätigen Sie diese mit Ihrer Unterschrift.

1. Bitte legen Sie Ihren Studentenausweis und einen Lichtbildausweis zur Personenkontrolle bereit!
2. Die angegebene Punkteverteilung gilt unter Vorbehalt.
3. Es sind keine Hilfsmittel erlaubt außer eines Stifts.
4. Die Lösung einer Aufgabe muss auf den vorgesehenen freien Raum auf dem Aufgabenblatt geschrieben werden; die Rückseite des Blatts kann mitverwendet werden. Wenn der Platz nicht ausreicht, können bei der Aufsicht zusätzliche Blätter angefordert werden.
5. Wenn Sie die Prüfung aus gesundheitlichen Gründen abbrechen müssen, so muss Ihre Prüfungsunfähigkeit durch eine Untersuchung in der Universitätsklinik nachgewiesen werden. Melden Sie sich in jedem Fall bei der Aufsicht und lassen Sie sich das entsprechende Formular aushändigen.
6. Die Bearbeitungszeit beträgt 90 min.
7. Sie können 95 Punkte erreichen, wenn Sie alle Aufgaben vollständig lösen. Allerdings zählen 90 Punkte bereits als volle Punktzahl, d.h. 5 Punkte sind Bonuspunkte.
8. Überprüfen Sie Ihr Exemplar der Klausur auf Vollständigkeit (17 Seiten inklusive Deckblatt und Hinweise) und einwandfreies Druckbild! **Vergessen Sie nicht, auf dem Deckblatt die Angaben zur Person einzutragen und diese Erklärung zu unterschreiben!**

Erklärung: Durch meine Unterschrift bestätige ich den Empfang der vollständigen Klausurunterlagen und die Kenntnisnahme der obigen Informationen.

Erlangen, July 20, 2021 .....

(Unterschrift)

1 Bayesian Reasoning

Note: When working with an upper case Boolean random variable X, you may abbreviate the event $X = \mathit{true}$ by the corresponding lower-case letter x. If you do that, make sure the distinction between upper and lower case letters is clear in your writing.

Problem 1.1 (Bayesian Rules) 10 pt

1. Assume that, in some area, the prevalence of SARS-CoV-2 infections is $1/10,000$. Moreover, assume that such an infection causes a cough half the time, and that in general on any given day 1 person out of $1,000$ is coughing. Apply Bayes' rule to determine the probability that someone who coughs is infected. 5 pt

2. Assume three random variables A, B, C such that A and B are conditionally independent given C. You know

   - the probability distribution of C,
   - the conditional probability distribution of A given C,
   - the conditional probability distribution of B given C.

   In terms of the above, give the formula for the probability distribution of C given the event $A = a, B = b$. 5 pt

This page was intentionally left blank for extra space

Problem 1.2 (Bayesian Networks) 15 pt

Consider the following situation:

- You have a rock in your yard, which can feel wet or not.
- Rain may cause humidity in the air.
- Any one of rain, humidity, and whether the lawn was sprinkled may cause the rock to be wet.
- There are no other causal relationships.

You want to model this situation using Boolean random variables W (wet rock), R (rain), H (humidity), and S (sprinkling).

2 pt

1. You do not know if it rained today and want to determine that by touching the rock. Which variables are the evidence, query, and hidden variables?

3 pt

2. Give a good variable ordering for forming a Bayesian network for this situation.

3 pt

3. Give the resulting network.

2 pt

4. Which of the variables are deterministic?

2 pt

5. Now assume your network is $W \to R \to H \to S$ (which *may or may not* be a correct solution to 3. above). How many entries do the conditional probability tables of that network have in total?

3 pt

6. Now assume a correct network is $W \to R \leftarrow H \leftarrow S$. Give the formula for

$$P(R|W = true, S = true)$$

in terms of the entries of the conditional probability table of that network.

This page was intentionally left blank for extra space

2 Markovian Reasoning

Problem 2.1 (Hidden Markov Models) 15 pt

Consider the following situation:

- You make weekly observations about your business with a client. Each week business is either good or bad.
- You know this is caused by the mood of your client, who feels either optimistic or pessimistic about the economy.
- You have previously obtained the following information:
  - when your client is optimistic, they remain optimistic next week 90% of the time,
  - when your client is optimistic, your business is good 70% of the time,
  - when your client is pessimistic, they remain pessimistic next week 25% of the time,
  - when your client is pessimistic, your business is bad 80% of the time,
  - your client was optimistic two weeks ago with probability 60%.

You want to model this situation as a hidden Markov model with Boolean random variables indexed by week number w.

2 pt

1. Give the state and evidence variables.

1 pt

2. Is the model stationary?

1 pt

3. What order does the model have?

4. Complete the following sentences:

3 pt

   (a) The transition model is given by the matrix

$$T = \begin{pmatrix} & \\ & \end{pmatrix} \quad \text{where} \quad T_{ij} = P(\quad = j| \quad = i).$$

3 pt

   (b) The sensor model is given by the matrix

$$M = \begin{pmatrix} & \\ & \end{pmatrix} \quad \text{where} \quad M_{ij} = P(\quad = j| \quad = i).$$

   To map Boolean values to matrix indices i, j, we use $true = 1$ and $false = 2$.

5 pt

5. Your business was good last week ($w = 1$) and bad this week ($w = 2$). Give the matrix form of the recursive filtering equation and state precisely which concrete values to plug in to obtain the probability distribution of your client's current mood at $w = 2$. (You do not have to actually compute the distribution.)

This page was intentionally left blank for extra space

Problem 2.2 (Markov Decision Processes) 10 pt

4 pt

1. Give an optimal policy π^* for the following MDP:
   - set of states: $S = \{0, 1, 2, 3, 4, 5\}$ with initial state 0
   - set of actions for $s \in S$: $A(s) = \{-1, 1\}$
   - transition model for $s, s' \in S$ and $a \in A(s)$: $P(s'|s, a)$ is such that
     - $s' = (s + a) \mod 6$ with probability 2/3,
     - $s' = (s + 3) \mod 6$ with probability 1/3.
   - reward function: $R(5) = 1$ and $R(s) = -0.1$ for $s \in S \setminus \{5\}$

2 pt

2. State the Bellman equation.

4 pt

3. Complete the following high-level description of the value iteration algorithm:
   - The algorithm keeps a table $U(s)$ for $s \in S$, that is initialized with
   - In each iteration, it uses the

     in order to
   - $U(s)$ will converge to the

This page was intentionally left blank for extra space

3 Learning

Problem 3.1 (Decision Trees) 15 pt

You observe the values below for 6 different football games of your favorite team. You want to construct a decision tree that predicts the result.

#	Day	Weather	Location	Opponent	Result
1	Monday	Rainy	Home	Weak	Win
2	Monday	Sunny	Home	Weak	Win
3	Friday	Rainy	Away	Strong	Loss
4	Sunday	Sunny	Home	Weak	Win
5	Friday	Cloudy	Home	Strong	Draw
6	Sunday	Sunny	Home	Strong	Draw

4 pt

1. Assume you choose attributes in the order
   $Opponent, Location, Weather, Day$.
   Give the resulting decision tree.

1 pt

2. How does the information-theoretic algorithm choose an attribute?

3 pt

3. Without using the above observations, give the formula for the information gain of the attribute *Opponent*.

3 pt

4. Using the above observations, give the results of

   - $I(P(Result)) =$
   - $P(Result = Loss | Opponent = Strong) =$

   You do not have to compute irrational logarithms.

2 pt

5. Give a minimal set A of attributes such that $A \succ Result$ holds for the above observations.

2 pt

6. Explain *why or why not* the determination $Day, Weather \succ Result$ holds for the above observations.

This page was intentionally left blank for extra space

Problem 3.2 (Neural Networks) 15 pt

Consider the neural network below where units $1, 2$ are inputs, unit 5 is output, weights are given by the labels on the edges, and units $3, 4, 5$ are perceptron units with activation function $T(x) = 1$ for $x > 0.5$ and $T(x) = 0$ otherwise.

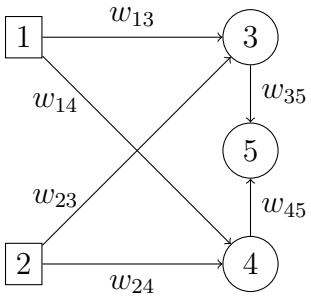

1 pt

1. Is it recurrent?

1 pt

2. How many hidden layers does it have?

3 pt

3. Give the formula for the activation a_3 of unit 3 in terms of the inputs a_1 and a_2 and the weights w_{ij}.

2 pt

4. Assume $w_{ij} = 1$ for all weights w_{ij} and $a_1 = a_2 = 1$. What is the resulting output a_5?

4 pt

5. Assume $a_1, a_2 \in \{0, 1\}$ and $w_{13} = w_{23} = 1$. Choose appropriate values for the other weights such that the network implements the XOR function, i.e., $a_5 = a_1 \, XOR \, a_2$.

4 pt

6. Complete the high-level description of the back-propagation algorithm *on the next page*.

To learn a target function $a_5 = f(a_1, a_2)$, do the following for each input to f:

- compute the

- determine the error between

  and propagate

- use the propagation results to update

4 Natural Language Processing

Problem 4.1 (Trigram Models) 5 pt

1 pt

1. How many trigrams does a language with 10 words have?

2 pt

2. Explain informally how we can obtain a trigram model for a language L.

2 pt

3. Name two applications of trigram models.

Problem 4.2 (Information Retrieval) 10 pt

Consider the following two texts

- d_1: "Information retrieval is hard."
- d_2: "Machine learning is very, very hard."

Let $D = \{d_1, d_2\}$.

Below we use alphabetical order for the vector components:

hard, information, is, learning, machine, retrieval, very

Simplify all results much as possible but without introducing approximate values.

3 pt

1. Give the vector $tf(\_, d_2)$.

3 pt

2. Give the vector $idf(\_, D)$.

2 pt

3. Let q be the query consisting of the word "retrieval". Give the value $tfidf(\text{retrieval}, q, D)$.

2 pt

4. How can we use $tfidf$ for choosing how to rank the texts in D for the query q?

This page was intentionally left blank for extra space

This page was intentionally left blank for extra space

This page was intentionally left blank for extra space