Name:

Birth Date:

Matriculation Number:

Exam
Artificial Intelligence 2

August 11., 2020

	To be used for grading, do not write here															
prob.	1.1	1.2	1.3	2.1	2.2	2.3	2.4	3.1	3.2	4.1	4.2	4.3	4.4	4.5	Sum	grade
total	6	10	2	6	4	6	9	12	3	10	2	4	8	3	85	
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 85 points if you fully solve all problems. You will only need 80 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 (24 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, August 11., 2020 .....
(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 85 Punkte erreichen, wenn Sie alle Aufgaben vollständig lösen. Allerdings zählen 80 Punkte bereits als volle Punktzahl, d.h. 5 Punkte sind Bonuspunkte.
8. Überprüfen Sie Ihr Exemplar der Klausur auf Vollständigkeit (24 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, August 11., 2020 .....

(Unterschrift)

Please consider the following rules; otherwise you may lose points:

- If you continue an answer on another page, please indicate the problem number on the new page and give a page reference on the old page.
- Always justify your statements (we would like to give poins for incorrect answers). Unless you are explicitly allowed to, do not just answer "yes", "no", or "42".
- If you write program code, give comments!

1 Bayesian Reasoning

Problem 1.1 (Basic Probability)
Let A, B, C be Boolean random variables. Which of the following equalities are always true? Justify each of your answers in one sentence. Each justification is worth half of the points for that equality. 6 pt

1. $P(b) = P(a, b) + P(\neg a, b)$
2. $P(a) = P(a|b) + P(a|\neg b)$
3. $P(a, b) = P(a) \cdot P(b)$
4. $P(a, b|c) \cdot P(c) = P(c, a|b) \cdot P(b)$
5. $P(a \vee b) = P(a) + P(b)$
6. $P(a, \neg b) = (1 - P(b|a)) \cdot P(a)$

Problem 1.2 (Bayesian Networks)

Consider the following Bayesian network with Boolean variables: 10 pt

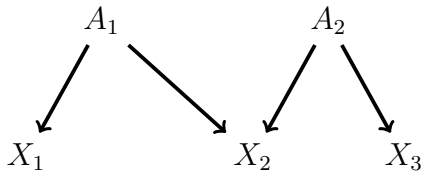

1. Give the definition of conditional independence. 2 pt
2. Which nodes in the network are conditionally independent with X_1 given A_1? Explain why. 3 pt
3. Give an example of two nodes from the network above that are stochastically independent. Explain why they are stochastically independent. 3 pt
4. What exactly (formal criterion) does an arrow between two nodes in a Bayesian network mean for the associated events? 2 pt

This page was intentionally left blank for extra space

Problem 1.3 (Arrows in Bayesian Networks)
Suppose that in a Bayesian network N we have variables C_1, E_1, C_2, and E_2, such that C_1 causes E_1 and C_2 causes E_2. 2 pt

1. How do we call an arrow going from C_1 to E_1?
2. How do we call an arrow going from E_2 to C_2?

2 Decision Theory

Problem 2.1 (Expected Utility)

6 pt

1. What is the formal(!) definition of *expected utility*? What is the meaning of every variable in the defining equation? 4 pt

2. How do we use expected utility to make decisions? 2 pt

Problem 2.2 (The Value of Information)

Chef Giordana runs a kitchen that provides food for a large organisation. A salad is sold for €6 and costs €4 to prepare. Therefore, the contribution per salad is €2. At present Giordana must decide in advance how many salads to prepare each day (40 or 60). Actual demand will also be 40 or 60 each day. So Giordana's payoff table looks as follows: 4 pt

Demand	Probability	40 salads	60 salads
40	0.4	€80	€0
60	0.6	€80	€120

Thus, the expected utility for making 40 salads is 80 and the expected utility for making 60 salads is 72. Based on these expected values without additional information, Giordana would choose to make 40 salads per day with an EU of €80 per day.

She is considering a new ordering system, where the customers must order their salad online the day before. With this new system Giordana will know for certain the daily demand 24 hours in advance. She can adjust production levels on a daily basis. How much is this system worth to her (per day)?

Task: Compute the concrete value in € and explain what you did.

This page was intentionally left blank for extra space

Problem 2.3 (Decision Network)

You try to decide on whether to take an umbrella to Uni. Obviously, it's useful to do so if it rains when you go back home, but it's annoying to carry around if it doesn't even rain. Here are the states you could end up in: 6 pt

- happy (or relieved) if it *doesn't rain* and you *did not bring an umbrella*,
- annoyed if it *doesn't rain* and you *brought an umbrella*,
- wet if it *rains* and you *did not bring an umbrella*,
- dry if it *rains* and you *brought an umbrella*.

You look at the weather forecast, which has two possible values: sunny and rainy.

You come up with this decision network:

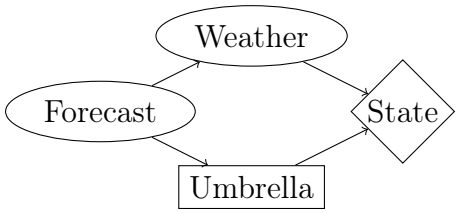

2 pt

1. Decision networks are extensions of Bayesian networks, which additional kinds of nodes do decision networks have? For each kind give an example from the network above. 4 pt
2. How would you compute whether or not to take an umbrella, assuming you know all of the probabilities $P(\text{state} = s|\text{forecast} = f, \text{umbrella} = u)$ for all
   - $s \in \{\text{happy}, \text{annoyed}, \text{wet}, \text{dry}\}$,
   - $u \in \texttt{Bool}$, and
   - $f \in \{\text{sunny}, \text{rainy}\}$.

This page was intentionally left blank for extra space

Problem 2.4 (Markov Decision Procedures)

9 pt

1. How do Markov decision procedures differ from (simple) decision networks? 3 pt
2. What do we use the value iteration and policy iteration for? How do they differ? 4 pt
3. What is the difference between *partially observable Markov decision procedures* and normal MDPs? 2 pt

This page was intentionally left blank for extra space

3 Markov Models

Problem 3.1 (Stock Market Predictions)

You bought SpaceY stock recently and try to predict whether to buy more or sell. The stock market is in one of two possible states; bull state or bear state. In a bull state, it will (in the long term) be advantageous to buy stock; in a bear state it will be more advantageous to sell. 12 pt

If the market is in a bull state, the probability it will still be in a bull state tomorrow is 60%. If it is in a bear state, the probability it will remain so tomorrow is 80%.

If the market is in a bull state, the probability that your stock will rise that day is 90%. If it is in a bear state, your stock will more likely fall (with 60% probability).

1. What are the observable and unobservable variables in this model? 1 pt
2. If we consider this as a hidden Markov model, what is its transition matrix T? Remember that we use transition matrices to compute the previous or future states. 6 pt
3. Explain what kind of probabilities *prediction*, *filtering* and *smoothing* compute in this scenario. Do not just give formulas. 4 pt
4. Give the underlying equations for the first two of these algorithms and explain what each variable in the equation represents.

This page was intentionally left blank for extra space

Problem 3.2 (Stationary)
Define what it means for a Markov model to be *stationary*, and why we are interested in stationarity. 3 pt

4 Learning

Problem 4.1 (Decision List)

Construct a decision list to classify the data below. The tests should be as small as possible (in terms of attributes), breaking ties among tests with the same number of attributes by selecting the one that classifies the greatest number of examples correctly. If multiple tests have the same number of attributes and classify the same number of examples, then break the tie using attributes with lower index numbers (e.g., select A_1 over A_2). 10 pt

Example	A_1	A_2	A_3	A_4	y
x_1	1	0	0	0	1
x_2	1	0	1	1	1
x_3	0	1	0	0	1
x_4	0	1	1	0	0

This page was intentionally left blank for extra space

Problem 4.2 (Information Theory)

Explain why it is possible (even common) that the learning curve (example given below) never gets to 100% correctness, even for large example sets. 2 pt

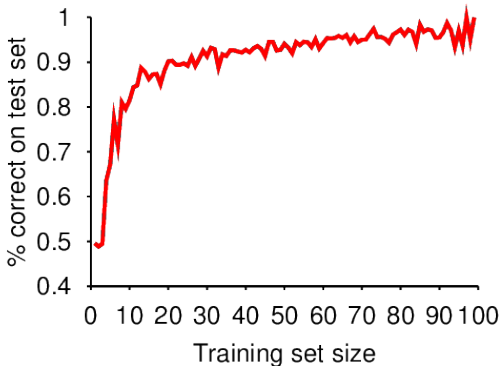

Problem 4.3 (Information Entropy)
Explain and define *information entropy.* 4 pt

Problem 4.4 (Sunbathing)

Eight people go sunbathing. Some of them got a sunburn, others didn't: 8 pt

Name	Hair	Lotion	Result
Sarah	Light	No	Sunburned
Dana	Light	Yes	None
Alex	Dark	Yes	None
Annie	Light	No	Sunburned
Julie	Light	No	None
Pete	Dark	No	None
John	Dark	No	None
Ruth	Light	No	None

2 pt

1. Which quantity does the decision tree learning algorithm use to pick the attribute to split on? Write down the formula for it.

   4 pt

2. Compute it for for the attributes Hair and Lotion. It is enough to give the formula and insert the correct values for the variables, you do not need to compute the final value.

   2 pt

3. Which one would the algorithm pick for the next step? Explain what happens next.

Note that *Name* is only an index, not a (meaningful) attribute!

This page was intentionally left blank for extra space

Problem 4.5 (Overfitting)
Explain what *overfitting* means and why we want to avoid it. 3 pt

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