Name:

Birth Date:

Matriculation Number:

Exam
Artificial Intelligence 1

July 17., 2018

	To be used for grading, do not write here												
prob.	1.1	1.2	2.1	2.2	2.3	3.1	3.2	3.3	4.1	5.1	5.2	Sum	grade
total	4	12	4	3	12	4	10	12	10	10	8	89	
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 89 points if you fully solve all problems. You will only need 125 points for a perfect score, i.e. -36 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 17., 2018 . . . .
(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 89 Punkte erreichen, wenn Sie alle Aufgaben vollständig lösen. Allerdings zählen 125 Punkte bereits als volle Punktzahl, d.h. -36 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 17., 2018 ....

(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 Prolog

Problem 1.1 (The Zip Function)

The zip function takes two lists with lengths that differ at most by 1, and outputs a list of lists containing one element from the first list and the element with the same index from the other list, possibly followed by a one-element list with the left-over argument. 4 pt

Create a `ProLog` predicate with 3 arguments: the first two would be the two lists you want to zip, and the third one would be the result. For instance:

```
?- zip([1,2,3],[4,5,6],L).       ?- zip([1,2],[3,4,5],L).
L = [[1, 4], [2, 5], [3, 6]].    L = [[1, 3], [2, 4], [5]].
```

Feel free to implement any helper functions.

Hint: Remember that you can pattern match a list L as [HEAD|TAIL].

Problem 1.2 (DFS in Prolog)

We want to implement DFS in ProLog using the following data structures for search trees: 12 pt

```
subtrees([]).
subtrees([(Cost,T)|Rest]) :- number(Cost),istree(T), subtrees(Rest).
istree(tree(_,Children)) :- subtrees(Children).
```

Write a Prolog predicate dfs such that dfs(G,T,X,Y) on a tree T returns the path to the goal G in X and the cost of the path in Y

2 Bayesian Reasoning

Problem 2.1 (Bayesian Rules)
Name four of the basic rules in Bayesian inference and explain each with a short sentence and formula. 4 pt

Problem 2.2 (Conditional Independence)

Define *conditional independence.* 3 pt

Problem 2.3 (Medical Bayesian Network 2)

Both Malaria and Meningitis can cause a fever, which can be measured by checking for a high body temperature. Of course you may also have a high body temperature for other reasons. We consider the following random variables for a given patient: 12 pt

- Mal: The patient has malaria.
- Men: The patient has meningitis.
- HBT: The patient has a high body temperature.
- F: The patient has a fever.

1. Draw the corresponding Bayesian network for the above data using the algorithm presented in the lecture, assuming the variable order Mal, Men, HBT, F. Explain rigorously(!) the exact criterion for whether to insert an arrow between two nodes.
2. Which arrows are causal and which are diagnostic? Which order of variables would be better suited for constructing the network?
3. How do we compute the probability the patient has malaria, given that he has a fever? State the query variables, hidden variables and evidence and write down the equation for the probability we are interested in.

This page was intentionally left blank for extra space

3 Decision Theory

Problem 3.1 (Expected Utility)
What is the formal(!) definition of *expected utility*? Explain every variable in the defining equation. 4 pt

Problem 3.2 (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. You look at the weather forecast, which hast three possible values: sunny, cloudy and rainy. 10 pt

1. Draw the decision network for bringing/leaving an umbrella depending on whether it does or doesn't rain later.

2. Explain *formally* how to compute whether or not to take an umbrella, assuming you know $P(rain = b|forecast = x)$ for all $b \in \texttt{Bool}, x \in \{\texttt{sunny}, \texttt{cloudy}, \texttt{rainy}\}$.

Problem 3.3 (Markov Decision Procedures)

12 pt

1. How do Markov decision procedures differ from (simple) decision networks?
2. How does the value iteration algorithm work? (Give an actual equation and explain its role in the algorithm)
3. What is the disadvantage of value iteration that is “fixed” by policy iteration?
4. How can we reduce *partially observable Markov decision procedures* to normal MDPs?

This page was intentionally left blank for extra space

4 Markov Models

Problem 4.1 (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. 10 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. Explain what kind of probabilities *prediction*, *filtering* and *smoothing* compute in this scenario.

2. 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

5 Learning

Problem 5.1 (Home Decisions)

Eight people go sunbathing. Some of them got a sunburn, others didn't: 10 pt

Name	Hair	Height	Weight	Lotion	Result
Sarah	Blonde	Average	Light	No	Sunburned
Dana	Blonde	Tall	Average	Yes	None
Alex	Brown	Short	Average	Yes	None
Annie	Blonde	Short	Average	No	Sunburned
Julie	Blonde	Average	Light	No	None
Pete	Brown	Tall	Heavy	No	None
John	Brown	Average	Heavy	No	None
Ruth	Blonde	Average	Light	No	None

Explain how the information-theoretic decision tree learning algorithm would proceed on this table (up to two iterations). Explicitly state how to compute the information gain (and what that means).

Note that you do not need to compute any actual values; if it is helpful for your explanation, you may guess any values you might want to use.

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

Problem 5.2 (Backpropagation)

Explain what *Backpropagation* means in the context of Neural Networks, when and why we need it, and how to do it using an example. 8 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