SOLUTIONS IN BLUE

Grading instructions in Red

i Frontpage

IN4080 Natural Language Processing

Fall 2021

Monday, December 13
15:00 AM - 19:00 PM (4 hours)

All questions should be answered!
Each question is assigned a weight which is indicated.
The maximum number of points for the whole set is 100 points.

Permitted materials: None
An on-screen calculator is available.

You may answer in English, Norwegian, Danish or Swedish.

1(a) Tokenization

One of the first steps in text processing is (word) tokenization. What is tokenization? With the sentence (*) from the Brown corpus as example, discuss the decisions a tokenizer have to take, and where different tokenizers may differ.

(*) "You shouldn't smoke so much " , he said , unconsciously imitating Victoria's holier-than-thou voice.

Fill in your answer here

Maximum marks: 5

In raw form a text is a sequence of characters, or even bytes. Tokenization splits it into chunks corresponding to words or, more generally, tokens. For languages like English, this can to a large degree be done by splitting on white space. (2 pts.)

Special care must be taken:

1. Punctuation signs are often not separated by white space in the raw text, e.g. "*You* and *voice*.
   - The most common is to split the punctuation sign off as a separate token: | " | You |
   - While some proposals keep it as part of the word
   - And some remove it.
2. How to handle contractions like ' shouldn't '. This is normally split into two tokens | should | n't |, where the representation of the last part may vary; some may normalize to | not |.
   ' Victoria's ' is similar, typically tokenized as | Victoria | 's |, while there are also approaches (e.g. the Brown corpus) which consider it as one token, a possevie form of *Vicoria*.
3. Whether to keep or split hyphenated expressions like ' holier-than-thou '

(1 pt for each)

1(b) Lemmatization

b) What is a lexeme? What is meant by lemmatization in NLP? What will (*) look like after lemmatization?

(*) "You shouldn't smoke so much " , he said , unconsciously imitating Victoria's holier-than-thou voice.

Fill in your answer here

Maximum marks: 10

A lexeme is an abstract unit of morphological analysis in linguistics, that roughly corresponds to a set of forms taken by a single word, e.g., the verb: {say, says, said, saying}. (3 pts.)

A lemma (plural lemmas or lemmata) is the canonical form, dictionary form, or citation form of a lexeme, for the example: 'say'. Lemmatization consists in replacing each token with the corresponding lemma. (3 pts.)

Token	Lemma
“	“
You	you
should	shall
n't	not
smoke	smoke
so	so
much	much
“	“
,	,
he	he
said	say
,	,
unconsciously	unconsciously/or unconsious
imitating	imitate
Victoria	Victoria
‘s	‘s
holier	holy
-	-
than	than
-	-
thou	thou
voice	voice
.	.

Alternatively:

- *Victoria’s* as one token with *Victoria* as the lemma.
- ‘holier-than-thou’ as one token.

(4 pts.)

2 Semantic relations

a) Which sematic relations are there between the words in each of the following pairs?

a. Horse – animal
b. Horse – cow
c. Horse – saddle
d. Big – small
e. Big – large

Fill in your answer here

Maximum marks: 5

a) Horse is hyponym to animal, animal is hypernym to horse
b) Similar
c) Related
d) Antonyms
e) Synonyms

(1 pt for each)

(a) BIO-format

We have sent two paragraphs from a longer text through the Stanford corenlp for Named Entity Recognition. The figure shows the result. (The break between sentence 3 and 4 is due to a typo in the source text we did not correct.)

A common format for representing NE-marked text is to use the so-called BIO-tags. Please present sentence 2 (“It was quickly … Friday evening.”) using the BIO-tag format and explain how the representation should be understood.

Fill in your answer here

Maximum marks: 10

Named Entity Recognition:

It	O
was	O
quickly	O
accepted	O
by	O
both	O
the	O
Progress	B-ORG
Party	I-ORG
and	O
the	O
Christian	B-IDE
Democrats	I-IDE
Friday	B-DATE
Evening	B-TIME
.	O

A complete solution should include the last line.
(Up to 5 pts.)

The tags starting with a B indicates that this token is the beginning of a named entity of the type indicated, e.g. *Progress: B-ORG* indicates the beginning of an ORG entity. The I-ORG tag indicates that this word is part of an ORG entity, but not its beginning, while O indicates that a word is not part of any entities. This mark-up indicates exactly which tokens constitute a NE.
(Up to 5 pts.)

(b) Evaluation

We have also run the same paragraph through Spacy. (For fairness sake, we should explain that we have used a small Spacy model; Spacy can do better with larger models.)

To compare the two, we will consider the Stanford corenlp result as the correct markup and the Spacy result as predicted outcome. Please calculate precision and recall for the various entity types. Explain how you find the numbers.

Also calculate the macro- and micro-averaged precision and recall across the entities.

For comparison, you may assume the following equivalences:

Corenlp	Spacy
PERSON	PERSON
ORGANIZATION	ORG
IDEOLOGY	NORP
DATE	DATE
TIME	TIME
MISC	MISC
NUMBER	CARDINAL
GPE	GPE
CITY	LOC

You can disregard the TITLE and P1Y entities, as we assume that Spacy does not aim at recognizing them.

You may ignore the additional detailed information in DATE, TIME, NUMBER.

If any expression ends up as 0/0, consider it undefined and exclude it from the (macro) averages.

Fill in your answer here

Maximum marks: 10

Solberg PERSON presented a long list of concrete projects set to benefit from the proposal that she said was formulated by her Conservative Party ORG .

It was quickly accepted by both the Progress Party ORG and the Christian Democrats NORP

Friday DATE evening TIME .

The Liberal Party ORG .

which has long promoted road tolls as a means of discouraging driving and raising money for public transport, remained a hold-out until finally going along as well just before midnight.

The Liberals and Progress parties had been most at odds among the four CARDINAL , and with both of them diving in public opinion polls lately, it was important that each could claim victory.

Most serious conflict yet

The road toll conflict has been among the most serious to face Solberg PERSON since taking over as prime minister of a minority conservative government in 2013 DATE , and threatened to topple her government.

Higher road tolls and far more extensive toll systems set up in cities like Stavanger GPE , Bergen LOC and Oslo GPE left motorists faced with paying tens of thousands CARDINAL of more kroner per year, so it became a pocketbook issue that nurtured the rise of protest parties that have attracted large numbers of voters away from the established parties.

We have chosen a strict interpretation where two entities have to contain the exact same tokens to be correct. For example
Spacy: Organization: the progressive party
Corenlp: ORG: progressive party
is considered incorrect because of the lacking 'the', and is hence counted as both a false positive and a false negative ORG.

We then get the following results where each line evaluates one label, and where
Tp – true positives
Fp – false positives
Fn – false negatives

Precision = tp/(tp+fp)
Recall = tp/(tp+fn)

Corenlp	Corenlp	tp	fp	fn	Precision	Recall
PERSON	PERSON	2	0	0	1	1
ORGANIZATION	ORG	1	2	2	1/3	1/3
IDEOLOGY	NORP	1	0	1	1	0.5
DATE	DATE	2	0	0	1	1
TIME	TIME	1	0	1	1	0.5
MISC	MISC	0	0	1	nan	0
NUMBER	CARDINAL	1	1	0	0.5	1
GPE	GPE	0	2	0	0	nan
CITY	LOC	1	0	2	1	1/3
Pooled/micro		9	5	7	9/14=0.64	9/16=0.56
Macro					(35/6)/8 = 35/48=0.73	(14/3)/8 = 7/12=0.58

To get the micro averages, we add the numbers in each column and calculate P and R from these sums.
For macro, we take the mean of the numbers in the respective columns.

Up to
3 pts for a correct understanding of the task
4 pts for the numbers for the various classes
2 pts for micro average
1 pt for macro average

A slightly different understanding of how to compare the chunks across the two models can still give a good score on the table and the micro and macro scores, given it is carried out consistently. If

- Spacy: Organization: the progressive party
- Corenlp: ORG: progressive party

is considered as a correct predicition, and similarly for all such examples, and this is the only mistake, subtract one point.

If the student evaluates the tags instead of the chunks, the whole point can give at most 5 pts altogether, including at most 1 pt. for micro average and 1 pt. for macro average.

4(a) Softmax

A central formula in classification is based on the softmax function:

$$P\left(C_j \middle| \vec{x}\right) = \frac{e^{\vec{w_j} \cdot \vec{x}}}{\sum_{i=1}^{k} e^{\vec{w_i} \cdot \vec{x}}}$$

Explain the formula, in particular

1. What does $P(C_i \mid \vec{x})$ express in a classification task?
2. What is C_i?
3. What is $\vec{w}_i$?
4. What is $\vec{x}$?
5. What is k?

(Observe that $\vec{x}$ can also be written as bold face $\mathbf{x}$ and similarly with $\vec{w}_i$.)

Fill in your answer here

Maximum marks: 10

- $P(C_i|\boldsymbol{x})$ is the probability of the class C_i For the observation represented by the vector $\boldsymbol{x}$. (4pts.)
- C_i is one class indexed by *i*. (1 pt.)
- $\boldsymbol{x}$ is a vector representing one observation (1 pt.)
- $\boldsymbol{w}_i$ is a vector of weights associated with the class *i*. It expresses how the various features in an observation $\boldsymbol{x}$ should be weighted for the class C_i. (3 pts.)
- k is the number of classes (1 pt.)

For $\boldsymbol{w}_i$ it is essential to express how this corresponds to the class *i*, that there are different weights for the different classes.

4(b) Recurrent Neural Nets

What is a recurrent neural net (RNN)? Sketch how a RNN can be used for POS-tagging. In particular, explain the softmax formula's place in the model.

Fill in your answer here

Maximum marks: 10

"A recurrent neural network (RNN) is any network that contains a cycle within its network connections, meaning that the value of some unit is directly, or indirectly, dependent on its own earlier outputs as an input." (Jurafsky og Martin). (3 pts.)

In its simplest form, an RNN consists of an input layer, one hidden layer, and an output layer. There are connections

- from the input layer to the hidden layer
- from the hidden layer to the hidden layer
- from the hidden layer to the ouput layer

The input consists of a sequence **x_1, x_2, …, x_n** which is fed to the network one input a time. The hidden layer at time *j* is calculated from the input at time *j* and the hidden layer a time *j-1*. The output at time *j* is calculated form the hidden layer at time *j*. This can be expressed by

- $\boldsymbol{h}_t = g(U\boldsymbol{h}_{t-1} + W\boldsymbol{x}_t)$
- $\boldsymbol{y}_t = f(V\boldsymbol{h}_t)$
- g and f are activation functions
- (There are also bias which we didn't include in the formulas)

(from lecture slides)

(3 pts.)

This can be applied to POS tagging as follows. Each input vector x_j represents a token, either one-hot-encoded or an embedding. Each output y_j is the corresponding tag. The activation function in the output layer, *f*, is the softmax yielding a probability distribution over POS-tags for the token given the input token and the preceding sequence. (4 pts.)

5(a) **Core dialogue concepts**

Here is a short example of dialogue:
Person 1: *did you manage to finish the third obligatory assignment due yesterday?*
Person 2: *I sent it last week already!*
Person 1: *Impressive!*
Could you send me your solutions?
Person 2: *Sure, I'll send you the PDF*

Answer the following questions based on this dialogue:
1) What are the speech acts (according to Searle's taxonomy) associated with each of those 5 utterances? *(2 points)*
2) Do you observe some conversational implicatures? Briefly explain. *(2 points)*
3) Does the common ground of those two persons evolve in the course of this short interaction? Explain in 2-3 sentences. *(2 points)*
Fill in your answer here

Maximum marks: 6

Question 1)

did you manage to finish the third obligatory assignment due yesterday? --> directive
I sent it last week already! --> assertive
Impressive! --> expressive
Could you send me your solutions? --> directive
Sure, I'll send you the PDF --> commissive

Grading: 0.4 points per correct answer (5 correct -> 2 points)

Question 2)

Yes, the second utterance ("I sent it last week already") is a conversational implicature. The utterance does not directly answer the preceding question. However, if we assume person 2 is cooperative and adheres to the maxim of relation ("be relevant"), we can nevertheless make sense of the utterance by understanding that, if person 2 has sent the obligatory assignment, it means the assignment was finished.

Grading: students do not have to give a very detailed explanation here, suffice it to say that it is a conversational implicature and that person 2 did not respond directly to the request from person 1.

Question 3)

Yes, the common ground evolves during the interaction. At the start of the dialogue, the common ground does not include the fact that person 2 has already submitted the assignment, but this element is added after person 2's statement. At the end of the dialogue, the common ground also contains the fact the promise to send the PDF.

Grading: 1 point for saying that common ground changes during the conversation, and 1 point for the explanation.

5(b) Chatbot design

Assume you wish to develop a chatbot to answer questions about the current time in various locations around the world. The chatbot should for instance be able to respond to queries such as "*What is the current time in Buenos Aires?*", "*What is the time difference between Boston and Oslo?*" or "*If it is 6:00 AM in Oslo, what time is it in Tokyo?*".

Your first task is to decide what kind of chatbot development strategy you wish to follow. We have covered four alternative approaches during the course: handcrafted chatbots, IR-based chatbots, sequence-to-sequence chatbots and NLU-based chatbots.

Answer the questions below:

1) Which type of approach (among the four approaches above) do you think is most suitable for this task? Motivate your choice in one or two paragraphs. *(5 points)*

2) Which data will you need to collect to train/develop your chatbot, based on the approach chosen above? What type of annotations would you need to add to this data, if any? Explain in a few sentences. *(2 points)*

3) Which system modules (machine learning models or rule-based components) would you need to integrate in your chatbot? Describe in one or two paragraphs the general processing pipeline of your chatbot. *(4 points)*

4) How would you evaluate the performance of the resulting chatbot? Describe in one or two paragraphs the evaluation procedure you would follow. *(3 points)*

Fill in your answer here

Maximum marks: 14

Question 1)

There are probably multiple possible answers, but the one that seems the most appropriate for me is to use an NLU-based chatbot that combines intent recognition with slot-filling (hours, places etc.). It's also possible to use a handcrafted system, although one would then need rules to detect the user intent and the possible slots.

An IR-based chatbot would be wrong here. There is no way a chatbot could derive the right answer for an utterance such as "what is the current time in Buenos Aires" from a fixed dialogue corpus. As for sequence-to-sequence, they would suffer from the same limitations as IR-chatbots (it would be technically possible to implement such a chatbot by implementing a "knowledge-grounded" seq2seq model which relies on a knowledge base to provide answers, but it's not something we have covered during the course).

Grading:

- The "worst" answer is IR-based chatbots, which are not suitable at all for such systems (0 points)

- Seq2seq models can give half of the points (2.5) provided that the student also mentions that the model must use an external (dynamically updated) source of knowledge to fill the answer with facts (such as the current time in Buenos Aires right now).

- A rule-based system is OK as long as the student mentions that the rules must recognize both the intent and the different slot values (if not, only 2 points)

- An NLU-based system is probably best. But the student must also mention both intent recognition and slot filling (-2 if they do not)

Question 2):

For IR-based and sequence-to-sequence chatbots, the data would correspond to a dialogue corpus. For rule-based systems, there is no absolute need for a dataset, although it is often useful to get inspiration to write rules that cover as many possible utterances as possible. For an NLU-based chatbot, the dataset would correspond to a labelled dataset where each utterance is associated with an intent. In addition, the utterances should be labelled with entities corresponding to the slots (for instance, "Buenos Aires" should be a location), at least if we are not reusing an existing NER model.

Grading: 2 points if the explanation is correct in relation to the approach chosen, otherwise 0. Also useful if the students also mention that the system must have access to some source of knowledge to find time zones etc, but it is not so important here.

Question 3):

At least three components would be necessary:

- A component (either rule-based or using a data-driven classification model) that detects the general user intent
- A component (based on handcrafted patterns, dictionaries, NER models or similar sequence labelling scheme) for detecting slots such as times and places
- A component for selecting/generating the response based on the intent, recognized slots, and an external knowledge source (for timezones etc.)

The two first components are typically part of what we call NLU, while the last one corresponds to response selection / NLG.

Alternatively, if the student has decided to go for a seq2seq model, we would need an external knowledge graph that can be "attended to" by the decoder when generating the response.

Grading: 1.5 points for each of these 3 components mentioned, and 0.5 for the description of how all the components are connected.

Question 4)

For the NLU-based model (with slot-filling), we can first evaluate the performance of the intent recognition and slot-filling using standard metrics (accuracy, precision, recall, F-score), since we would have a labelled dataset at our disposal. This could also be done for handcrafted systems, provided we have collected a labelled corpus that can be employed for such an evaluation.

But that is not all. We also need to evaluate the quality of the system responses. For this, a human evaluation would be preferable, where human annotators provide scores on the quality of the responses, often among several axes. There are also some automated dialogue metrics, although they would be relatively difficult to apply in our case. Since we already know which answer is the correct one (there can be only one correct time in Buenos Aires at a given moment), one could also design an automated metric tailored to this task, and that checks whether the system response contains the correct time.

Grading: 1.5 points if the student mentions the evaluation of components such as intent recognition and slot-filling, and 1.5 points if they also mention a human evaluation of the entire system.

5(c) Speech recognition

Answer the following questions:

1) Is is possible to have a Word Error Rate (WER) that is larger than 100%? Explain. *(2 points)*

2) During the lecture on speech processing, we mentioned that one challenge faced when training speech recognition models was the lack of explicit alignments between the speech inputs and the output transcriptions. What did we mean by that? Explain in 2-3 sentences. *(2 points)*

Fill in your answer here

Maximum marks: 4

Question 1)

Yes, it's possible to have a WER larger than 100%, for instance if there are many insertions (the edit distance between the recognition hypothesis and the gold standard utterance may be larger than the number of words in the gold standard).

Grading: 1 point for the correct answer, 1 point for the explanation

Question 2)

When training speech recognition models, we are typically provided with speech recordings together with their gold standard transcriptions. But there is no 1:1 correspondence between the speech inputs (audio frames of short duration, such as 50 ms) and the transcribed words or phonemes, which are substantially fewer. We don't initially know to which part of a phoneme/letter belongs a given audio frame. So we typically need to *infer* this alignment when training a speech recognition model (there are many ways to implement this kind of inference, but that is beyond the question).

Grading: 1 point if students correctly explain what input and output in speech recognition are, and 1 point if it mentions that there is no 1: 1 comparison between the audio stream and the words in the transcription, so we have to infer a latent mapping.

5(d) Reinforcement learning

Assume you wish to develop a talking robot that can respond to various requests. When a human is perceived around the robot, the first first step is to engage the human, for instance with a greeting, accompanied or not by gestures. We wish to apply reinforcement learning to determine the best engagement strategy among two options: *SayHi* and *SayHiWithGestures*.

Formally speaking, we can frame this problem as an MDP with

- 2 states: *HumanNotEngaged* (which is the starting state) and *HumanEngaged* (which is a final state).
- 3 actions: *SayHi*, *SayHiWithGestures*, and *AskHowCanIHelpYou*. The first two actions (*SayHi* and *SayHiWithGestures*) are only available for the starting state *HumanNotEngaged*, while the action *AskHowCanIHelpYou* is only available for the final state *HumanEngaged*.

The transition model $P(s'|s,a)$ for this MDP is as follows:

- If the robot executes action *SayHi* in the state *HumanNotEngaged*, we have a probability 0.5 of reaching the state *HumanEngaged*, and a probability 0.5 of staying in the state *HumanNotEngaged*.
- If the robot executes action *SayHiWithGestures* in the state *HumanNotEngaged*, we have a probability 0.7 of reaching the state *HumanEngaged*, and a probability 0.3 of staying in the state *HumanNotEngaged*.

The reward function $R(s,a)$ of this MDP is as follows:

- The reward of executing action *SayHi* in the state *HumanNotEngaged* is -1
- The reward of executing action *SayHiWithGestures* in the state *HumanNotEngaged* is -2, since the physical gestures "cost" more to the agent in terms of energy and mechanical wear.
- Finally, the reward of executing action *AskHowCanIHelpYou* in the state *HumanEngaged* is +10. You can consider this action as the final one in this MDP, which means that the expected cumulative reward Q(*HumanEngaged, AskHowCanIHelpYou)* = 10.

Questions:

1) Compute the expected cumulative rewards Q(*HumanNotEngaged*, *SayHi*) and Q(*HumanNotEngaged*, *SayHiWithGestures*) based on the MDP described above. To compute those Q-values, you need to use Bellman's equation and refine your estimates through several iterations. You can stop after 5 iterations and use a discount factor γ=0.9. For the first iteration, you can initialize the Q-values to zero. *(10 points)*

2) Based on those Q-values, which engagement strategy should the robot chose? *(2 points)*

Tip: Bellman's equation is:

$$Q(s,a) = R(s,a) + \gamma \sum_{s' \in S} P(s'|s,a)\ max_{a'} Q(s',a')$$

Fill in your answer here

Maximum marks: 12

Question 1)

We need to compute two Q-values (indicating the cumulative expected reward of an action in a particular state):

$Q(HumanNotEngaged, SayHi)$

$$= R(HumanNotEngaged, SayHi) + 0.9\left(0.5 * max_{a_1} Q(HumanNotEngaged, a_1) + 0.5 * max_{a_2} Q(HumanEngaged, a_2)\right)$$

We know that only one action is possible in the *HumanEngaged* state, with a Q-value of 10, so we can simplify the formula as:

$$= -1 + 0.9\left(0.5 * max_{a_1} Q(HumanNotEngaged, a_1) + 0.5 * 10\right)$$

$$= 3.5 + 0.45 * \ max_{a_1} Q(HumanNotEngaged, a_1)$$

We can write the same formula for the other action, *SayHiWithGestures*:

$$Q(HumanNotEngaged, SayHiWithGestures)$$

$$= R(HumanNotEngaged, SayHiWithGestures) + 0.9\left(0.3 * max_{a_1} Q(HumanNotEngaged, a_1) + 0.7 * max_{a_2} Q(HumanEngaged, a_2)\right)$$

$$= -2 + 0.9\left(0.3 * max_{a_1} Q(HumanNotEngaged, a_1) + 0.7 * 10\right)$$

$$= 4.3 + 0.27 * \ max_{a_1} Q(HumanNotEngaged, a_1)$$

So, for the first iteration, starting with Q-values initialized to 0, we get:

$$Q_{t_1}(HumanNotEngaged, SayHi) = 3.5 + 0.45 * 0 = 3.5$$

$$Q_{t_1}(HumanNotEngaged, SayHiWithGestures) = 4.3 + 0.27 * 0 = 4.3$$

For the second iteration, we get:

$$Q_{t_2}(HumanNotEngaged, SayHi) = 3.5 + 0.45 * \max([3.5, 4.3])$$

$$= \ 3.5 + 0.45 * 4.3 = 5.44$$

$$Q_{t_2}(HumanNotEngaged, SayHiWithGestures) = 4.3 + 0.27 * 4.3 = 5.46$$

We continue with the second, third, fourth and fifth iteration:

$Q_{t_3}(HumanNotEngaged, SayHi) = 3.5 + 0.45 * 5.46 = 5.96$

$Q_{t_3}(HumanNotEngaged, SayHiWithGestures) = 4.3 + 0.27 * 5.46 = 5.77$

$Q_{t_4}(HumanNotEngaged, SayHi) = 3.5 + 0.45 * 5.96 = 6.18$

$Q_{t_4}(HumanNotEngaged, SayHiWithGestures) = 4.3 + 0.27 * 5.96 = 5.91$

$Q_{t_5}(HumanNotEngaged, SayHi) = 3.5 + 0.45 * 6.18 = 6.28$

$Q_{t_5}(HumanNotEngaged, SayHiWithGestures) = 4.3 + 0.27 * 6.18 = 5.97$

So the correct answer at the end of the fifth iteration would be 6.28 for Q(HumanNotEngaged, SayHi) and 5.97 for Q(HumanNotEngaged, SayHiWithGestures).

Grading: 4 points if the student is able to use Bellman to express the Q-formula for the two possible actions, 4 points if they are able to refine the values through 5 iterations, and 2 points if they arrive at the correct answer.

Comment: Some students tried to calculate the true expected cumulative reward. That can be done as follows:

To simplify notation set $x = Q(HumanNotEngaged, SayHi)$ and $y = Q(HumanNotEngaged, SayHiWithGestures)$

We see that $max_{a_1} Q(HumanNotEngaged, a_1) = \max(\{x, y\})$

Suppose $x > y$:

Then $x = 3.5 + 0.45\,x$, hence $0.55x = 3.5$ and $x = \frac{3.5}{0.55} = 6.36$

and $y = 4.3 + 0.27x = 4.3 + 0.27 * 6.36 = 6.02$

(If we instead assume $y > x$, we get a result where $x > y$, hence a contradiction)

So the expected cumulative rewards will be 6.36 for Q(HumanNotEngaged, SayHi) and 6.02 for Q(HumanNotEngaged, SayHiWithGestures).

If the student wrote the correct initial Q-formulae and tried (without success) to solve the problem analytically, they should get 7 points.

Question 2:

The agent should always choose the action that provides the maximum expected cumulative reward Q, so the best action in this case would be to say hi (without gestures).

Grading: 2 points if they explain that the system must select actions that give the maximum Q value, or 0.

6 Data privacy

In the field of data privacy, what is the difference between a direct identifier and a quasi-identifier? Explain in a few sentences, and illustrate those types of personal identifiers with a few examples.

Fill in your answer here

Maximum marks: 4

A direct identifier is a type of personal information that can univocally disclose the identity of a given person, such as a full person name, a social security number, a bank account, a mobile phone number, etc.

A quasi-identifier, on the other hand, is a personal information that does not typically reveal the identity of the person when considered in isolation but may do so when combined with other quasi-identifiers and some background knowledge. Examples of such quasi-identifiers: gender, city of residence, birth date, age, ethnicity, etc.

Grading: 2 points for correct explanations and 2 points for the examples.