Name: Matrikelnummer:

Chair for Algorithms
and Data Structures
Prof. Dr. Hannah Bast
Allthetu Tors

Information Retrieval
WS 2021/2022

http://ad-wiki.informatik.uni-freiburg.de/teaching

Exam

Friday, 4th of March 2022, 14:00 - 16:15 h, HS 026

General instructions:

There are five tasks, of which you can select *four tasks of your choice*. Each task is worth 25 points. If you do all five tasks, we will only count the best four. That is, you can reach a maximum number of 100 points.

You need 50 points to pass the exam. You have 2 hours 15 minutes of time overall. If you do four tasks, this is 30 minutes per task on average plus 15 minutes buffer time.

You are allowed to use one DIN A4 sheet of paper with any contents of your choice related to the lectures and the exercises. You can write on the front and on the back and it can be hand-written or a printout, but it must be from yourself, not from someone else. You are not allowed to use any computing devices or mobile phones, in particular nothing with which you can communicate with others or connect to the Internet or parallel universes.

You may write down your solutions in either English or German.

Please write your solutions on this hand-out, below the description of the tasks! You can also use the back side of the pages. Please write your name and Matrikelnummer on the top of this cover sheet in the framed box. If you need additional pages, please write your name and Matrikelnummer on each of them, too.

Important:

For the programming tasks: You can use Python, Java, or C++. None of your functions must be longer than TWENTY lines.

For all other tasks: Do not simply write down the final result. It should also be clear how you derived it.

Good luck!

T1	T2	T3	T4	T5	Points	Grade

Task 1 (Ranking, Evaluation and List Intersection, 25 points)

1.1 (5 points) Consider the following collection of four documents D_1, ..., D_4. Write down the term-document matrix for this collection with *tf.idf* scores.

D_1: bla bla bla
D_2: bli blu blo
D_3: bla blo
D_4: blu blu

1.2 (5 points) Consider a query "bla bli". Compute the dot-product similarities of the four documents above (using your *tf.idf* scores from 1.1) to that query.

1.3 (5 points) Rank the four documents by the scores from 1.2 (highest score first). Assume that documents D_1 and D_3 are relevant. Compute the following metrics: P@2, P@R, AP (average precision).

1.4 (10 points) Given an array A with integer values in ascending order, we want to locate element x in A with galloping search (an exponential search followed by a binary search). Prove that, if x is contained in A, the algorithm needs $O(\log d)$ time, where d is the position of x in A.

Task 2 (Zipf's law, Compression, and UTF-8, 25 points)

2.1 (8 points) Zipf's law states the frequency F of the n-th most frequent term in a text collection as a function of n. State the law and prove that F shows as a straight line with negative slope in a log-log plot.

2.2 (7 points) Write the following numbers in Elias-Delta encoding: 1, 7, 16. For each code, underline the part that comes from the Elias-Gamma encoding.

2.3 (5 points) Consider the following encoding: $a = 000$, $b = 001$, $c = 01$, $d = 10$, $e = 11$. Find a distribution over $\{a, b, c, d, e\}$ such that this encoding is entropy-optimal, with a proof that it indeed is.

2.4 (5 points) Give an example of a two-byte UTF-8 sequence that is a *valid* encoding of a single Unicode character. Give an example of a two-byte UTF-8 sequence that is an *invalid* encoding of a single Unicode character. With explanation! Specify each sequence as a bit sequence and as a sequence of four hexadecimal digits.

Task 3 (SPARQL and fuzzy prefix search, 25 points)

3.1 (5 points) Assume we have a knowledge base which includes relations for *nationality* (relating persons to countries), *place_of_birth* (relating persons to their place of birth), *language_spoken* (relating countries to languages), *olympic_discipline* (relating persons to their olympic disciplines). Express the query Olympians from German-speaking countries with their place of birth and their olympic discipline in SPARQL.

3.2 (10 points) Write a function *ped(x, y)* that computes the PED between two strings x and y in $O(|x| \cdot |y|)$ time.

3.3 (10 points) Let x and y be non-empty strings, and let $\delta \in \mathbb{N}$ such that $|x| + \delta < |y|$ and $\mathrm{PED}(x, y) \leq \delta$. Proof that under these conditions $\mathrm{PED}(x, y) = \mathrm{PED}(x, y')$, where y' is the prefix of length $|x| + \delta$ of y.

Task 4 (Web Apps and Naive Bayes, 25 Punkte)

4.1 (5 Points) Write a valid HTML page with heading *Hardle*, three text input fields with ids *f1*, *f2*, *f3*, and a button labeled *Guess*. The HTML should include a JavaScript file *script.js*, to be written in Task 4.2.

4.2 (10 Points) Write *script.js* with the following functionality. Whenever the user presses the *Guess* button, send the content of input field *f1* to a backend at the relative URL */backend*. The backend returns a JSON object with 3 integer attributes *correct_positions*, *correct_letters* and *tries_left*. Show the value of *correct_positions* and *correct_letters* in input fields *f2* and *f3*. If *correct_positions* matches the length of the content of field *f1*, show *You win* in field *f1*. If field *tries_left* is 0, show *You lose*. You may use *jQuery*.

4.3 (10 Points) Consider each number from 0..15 as a document with four words, where the words are 0 or 1 and the ith word stands for the ith bit in the binary representation of the number. For example, document 1100 is the number 3. Each document is labeled with class A if the number is even, and with class B if the number is odd. Train a Naive Bayes classifier on these 16 labeled documents. Determine all probabilities p_c and p_{wc} (and write down the intermediate steps). Show that using these probabilities, Naive Bayes predicts class A for the number 1.

Task 5 (Latent Semantic Indexing and Linear Classifiers, 25 points)

5.1 (10 points) Write a function *top(q, U, S, V, k)* that returns the index of the top-ranked document for a given query vector q, the *numpy* matrices U, S, V from the singular value decomposition, and the dimension k of the approximation. Use the function *transpose* for transposition and *dot* for matrix-matrix or matrix-vector products. In your code, indicate (using a different color) the dimensions of the vectors and matrices and make sure that they match.

5.2 (5 points) Let $H = \{x : 2 \cdot x_1 - 2 \cdot x_2 + x_3 = 7\}$ be a 2D plane in 3D space. Compute the distance of the point $x = (1, 2, 3)$ to H. If you know the formula for the distance computation, you can just use it. Otherwise, derive the formula by writing x as the sum of a point on H and a multiple of the unit normal vector of H; the multiple is then the distance.

5.3 (10 points) Let $\sigma(t) = 1/(1 + e^{-t})$ be the sigmoid function. Prove that $\sigma'(t) = \sigma(t) \cdot \sigma(-t)$.