

Name:

Matrikelnummer:

Chair for Algorithms
and Data Structures
Prof. Dr. Hannah Bast
Claudius Korzen

Information Retrieval
WS 2018/2019

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



Exam

Tuesday, February 12, 2019, 14:00 - 16:30, Lecture Hall 101-026

General instructions:

There are six tasks, of which you can select *five tasks of your choice*. Each task is worth 20 points. If you do all six tasks, we will only count the best five, that is, you can reach a maximum number of 100 points.

You need 50 points to pass the exam. You have 150 minutes of time overall. If you do five tasks, this is 30 minutes per task on average.

You are allowed to use any amount of paper, books, etc. 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	T6	Points	Grade
----	----	----	----	----	----	--------	-------

Task 1 (List intersection, 20 points)

In the following tasks, the two list lengths are denoted by k and n , ordered such that $k \leq n$.

1.1 (10 points) Write a function *intersect(list1, list2)* that takes two sorted lists and computes their intersection by binary searching each of the elements from the shorter list in the longer list. You must implement the binary search yourself. The running time must be $O(k \cdot \log n)$.

1.2 (5 points) We have shown that the running time of the galloping search algorithm is $O(k \cdot \log(1 + n/k))$. Prove that this running time is asymptotically never worse than that of the simple linear-time algorithm. You can use without proof that $e^x \geq 1 + x$ for all real x .

1.3 (5 points) You are given two sorted lists and you want to compute their *union* as fast as possible. You care about the running time in *practice*, not in theory. How would you decide whether to use the simple linear-time algorithm or the galloping search algorithm? Explain the reasons for your choice.

Task 2 (Encodings and UTF-8, 20 points)

2.1 (5 points) Let us assume we want to encode the numbers 1, 2, 3, 4, 5, 6, 7, 8 and they all occur with equal probability. Specify a prefix-free entropy-optimal encoding and prove that it is entropy-optimal.

2.2 (5 points) Shannon's source coding theorem says that there always is a prefix-free encoding with expected code length of at most $H(X) + 1$, where $H(X)$ is the entropy of the probability distribution of the symbols to be encoded. Under exactly which condition can we find a prefix-free encoding with expected code length of exactly $H(X)$, that is, without the +1? You should state the condition as simple as possible.

2.3 (5 points) The Greek lowercase α has a UTF-8 codepoint of 945, while the greek uppercase A has a codepoint of 913. Write down the UTF-8 codes of the two characters in binary. Write them down one below the other.

2.4 (5 points) Write a function *greek_lower_case(bytes)* that takes a UTF-8 encoded string (in the form of a byte array) that is composed entirely of Greek letters and converts all letters to lower case. You can assume that, just like in the ASCII code, the lowercase letters as well as the uppercase letters have consecutive codepoints. For example, β has codepoint 946, γ has codepoint 947, and so on.

Task 3 (Fuzzy prefix search, 20 points)

3.1 (5 points) Write down the 2-gram index for the following list of names: *abba*, *bab*, *aba*. Use padding from the left with a single \$.

3.2 (5 points) Let $x = aaab$. Use your 2-gram index from Task 3.1 to find all words y from the list from Task 3.1 such that $\text{PED}(x, y) \leq 1$. It is not enough if you just write down the result, it has to be clear how you used the 2-gram index. You can do the required PED calculations in your head and just write down the result (you don't have to write down the dynamic programming tables).

3.3 (10 points) Write a function *fuzzy-prefix-search*(x , *index*) that for a given string and a given 2-gram index returns the number of PED computations needed to compute all words y with $\text{PED}(x, y) \leq 1$ using the 2-gram index. Use a hash map for counting the number of matching 2-grams for each word. The running time should be linear in the total size of the inverted lists of the 2-grams of x .

Task 4 (Search web app, 20 points)

4.1 (10 points) Write a function *search(query, index, titles)* that for a query (given as an array of keywords), an inverted index (given as a map from words to arrays of pairs of document id and score), and the document titles (given as an arrays of strings), returns the title of the matching document with the highest score. A matching document does not necessarily have to contain all query words. Use a hash map to compute the various document scores. Do not assume that there is a function for merging sorted lists.

4.2 (10 points) Write a function *search_server(socket, index, titles)* that for a given server socket and an index and document titles like in Task 4.1, listens to HTTP requests, interprets them as keyword queries and returns the best matching title using your function from Task 4.1. You can assume that *server_socket* has a method *accept* that returns a client socket, that the client socket has methods *receive* and *send* for reading and writing an arbitrary string in one slurp, that all requests are properly formed GET requests, and that no errors of any kind occur. You can also assume that all your titles contain only ASCII characters. The response should include header lines stating the status, the response length and the response type.

Task 5 (Latent Semantic Indexing, 20 points)

Consider the following term-document matrix (three terms, five documents):

$$A = \begin{pmatrix} 1 & 1 & 1 & 2 & 2 \\ 2 & 1 & 1 & 3 & 1 \\ 1 & 1 & 2 & 2 & 1 \end{pmatrix}$$

5.1 (5 points) Determine the rank of matrix A . Don't just write down the result, but also provide a proper argument.

5.2 (5 points) Change *at most two* entries of A such that its rank becomes 2. Change *at most four* entries of A such that its rank becomes 1. Don't just write down the result, but also provide a proper argument.

5.3 (10 points) Write a function *is_row_orthonormal(matrix)* that determines if the product of the given matrix with its transpose is the identity matrix. Your function should not materialize the transpose and it should not use any special linear-algebra functions, just simple for-loops. You can assume that there are functions which give you the number of rows and columns of the matrix.

Task 6 (SPARQL and POS-Tagging, 20 points)

6.1 (5 points) Write down the SPARQL query that produces a table of all female astronauts from Asia with the following columns: name of the astronaut, country of nationality, birth date. You can assume that you have a knowledge base with the appropriate entities and relations (the exact names of the relations are up to you). The body of your query should contain at least four triples and the corresponding subgraph should be connected.

6.2 (5 points) Write down the equivalent query in SQL, assuming that you have one two-column table for each of the relations you used in your query from Task 6.1.

6.3 (10 points) Consider a HMM with two observable states *GOOD* and *BAD* (the mood) and two hidden states *SUNNY* and *RAINY* (the weather). Given a sequence o_1, \dots, o_n consider the following greedy algorithm to find a sequence h_1, \dots, h_n of hidden states: In the first step, find h_1 which maximizes $P(h_1) \cdot P(o_1|h_1)$. In subsequent steps, find h_{i+1} which maximizes $P(h_{i+1}|h_i) \cdot P(o_i|h_i)$. Find an example problem instance (all the probabilities and a sequence of observations) such that this greedy algorithm finds a different sequence of hidden states than the Viterbi algorithm (with explanation, of course). Hint: a sequence of two observations is enough.