

Name:

Matrikelnummer:

Chair for Algorithms
and Data Structures
Prof. Dr. Hannah Bast
Natalie Prange

Information Retrieval WS 2022/2023

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



Exam

Tuesday, 28th of February 2023, 11:00 - 13:15 h, in building 101 SR 00-010/14, SR 01-009/13 and SR 02-016/18

General instructions: There are five tasks. Each task is worth 20 points, 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. This is 20 minutes per task plus 35 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 handwritten 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 handout, 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 should use Python. None of your functions must be longer than *fifteen* 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 (Exponential Search, Elias-Gamma, Lagrangian multipliers, 20 points)

1.1 (5 points) Implement a function *exponential_search*(A, x), which takes an array A of integers sorted in ascending order and an integer x . You can assume that x is contained in A . Let j be the smallest index for which $A[j] = x$. The function should compute an index i with $j \leq i \leq 2j$ and the function should run in $O(\log(1 + j))$ time.

1.2 (5 points) Write down the Elias-Gamma codes for $x = 3, 4, 7, 8, 15, 16$ with the leading zeroes underlined. Then write down the code length (in number of bits) for each of these numbers. Then write down the general formula for the code length for an arbitrary number $x \in \mathbb{N}$ and briefly explain why it's correct.

1.3 (10 points) Let p_1, \dots, p_n be a probability distribution over n symbols. Determine the probability distribution for which the entropy is maximized, using the method of Lagrange multipliers. You can assume without proof that the assumptions for applying the method are satisfied and that the point where all partial derivatives of the Lagrangian function are zero is a global maximum.

Task 2 (Q-Grams and Edit Distance, 20 points)

2.1 (5 points) Write down the edit distance between the strings $x = \textit{alfa}$ and $y = \textit{alpha}$ and the sequence of operations needed to get from x to y (you don't have to write down the table). Then pad the words with dollar signs on both ends, as appropriate for a 3-gram index, and write down the set of 3-grams that the two padded strings have in common.

2.2 (5 points) Calculate the minimal possible edit distance between two strings of length 5 that have three 3-grams in common when using padding on both sides. Write down the intermediate steps for the calculation.

2.3 (10 points) The general algorithm for computing the edit distance between two strings of length n and m runs in $O(n \cdot m)$. Write a function $\textit{one_delete}(x, y)$ that runs in time $O(m + n)$ and returns *True* if the string x can be transformed into the string y with exactly one *delete* operation, and *False* otherwise.

Task 3 (Ranking and Linear Classifiers, 20 points)

3.1 (8 points) You are given the following term-document matrix (3 terms, 5 documents):

$$A = \begin{bmatrix} 1 & 2 & 0 & 1 & 5 \\ 4 & 1 & 0 & 2 & 3 \\ 1 & 1 & 2 & 1 & 0 \end{bmatrix}$$

State a query $q \in \mathbb{N}^3$ such that $q^T A$ returns the following document ranking in descending order (ties will be decided in favor of lower document numbers): 1, 3, 4, 2, 5. Which documents should be relevant for q such that the ranking gets a P@R score of 0.5 and a P@4 score of 0.25? Calculate the average precision of your solution.

3.2 (5 points) Let $H = \{x: w \bullet x = b\}$ be a hyperplane in \mathbb{R}^d with $w \in \mathbb{R}^d, w \neq \mathbf{0}$ and $b \in \mathbb{R}$. For which b does H contain the origin (with proof)? Now assume that $d = 4$, $w = (\frac{1}{2}, \frac{1}{2}, \frac{1}{2}, \frac{1}{2})$, and $b = 0$. Compute the distance of the point $p = (1, 1, 1, 1)$ to H .

3.3 (7 points) The *F1 score* of the predictions for a class is defined as $2/(1/P + 1/R)$, where P is the precision and R is the recall. Prove that $F1 = 1$ if and only if $P = 1$ and $R = 1$. Draw four two-dimensional points, each with a label from one of two classes, such that no linear classifier can achieve $F1 = 1$ for both classes. Briefly explain why your points have this property.

Task 4 (Web Apps and Naive Bayes, 20 points)

4.1 (10 points) Write a function $predict(doc, wordIds, pwc, pc)$ that expects a document (as an array of strings), the word IDs (as a map where $wordIds[<word>]$ gives the ID), the p_{wc} learned during training of Naive Bayes (as an array of k arrays, one for each of the k classes, and all entries are non-zero) and the p_c (as a one-dimensional array). The function should return the most likely class and its probability. Do not use smoothing. You can ignore rounding problems.

4.2 (5 points) Write an HTML page with heading *Enter Query*, an input field, a paragraph, and a button labeled *Predict*. The input field and the paragraph should have an ID. The HTML should include a JavaScript file *script.js*, to be written for Task 4.3.

4.3 (5 points) Write *script.js* with the following functionality: If the user presses the button, send the content of the input field to a server using a GET request to a URL of your choice. Assume that the backend implements the method from Task 4.1 and returns a single JSON object with fields *class* and *probability*. Replace the content of the paragraph with the predicted class and its probability, separated by a space.

Task 5 (Latent Semantic Indexing and SPARQL, 20 points)

5.1 (10 points) Write a function $top(q, A, U, k)$ that returns the index of the top-ranked document using the document expansion method (the third variant discussed in the lecture) and takes the following parameters, each as a *numpy.array*: the query vector q (as one column), the original term-document matrix A , the U from the SVD of A , and the dimension k of the approximation. Use the function *transpose* for transposition and *dot* for a matrix-matrix or matrix-vector product.

5.2 (5 points) Let A be an $m \times n$ matrix with rank r , and let $A = U \cdot S \cdot V$ be the singular value decomposition, where U is an $m \times r$ column-orthonormal matrix, S is a diagonal $r \times r$ matrix, and V is an $r \times n$ row-orthonormal matrix. Prove that $V = S^{-1} \cdot U^T \cdot A$.

5.3 (5 points) Assume that we have a knowledge graph which includes relations for *occupation* (relating persons to their occupation), *place_of_birth* (relating persons to their place of birth, which you can assume to be unique) and *educated_at* (relating persons to the university they studied at). Express the following question as a SPARQL query: YouTubers who studied at the same university as MrBeast, each with their place of birth.