

Name:

Matriculation number:

↑ Please write **clearly** and use capitalized BLOCK LETTERS ↑

Chair of Algorithms  
and Data Structures  
Prof. Dr. Hannah Bast  
Sebastian Walter

**Information Retrieval /  
Databases and Information Systems  
WS 2023/2024**

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



## Exam

Monday, 18th of March 2024, 14:00 - 16:30 h, in building 101

There are *five* tasks. Each task is worth 20 points, that is, you can reach a maximum of 100 points. You need 50 points to pass the exam.

You have 2 hours 30 minutes of time overall. This is 20 minutes plus 10 minutes buffer per task.

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.

**Important:** For the programming tasks you should use Python. Each programming task specifies a maximum number of lines you are allowed to write. For all other tasks, do not simply write down the final result, but it should also be clear how you derived it.

**Multiple-choice tasks:** Some of the tasks consist of a set of questions, to each of which the answer is either *true* or *false*. A correct answer gives plus 2 points, a wrong answer gives minus 2 points, and no answer gives 0 points. So if you don't know the answer, do not guess but write *no answer* or simply nothing. No explanations are needed. Altogether, you cannot get less than 0 points for such tasks.

**How and what to hand in:** Please write, [in the blue box at the top](#), your name and your matriculation number. Please write your solutions on the exam: use the front side of the sheet on which the question is printed and then the back side of the *previous* sheet. If you wish to submit additional paper (which should be the exception), write your name and matriculation number on every additional sheet.

**Good luck!**

A1

A2

A3

A4

A5

Points

Grade

**Task 1** (Inverted Index and Ranking, 20 points)

**1.1** (6 points) Consider the following collection of four documents  $D_1, \dots, D_4$ . Write down the inverted index with *tf* scores. For which  $C$  and  $\alpha$  does Zipf's Law ( $F_n = C \cdot n^{-\alpha}$ ) hold exactly for this document collection?

$D_1$  : b c a c d c

$D_2$  : b c a c d c b c

$D_3$  : a d a c c b

$D_4$  : c a a c c

**1.2** (10 points) Write a function  $merge(A: list[int], B: list[int]) \rightarrow list[int]$  that merges two sorted lists of integers into a single sorted list of integers, but without removing duplicates. For example,  $merge([1, 3, 4], [2, 3, 5]) = [1, 2, 3, 3, 4, 5]$ . The function should run in time  $O(|A| + |B|)$  and should have at most 15 lines.

**1.3** (4 points) Given the list of result IDs  $[6, 7, 2, 4, 5]$  and the set of relevant IDs  $\{1, 2, 3, 4\}$ . Write down the following values:  $P@2$ ,  $P@R$ ,  $AP$ . For  $rel_1 = 0$ ,  $rel_2 = 0$ , and  $rel_3 = 4$ , write down  $DCG@3$ .

**Task 2** (Database Basics, 20 points)

**2.1** (12 points) You are given the following database schema, where columns marked PK are part of a table's primary key and columns marked FK are foreign keys to other tables.

movies(id INT PK, title TEXT, release\_year INT)

persons(id INT PK, name TEXT)

actors(movie\_id INT FK(movies.id) PK, person\_id INT FK(persons.id) PK)

streaming(movie\_id INT FK(movies.id) PK, platform TEXT PK, views INT)

Write a corresponding SQL query for each of the following queries:

1. Titles of all movies released in the 20th century
2. Titles of the top 10 movies ordered by total number of views on "Netflix" and "Amazon Prime"
3. Title and comma-separated list of actors for all movies with at least 10 actors

**2.2** (8 points) The two query plans below each execute the same query over two tables  $A$  and  $B$ . Compute the cost estimate for each plan if  $A$  has 10 rows and 4 columns and  $B$  has 20 rows and 2 columns. Give an example of a shape for  $A$  and  $B$  such that plan 2 is cheaper than plan 1.

Assume the following cost estimates:  $select \rightarrow n \cdot k$ ,  $project \rightarrow n \cdot k'$ , and  $cartesian-product \rightarrow n_1 \cdot n_2 \cdot (k_1 + k_2)$ , where  $n$ ,  $n_1$ ,  $n_2$  and  $k$ ,  $k_1$ ,  $k_2$  refer to the number of rows and columns of the respective input tables and  $k'$  refers to the number of output columns. Assume that  $select$  selects half of the rows. It doesn't matter what exactly  $\phi_1$  and  $\phi_2$  are.

Query plan 1:

1.  $A' = select(A, \phi_1)$
2.  $B' = select(B, \phi_2)$
3.  $C = cartesian-product(A', B')$
4.  $R = project(C, \text{all cols of } C, true)$

Query plan 2:

1.  $C = cartesian-product(A, B)$
2.  $C' = select(C, \phi_1)$
3.  $C'' = select(C', \phi_2)$
4.  $R = project(C'', \text{all cols of } C'', true)$

**Task 3** (Advanced SQL, SPARQL, and Knowledge Graphs, 20 points)

**3.1** (7 points) The following SQL query is the result of transforming a SPARQL query to SQL for a database that stores all the triples in one table *triples*. Write down a SPARQL query that leads to this SQL query. Write down in natural language what this query computes.

```
SELECT t3.subject, t3.object
FROM triples AS t1, triples AS t2, triples AS t3
WHERE t1.predicate = "has_mother" AND
      t2.predicate = "has_father" AND
      t1.subject = t2.subject AND
      t3.predicate = "married_to" AND
      t3.subject = t1.object AND
      t3.object = t2.object;
```

**3.2** (6 points) For each of the following queries, write down whether it computes the three Oscars that Meryl Streep has won, together with the corresponding movie. For each query write either *true* or *false*. No explanation is needed. A correct answer gives plus 2 points, a wrong answer gives minus 2 points. If you don't know the answer, write *no answer* or nothing. You cannot get less than 0 points for this task. Hints: *SELECT \** selects all variables from the query, *P166* is *won-award*, *Q873* is *Meryl Streep*, *P1686* is *for which film*, and *P31 Q19020* means *is an Oscar*.

```
SELECT * WHERE {
  wd:Q873 wdt:P166 ?a .
  ?a wdt:P1686 ?b .
  ?a wdt:P31 wd:Q19020 .
}

SELECT * WHERE {
  wd:Q873 p:P166 ?a .
  ?a ps:P166 ?b .
  ?a pq:P1686 ?b .
  ?a wdt:P31 wd:Q19020 .
}

SELECT * WHERE {
  wd:Q873 p:P166 ?a .
  ?a ps:P166 ?b .
  ?a pq:P1686 ?c .
  ?b wdt:P31 wd:Q19020 .
}
```

**3.3** (7 points) Manually compute the result of the SQL query given on the left, using the table *movies* given on the right. Write down in natural language what this query computes.

```
SELECT year, title, score
FROM movies NATURAL JOIN (
  SELECT year, (MAX(score) AS score)
  FROM movies
  GROUP BY year
)
ORDER BY score DESC;
```

title	score	year
Spider-Man	8.3	2021
Everything Everywhere	7.8	2022
Oppenheimer	8.4	2023
Don't Look Up	7.2	2021
Dune	8.0	2021
Top Gun	8.2	2022

**Task 4** (Fuzzy Search, Q-Gram Index and Web Applications, 20 points)

**4.1** (5 points) For the query prefix  $x = bree$  and the words  $y_1 = free$ ,  $y_2 = dreisam$  and  $y_3 = spree$ , write down the left-padded 3-grams for the 3-gram index. For each word  $y_i$  with  $i \in \{1, 2, 3\}$ , determine how many 3-grams it has in common with the query prefix. Then for each  $y_i$ , say whether  $PED(x, y_i)$  needs to be computed explicitly in order to find out if  $PED(x, y_i) \leq 1$  or whether it is enough to look at the number of common 3-grams.

**4.2** (8 points) For each of the following statements, say whether it is *true* or *false*. No explanation is needed. A correct answer gives plus 2 points, a wrong answer gives minus 2 points. If you don't know the answer, write *no answer* or nothing. You cannot get less than 0 points for this task. The "algorithm for fuzzy prefix search" refers to the algorithm from the respective lecture.

1. For all strings  $x$  and  $y$ , it holds that  $PED(x, y) \leq ED(x, y)$ .
2. For all strings  $x$  and  $y$ , it holds that  $PED(x, y) \geq abs(|x| - |y|)$ .
3. Without any padding, the algorithm for fuzzy prefix search is still correct but less efficient.
4. With padding on both sides, the algorithm for fuzzy prefix search is still correct.

**4.3** (7 points) Write a simple HTML page with heading *Search*, an input field and a button, which upon a click of the button, sends the contents of the input field back to the server. Use forms, not JavaScript. Given your HTML, the input text *example*, and assuming that the page is accessed via *localhost:8000/search.html* how does the first line of the sent request look like when clicking the button?

**Task 5** (Word Embeddings, Logistic Regression, Language Models, 20 points)

**5.1** (7 points) Write a function  $\text{cos\_sim}(emb1: \text{list}[\text{float}], emb2: \text{list}[\text{float}]) \rightarrow \text{float}$  that computes the cosine similarity between two given word embeddings  $emb1$  and  $emb2$ . You can assume that the embeddings have the same dimension. You can use basic arithmetic operations and functions from the *math* module, like *math.sqrt*, but you should not rely on the module *torch*. Your function should have at most 10 lines.

**5.2** (8 points) Write a function  $\text{epoch}(X: \text{tensor}, y: \text{tensor}, w: \text{tensor}) \rightarrow \text{tensor}$  that does one epoch of logistic regression, where *tensor* is an abbreviation for *torch.Tensor*. The arguments are as follows:  $X$  is a 2D tensor with input vectors as rows,  $y$  is a 1D tensor with the input labels,  $w$  is a 1D tensor with the current weights, and the function should return a 1D tensor with the new weights after the epoch. You can assume that all values of all tensors involved are of type *float* and that the bias is already part of  $X$  and  $w$ . You should take a batch size of 1 and a learning rate of 0.1. Your function should have at most 10 lines.

**5.3** (5 points) Express  $\text{softmax}(x, 0)$  in terms of  $x \in \mathbb{R}$  and the sigmoid function  $\sigma$ . Do not just write down the result, but also how you got there.