



Computer Engineering II

Exercise Sheet 7

Quiz

1 Quiz Questions

- A candidate key is a column or set of columns of a table that can uniquely identify any database record (row) without ambiguity. Specify all the candidate keys of a given table.
- Given a set of SQL queries, determine which statement will produce an error and which will yield a valid response from a SQL interpreter.
- Determine how many rows the result of a given query will produce.
- Determine the result of a right join statement.

Basic

2 Database Design

Swiss cat and dog owners are required by law to microchip their pets. This greatly improves the chances of successful pet discovery in case the pet has run away or was stolen. The Animal Identity Service (ANIS) is a company that operates the Swiss animal database and allows veterinarians, animal shelters and police stations access if they are trying to find owners of lost pets. Unfortunately, the ANIS website was vulnerable to SQL injection attacks and an evil attacker managed to delete the entire database including all backups. Let's help them reconstruct their lost data!

- For each animal ANIS wishes to provide name and microchip id
- Every animal has exactly one owner. For each owner ANIS wishes to provide a name, an address, a phone number, and an email address
- To ensure the safety of a found animal at a shelter, ANIS provides additional information by recording any number of allergies the animal may have. We store a unique identifier for every possible allergy including its name and a number from 1 to 5 which indicates how strongly the allergy in question afflicts the animal.
- To help protect against animal microchip tampering and replacement, ANIS provides information about the type (cat, dog or other) and about any possible breed (Rottweiler, Norwegian Forest Cat, etc.) of the animal in question. For every distinct animal type or breed we store a unique identifier including the type or breed's name.

- a) Draw an Entity Relationship Diagram for the dataset described above. Make sure that you specify all cardinality constraints and that the resulting tables don't contain redundant entity sets, relationships, or attributes. Please include all made assumptions in your answer. Don't forget to assign primary key and foreign key constraints and mark attributes as unique wherever necessary
- b) Convert your Entity Relationship Diagram into relations and write SQL statements to create these relations. You may make reasonable choice of data types. Don't forget to assign primary key and foreign key constraints. If you're unsure how to apply foreign key constraints, or it's unclear what many to many relationships are, this resource may help: <http://www.databaseprimer.com/pages/table-relationships/>. For instant feedback of your SQL statements you may want use an SQL interpreter¹.

3 Database Queries

Take a look at the database schema and use the SQL query interface on <http://pc-10129.ethz.ch/sqlquery/> to answer the following questions:

1. Display the title and id of the first five entries in the movie table.
2. Display the first two rows of the movie table after sorting it by title in alphabetically reversed order.
3. How many movies came out since the year 2000?
4. What's the tomatometer score of the movie titled 'The Matrix'?
5. How many movies have a higher tomatometer score than 'The Matrix'?
6. Display in one row the overall maximum, minimum and average tomatometer score.
7. Which movie titles start with a capital X?
8. How many movie titles contain the string 'fight'? How many movies do you find if you ignore the case?

Advanced

4 Advanced Database Queries

You will need to use table joins to answer the following queries. To find the first movie 'Steve Buscemi' was cast in, you could use a query like:

```
SELECT movie.title,movie.year FROM cast_info
      JOIN person ON cast_info.person_id = person.id
      JOIN movie ON movie.id = cast_info.movie_id
WHERE person.name = 'Buscemi, Steve'
ORDER BY year ASC;
```

¹<https://kripken.github.io/sql.js/GUI/>

If you're unsure which tables you need to join, take a look at the schema at <http://pc-10129.ethz.ch/sqlquery/schema>.

1. List all actresses who were cast in 'The Matrix'.
2. How many female directors can you find in the database?
3. List all directors who have directed more than 20 projects and have also appeared at least once as an actor in any project (movie, tv show, etc).
4. Which movie is described by the most keywords?
5. What's the average, maximum and minimum number of keywords per movie?
6. List all actors with an average tomatometer score over 85% who were cast in at least 30 movies and who appear in at least one movie made since 2000?
7. List all directors who have achieved to direct at least 10 movies with a tomatometer score of over 90%