Weekly exercises IN3020&4020 (2025)
Sheet 1: Relational Data Model

1 The relational model revision

1. Briefly recall the main concepts and ideas of the relational data model:
   - relation: relation name, attribute names, signature, relation schema, tuples, relation state, attribute values, domain
   - relation properties: tuples are unordered and unique, attribute names in the signature may be seen as ordered or unordered

2. Give the ideas of the following notions and explain the differences between them:
   - Super key
   - Candidate key
   - Primary key
   - Foreign key

3. Why are the following statements false?
   - A relation cannot have multiple super keys.
   - A candidate key can include excess attributes as long as the key still uniquely identifies tuples.
   - A foreign key cannot refer to a candidate key in the same relation.

4. (*) In the lectures, we defined: 'Relation with schema $R[A_1, A_2, \ldots, A_n]$ *satisfies* the key constraint if the key subset of the constraint is a super key in the relation'. Explain what is the problem if we replace 'super key' by 'candidate key' in this definition (in this case, a problem is an undesirable, counter-intuitive behaviour).

5. Can you think of other types of constraints that we may put on data, that are not expressible via primary and foreign keys? Hint: look for "Domain Constraints" the Web.

2 PostgreSQL Practical Review

Note: To use the PostgreSQL databases in this course, you should log in to a Linux terminal at IFI or log in remotely. Then,
`psql -h dbpg-ifi-kurs02 -U [username] -d in3020_[username]`
will log you in to you personal database. Here you will be able to create tables, insert, update, etc., while
`psql -h dbpg-ifi-kurs02 -U [username] -d fdb`
will log you in to the movie database. This database is read only.

1. Log in to your PostgreSQL environment; information on how this can be done is found at `https://www.uio.no/studier/emner/matnat/ifi/IN2090/h21/programvare/postgres.pdf`

2. Create the table `Ships` with the following fields (i.e., attributes):

   - `ship_id`
   - `name`

   Make attribute `ship_id` the primary key.

3. Insert the following data into the table `Ships`:

   - (1, Pequod)
   - (2, Titanic)
   - (3, Le pharaon)

4. Create a table `Persons` with the attributes:

   - `name`
   - `ssn`
   - `street address`
   - `ZIP code`

   Use reasonable datatypes.

5. What is a reasonable primary key for `Persons`?

6. Alter table`Persons` and add a primary key.

7. Insert the following data into `Persons`:
   - (Ishmael, 111111110, 20 Cooper Square Manhattan, 10001)
   - (Captain Ahab, 111111112, 16 Nantucket Avenue, 02554)
   - (Queequeg, 111111113, Rocovoco, 99991)
   - (Molly Brown, 111111114, Hannibal Missouri, 63401)
   - (Edmond Dantes, 111111115, Marseille, 13003)

8. Create a table `Crew` with attributes
   - `ssn`
   - `ship_id`

9. Let the combination of `ssn` and `ship_id` be the primary key.
10. Add foreign keys to `Persons` and `Ships`.
11. Insert the following data into `Crew`:
    - (111111110, 1)
    - (111111112, 1)
    - (111111113, 2)
    - (111111114, 3)

12. Use SQL to find all persons with an address containing Manhattan.

13. Print out the name of the persons that were crew on the ship `Pequod`. Order the names alphabetically.

14. Use SQL to find how many crew members there were on the ship `Titanic`.