

Homework 4 | Database Design

Updates made to the assignment after its release are *highlighted in red*.

Objectives: To translate from entity-relationship diagrams to a relational database, and to understand functional dependencies and normal forms.

Due date: Friday, May 2 @ 9:00pm

Median completion time (23sp): 5 hours

Resources

For this assignment, you will need:

- Pen and paper, or any drawing tools you prefer (e.g., PowerPoint, draw.io).
- SQLite, Postgres, or any other relational database.

Problem Set

Question 2: Geography E/R Diagram (10 points)

Draw an E/R diagram for a mapping application. Your diagram must contain the following list of entities, along with their listed attributes. Your diagram should also contain the relationships described below, modeled as edges between entities; note that edges between entities can be labeled with constraints.

Make sure to choose the correct primary key(s) based on the information below.

Entities

- `countries`: name, area, population, GDP ("gross domestic product")
 - a country's name uniquely identifies the country
- `cities`: name, population, longitude, latitude
 - a city is uniquely identified by its (longitude, latitude), not by name!
 - eg: there are 41 different cities and towns named Springfield in the US
- `rivers`: name, length
- `seas`: name, max depth
- rivers and seas are uniquely identified by their name *within the set of all water entities*

- e.g., if there were a river named "Ganges" would be a unique water entity, no other river or sea could have that name

Relationships

- each city belongs to exactly one country
- each river crosses one or several countries
- each country can be crossed by zero or several rivers
- each river ends in either a river or a sea

If you choose to draw your diagrams on paper, take *high quality pictures* and then use <https://imagnetopdf.com> to convert the image to a pdf.

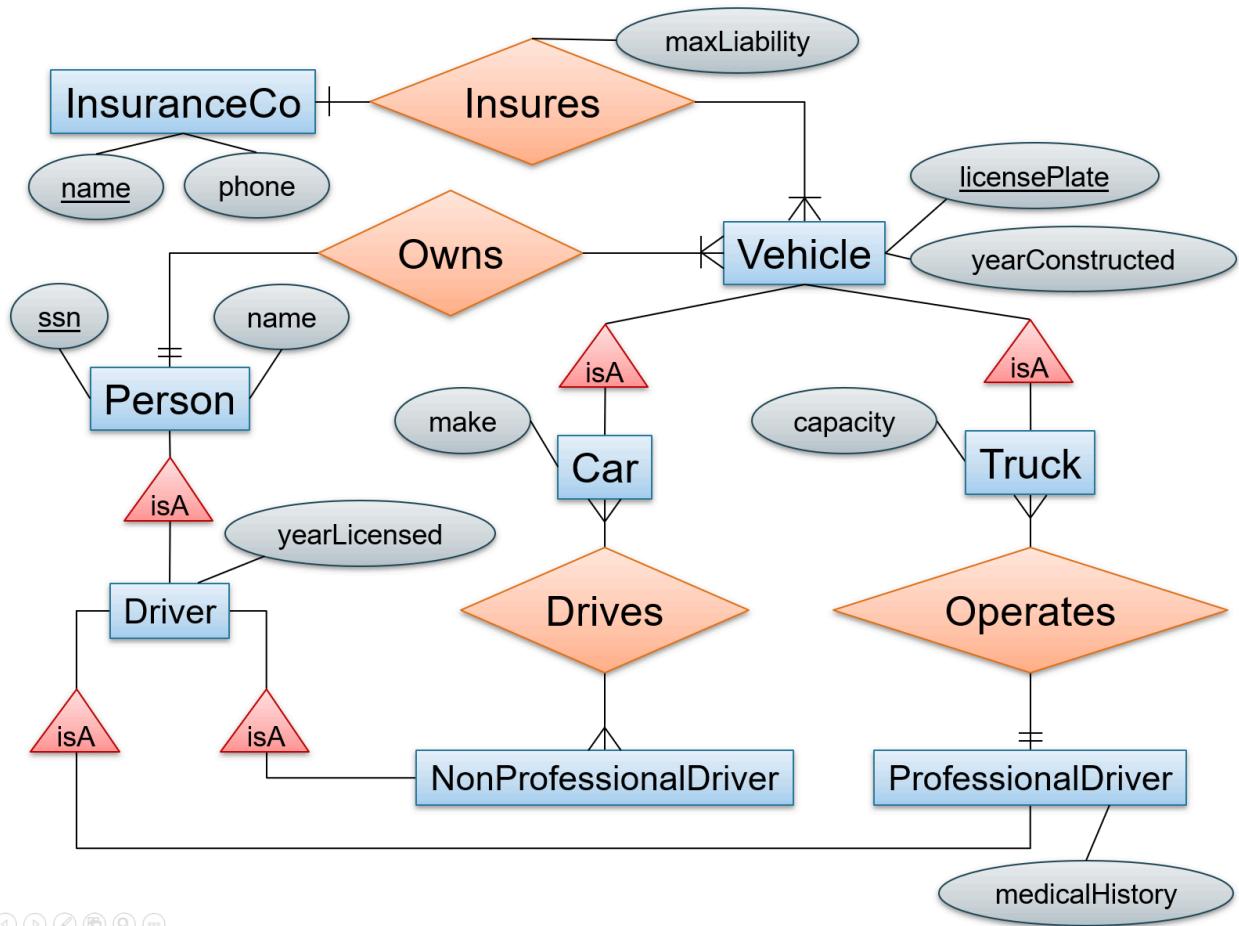
Alternatively, you can draw your diagrams using a favorite drawing tool, such as PowerPoint, Keynote, or draw.io. We do not recommend Google Slides because it lacks a few shapes that you might need (eg, rounded arrows), but you can work around those omissions if you are committed (eg, you can use a crescent and a line, or simply type a constraint label such as "=1".)

Name your file **geography.pdf** and upload to Gradescope.

Question 3: E/R to Schema (20 points)

Consider the following E/R diagram. All attributes in this diagram are alphabetic strings, with the following exceptions:

- `licensePlate` and `medicalHistory` have letters and digits
- `ssn` only contains digits
 - In other words, can be represented as integers
- `maxLiability` is a real (floating point) number
- `yearConstructed`, `yearLiscensed`, `phone`, `capacity` are integers



1. (10 points)
Translate the above E/R diagram into a series of SQL `CREATE TABLE` statements, including all key constraints (ie, primary and foreign keys). We will run these statements, so ensure that the statements are syntactically correct (ie, check them using SQLite, SQLServer, or another relational database).
2. (5 points)
Which relation in your schema represents the relationship `insures` in the E/R diagram? Why is that your representation?
3. (5 points)
Compare the representation of the relationships `drives` and `operates` in your schema, and explain why they are different.

Question 4: Mr. Frumble: Relationship Discovery & Normalization (35 points)

[Mr. Frumble](#) (a great character for small kids who always gets into trouble) designed a simple database to record the projected monthly sales in his small store. He never took a database class, so he came up with the following schema:

```
Sales (name, discount, month, price)
```

He quickly realized it was difficult to update. Fortunately, he has hired *you* to fix his data management problems! He gives you his file [mrFrumbleData.csv](#) and says: "Fix it for me!".

You decide to help Mr. Frumble by normalizing his database. Unfortunately, he has already dashed off (*hats don't chase themselves, doncha know?*) and you are unable to ask him what functional dependencies make sense in his business. Instead, you will reverse engineer them from his data instance.

1. Create a table in a DBMS of your choosing, and load Mr. Frumble's data into that table; use SQLite or any other relational DBMS of your choosing.

Submit your `CREATE TABLE` statement. Do not include the data import statements.

2. Find all non-trivial functional dependencies in his schema.

This is a *reverse engineering task*, so expect to proceed in a trial and error fashion. We recommend first searching for simple dependencies (eg, `name → discount`), and only searching for more-complex dependencies (eg, `name, discount → month`) as needed. For each candidate FD you consider, you will need to write a SQL query to verify whether it does or does not hold.

A few notes:

- a. Your query's answer should always be short
 - i. A good rule of thumb is no more than ten results
 - ii. Remember that 0 results can be instructive as well
- b. You should be able to determine whether a candidate FD does or does not hold by examining the query's answer. Don't miss relevant dependencies but, if you can, try to minimize the number of queries you write.
 - i. For example, if you have queries demonstrating `A → B` and `C → D`, you do not need to write a query for `AC → BD` (but don't forget to find all non-trivial FDs!).
- c. Write a SQL query for each functional dependency you find, along with a comment describing the FD.

- d. Also write a SQL query for at least one functional dependency that *does not exist in the dataset* along with a comment mentioning that the functional dependency does not exist.
- e. Not only should your query's results be short, the queries themselves should be short.

Submit one SQL query for each FD you identify. For example: if you find 3 functional dependencies, submit 4 FDs and 4 queries. You should submit 3 queries showing that your 3 FDs hold over this instance, and a fourth FD along with a query demonstrating how it does not hold on this instance.

3. Decompose Mr. Frumble's single table into Boyce-Codd Normal Form (BCNF), and create SQL tables for the decomposed schema. Create keys and foreign keys where appropriate.

Submit the SQL commands for creating the tables.

4. Submit the SQL queries that *load data* into your new tables. In a SQL comment, count the size of each of your new tables after they've been loaded (ie, the result of `SELECT COUNT(*) FROM MyNewTable`).

Question 5: Further Relationship Discovery (15 points)

Recall that functional dependencies are derived from relationships in existing datasets (as we saw in the previous part with Mr. Frumble) but can also be derived from our knowledge of the domain. For each of the following functional dependencies, please state whether you believe it *holds* or *does not hold* and then state to which domain your answer applies. Lastly, if you believe it holds, please give an example tuple demonstrating this fact; if you believe it does not hold, please give a tuple demonstrating that instead.

Three examples follow:

- `Species -> NumberOfEyes` **holds** over the domain of all known species on earth; if you know that the species is `Black Widow Spider`, then you know that it has 8 eyes.
- `NumberOfEyes -> Species` **does not hold** over the domain of all known species on earth; if you know that someone has 2 eyes, then it could be a `Human` or a `Husky Dog`. Thus, the number of eyes does not determine the species
- On the other hand, `NumberOfEyes -> Species` **holds** over the domain of all UW students; if you know that someone has 2 eyes and is also a UW student, then their species is `Human`.

1. (5 points)
`UnitedStatesZipCode -> State`
2. (5 points)
`TelephoneAreaCode -> State`
3. (5 points)
`IPAddress -> State`

Question 6: Functional Dependency Theory (20 points)

Consider the following two relational schemas and sets of functional dependencies:

1. (10 points)
 $\mathbf{R}(A, B, C, D, E)$ with functional dependencies $D \rightarrow B, CE \rightarrow A$.
2. (10 points)
 $\mathbf{S}(A, B, C, D, E)$ with functional dependencies $A \rightarrow E, BC \rightarrow A, DE \rightarrow B$.

Decompose both of the above schemas into BCNF. At each step of the decomposition, explain which bad FD you are correcting. For each table in your final answer, please note which (if any) of the original FDs still apply to that table.

Submission Instructions

You should answer the questions on [Gradescope](#) under HW4. Please follow the prompts on Gradescope to ensure that you answer all the questions; this document provides additional explanation and materials. You can resubmit an arbitrary number of times until the due date, which will save your progress without submitting all answers.