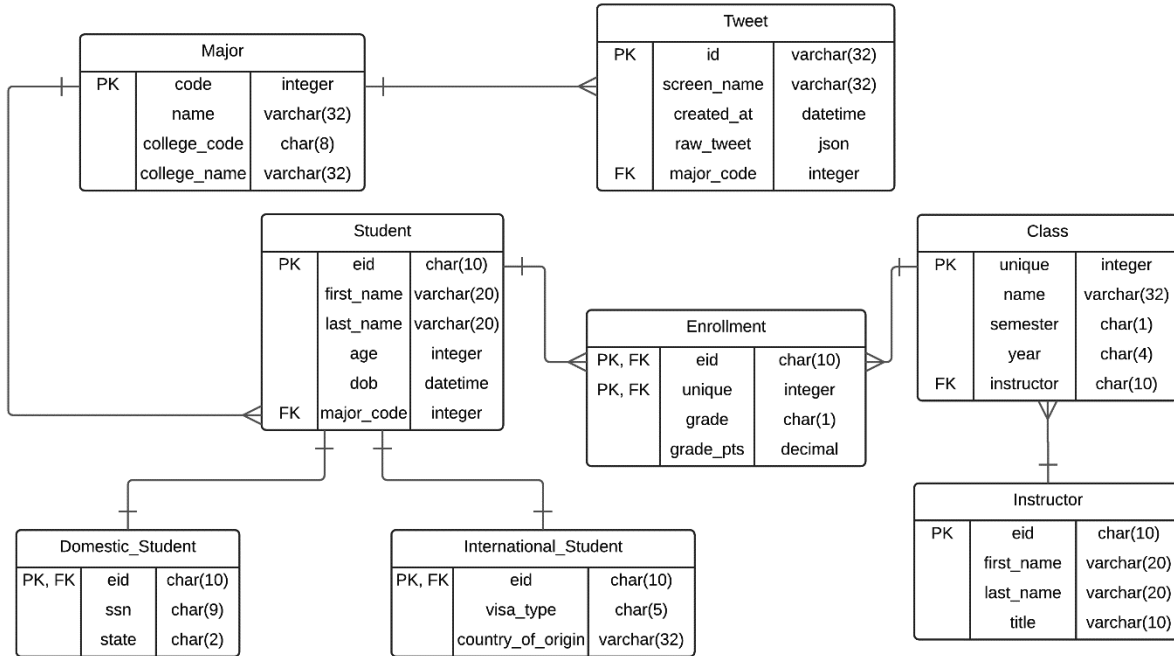


Practice Exam - CS 327E Fall 2016

This practice exam is based on latest rendition of the *utexas* database, a simple database of UT students, instructors, classes, enrollments, and most recently, majors and tweets.

Below is the logical ERD for this database:



Logical ERD - utexas Database - CS 327E Fall 2016

For convenience, each table in the database is also listed below. Note that the underlined fields are the primary keys for each table, while the **bold** fields are the foreign keys for each table.

Major(code, name, college_code, college_name)
 Tweet(id, screen_name, created_at, raw_tweet, **major_code**)
 Student(eid, first_name, last_name, age, dob, **major_code**)
 Domestic_Student(eid, ssn, state)
 International_Student(eid, visa_type, country_of_origin)
 Instructor(eid, first_name, last_name, title)
 Class(unique, name, semester, year, **instructor**)
 Enrollment(eid, unique, grade, grade_pts)

Answer the questions about this database and on the following pages. You may remove this page from the practice exam for reference if that is convenient.

Part 1: Database Design:

Question 1. Is this schema in 1NF? Recall that to be in 1NF, every cell must have a scalar value. If the answer is no, identify which fields violate 1NF. Use the notation *table_name.column_name* to specify any violations found.

Question 2. Is this schema in 2NF? Recall that to be in 2NF, the schema must be in 1NF and every non-key field must be dependent on the **entire** key. Use the notation *table_name.column_name* to specify any violations found.

Question 3. Is this schema in 3NF? Recall that to be in 3NF, the schema must be in 2NF and every non-key field must be dependent **only** on the key. Use the notation *table_name.column_name* to specify any violations found. If you found any new violations, please provide a brief explanation.

Question 4. One weakness of the current schema is that a student can only have one major. Suppose we wanted to represent double-majors, that is allowing a student to have up to two majors. Make the necessary changes to the schema to allow for double-majors. Specify the new table definitions below using the same table notation as above. Note: you only need to provide definitions for any new or changed tables.

Question 5. Suppose we wanted to represent the pre-requisites for a class. That is, a class can have zero or more pre-requisite classes. If a class has a pre-requisite, then it also has minimum grade to satisfy the pre-requisite. Make the changes to the schema to model pre-requisites classes and specify the new table definitions below. Again, please use the same table notation as above for your answer. Note: you only need to provide definitions for any new or changed tables.

Part 2: SQL:

The following questions are asking you to translate the description to a valid SQL statement.

Question 1. Add a new Class to the database with the following details: unique = 98765, name = Data Science, semester = S, year = 2017, instructor = jeffwu. Assume that this instructor already exists in the Instructor table.

Question 2. Change the Data Science class to name = Data Science and Prediction and semester to null.

Question 3. Delete the Data Science class from the database.

Question 4. Delete all tweets for major_code = CS and created_at = 12/01/2016.

Question 5. List the distinct class names in the database. Sort the results by class name.

Question 6. List the eids and names of all domestic students who are from out-of-state. Sort the results by eid.

Question 7. List the eids and names of all the Math majors who are currently enrolled in the class = Elements of Databases. Assume that "currently enrolled" means a null grade in the Enrollment table and Math majors have the major_code 'math'. Sort the results by eid.

Question 8. Based on your answer in Part 1, find the most common pre-requisite classes in the database. Your answer should return the unique of the pre-requisite class as well as the number of classes that use it as a pre-requisite. Return the top 10 results.

Question 9. List the eids and names of all instructors who have **not** taught a class. The results should be sorted by the eid of the instructor.

Question 10. Find the instructors who have taught more than 3 classes in 2016. Return the eids and number of classes taught. Sort the results by the number of classes in descending order.

Question 11. Create a view that shows the number of enrollments for each student in 2016. The view should return the eid and enrollment count.

Question 12. Using the view from the previous question, return the minimum, maximum, and average number of enrollments for 2016. The query should not count those students who have zero enrollments this year.