

## 67-262 Database Design and Development, Fall 2016

---

### Course Logistics

Class: Tuesday / Thursday, 1:30 - 2:50 / DH 1112  
Labs: Tue 7pm; Wed 1pm and 7pm; Thu 7pm  
Information Systems room, Wean 5336  
Website: [www.andrew.cmu.edu/~sraja/67262](http://www.andrew.cmu.edu/~sraja/67262)  
Instructor: Raja Sooriamurthi, Ph.D.  
Office: Porter Hall 224D  
Email: raja@cmu.edu

*Data is the new oil.  
It's valuable, but if unrefined it cannot really be used.  
It has to be changed into gas, plastic, chemicals, etc. to create  
a valuable entity that drives profitable activity; so must data  
be broken down, analyzed for it to have value.  
... Clive Humby*

### Course Description

We currently live in a data driven age. Data has emerged as the new oil that drives an organization: The successful operation of modern organizations relies on the effective use of their operational data. Database management systems (DBMS) are the engines of this data driven world.

Data collected and used by an organization is broadly divided into two types (i) line of business data and (ii) customer behavior data. Traditionally data management has focused on line of business data. For example, when a ride request is made to a ride sourcing company (Uber, Lyft, etc.), what data is needed to meet that request? When a purchase is made in a grocery store what is the flow of data during that transaction? Line of business data is used to support the core business processes of the organization. Alternatively, based on the purchase patterns of a shopper or the volume or location of ride requests, how can a grocery store or a ride sourcing company make their operation more effective? The answer to this question is based on customer behavior data. Whatever type of data it may be, many fundamental questions are the same: How do you gather, organize, curate, and process data to help run the organization or to extract actionable information to increase their effectiveness?

In this course we will study tools and techniques for managing data with database systems. At the highest level we will study two questions (i) how to use a database system and (ii) how to build a database system. For the past three decades the relational model has been the predominant model of data management. Most of the course will focus on the classic relational model. In the past several years, driven by the evolving functional and non-functional (quality) needs of an organization, alternatives to the classic relational model have emerged. We will also examine illustrative samples of these popular alternatives known as non-relational or NoSQL models.

We will sample course themes from the following:


1. Relational Model: table, rows, columns, keys
2. SQL 1: SELECT, FROM, WHERE, ORDER BY, DISTINCT, LIKE, BETWEEN, IN
3. SQL 2: JOIN, GROUP BY, HAVING, sub-queries, ANY, ALL, UNION
4. Programming interface to a database (e.g., psycopg2, SQLAlchemy, Flask)
5. Concurrency and Recovery: transactions
6. Integrity constraints: entity, referential, check, functions, triggers
7. Security: views, authorization, SQL injection
8. DB Design: functional dependencies, normalization, ER models
9. Non-relational models (e.g., Redis, MongoDB)

The exact themes we discuss and the depth to which we discuss them will depend on the pace of the course. As the course progresses the online schedule of the course, available off of the course website, will be updated.

## 67-262 Database Design and Development, Fall 2016

### Text Book

There are a wide range of books on various aspects of database systems (from abstract theory to the details of a specific implementation). In addition to class discussions, we will use the required textbook *Database Systems: A Practical Approach to Design, Implementation, and Management (6<sup>th</sup> edition)* by Connolly and Begg. This is a comprehensive book with an applied perspective. It walks the reader through multiple aspects of a database driven system from requirements, to design, to implementation, to security and beyond. While there are numerous resources available on the web, a textbook cohesively ties various concepts together with running examples, review questions, and problems. Periodic readings will be assigned from this text. Hard copies of the book are available in the CMU bookstore (as well as, naturally, online stores). Depending on your reading preference, you can also consider renting (or purchasing) an e-version of the book from the publisher, Pearson, at <https://www.vitalsource.com/referral?term=9780133869750>



I find the e-version of the book to be well done: hyperlinked, each theme (few paragraphs) on individual pages, easy navigation, exclusive online appendixes included with the main text, web access as well as a mobile app to support offline reading etc. I personally prefer the e-version to the hard copy.

### Learning Objectives

Upon successful completion of this course, students will have achieved the following learning objectives.

- Given an organizational need, identify the data needed to meet that need
- Build an ER data model and implement it with SQL
- Manipulate and extract information from a database using SQL
- Understand advanced database concepts and alternatives to the classic relational model

Performance in the course will be assessed along the adjacent components.

Component	Weight
Labs	10
Attendance + Quizzes	5
Assignments	15
Project	20
Exam 1	15
Exam 2	15
Final Exam	20

### Class Policies

**Attendance and Preparation for Class:** To fully engage in classroom discussions, you are expected to attend all class sessions and come prepared for each class. Class participation contributes towards the final grade assessment. There will be in-class assignments and occasionally short quizzes at the beginning of class. Students who have an unexcused absence or tardiness will not be able to make up these assignments and quizzes. Unexcused absences can reflect upon your grade. In the event of a situation requiring you to be absent (e.g., job interview) please contact the professor in advance.

**Laptops:** This is a technology-oriented course but there is a time and place to use technology. As the need arises, we will have hands on class sessions where you will need to use your laptop. But in other instances, laptops and other devices (iPads, smart phones etc.) tend to hinder classroom participation and discussions. Hence, unless explicitly stated otherwise, please close or turn off all such devices when in class. In this context, amongst other articles discussing this issue, you may find the Washington Post op-ed piece of David Cole, *Laptops & Learning*, to be interesting (<http://www.washingtonpost.com/wp-dyn/content/article/2007/04/06/AR2007040601544.html>).

## *67-262 Database Design and Development, Fall 2016*

---

*Classroom Etiquette:* Arriving late to class or leaving early disrupts the instructor and the learning environment. Please plan ahead and do everything you can to avoid these situations. Exiting and re-entering the classroom during session is a significant distraction interrupting class flow. If there is an emergency exit and re-enter with a minimum of disruption to the class.

*Labs.* In this course you will gain knowledge and skill in designing, developing, and working with database systems. As with all skills, expertise comes with hands on practice. Practice in this course is provided via guided lab exercises and assignments. It is expected that you will attend your assigned lab session each week. TAs will be available to provide immediate assistance as needed. Each lab will have a deliverable. Ensure that you have your deliverable assessed by your TA before concluding your lab session.

*Assignments and Flex days:* Part of professional behavior is submitting deliverables on time. Due dates of all deliverables (assignments, projects etc.) will be specified when issued and it is expected that assignments will be submitted on time. Equipment failure, last minute printing issues etc. are not valid reasons for missing a deadline. At the same time 'life happens' — you may have to travel for an interview, may fall sick, it may be an extremely busy week etc. To accommodate such situations, each student has 3 flex days. Unless explicitly specified otherwise, you may apply at the max 2 flex days (48 hours) for submitting an assignment beyond the due date. After that, submissions will not be accepted. Please email the professor when you avail of a flex day.

*Academic Integrity:* Unless explicitly stated otherwise, all work needs to be individually done. While it is fine to discuss general ideas, all submitted work must be your own. Sharing of work with another student or using the work of another's when completing your own will result in a grade of zero. Any case of suspected cheating will be brought to the Dean's attention. If you referred to external sources or consulted with others be sure to clearly indicate so. Be sure to familiarize yourself with the University policies on academic integrity <http://www.cmu.edu/policies/student-and-student-life/academic-integrity.html>.

*Reassessment:* If you would like a component of the course (assignment, exam etc.) to be reevaluated, submit your request in writing (email will suffice) explaining in detail why you feel your response needs to be re-assessed. Any reassessment requests need to be submitted within two weeks of the assignment or exam being returned.

*For Students with Learning Disabilities:* If you wish to request an accommodation due to a documented disability, please inform your instructor and contact: Disability Resources, 102 Whitfield Hall 412.268.2013, or by email at: [lpowell@andrew.cmu.edu](mailto:lpowell@andrew.cmu.edu).

## ***67-262 Database Design and Development, Fall 2016***

---

### ***Take care of yourself***

Do your best to maintain a healthy lifestyle this semester by eating well, exercising, avoiding drugs and alcohol, getting enough sleep, and taking time to relax. Despite what you might hear, using your time to take care of yourself will actually help you achieve your academic goals more than spending too much time studying.

All of us benefit from support and guidance during times of struggle. There are many helpful resources available on campus. An important part of the college experience is learning how to ask for help. Take the time to learn about all that's available and take advantage of it. Ask for support sooner rather than later – this always helps.

If you or anyone you know experiences any academic stress, difficult life events, or difficult feelings like anxiety or depression, we strongly encourage you to seek support. Consider reaching out to a friend, faculty or family member you trust for assistance connecting to the support that can help. Counseling and Psychological Services (CaPS) is here for you: call 412-268-2922 and visit their website at <http://www.cmu.edu/counseling/>. Over 25% of students reach out to CaPS some time during their time at CMU.

If you or someone you know is feeling suicidal, call someone immediately, day or night:

*CaPS: 412-268-2922*

*Re:solve Crisis Network: 888-796-8226*

*If the situation is life threatening, call the Police:*

*On campus: CMU Police: 412-268-2323*

*Off campus: 911*

<http://www.cmu.edu/teaching/designteach/design/syllabus/syllabussupport.html>

***Let's have a fun and productive course!***