

CMSI 3520

DATABASE SYSTEMS

<https://dondi.lmu.build/fall2022/cmsi3520>

Fall 2022—Pereira 140
MW 1:45–3:25pm, 4 semester hours
Office Hours: MTW 4:30–6pm,
or by appointment (*don't hesitate to ask!*)

John David N. Dionisio, Ph.D.
email: dondi@lmu.edu
Doolan 102; (310) 338-5782

Objectives and Outcomes

This course introduces the computer science sub-field of *databases*, which is concerned with the theory, design, and implementation of systems that manage large amounts of data. Long after the course concludes, my hope is that you will:

1. Know and understand how databases are designed, implemented, and deployed.
2. Be acquainted with database system theory and algorithms.
3. Apply this knowledge by designing and implementing a database application library for a particular domain.

In addition to the course-specific content, you are also expected to:

4. Follow disciplinary best practices throughout the course.

Prerequisites/Prior Background

Although there are no absolute prerequisites to this course, students will benefit greatly from having taken programming language (386/3801) and operating systems (387/3510) courses. Intermediate to advanced programming proficiency in any language will be helpful, as well as familiarity and experience with the command line interaction style.

Materials and Texts

This course does not have a preassigned textbook, with materials consisting primarily of assorted websites, articles, videos, and sample code to be made available online. However, the following can serve as foundational reading:

- Ramez Elmasri and Shamkant B. Navathe. *Fundamentals of Database Systems*, 7th Edition, Pearson, 2016.

Having the latest exact edition isn't critical, so no problems if an older edition is more cost-effective for you.

Course Work and Grading

Your final grade will be based on the percentage of the points you get for the following deliverables against the total number of possible points:

Database fiddle tutorial	50
Key-value database code-along	50
File database mini-stack	100
Relational database mini-stack	100
Document database mini-stack	100
Graph database mini-stack	100
Full database SDK	200
Total	700 points

Percentages $\geq 90\%$ get an A– or better; $\geq 80\%$ get a B– or better; $\geq 70\%$ get a C– or better. I may nudge grades upward based on qualitative considerations such as degree of difficulty, effort, class participation, time constraints, and overall attitude toward the course.

Term Portfolio

Your accumulated assignments for the semester comprise the *term portfolio*—the final, definitive artifact that demonstrates the course's outcomes. It is how you show whether you have, indeed, accomplished the objectives of this course.

An assignment's number is its due date in *mmdd* format, and it is always due by 11:59:59.999pm of that date. Point values are based on the state of your assignments at that moment.

Your portfolio for this course will consist of the following types of deliverables:

- Tutorial and code-along—introductory forays into basic database operations and implementation on an online “fiddle” system and a code-along in-memory key-value database (*1a–1c, 3c*)
- “Mini-stacks”—case studies where data sets are loaded and managed using a range of database paradigms: modeling, CRUD, and the beginnings of a persistence layer API (*1a–1d, 3a–3d*)

- Full database SDK—a fully-functional persistence layer implementation consisting of a data model, CRUD implementations, and a small demonstration application in a chosen data set's domain (2b, 3a–3d)

Outcomes 4d–4f apply to all assignments; the full range of 4a–4f applies to all assignments that involve programming.

For maximum enthusiasm and interest level, you will be allowed to select a particular application domain (assuming there is a robust data set available for it) around which your assignments can revolve. The hope is that such a selection will allow you to formulate database queries and implement database operations that you will find to be personally interesting and fulfilling.

Version Control

Version control is an indispensable part of today's computer science landscape in industry, the academe, and the open source community. We use version control heavily in this course: make sure that you get the hang of it.

None of the assignments can be completed (well) overnight; they should be the result of steady progress from the moment they are assigned to the date they are due. “One and done” submissions will negatively affect the final score.

Workload Expectations

In line with the *LMU Credit Hour Policy*, the workload expectation for this four-credit-hour course is a minimum of $4 \times 3 = 12$ hours of work by an average student per week, including the time that we spend together in the classroom.

Attendance

Attendance at all sessions is expected in person, but not absolutely required. If you must miss class, it is your responsibility to notify me about this and keep up with the course.

The last day to add or drop a class without a grade of W is September 2. The withdrawal or credit/no-credit deadline is November 4.

Academic Honesty

Loyola Marymount University is a community dedicated to academic excellence, student-centered education, and the Jesuit and Marymount traditions. As such, the University expects all members of its community to act with honesty and integrity

at all times, especially in their academic work. Academic honesty requires that all members of the LMU community act with integrity, respect their own intellectual and creative work as well as that of others, acknowledge sources consistently and completely, act honestly during exams and on assignments, and report results accurately. As an LMU Lion, by the Lion's Code, you are pledged to join the discourse of the academy with honesty of voice and integrity of scholarship.

Academic dishonesty will be treated as an extremely serious matter, with serious consequences that can range from receiving no credit for assignments/tests to expulsion. It is never permissible to turn in any work that has been copied from another student or copied from any source (including the Internet) without properly acknowledging/citing the source. It is never permissible to work on an assignment, exam, quiz, or any project with another person unless your instructor has indicated so in the written instructions/guidelines. It is your responsibility to make sure that your work meets the standard of academic honesty set forth in the “Academic Honesty Policy” found at:

<https://academics.lmu.edu/honesty>

Special Accommodations

The Disability Support Services (DSS) Office offers resources to enable students with ADD/ADHD; physical, learning, and psychiatric disabilities; and those on the autism spectrum to achieve maximum independence while pursuing their educational goals. Staff specialists interact with all areas of the University to eliminate physical and attitudinal barriers. Students must provide documentation for their disability from an appropriate licensed professional. Services are offered to students who have established disabilities under state and federal laws. DSS personnel also advise students, faculty, and staff regarding disability issues. Students who need reasonable modifications, special assistance, academic accommodations or housing accommodations should direct their request to the DSS Office as soon as possible. All discussions will remain confidential. The DSS Office is located on the 2nd floor of Daum Hall and may be reached by email at dsslmu@lmu.edu or phone at (310) 338-4216. Please visit <http://www.lmu.edu/dss> for additional information.

Topics and Important Dates

Correlated outcomes are shown for each topic. Specifics may change as the course progresses. University dates (*italicized*) are less likely to change.

August	Overview; introduction to data sets and database application tiers; hands-on with a database fiddle (<i>1a–1c, 3c</i>)
September	In-memory database code-along (<i>1d, 3c, 3d</i>); file databases (<i>2b</i>); database server setup, initialization, startup and shutdown, and general use (<i>3a–3d</i>); introduction to relational databases and SQL (<i>1b–1c</i>)
<i>September 2</i>	<i>Last day to add or drop a class without a grade of W</i>
<i>September 5</i>	<i>Labor Day; no class</i>
October	Advanced SQL: aggregation, subqueries, JSON fields (<i>1c</i>); document-centric databases (<i>1d, 2b, 3a–3d</i>)
November	Graph databases (<i>1d, 2b, 3a–3d</i>); relational database theory: algebra and calculus; functional dependencies and normalization; integrity constraints (<i>2a</i>); indexing and transactions (<i>2b</i>)
<i>November 4</i>	<i>Withdraw/credit/no-credit deadline</i>
<i>November 23–25</i>	<i>Thanksgiving; no class</i>
December	Database SDK work sessions (<i>1a–1d, 2b, 3a–3d, 4a–4e</i>); miscellaneous topics (<i>time permitting</i>)
<i>December 12</i>	<i>Last set of term portfolio assignments due</i>

You can view my class calendar and office hour schedule in any iCalendar-savvy client. Its subscription link can be found on the course web site (it's too long to provide in writing).

Tentative Nature of the Syllabus

This syllabus and its contents are subject to revision; students are responsible for any changes or modifications announced or distributed in class, emailed to students' LMU Lion accounts, or posted on LMU's course management system.

Course Evaluations

Student feedback provides valuable information for continued improvement. All students are expected to fairly and thoughtfully complete a course evaluation for this course. This semester, course evaluations will be administered online through the Blue™ evaluation system. You will receive an email notification at your Lion email address when the evaluation form is available. You may also access the evaluation form on Brightspace (<https://brightspace.lmu.edu>) dashboard during the evaluation period. Your responses will be anonymous and will not be linked to you in any way.

Course Outcomes

1 Know and understand how databases are designed, implemented, and deployed.

1a	<i>Know and understand the structure of modern database applications.</i>	Modern database applications include not only the database layer itself, but other layers for logic beyond pure data access.
1b	<i>Know and understand the relational database model.</i>	Outcomes 1b and 1c include the ability to interact directly with a “bare” database system—for this course, this is PostgreSQL. Relevant activities include creating tables with appropriate keys and foreign keys, loading these tables with data, performing a variety of queries, and altering a database schema.
1c	<i>Be proficient at database definition and manipulation with SQL.</i>	
1d	<i>Know about alternatives to the relational database model.</i>	For this course, this includes using files directly as a database as well as document-centric and graph databases.

2 Be acquainted with database system theory and algorithms.

2a	<i>Know the central concepts behind relational database theory.</i>	In addition to the data model itself, relational database theory includes the relational algebra, relational calculus, functional dependencies, normalization, and integrity constraints.
2b	<i>Be aware of the primary implementation and performance issues that database systems face.</i>	Issues include transaction management, security, storage, indexing, and query processing & optimization.

3 Apply this knowledge by designing and implementing a database application library for a particular domain.

3a	<i>Install, set up, and manage a database server.</i>	The intent of this outcome is to make you as comfortable with installing and running industrial-strength database servers as you are with more conventional types of applications.
3b	<i>Load data from flat files into a database server.</i>	These outcomes are demonstrated by writing code. Thus, some specific set of technologies, languages, and libraries unavoidably must be learned and used in order to accomplish these. However, it must also be understood that these concepts are general and technology-independent: when called for, one should be able to transfer this knowledge to other platforms.
3c	<i>Perform create, read, update, and delete operations (CRUD).</i>	
3d	<i>Design and implement a programming layer that will perform CRUD as function calls for a particular domain or set of use cases.</i>	

4 Follow disciplinary best practices throughout the course.

4a	<i>Write syntactically correct, functional code.</i>	Code has to compile. Code has to work. No errors, no bugs. Use unit tests as much as possible.
4b	<i>Use programming best practices, demonstrating principles such as DRY, proper separation of concerns, correct scoping of variables and functions, etc.</i>	This is the basis of good software design. It makes software easier to maintain, improve, and extend. Heed feedback well. <i>What you learn here will apply to future work in this field, in school and beyond.</i>
4c	<i>Write code that is easily understood by programmers other than yourself.</i>	This outcome involves all aspects of code readability and clarity for human beings, including but not limited to spacing & indentation, proper naming, presenting code in a manner that is consistent with its structure, documentation & comments when appropriate, and adherence to conventions or standards.
4d	<i>Use available resources and documentation to find required information.</i>	The need to look things up never goes away. Remember also that the course instructor counts as an “available resource,” so this outcome includes asking questions and using office hours.
4e	<i>Use version control effectively.</i>	In addition to simply using version control correctly, effective use also involves appropriate time management, commit frequency, and descriptive commit messages.
4f	<i>Meet all designated deadlines.</i>	