



Data Analytics I: PostgreSQL

Course Syllabus

Copyright © 2018-19 NCLab Inc.

This course can be accessed at: www.nclab.com/learn-sql

What you will learn

“You must ask the right question.” - Isaac Asimov

The language used to create, manage and query (question) databases is SQL, or Structured Query Language. This course will train you in SQL (specifically PostgreSQL). At the end of this course, you will be prepared to “ask the right questions” and perform analysis on meaningful data.

Once you are familiar with SQL commands and structures, you will analyze data using several tools. Data analysis and visualization continues in Data Analytics II: Introduction to Python for Data Science and Data Analytics II: Regression Analysis with Python.

After completing Data Analytics I and II, you will be prepared to use databases, SQL, Python, and statistics to uncover insights, communicate critical findings, and create data-driven solutions.

Recommended background

For this course, you should be comfortable with basic math and algebra operations, number systems, and data functions such as average and sum.

A background in coding is helpful but not required. You will learn how to work with SQL commands within this course. In Data Analytics II, you will learn enough of the Python language to write scripts and use specific Python libraries. If you wish to learn more

about programming, coding, or the Python language, try NCLab’s Karel, Visual Introduction to Python, or any one of several advanced Python self-paced courses.

Contact us

If you have any questions or concerns as you go through the course, please contact NCLab at support@nclab.com.

Course structure and length

The course is self-paced, and you will practice each skill or concept as you go. Automatic feedback is built into the course for both practices and quizzes.

The course is divided into eight Units, and each Unit is composed of five Sections. Each Section consists of 7 instructional/practice levels, a quiz, and a master (proficiency) level. You can return to any level or quiz for review.

Unit 1	Unit 2	Unit 3	Unit 4	Unit 5	Unit 6	Unit 7	Unit 8
--------	--------	--------	--------	--------	--------	--------	--------

Unit 1 (Each unit is composed of five sections)																	
Section 1 Section 5										
1.1	1.2	1.3	1.4	1.5	1.6	1.7	1.8	1.9	5.1	5.2	5.3	5.4	5.5	5.6	5.7	5.8	5.9
Instruction and practice							Quiz	Master	Instruction and practice						Quiz	Master	

While learning skills in the Data Analytics course, you can practice queries and create portfolio artifacts with NCLab’s PostgreSQL and Python apps. Use a project idea from NCLab or create your own. NCLab provides several practice databases that can be queried with the PostgreSQL app.

Data Analytics I is designed as a semester length course (3 credit hours). Since the course is self-paced, the amount of time required to complete the course will vary from student to student. You are responsible for learning both the tutorial content and the skills acquired through practice. Upon successful completion of each Unit, you will receive a certificate from NCLab listing the skills learned. NCLab partners with certain colleges who can provide certification at the college level. Official certification will require a final exam.

Unit 1 Introduction to Data, Databases and SQL

Section 1

- Define and explain the purpose of data and the most widely used data types.
- Understand the hierarchy of units used to calculate data size.
- Calculate the size of your own data.
- Define and explain the purpose of databases.
- Understand relational vs. non-relational databases.
- Use design and ethics principles to guide database use, including referential integrity and ACID.
- Learn about the Structured Query Language (SQL) and its history.
- Understand the main differences between various SQL flavors: PostgreSQL, MySQL, and SQLite.
- Learn further details about PostgreSQL, which will be used in this course.

Section 2

- Examine the structure of a PostgreSQL database.
- Use the SELECT statement to display all columns of a table.
- Understand that SQL is case-insensitive.
- Understand the purpose of the keyword NULL and SQL's three-valued logic.
- Order query results using ORDER BY.
- Limit the number of results using the keyword LIMIT.

Section 3

- Use the SELECT statement to only display one column of interest.
- Use the SELECT statement to display two or more columns.
- Change the order of the displayed columns if needed.
- Skip some rows in the output using the keyword OFFSET.
- Order multiple columns at the same time.

Section 4

- Filter the results of queries using the keyword WHERE.
- Access a table in a schema by typing schema.table.
- Remove duplicate results with SELECT DISTINCT.

Section 5

- Use aggregate functions:
 - Function COUNT() to count results.
 - Functions MIN(), MAX() to find minimum and maximum values.
 - Functions SUM(), AVG() to calculate the sum and average of values, respectively.
- Narrow down queries using WHERE ... LIKE ... and WHERE ... BETWEEN ...
- Combine multiple conditions in the search using the keywords AND, OR, NOT.

Unit 2 SQL Queries

Section 6

- Practice basic SQL queries by exploring the Northwind database and in particular the table Employees.

Section 7

- Practice basic SQL queries by exploring the table Products of the Northwind database.

Section 8

- Practice basic SQL queries by exploring the table Products of the Northwind database.
- Use SQL for calculations
- Use the WHERE ... IN ... clause.

Section 9

- Practice basic SQL queries by exploring the table Order Details of the Northwind database.
- Create new columns and name them using the keyword AS
- Group results for better readability using the GROUP BY ... and GROUP BY ... HAVING ... clauses.

Section 10

- Practice basic SQL queries by exploring the Customers table of the Northwind database.

Unit 3 Creating and Managing Tables

Section 11

- Explain the different data types in the SQL standard and in PostgreSQL.
- Create and drop schemata.
- Learn about the two schemata in NCLab where you can store your own data.
- Create tables.

Section 12

- Insert complete rows (without using column names).
- Insert incomplete rows (without using column names).
- Insert incomplete rows (using the names of columns).

Section 13

- Insert multiple rows at once.
- Create a new table by copying an existing table.
- Create a new empty table which has the same structure as an existing table.
- Create a new table by copying selected rows from an existing table.

Section 14

- Insert selected rows from a table into an existing table.
- Use the powerful statement ALTER TABLE to modify tables, such as:
 - rename tables,
 - add / rename / drop (= delete) columns,
 - change data types of columns, etc.
- Delete all rows from a table, selected rows, or entire tables.
- Delete selected rows from a table.
- Delete entire tables.

- Understand that the result of the VALUES clause and of the SELECT statement is a TABLE.

Section 15

- Define constraints and specify default values.

Unit 4 Joining Tables

Section 16

- Combine data from two different tables using the inner join operation.
- Initially, the tables are matched based on a shared column of the same name.
- Write inner joins in three simple steps:
 - Basic SELECT query followed by the columns one wants to display,
 - Add INNER JOIN followed by the second table name,
 - Add USING followed by the name of the shared column in parentheses.
- Join three and more tables, and that joining multiple tables is equally simple as joining just two.

Section 17

- Review the structure of the tables in the schema World.
- Perform an inner join of tables based on columns of different names, using the keyword ON.
- Abbreviated table names using the keyword AS.
- Understand that some keywords (for example, AS) can be omitted.
- Write inner joins using an alternative (implicit) syntax without the keywords INNER JOIN and USING/ON.

Section 18

- Practice inner joins by solving practical tasks related to the schema World.

Section 19

- Review the schema Northwind, and then perform calculations which require combining data from various tables in this schema.
- Combine inner joins with filtering, grouping and other basic SQL techniques.

Section 20

- Practice inner joins performing calculations which combine data from various tables in the schema Northwind.
- Combine inner joins with filtering, grouping and other basic SQL techniques.

Unit 5 Conditional Expressions

Section 21

- Understand the difference between a statement and an expression.
- Insert "intelligent" conditional expressions into your queries using the keyword CASE.
- Use both the "searched version" of the CASE expression which is more versatile, and the "simple (switch) version".
- Combine a CASE expression with an inner join.

Section 22

- Learn about additional important applications of conditional expressions including:
 - How to use them to split values into different columns.
 - How to count non-NULL values in the columns of a SELECT statement.
 - How to use conditional expressions in aggregate functions COUNT, SUM, MIN, MAX, AVG.
 - Understand the difference between using conditional expressions and the WHERE clause.

Section 23

- Simplify filtering values in aggregate functions using the FILTER clause.
- Learn how integer division works in SQL, and how to prevent potential problems which it may cause.

Section 24

- Learn about various situations which can cause an error in SQL, and how to prevent them using suitable CASE expressions. In particular, learn how to prevent division-by-zero errors.
- Learn how SQL responds when NULL values are combined with numbers and text strings.
- Use CASE expressions in the ORDER BY clause.

Section 25

- Replace NULL with selected values using the function COALESCE
- Replace selected values with NULL using the function NULLIF.
- Learn about the relation between these functions and the conditional expression CASE.
- Replace values in tables by combining the COALESCE and NULLIF functions.

Unit 6 Text Strings

Section 26

- Using the Museum of Modern Art (MoMA) database, review pattern matching and text string operations.
- Understand the advantages of performing case-insensitive searches.
- Use the functions LOWER or UPPER to make a text search case-insensitive.
- Refine searches using the PostgreSQL clause ILIKE, and the shortcuts ~~, ~~, !~, and !~*.
- Search for the special characters % and _, and for the escape character \.
- Become familiar with the official PostgreSQL documentation.

Section 27

- Search for single quotes ' and use the clause SIMILAR TO with entry-level regular expressions.
- Understand that SIMILAR TO is case-sensitive but can be made case-insensitive in combination with LOWER or UPPER (same as LIKE).
- Use regular expressions with groups of characters or text strings enclosed in parentheses and separated by the | character.

- Escape characters when using SIMILAR TO to search for parentheses
- Use the character class wildcard [...] in regular expressions.

Section 28

- Use regular expressions with a pattern which is repeated (exactly) N times.
- Use regular expressions with a pattern which is repeated at least N times.
- Understand that repeating patterns may contain wildcards and other repeating patterns.
- Use the shortcuts *, + and ? when working with repeating patterns.

Section 29

- Use the ASCII table and its Unicode extension, as well as the functions ASCII and CHR.
- Understand that both the ASCII and CHR functions work with Unicode.
- Use Unicode and the web when searching for text strings with special characters.
- Use the concatenation operator || and the cast operator :: .
- Use the functions SUBSTRING, POSITION, LEFT, RIGHT, REVERSE and INITCAP.
- Use at least two different methods to search for substrings.

Section 30

- Modification of text strings via search-and-replace. There are
- many applications for this, related to cleaning, adjusting, extracting, combining and presenting data.
- First we will show you two entry-level functions OVERLAY and REPLACE.
- Then we will bring your search-and-replace skills to a new level with POSIX regular expressions.
- You will learn how to use POSIX regular expressions in the SELECT statement, in the functions
- REGEXP_REPLACE and SUBSTRING, and in the WHERE clause.

Unit 7 Sets and Subqueries

Section 31

- Distinguish sets, relations and bags in SQL

- Understand that tables are relations and not ordered.
- Understand why ORDER BY should be used before the clauses LIMIT and OFFSET.
- Find and count duplicate values in columns and duplicate rows in tables.

Section 32

- Use the UNION and UNION ALL clauses in queries.
- Understand that the UNION clause produces a set with no duplicates, while the UNION ALL clause produces a bag which may have duplicates.
- Understand table compatibility with UNION and UNION ALL.
- Combine rows with UNION and UNION ALL.
- Write queries using UNION for qualitative analysis and UNION ALL for calculations.

Section 33

- Write subqueries to access data in multiple tables at the same time.
- Use subqueries in combination with set operations.
- Design subqueries to be simple and not reach into the same table as the main query.
- Use subqueries with the FROM and WHERE ... IN ... clauses.
- Understand how to alias results.
- Use subqueries with aggregate functions.

Section 34

- Use the set theory clauses INTERSECT and INTERSECT ALL.
- That the intersection operation is associated with the keyword AND.
- Understand how set theory clauses in SQL affect computing complexity.
- Use the INTERSECT clause in subqueries.
- Understand the similarities and differences between set theory clauses, subqueries and joins.

Section 35

- Use set difference in queries.
- Use the related clauses EXCEPT and EXCEPT ALL.
- Practice subqueries together with INNER JOINS.

Unit 8 Advanced Joins

Section 36

- Review the INNER JOIN.
- Use the LEFT(OUTER)JOIN in queries.
- Understand how LEFT(OUTER) JOIN relates to the INNER JOIN.
- Use the LEFT JOIN to find orphaned or missing data.
- Use subqueries to work with orphaned or missing data.

Section 37

- Use the RIGHT (OUTER) JOIN and the FULL (OUTER) JOIN in queries.
- Understand when to use the INNER, LEFT, RIGHT, and FULL JOIN.
- Rewrite a RIGHT JOIN as a LEFT JOIN.

Section 38

- Understand the CROSS JOIN as the Cartesian product of the rows of the two tables.
- Write queries using the CROSS JOIN and compare with the INNER JOIN, which is a special case of the CROSS JOIN.
- Write queries with self joins (CROSS JOIN of a table with itself).
- Use self-joins to establish a hierarchy among items in a table, such as among employees.
- Create complex queries.

Section 39

- Work with semi-joins between two tables.
- Understand how matching values work in left and right semi-joins.
- Work with anti-joins, and understand how they return values that do not match.
- Practice INNER JOIN, CROSS JOIN, and subqueries.

Section 40

- Create functions using the CREATE FUNCTION and CREATE OR REPLACE FUNCTION clauses.
- Specify the returned data type using RETURNS.
- Use names or numbers (\$1,\$2etc.) for function parameters.
- Use the keywords IN and OUT to specify input and output parameters.
- Enclose the body of the function in double-dollars \$\$...\$\$ or quotes '...'.
- Use and understand the flags IMMUTABLE, STABLE and VOLATILE.
- Use the flags RETURNS NULL ON NULL INPUT (STRICT), and CALLED ON NULL INPUT.
- Use the keyword SETOF to return a column of values, or a table.
- Create composite data types, and use them in functions.