CS 61:
Database Systems

Basic SQL

Adapted from Silberschatz, Korth, and Sundarshan unless otherwise noted
Augmented by Claude Opus

Agenda

1. Relational algebra
2. Intro to SQL SELECT statement
3. NYC Open Data

A relational database contains related relations ☺

College schema

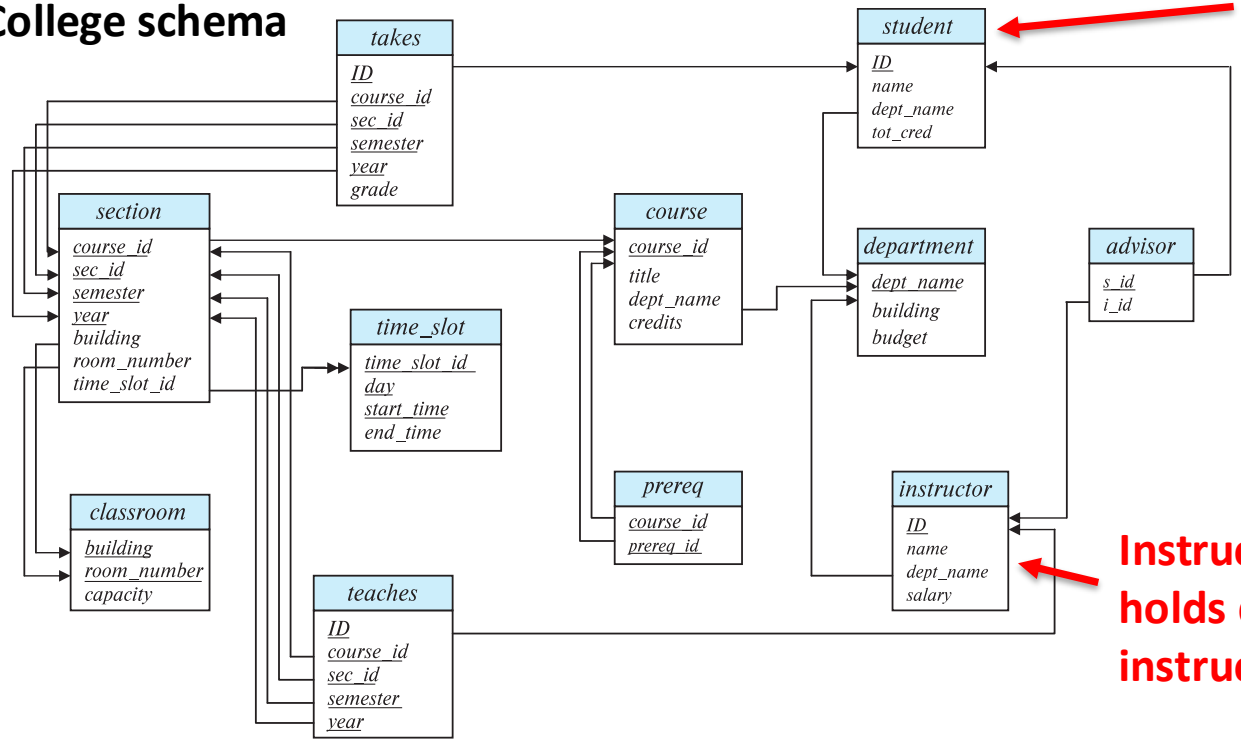

Student relation holds data about instructors

Instructor relation holds data about instructors

Relations

- Each relation holds data about people, places, things or events (nouns)
- Relations (tables) consist of rows and columns
- Each row (tuple) represents one person, place, thing, or event
- Each column represents one attribute about a person, place, thing, or event (e.g., name)
- A column (FK) can refer to a column (PK) in another table, creating a relationship between tables

A relational database contains related relations ☺

College schema

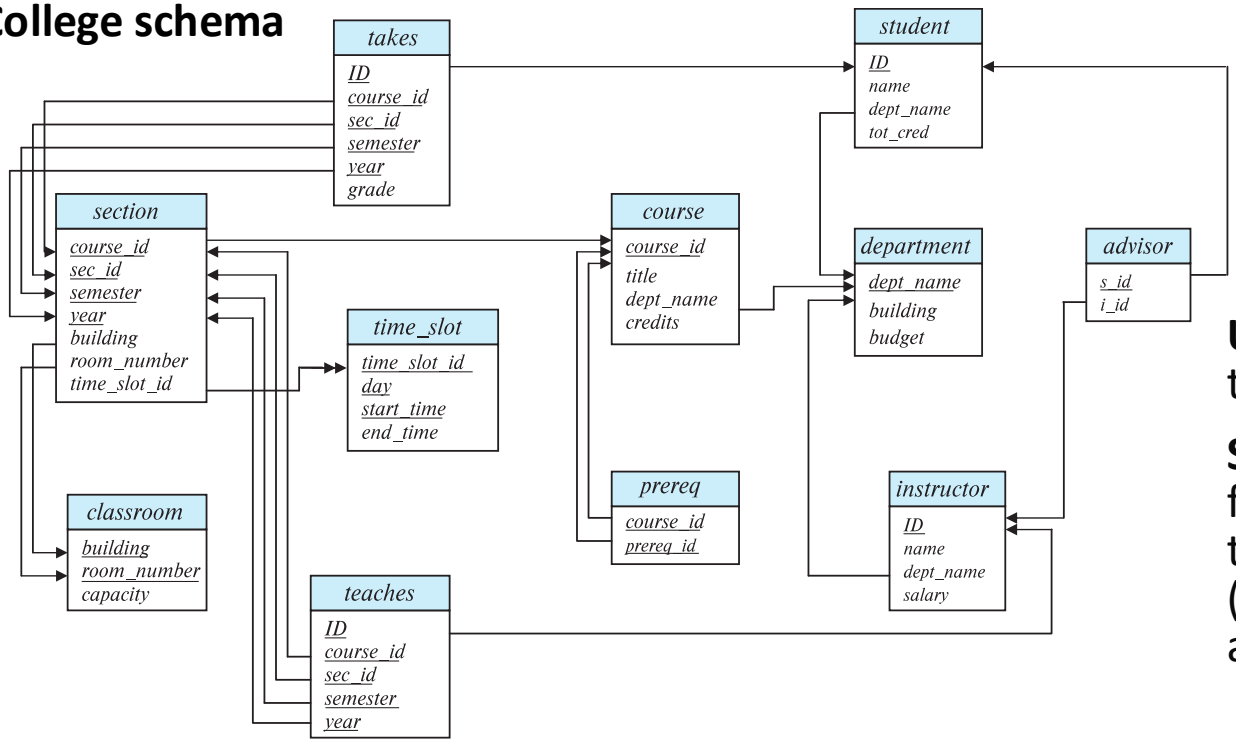

Unique Identifiers: Every table has a primary key

Schema Enforcement: A table follows a set schema, which is the definition of the structure (column names, data types, and constraints)

Set-Based Theory: In formal database theory, a relation is treated as a set, meaning it cannot contain duplicate rows (although SQL can allow duplicate rows unless constrained not to)

No Specific Order: Formally, the rows and columns in a relation are unordered, though in practice, databases return data in specific orders

Atomic Values: Each cell contains only one, indivisible value (atomic value), preventing nested lists or multiple values in a single cell

Relational algebra allows us to work with data in relations (tables)

Mathematically

- Let A_1, A_2, …, A_n be a set of *n* **attributes**
- Let $R = (A_1, A_2, \ldots, A_n)$ be the set of attributes in the schema of relation ***r***
- A relation instance *r* defined over schema *R* is denoted by *r* (*R*)

  Example:

  R = (*ID, name, dept_name, salary*)

  instructor(R) = instructor(ID, name, dept_name, salary)

Implementation

- The current values (**relation instances**) of a relation are specified by a table
- An element ***t*** of relation ***r*** is called a **tuple** and is represented by a table *row*
- Duplicates tuples (rows) are not allowed in a relation (but are in table!)

  If t_1 and t_2 are tuples in r, then $t_1 \neq t_2$

Project: returns a subset of attributes from relation *r*

Project notation: $\Pi_{A_1, A_2, A_3 \ldots A_k}(r)$

project

$\Pi_{ID,\ name,\ salary}(instructor)$

relation *r*

attributes

What ***columns*** of the relation do we want

instructor relation

ID	*name*	*dept_name*	*salary*
22222	Einstein	Physics	95000
12121	Wu	Finance	90000
32343	El Said	History	60000
45565	Katz	Comp. Sci.	75000
98345	Kim	Elec. Eng.	80000
76766	Crick	Biology	72000
10101	Srinivasan	Comp. Sci.	65000
58583	Califieri	History	62000
83821	Brandt	Comp. Sci.	92000
15151	Mozart	Music	40000
33456	Gold	Physics	87000
76543	Singh	Finance	80000

result

dept_name left out

ID	*name*	*salary*
10101	Srinivasan	65000
12121	Wu	90000
15151	Mozart	40000
22222	Einstein	95000
32343	El Said	60000
33456	Gold	87000
45565	Katz	75000
58583	Califieri	62000
76543	Singh	80000
76766	Crick	72000
83821	Brandt	92000
98345	Kim	80000

Project returns a relation

Select: returns tuples from relation *r* that satisfy predicate *p*

Select notation: $\sigma_p(r)$

select

$\sigma_{dept\_name=\text{“Physics”}}(instructor)$

relation *r*

predicate *p*

What *rows* of the relation do we want

instructor relation

ID	*name*	*dept_name*	*salary*
22222	Einstein	Physics	95000
12121	Wu	Finance	90000
32343	El Said	History	60000
45565	Katz	Comp. Sci.	75000
98345	Kim	Elec. Eng.	80000
76766	Crick	Biology	72000
10101	Srinivasan	Comp. Sci.	65000
58583	Califieri	History	62000
83821	Brandt	Comp. Sci.	92000
15151	Mozart	Music	40000
33456	Gold	Physics	87000
76543	Singh	Finance	80000

Select returns a relation

result

ID	*name*	*dept_name*	*salary*
22222	Einstein	Physics	95000
33456	Gold	Physics	87000

- In selection predicate can use:
  =, ≠, >, ≥, <, ≤
- Can combine several predicates:
  ^ (and), v (or), ¬ (not)

Example:

$\sigma_{dept\_name=\text{“Physics”} \wedge salary > 90{,}000}(instructor)$

The result of an operation is a relation, so we can combine them into an expression

Relational algebra expression

project

Select returns relation

relation *r*

$$\Pi_{name}(\sigma_{dept\_name = \text{“Physics”}}(instructor))$$

attribute

predicate *p*

instructor relation

ID	*name*	*dept_name*	*salary*
22222	Einstein	Physics	95000
12121	Wu	Finance	90000
32343	El Said	History	60000
45565	Katz	Comp. Sci.	75000
98345	Kim	Elec. Eng.	80000
76766	Crick	Biology	72000
10101	Srinivasan	Comp. Sci.	65000
58583	Califieri	History	62000
83821	Brandt	Comp. Sci.	92000
15151	Mozart	Music	40000
33456	Gold	Physics	87000
76543	Singh	Finance	80000

result

name
Einstein
Gold

“Find the name of instructors in the Physics department”

- Select eliminated rows we do not want
- Project eliminated columns we do not want
- Operation returns a relation (here with 2 tuples)

Agenda

1. Relational algebra
2. Intro to SQL SELECT statement
3. NYC Open Data

SQL SELECT command has three parts, like relational algebra expression

- SQL SELECT returns a relation with specified attributes from one or more relations with tuples matching provided criteria
- A typical SQL Select query has the form:

SELECT $A_1, A_2, \ldots, A_n$

FROM $r_1, r_2, \ldots, r_m$

WHERE P

Specify the attributes (columns) we want (like Project)
Use * if we want all attributes

From the relations (tables) we want, today only one

Return tuples meeting some requirement (like Select)

SQL SELECT command has three parts, like relational algebra expression

Generic SELECT

SELECT $A_1, A_2, \ldots, A_n$
FROM $r_1, r_2, \ldots, r_m$
WHERE P

Example SELECT

SELECT name
FROM instructor
WHERE dept_name = 'Physics';

Resulting relation

name
Einstein
Gold

Equivalent relational algebra

$$\Pi_{name}(\sigma_{dept\_name = \text{“}physics\text{”}}(instructor))$$

SQL Select command

- Three main parts:
  1. **SELECT** attributes – what columns we want (instructor name)
  2. **FROM** – relation to use (instructor relation)
  3. **WHERE** – criteria for selecting tuples (dept_name = ‘Physics’)
- SQL command capitalization does not matter (Select == select == SELECT)
- Convention is to capitalize SQL commands, but not required
- Relation/attribute names in MySQL not case sensitive but are in some databases
- Use a single quote for strings in SQL (this may bite you, if you get an error: Unknown column “Physics”, it did! Make it 'Physics' (single quote) instead

Agenda

1. Relational algebra
2. Intro to SQL SELECT statement
3. NYC Open Data

NYC is made up of five boroughs

Five NYC boroughs

Manhattan

The Bronx

Queens

Brooklyn

Staten Island

New York makes an incredible amount of data publicly available in NYC Open Data

https://data.cityofnewyork.us/browse?sortBy=most_accessed

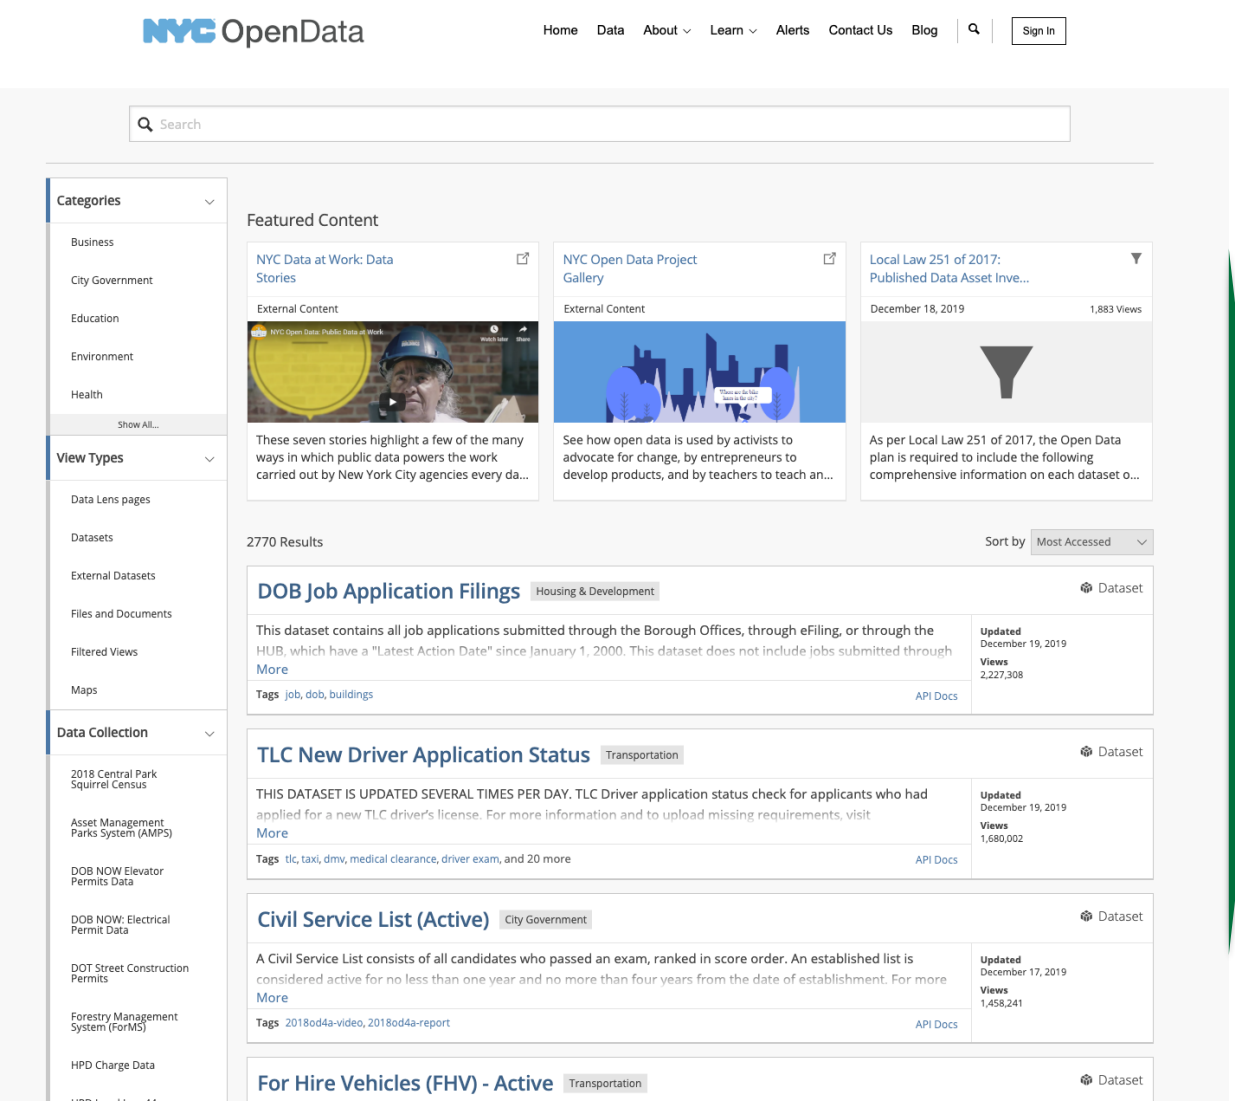

Lots of data collected by NYC (and other cities) is freely available

One data set NYC publishes contains all restaurant health inspections

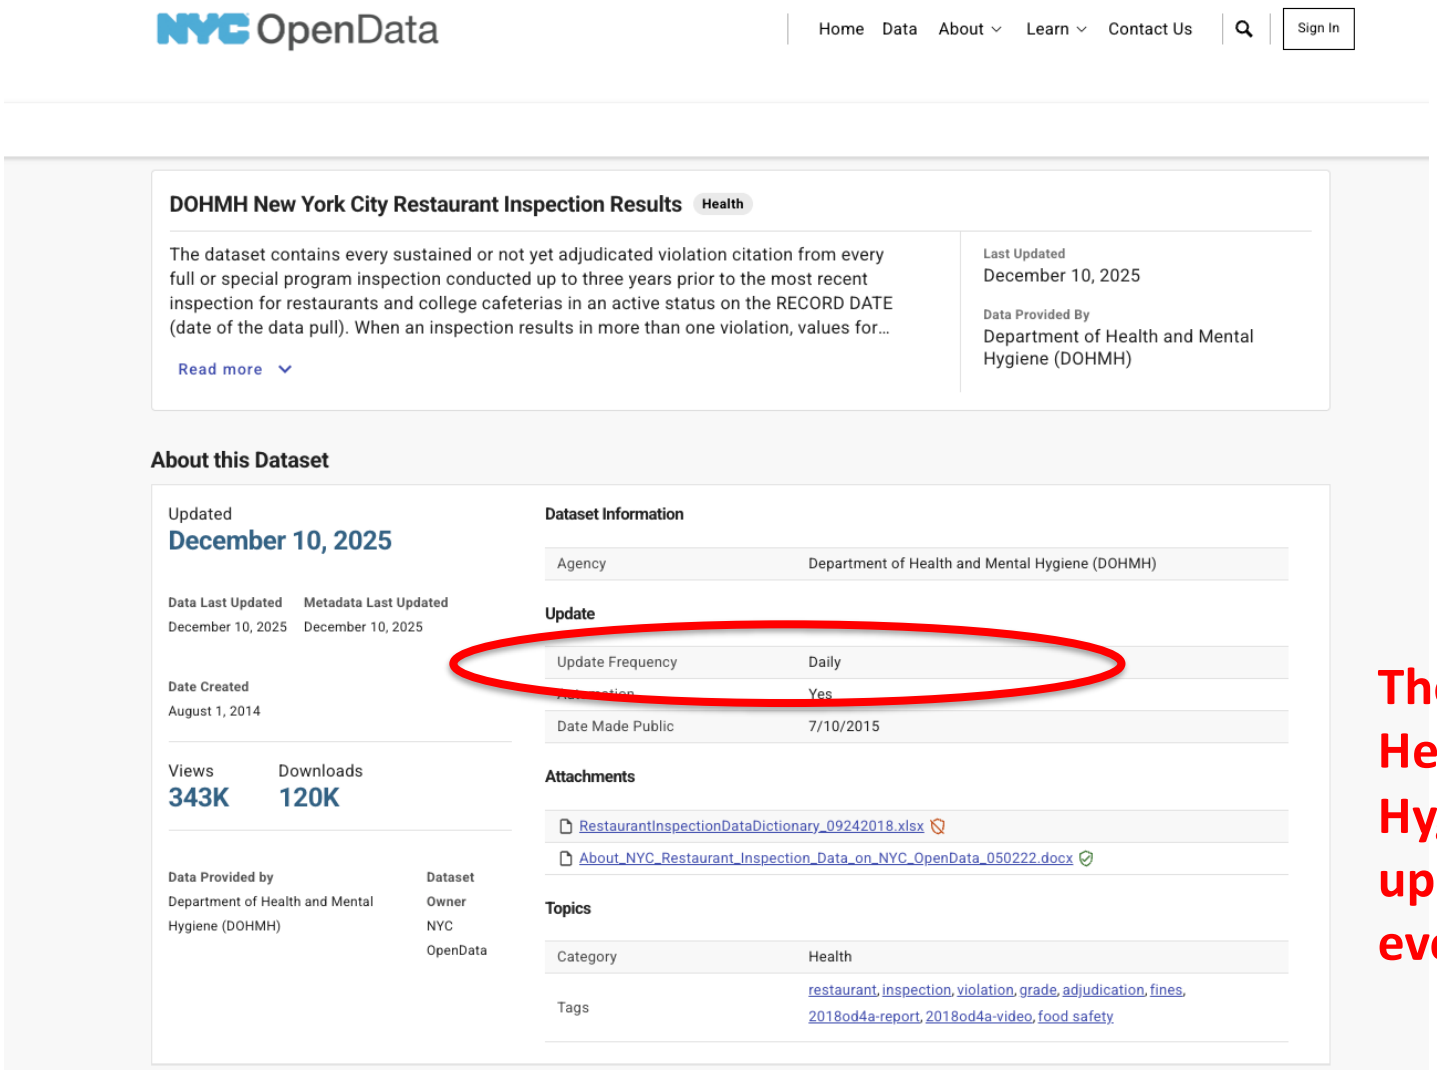

The Department of Health and Mental Hygiene (DOHMH) updates this data set everyday!

Source: https://data.cityofnewyork.us/Health/DOHMH-New-York-City-Restaurant-Inspection-Results/43nn-pn8j

NYC provides a “data dictionary” that describes each column

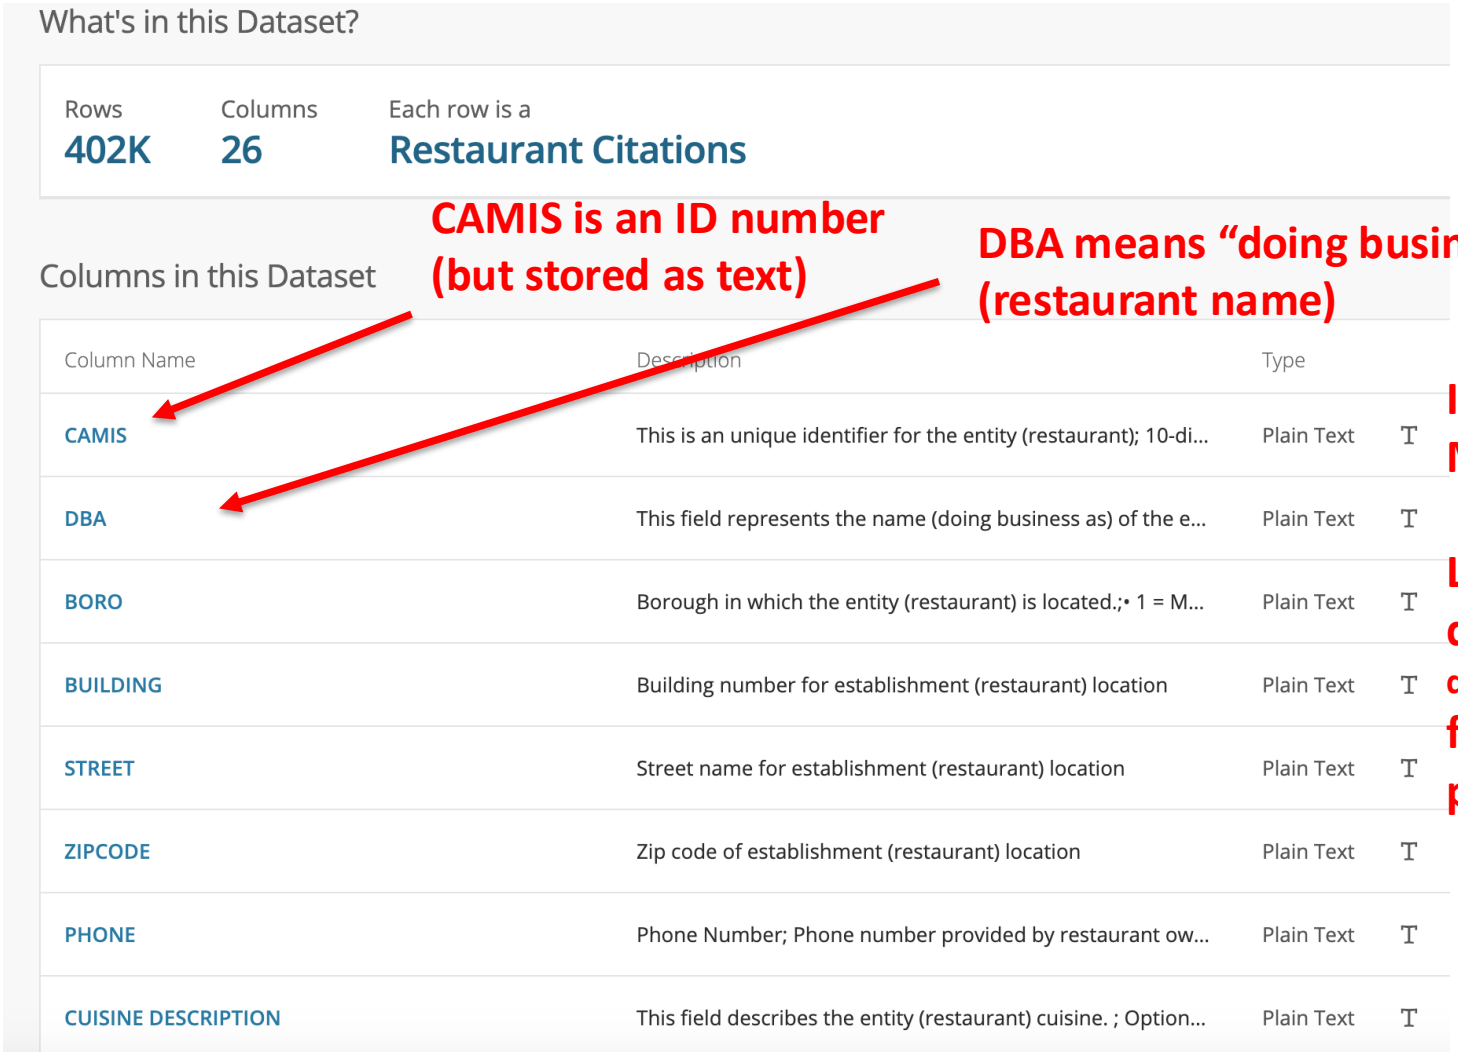

What's in this Dataset?

Rows	Columns	Each row is a
402K	26	Restaurant Citations

CAMIS is an ID number (but stored as text)

DBA means “doing business as” (restaurant name)

Columns in this Dataset

Column Name	Description	Type
CAMIS	This is an unique identifier for the entity (restaurant); 10-di...	Plain Text
DBA	This field represents the name (doing business as) of the e...	Plain Text
BORO	Borough in which the entity (restaurant) is located.;• 1 = M...	Plain Text
BUILDING	Building number for establishment (restaurant) location	Plain Text
STREET	Street name for establishment (restaurant) location	Plain Text
ZIPCODE	Zip code of establishment (restaurant) location	Plain Text
PHONE	Phone Number; Phone number provided by restaurant ow...	Plain Text
CUISINE DESCRIPTION	This field describes the entity (restaurant) cuisine. ; Option...	Plain Text

I loaded this data into MySQL

Load it into your local database with database_before_day2.sql from the course web page

Source: https://data.cityofnewyork.us/Health/DOHMH-New-York-City-Restaurant-Inspection-Results/43nn-pn8j

Explore data via command line

First load data “database_before_day2.sql” from today’s link on the course website Schedule page

Connect from command line:
mysql –u root -p

```
d84607y@MacBook-Pro-5 ~ % sudo mysql -u root -p
Enter password:
Welcome to the MySQL monitor.  Commands end with ; or \g.
Your MySQL connection id is 17
Server version: 9.5.0 MySQL Community Server - GPL

Copyright (c) 2000, 2025, Oracle and/or its affiliates.

Oracle is a registered trademark of Oracle Corporation and/or its
affiliates. Other names may be trademarks of their respective
owners.

Type 'help;' or '\h' for help. Type '\c' to clear the current input statement.

mysql> SHOW SCHEMAS;
+--------------------+
| Database           |
+--------------------+
| college            |
| information_schema |
| lab2_db            |
| mysql              |
| nyc_data           |
| nyc_inspections    |
| performance_schema |
| sql-murder-mystery |
| sys                |
+--------------------+
9 rows in set (0.060 sec)
```

See what schemas are available with:
SHOW DATABASES; (or SHOW DATABASES;)

Source: https://data.cityofnewyork.us/Health/DOHMH-New-York-City-Restaurant-Inspection-Results/43nn-pn8j

Explore data via command line

First load data “database_before_day2.sql” from today’s link on the course website Schedule page

```
mysql> SHOW SCHEMAS;
+--------------------+
| Database           |
+--------------------+
| college            |
| information_schema |
| lab2_db            |
| mysql              |
| nyc_data           |
| nyc_inspections    |
| performance_schema |
| sql-murder-mystery |
| sys                |
+--------------------+
9 rows in set (0.060 sec)

mysql> USE nyc_data;
Reading table information for completion of table and column names
You can turn off this feature to get a quicker startup with -A


Database changed
mysql> SHOW TABLES;
+-----------------------+
| Tables_in_nyc_data    |
+-----------------------+
| restaurant_inspections |
+-----------------------+
1 row in set (0.006 sec)
```

Choose one of the schemas (databases) to use

See what tables the nyc_data schema has
Here only one table called *restaurant_inspections*

Source: https://data.cityofnewyork.us/Health/DOHMH-New-York-City-Restaurant-Inspection-Results/43nn-pn8j

Explore data via command line

First load data “database_before_day2.sql” from today’s link on the course website Schedule page

```
mysql> DESCRIBE restaurant_inspections;
+----------------------+--------+------+-----+---------+-------+
| Field                | Type   | Null | Key | Default | Extra |
+----------------------+--------+------+-----+---------+-------+
| CAMIS                | bigint | YES  |     | NULL    |       |
| DBA                  | text   | YES  |     | NULL    |       |
| BORO                 | text   | YES  |     | NULL    |       |
| BUILDING             | text   | YES  |     | NULL    |       |
| STREET               | text   | YES  |     | NULL    |       |
| ZIPCODE              | text   | YES  |     | NULL    |       |
| PHONE                | bigint | YES  |     | NULL    |       |
| CUISINE DESCRIPTION  | text   | YES  |     | NULL    |       |
| INSPECTION DATE      | text   | YES  |     | NULL    |       |
| ACTION               | text   | YES  |     | NULL    |       |
| VIOLATION CODE       | text   | YES  |     | NULL    |       |
| VIOLATION DESCRIPTION| text   | YES  |     | NULL    |       |
| CRITICAL FLAG        | text   | YES  |     | NULL    |       |
| SCORE                | int    | YES  |     | NULL    |       |
| GRADE                | text   | YES  |     | NULL    |       |
| GRADE DATE           | text   | YES  |     | NULL    |       |
| RECORD DATE          | text   | YES  |     | NULL    |       |
| INSPECTION TYPE      | text   | YES  |     | NULL    |       |
| Latitude             | double | YES  |     | NULL    |       |
| Longitude            | double | YES  |     | NULL    |       |
| Community Board      | text   | YES  |     | NULL    |       |
| Council District     | text   | YES  |     | NULL    |       |
| Census Tract         | text   | YES  |     | NULL    |       |
| BIN                  | text   | YES  |     | NULL    |       |
| BBL                  | bigint | YES  |     | NULL    |       |
| NTA                  | text   | YES  |     | NULL    |       |
| Location             | text   | YES  |     | NULL    |       |
+----------------------+--------+------+-----+---------+-------+
```

See what attributes (fields) the table contains

See:

- **Attribute names**
- **Types**
- **If can be NULL**
- **Key (more later)**
- **Default values**
- **Extra info**

Explore data using MySQL Workbench

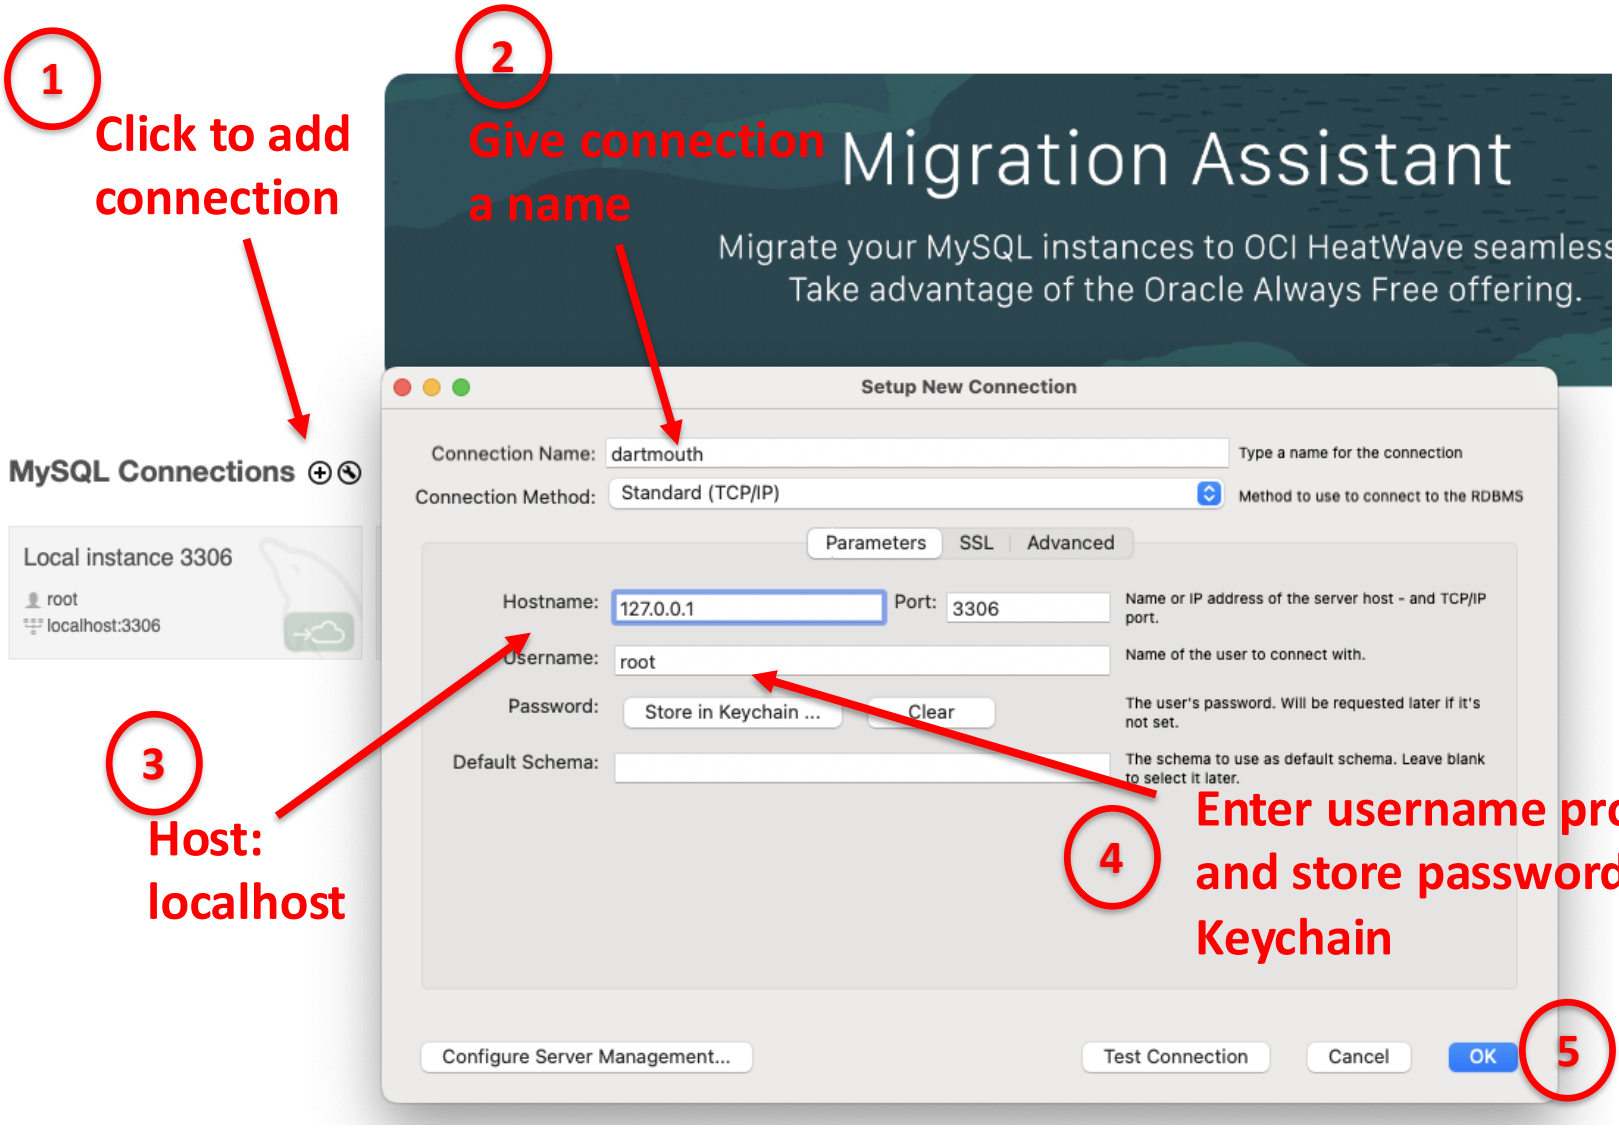

Explore data using MySQL Workbench

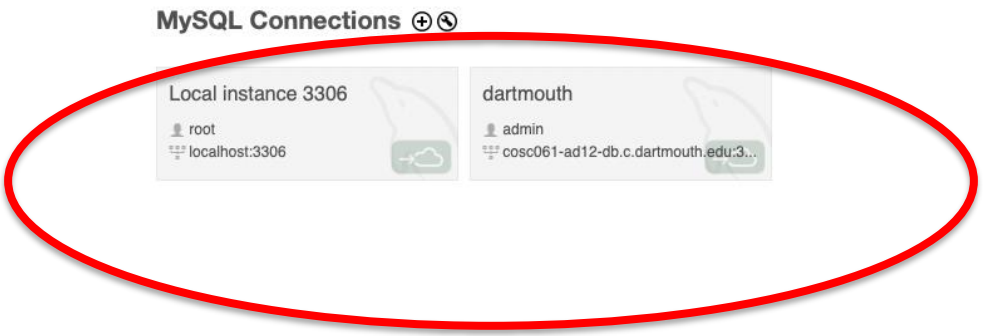

From main screen, choose the database you'd like to work with

Explore data using MySQL Workbench

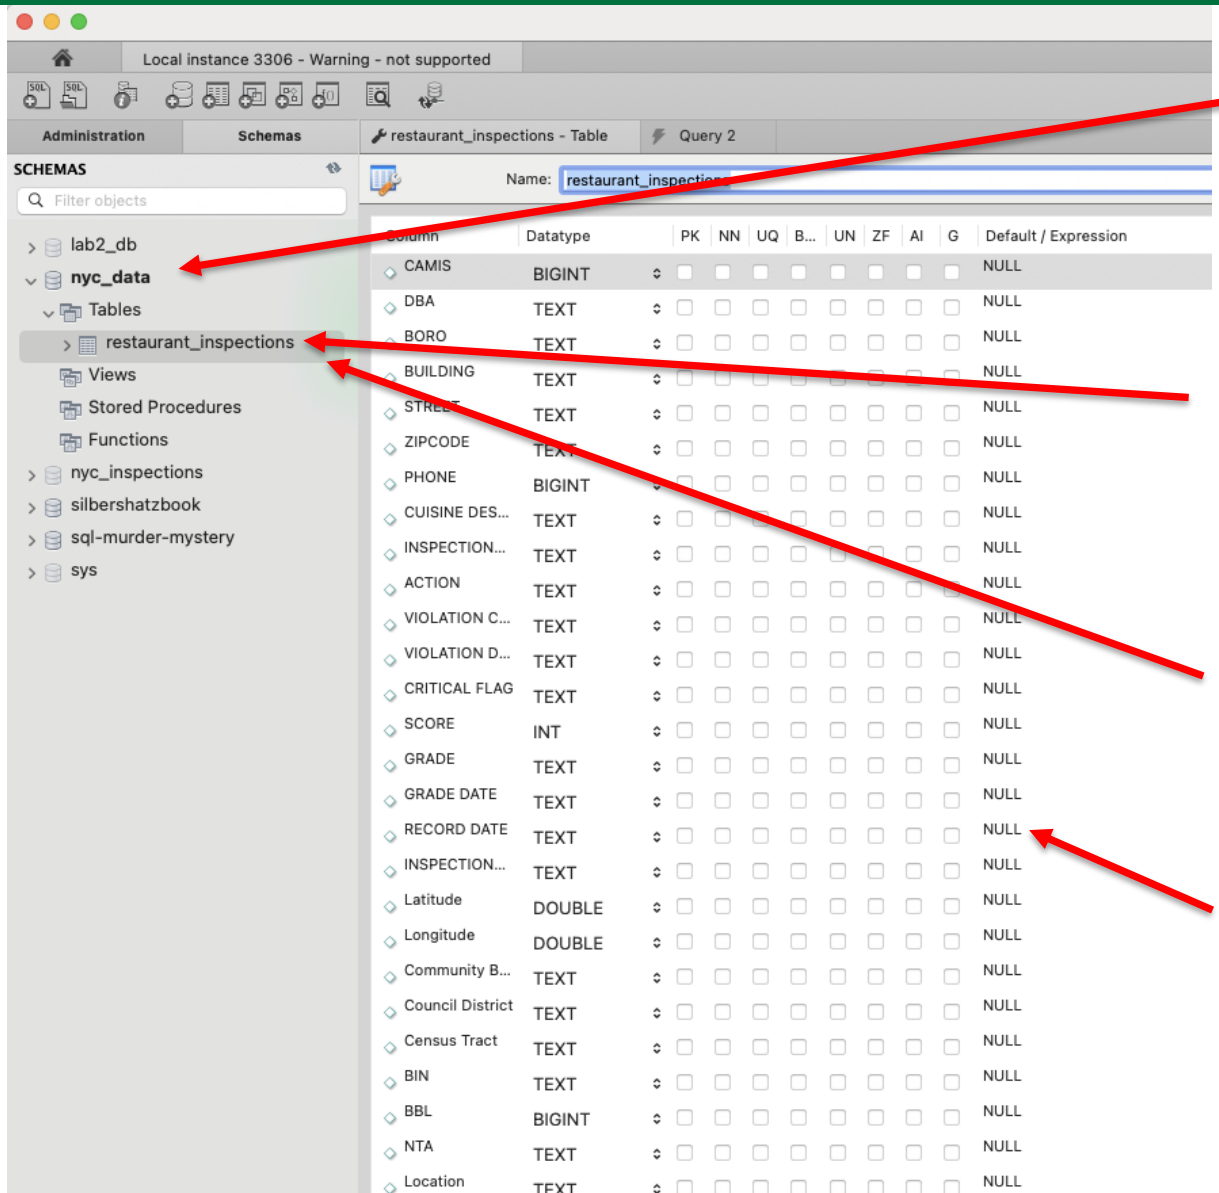

See what schemas (databases) are available

Expand schema to see tables
Here only one table called *restaurant_inspections*

Click wrench icon to see what attributes the table contains

See attribute names, types, and details

Explore data using MySQL Workbench

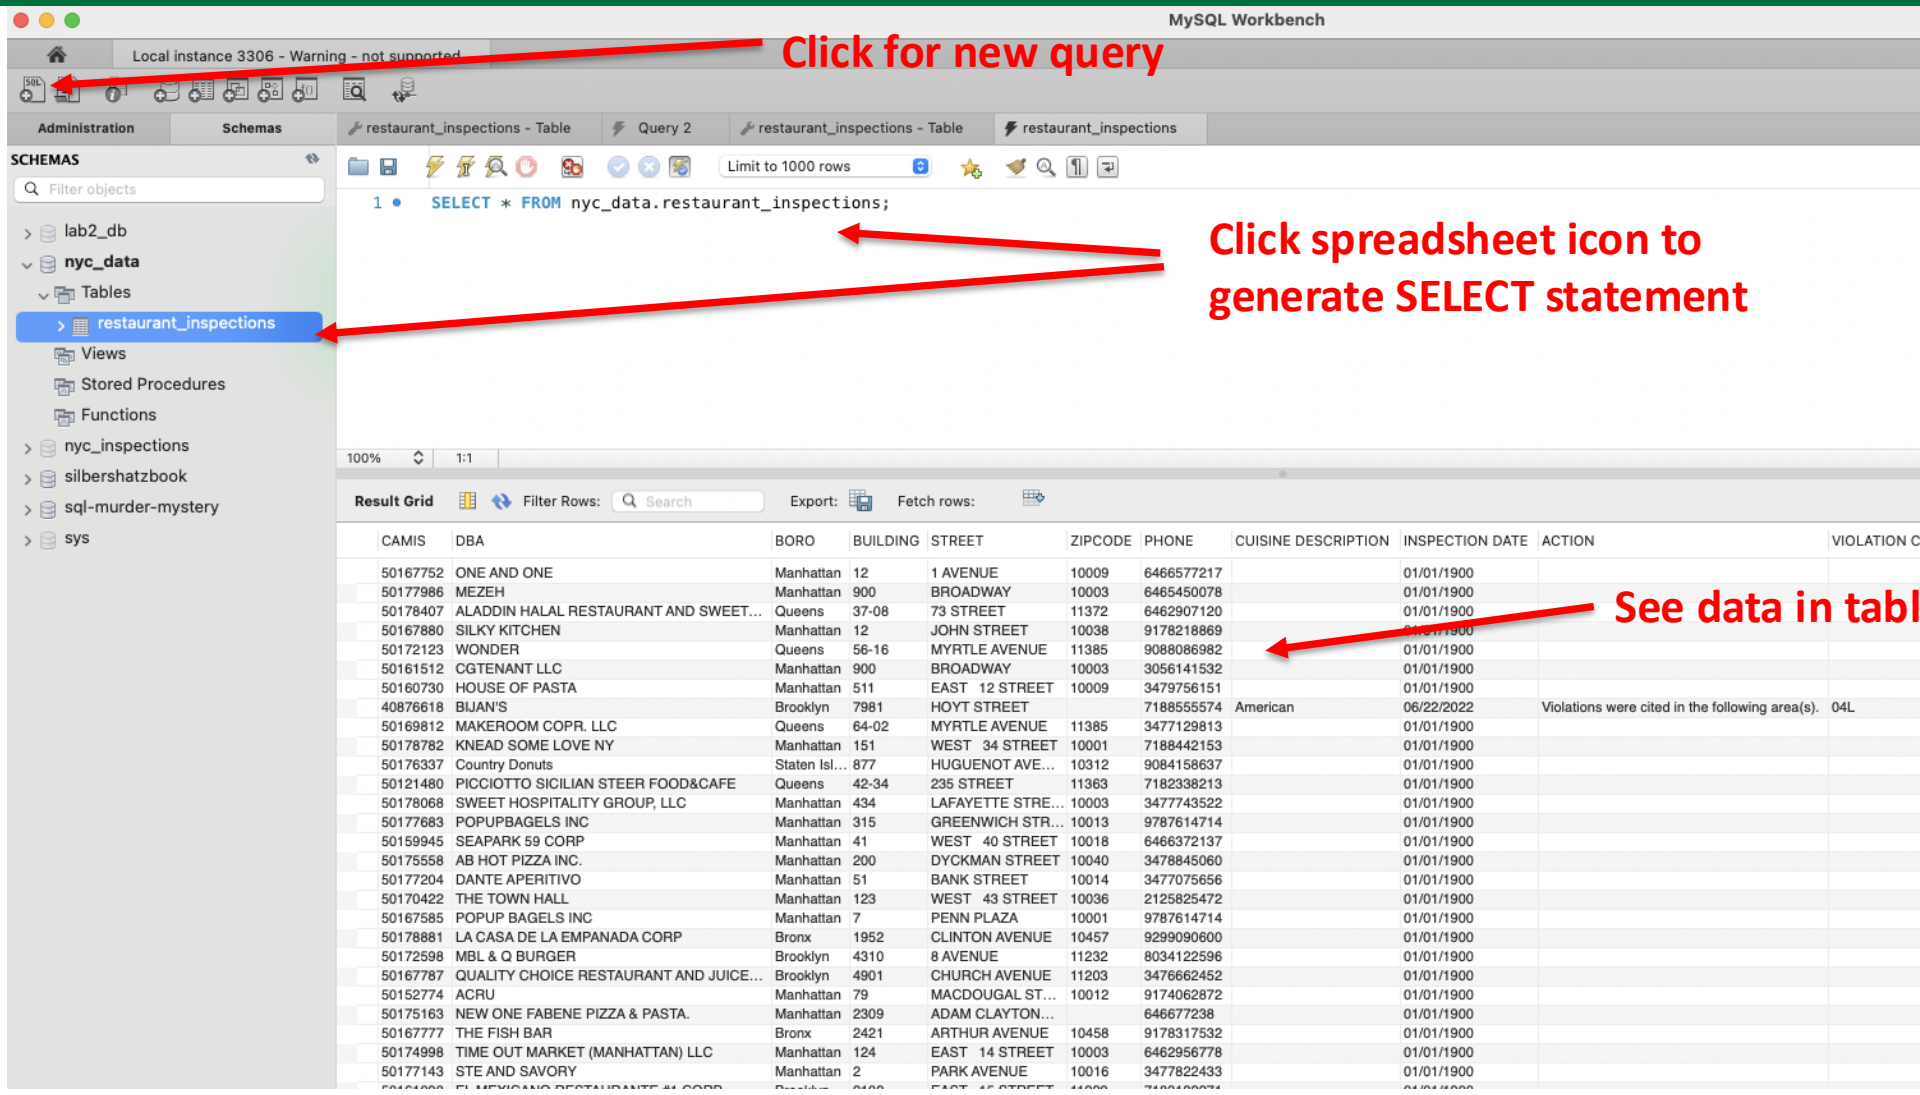

CAMIS	DBA	BORO	BUILDING	STREET	ZIPCODE	PHONE	CUISINE DESCRIPTION	INSPECTION DATE	ACTION	VIOLATION C
50167752	ONE AND ONE	Manhattan	12	1 AVENUE	10009	6466577217		01/01/1900		
50177986	MEZEH	Manhattan	900	BROADWAY	10003	6465450078		01/01/1900		
50178407	ALADDIN HALAL RESTAURANT AND SWEET...	Queens	37-08	73 STREET	11372	6462907120		01/01/1900		
50167880	SILKY KITCHEN	Manhattan	12	JOHN STREET	10038	9178218869		01/01/1900		
50172123	WONDER	Queens	56-16	MYRTLE AVENUE	11385	9088086982		01/01/1900		
50161512	CGTENANT LLC	Manhattan	900	BROADWAY	10003	3056141532		01/01/1900		
50160730	HOUSE OF PASTA	Manhattan	511	EAST 12 STREET	10009	3479756151		01/01/1900		
40876618	BIJAN'S	Brooklyn	7981	HOYT STREET		7188555574	American	06/22/2022	Violations were cited in the following area(s).	04L
50169812	MAKEROOM COPR. LLC	Queens	64-02	MYRTLE AVENUE	11385	3477129813		01/01/1900		
50178782	KNEAD SOME LOVE NY	Manhattan	151	WEST 34 STREET	10001	7188442153		01/01/1900		
50176337	Country Donuts	Staten Isl...	877	HUGUENOT AVE...	10312	9084158637		01/01/1900		
50121480	PICCIOTTO SICILIAN STEER FOOD&CAFE	Queens	42-34	235 STREET	11363	7182338213		01/01/1900		
50178068	SWEET HOSPITALITY GROUP, LLC	Manhattan	434	LAFAYETTE STRE...	10003	3477743522		01/01/1900		
50177683	POPUPBAGELS INC	Manhattan	315	GREENWICH STR...	10013	9787614714		01/01/1900		
50159945	SEAPARK 59 CORP	Manhattan	41	WEST 40 STREET	10018	6466372137		01/01/1900		
50175558	AB HOT PIZZA INC.	Manhattan	200	DYCKMAN STREET	10040	3478845060		01/01/1900		
50177204	DANTE APERITIVO	Manhattan	51	BANK STREET	10014	3477075656		01/01/1900		
50170422	THE TOWN HALL	Manhattan	123	WEST 43 STREET	10036	2125825472		01/01/1900		
50167585	POPUP BAGELS INC	Manhattan	7	PENN PLAZA	10001	9787614714		01/01/1900		
50178881	LA CASA DE LA EMPANADA CORP	Bronx	1952	CLINTON AVENUE	10457	9299090600		01/01/1900		
50172598	MBL & Q BURGER	Brooklyn	4310	8 AVENUE	11232	8034122596		01/01/1900		
50167787	QUALITY CHOICE RESTAURANT AND JUICE...	Brooklyn	4901	CHURCH AVENUE	11203	3476662452		01/01/1900		
50152774	ACRU	Manhattan	79	MACDOUGAL ST...	10012	9174062872		01/01/1900		
50175163	NEW ONE FABENE PIZZA & PASTA.	Manhattan	2309	ADAM CLAYTON...		646677238		01/01/1900		
50167777	THE FISH BAR	Bronx	2421	ARTHUR AVENUE	10458	9178317532		01/01/1900		
50174998	TIME OUT MARKET (MANHATTAN) LLC	Manhattan	124	EAST 14 STREET	10003	6462956778		01/01/1900		
50177143	STE AND SAVORY	Manhattan	2	PARK AVENUE	10016	3477822433		01/01/1900		

First get a feel for the data by selecting all attributes

SELECT command (query)

```
USE nyc_data;

SELECT * FROM restaurant_inspections LIMIT 100;
```

Tell MySQL which schema (database) to use

- Commands end with ;
- Can run multiple commands, like a program

* means return all attributes (columns)

Table "restaurant_inspections" has results of 291,534 health inspections

- Each restaurant may have been inspected multiple times over the years
- Only active restaurants listed in this dataset

No WHERE clause so all tuples (rows) match select criteria

Only return the first 100 rows

Source: https://data.cityofnewyork.us/Health/DOHMH-New-York-City-Restaurant-Inspection-Results/43nn-pn8j

First get a feel for the data by selecting all attributes

SELECT command (query)

USE nyc_data;

SELECT * **FROM** restaurant_inspections **LIMIT** 100;

SQL File 12* | SQL File 10*

Limit to 1000 rows

```
1  USE nyc_data;
2
3  SELECT * FROM restaurant_inspections LIMIT 100;
4
```

100% 1:4

Result Grid | Filter Rows: Search | Export: | Fetch rows:

CAMIS	DBA	BORO	BUILDING	STREET	ZIPCODE	PHONE	CUISINE DESCRIPTION	INSPECTION DATE	ACTION	VIOLATION COD
50167752	ONE AND ONE	Manhattan	12	1 AVENUE	10009	6466577217		01/01/1900		
50177986	MEZEH	Manhattan	900	BROADWAY	10003	6465450078		01/01/1900		
50178407	ALADDIN HALAL RESTAURANT AND SWEET...	Queens	37-08	73 STREET	11372	6462907120		01/01/1900		
50167880	SILKY KITCHEN	Manhattan	12	JOHN STREET	10038	9178218869		01/01/1900		
50172123	WONDER	Queens	56-16	MYRTLE AVENUE	11385	9088086982		01/01/1900		
50161512	CGTENANT LLC	Manhattan	900	BROADWAY	10003	3056141532		01/01/1900		
50160730	HOUSE OF PASTA	Manhattan	511	EAST 12 STREET	10009	3479756151		01/01/1900		
40876618	BIJAN'S	Brooklyn	7981	HOYT STREET		7188555574	American	06/22/2022	Violations were cited in the following area(s).	04L
50169812	MAKEROOM COPR. LLC	Queens	64-02	MYRTLE AVENUE	11385	3477129813		01/01/1900		
50178782	KNEAD SOME LOVE NY	Manhattan	151	WEST 34 STREET	10001	7188442153		01/01/1900		
50176337	Country Donuts	Staten Isl...	877	HUGUENOT AVE...	10312	9084158637		01/01/1900		
50121480	PICCIOTTO SICILIAN STEER FOOD&CAFE	Queens	42-34	235 STREET	11363	7182338213		01/01/1900		
50178068	SWEET HOSPITALITY GROUP, LLC	Manhattan	434	LAFAYETTE STRE...	10003	3477743522		01/01/1900		
50177683	POPUPBAGELS INC	Manhattan	315	GREENWICH STR...	10013	9787614714		01/01/1900		
50159945	SEAPARK 59 CORP	Manhattan	41	WEST 40 STREET	10018	6466372137		01/01/1900		

Source: https://data.cityofnewyork.us/Health/DOHMH-New-York-City-Restaurant-Inspection-Results/43nn-pn8j

SELECT can specify the rows we want and sort them using WHERE and ORDER BY

SELECT command (query)

SQL File 10*

Limit to 1000 rows

```
1 • use nyc_data;
2 • SELECT * FROM restaurant_inspections
3   WHERE DBA is not NULL
4   ORDER BY DBA asc
5   LIMIT 100;
```

- Commands on multiple lines are ok (encouraged!)
- "where is not null" returns tuples where attribute is not null (can also say "is null" for nulls)
- "order by" sorts either asc (default) or desc
- Note the sorted DBA names start with :
- Also note lower case SQL commands are ok

100% 11:5

Result Grid Filter Rows: Search Export: Fetch rows:

CAMIS	DBA	BORO	BUILDING	STREET	ZIPCODE	PHONE	CUISINE DESCRIPTION	INSPECTION DATE	ACTION
50134736	:3 COFFEE	Manhattan	29	WEST 21 STREET	10010	3472658602	Coffee/Tea	05/01/2023	Violations were c
50134736	:3 COFFEE	Manhattan	29	WEST 21 STREET	10010	3472658602	Coffee/Tea	05/01/2023	Violations were c
50134736	:3 COFFEE	Manhattan	29	WEST 21 STREET	10010	3472658602	Coffee/Tea	11/17/2025	Violations were c
50134736	:3 COFFEE	Manhattan	29	WEST 21 STREET	10010	3472658602	Coffee/Tea	11/17/2025	Violations were c
50134736	:3 COFFEE	Manhattan	29	WEST 21 STREET	10010	3472658602	Coffee/Tea	07/22/2024	Violations were c
50134736	:3 COFFEE	Manhattan	29	WEST 21 STREET	10010	3472658602	Coffee/Tea	07/22/2024	Violations were c
50134736	:3 COFFEE	Manhattan	29	WEST 21 STREET	10010	3472658602	Coffee/Tea	11/17/2025	Violations were c
50134736	:3 COFFEE	Manhattan	29	WEST 21 STREET	10010	3472658602	Coffee/Tea	07/22/2024	Violations were c
50134736	:3 COFFEE	Manhattan	29	WEST 21 STREET	10010	3472658602	Coffee/Tea	11/17/2025	Violations were c
50178828	'ESSEN	Manhattan	290	MADISON AVENUE	10017	2126890800		01/01/1900	

Source: https://data.cityofnewyork.us/Health/DOHMH-New-York-City-Restaurant-Inspection-Results/43nn-pn8j

Rename attributes (and tables) using the AS operator

Limit to 1000 rows

```
1  USE nyc_data;
2
3  SELECT dba as RestaurantName, `inspection date` as InspectionDate, Score,
4      score/100 AS ScorePercent, 0 as AdjustedScore
5      FROM restaurant_inspections
6      WHERE Score is not NULL;
```

I prefer on word TitleCase names (no spaces)

Use backticks ` for multiple words

AS changes column name in returned relation

100% 29:6

Result Grid Filter Rows: Search Export: Fetch rows:

RestaurantName	InspectionDa...	Score	ScorePercent	AdjustedScore
CAFE GRUMPY	05/22/2024	51	0.5100	0
TULCINGO RESTAURANT	03/04/2025	16	0.1600	0
SAGE	06/13/2023	26	0.2600	0
MORI MORI	09/12/2025	8	0.0800	0
AFRICA KINE	04/30/2024	25	0.2500	0
CARIBBEAN KRAVE	05/09/2025	13	0.1300	0
CITY ONE	08/31/2023	15	0.1500	0
TANGRA	03/31/2022	22	0.2200	0
65 MARKET PLACE	11/30/2022	27	0.2700	0
BARAKA BUFFET	05/08/2025	8	0.0800	0
BLUJEEN	02/05/2020	29	0.2900	0
HARLEM HOPS	06/02/2022	23	0.2300	0
BABA COOL	01/11/2024	34	0.3400	0
DONS BOGAM	08/18/2022	12	0.1200	0
IHOP	08/14/2025	27	0.2700	0

Can do math in the select (e.g., score/100) or create an attribute with a given value (0 here)

Load this health inspection data into your local MySQL installation

1. Download “database_before_day2.sql” from today’s link on the course website Schedule page (right click, choose “Save as”)
2. Open MySQL Workbench
3. Connect to the database on your localhost
4. Click File -> Open SQL Script… (or just double click downloaded file)
5. Choose downloaded file from step 1
6. Run the script
7. This will create a database schema with one table holding all NYC restaurant health inspections (only for restaurants currently open)
8. Do exercises on next slide

Practice 1: Find details about the restaurants at the Met

- The Metropolitan Museum of Art is located at 1000 5 Avenue in Manhattan
- Find all restaurants in the Met
  - There are 5 restaurants at the Met, each with a different name!
  - Think about how to find them all (address perhaps?)
- List all inspections for any of the Met's restaurants
- Use DISTINCT to find all restaurant names, each listed one time

Practice 2: use SELECT to answer questions about Ray's Pizza locations

Exercises

1. New Yorkers sometimes joke that there are many "Ray's Pizza" variants ("Original Ray's", "Famous Rays", "Famous Original Ray's"…), find all inspections for each the Ray's Pizza variants
   - Use "LIKE" instead of = in a WHERE clause (WHERE DBA LIKE '%Ray%')
   - LIKE works with wildcards
     - "%" matches zero or more characters
     - "_" matches any one character
2. What are the "gotchas" with Ray's name?
3. How inspections many of Ray's are in the Bronx boro? Queens? Manhattan?
   - WHERE clauses can use "AND", "OR", "NOT"
4. Limit your results to only Ray's Pizza inspections that raised a "critical flag"
   - Use ` character (same key as tilde ~, by the 1 key) around attributes that have more than one word, e.g., `critical flag`
5. Would you eat at the Columbus Ave Ray's Pizza store?