

Post-class Exercises: Nested Queries

Schema and data

We will work on the **DineSafe** database, which is a collection of food safety inspections carried out between 2010 and 2012 in Toronto.

- Find out the establishment that has received the highest amount of fine for a single infraction. Report the name and the address of the establishment, as well as the amount of the fine. If there are ties, report for each of them.
 - Find out every establishment that has ever paid any fine because of some infraction, ordered by their names alphabetically.
 - Find out the establishment which a total amount of fine over \$3000 in the history.
- Fine the top 5 most popular type of establishment and the number of such establishment on Bloor Street West (Its address contains "BLOOR ST W").
- Find the top 10 establishment that has the most infractions in one single inspection. Report the name and the address of the establishment, as well as the number of infractions in that inspection. If there are ties, report for each of them.
- In this query, we want to study if the minimal annual inspections are properly conducted. We will break this complicated query into some small steps, each of them involves using a nested query itself. You will find it handy to handle these complicated queries using views if possible.
 - Find out the establishment that was never inspected in year 2010. Report their names and addresses, ordered by their names alphabetically. Hint: for PostgreSQL, use `extract(year from date)` to extract the year from the date; for SQLite, need to use `STRFTIME('%Y', date)`.
 - For each establishment, find out the actual inspections happened in year 2010. Hint: don't forget to include those with zero inspection. Verify your result by checking whether the number of returned rows equals to the number of rows in the establishment relation.
 - Find out every establishment that has ever failed to meet the minimal number of required annual inspections in year 2010, ordered by name alphabetically. For each of them, report its name, address, minimal number of required annual inspections, and actual inspections conducted in 2010. Hint: start your work from the results in Q4(b).
- According to DineSafe Inspection and Disclosure System, a "Conditional Pass" notice will be issued when one or more significant infractions are observed during an inspection, and a second re-inspection will be conducted within 24-48 hours. Find out the number of times that an establishment was issued a "Conditional Pass" notice at first inspection, but no re-inspection was conducted within 24-48 hours (holes in execution!). Hint: for PostgreSQL, use `extract(day from date)` to extract the day from the date; for SQLite, need to use `julianday(R2.date)`.
- In the database, you may have found several establishment have the same name but in different locations. These are franchise stores or chain stores, e.g., McDonald's.
 - Find the name of the top 5 largest chains in the database. For each of them, report its name and the number of stores it owns in Toronto, ordered by the number in descending order.
 - For each of the chain brand in (a), report the average minimal annual inspections per store, in the same order as in (a).

- (c) Find out if any of the chain brand in (a), has received any inspections with a “Closed” status in the history. If yes, report the name, the number of “Closed” inspections for each of such chain brand.
7. (a) Find the pass rate of “SUSHI WORLD” on COLLEGE ST. Report the total number of passed inspections, total number of previous inspections, and the pass rate in its history.
- (b) **A very hard one:** For each establishment, find its pass rate among all its previous inspections in the database. Report the name, address, total number of passed inspections, total number of previous inspections, and the pass rate. Order your results in an descending order of the pass rate, followed by the total number of inspections (check some of your favorite restaurants!). Verify that the number of rows in your results equals to the number of establishment in the database.
8. Which rule has been violated the most?