

Lecture 10-3: Advanced Query Techniques

DATA 503: Fundamentals of Data Engineering

Lucas P. Cordova, Ph.D.

2026-03-16

This lecture covers advanced SQL query techniques: subqueries in WHERE, FROM, and SELECT clauses, correlated subqueries, EXISTS, LATERAL joins, Common Table Expressions, cross tabulations, and CASE expressions. We apply these patterns to Census, survey, and temperature data. Based on Chapter 13 of Practical SQL, 2nd Edition.

Table of contents

1	Part 1: Setting Up	2
2	Part 2: Subqueries	2
3	Part 3: EXISTS and IN with Subqueries	7
4	Part 4: LATERAL Joins	9
5	Part 5: Common Table Expressions (CTEs)	11
6	Part 6: Cross Tabulations	13
7	Part 7: CASE Expressions	16
8	Part 8: Summary	18

1 Part 1: Setting Up

1.1 Database Setup

1.2 Create and Populate the Database

Create a fresh database for this lecture:

```
1 CREATE DATABASE advanced_queries;
```

Connect to it, then run the setup script provided: [[advanced_queries.sql](#)]

The script creates and loads the following tables:

Table	Rows	Description
us_counties_pop_est_2019	3,142	County population estimates
cbp_naics_72_establishments	2,764	Food/accommodation businesses per county
employees	6	Small HR demo table
teachers	6	Small demo table (from Ch 7)
ice_cream_survey	200	Office ice cream flavor preferences
temperature_readings	730	Two years of daily temps at two stations

All CSVs are in the `Chapter_13/` folder alongside these slides. Download them before class.

1.3 Verify Your Setup

```
1 SELECT count(*) FROM us_counties_pop_est_2019; -- 3,142
2 SELECT count(*) FROM cbp_naics_72_establishments; -- 2,764
3 SELECT count(*) FROM employees; -- 6
4 SELECT count(*) FROM ice_cream_survey; -- 200
5 SELECT count(*) FROM temperature_readings; -- 730
```

If any count is wrong, re-run the setup. We need all of these today.

2 Part 2: Subqueries

A query inside a query. It's queries all the way down.

2.1 Subqueries in WHERE

2.2 Scalar Subquery in WHERE

The most common pattern: use a subquery to compute a threshold, then filter on it.

```
1 SELECT county_name,  
2         state_name,  
3         pop_est_2019  
4 FROM us_counties_pop_est_2019  
5 WHERE pop_est_2019 >= (  
6     SELECT percentile_cont(.9) WITHIN GROUP (ORDER BY pop_est_2019)  
7     FROM us_counties_pop_est_2019  
8 )  
9 ORDER BY pop_est_2019 DESC;
```

Run this. The inner query computes the 90th percentile of county populations. The outer query returns only counties at or above that threshold.

The subquery executes first, returns a single value, and the outer query uses it like a constant. This is a **scalar subquery** – it returns exactly one value.

2.3 Subquery with DELETE

Subqueries work with DML too. Let's create a copy and prune it:

```
1 CREATE TABLE us_counties_2019_top10 AS  
2 SELECT * FROM us_counties_pop_est_2019;  
3  
4 DELETE FROM us_counties_2019_top10  
5 WHERE pop_est_2019 < (  
6     SELECT percentile_cont(.9) WITHIN GROUP (ORDER BY pop_est_2019)  
7     FROM us_counties_2019_top10  
8 );  
9  
10 SELECT count(*) FROM us_counties_2019_top10;
```

Run this. You started with 3,142 counties and now have ~315 (the top 10%).

In a pipeline context: this is how you build filtered summary tables. Copy the source, delete what you don't need, keep the rest for downstream consumers.

2.4 Quick Quiz: Subquery Execution

When does the subquery in a WHERE clause execute?

- A. Once before the outer query starts
- B. Once per row of the outer query
- C. After the outer query finishes
- D. It depends

...

A (for scalar subqueries). The database computes the inner value once, then scans the outer table. Correlated subqueries (coming up) are different – they run once *per row*.

2.5 Try It: Above-Average Counties

Write a query that returns all counties where `pop_est_2019` is above the national average. Show county name, state, and population. How many are there?

```
1 -- Your query here
...
1 SELECT county_name, state_name, pop_est_2019
2 FROM us_counties_pop_est_2019
3 WHERE pop_est_2019 > (
4     SELECT avg(pop_est_2019) FROM us_counties_pop_est_2019
5 )
6 ORDER BY pop_est_2019 DESC;
```

Far fewer than half! The mean is pulled up by massive counties (LA, Cook, Harris), so most counties are below average. Median vs. mean matters.

2.6 Subqueries in FROM (Derived Tables)

2.7 Using a Subquery as a Table

A subquery in the FROM clause creates a **derived table** – a temporary result set you can query like a regular table:

```

1 SELECT round(calcs.average, 0) AS average,
2         calcs.median,
3         round(calcs.average - calcs.median, 0) AS median_average_diff
4 FROM (
5     SELECT avg(pop_est_2019) AS average,
6            percentile_cont(.5)
7            WITHIN GROUP (ORDER BY pop_est_2019)::numeric AS median
8     FROM us_counties_pop_est_2019
9 ) AS calcs;

```

Run this. The inner query computes two aggregates. The outer query calculates the difference. The AS calcs alias is required – PostgreSQL demands that derived tables have names.

2.8 Joining Two Derived Tables

This is powerful for combining aggregations from different tables:

```

1 SELECT census.state_name AS st,
2        census.pop_est_2018,
3        est.establishment_count,
4        round((est.establishment_count / census.pop_est_2018::numeric) * 1000, 1)
5        AS estabs_per_thousand
6 FROM
7     (SELECT st, sum(establishments) AS establishment_count
8      FROM cbp_naics_72_establishments
9      GROUP BY st) AS est
10 JOIN
11     (SELECT state_name, sum(pop_est_2018) AS pop_est_2018
12      FROM us_counties_pop_est_2019
13      GROUP BY state_name) AS census
14 ON est.st = census.state_name
15 ORDER BY estabs_per_thousand DESC;

```

Each derived table aggregates to the state level separately, then they join on state name. This is a common ETL pattern: aggregate separately, join at the grain you need.

2.9 Quick Quiz: Derived Table Rules

What happens if you forget the AS alias on a derived table?

- A. PostgreSQL infers a name automatically
- B. You get a syntax error
- C. It works but you can't reference the columns
- D. The query runs but returns wrong results

...

B. PostgreSQL requires every derived table to have an alias. `AS calcs`, `AS est`, `AS census` – pick a name. Without it, the parser rejects the query.

2.10 Subqueries in SELECT

2.11 Column-Level Subqueries

You can put a subquery directly in the SELECT list to add a computed column:

```

1 SELECT county_name,
2         state_name AS st,
3         pop_est_2019,
4         (SELECT percentile_cont(.5) WITHIN GROUP (ORDER BY pop_est_2019)
5          FROM us_counties_pop_est_2019) AS us_median
6 FROM us_counties_pop_est_2019;
```

Every row gets the same median value alongside its own population. This lets you compare each county to the national benchmark without a JOIN.

2.12 Subquery in a Calculation

Take it further – compute the difference from the median inline:

```

1 SELECT county_name,
2         state_name AS st,
3         pop_est_2019,
4         pop_est_2019 - (SELECT percentile_cont(.5) WITHIN GROUP (ORDER BY pop_est_2019)
5                          FROM us_counties_pop_est_2019) AS diff_from_median
6 FROM us_counties_pop_est_2019
7 WHERE (pop_est_2019 - (SELECT percentile_cont(.5) WITHIN GROUP (ORDER BY pop_est_2019)
8                       FROM us_counties_pop_est_2019))
9        BETWEEN -1000 AND 1000;
```

This finds counties whose population is within 1,000 of the national median. Notice the subquery appears twice – once in SELECT, once in WHERE. That redundancy is why CTEs exist (we'll get there).

2.13 Try It: Every County vs. Its State Average

Write a query that shows each county's population alongside its state's average population (using a correlated subquery in the SELECT list). Show `county_name`, `state_name`, `pop_est_2019`, and `state_avg`.

Hint: the subquery references the outer table's `state_name`.

```
1 -- Your query here
2
3 . . .
4
5 1 SELECT county_name,
6     state_name,
7     pop_est_2019,
8     (SELECT round(avg(pop_est_2019), 0)
9     FROM us_counties_pop_est_2019 AS inner_t
10    WHERE inner_t.state_name = outer_t.state_name) AS state_avg
11 FROM us_counties_pop_est_2019 AS outer_t
12 ORDER BY state_name, pop_est_2019 DESC;
```

This is a **correlated subquery** – it runs once per row of the outer query because it references `outer_t.state_name`. More expensive than scalar subqueries, but very expressive.

3 Part 3: EXISTS and IN with Subqueries

Checking membership across tables – the SQL equivalent of “are you on the list?”

3.1 EXISTS, IN, and Correlated Subqueries

3.2 Setup: Employees and Retirees

```
1 CREATE TABLE retirees (
2     id int,
3     first_name text,
4     last_name text
5 );
6
7 INSERT INTO retirees
8 VALUES (2, 'Janet', 'King'),
9         (4, 'Michael', 'Taylor');
```

Now we have two employees who retired. Let's find them – and find who's still active.

3.3 Using IN with a Subquery

```
1 SELECT first_name, last_name
2 FROM employees
3 WHERE emp_id IN (
4     SELECT id FROM retirees
5 )
6 ORDER BY emp_id;
```

IN checks if `emp_id` matches any value in the subquery result. Simple and readable.

3.4 Using EXISTS (Correlated)

```
1 SELECT first_name, last_name
2 FROM employees
3 WHERE EXISTS (
4     SELECT id
5     FROM retirees
6     WHERE id = employees.emp_id
7 );
```

EXISTS returns TRUE if the subquery finds at least one row. The subquery references `employees.emp_id` from the outer query – making it correlated.

For small tables, IN and EXISTS perform identically. For large tables, EXISTS can be faster because it stops scanning as soon as it finds one match.

3.5 NOT EXISTS: The Anti-Join

```
1 SELECT first_name, last_name
2 FROM employees
3 WHERE NOT EXISTS (
4     SELECT id
5     FROM retirees
6     WHERE id = employees.emp_id
7 );
```

Run this. Returns employees who are NOT retirees. This is an **anti-join** – one of the most useful patterns in data engineering. “Give me everything in table A that has no match in table B.”

Pipeline use case: find new records that haven’t been processed yet. `WHERE NOT EXISTS (SELECT 1 FROM processed WHERE processed.id = raw.id).`

3.6 Quick Quiz: IN vs EXISTS

For a table with 10 million rows, which is typically faster?

- A. IN (always)
- B. EXISTS (always)
- C. EXISTS (usually, because it short-circuits)
- D. They're always identical

...

C. EXISTS stops at the first match. IN materializes the full subquery result into a list. For large datasets, that difference matters. But the optimizer is smart – sometimes it rewrites one to the other. Profile, don't guess.

4 Part 4: LATERAL Joins

The most powerful subquery pattern you've never heard of.

4.1 LATERAL Subqueries

4.2 LATERAL in FROM: Inline Calculations

LATERAL lets a subquery in the FROM clause reference columns from preceding tables. Think of it as a “for each row, compute this”:

```
1 SELECT county_name,
2         state_name,
3         pop_est_2018,
4         pop_est_2019,
5         raw_chg,
6         round(pct_chg * 100, 2) AS pct_chg
7 FROM us_counties_pop_est_2019,
8      LATERAL (SELECT pop_est_2019 - pop_est_2018 AS raw_chg) rc,
9      LATERAL (SELECT raw_chg / pop_est_2018::numeric AS pct_chg) pc
10 ORDER BY pct_chg DESC;
```

Run this. Each LATERAL subquery can reference columns from earlier in the FROM clause – including results from *other* LATERAL subqueries. `pct_chg` uses `raw_chg` from the first LATERAL. Without LATERAL, you'd need nested subqueries or repeat the calculation.

4.3 LATERAL with JOIN: Top-N Per Group

The killer feature: get the top N related rows for each row in the driving table.

```
1 ALTER TABLE teachers ADD CONSTRAINT id_key PRIMARY KEY (id);
2
3 CREATE TABLE teachers_lab_access (
4     access_id bigint PRIMARY KEY GENERATED ALWAYS AS IDENTITY,
5     access_time timestamp with time zone,
6     lab_name text,
7     teacher_id bigint REFERENCES teachers (id)
8 );
9
10 INSERT INTO teachers_lab_access (access_time, lab_name, teacher_id)
11 VALUES ('2022-11-30 08:59:00-05', 'Science A', 2),
12         ('2022-12-01 08:58:00-05', 'Chemistry B', 2),
13         ('2022-12-21 09:01:00-05', 'Chemistry A', 2),
14         ('2022-12-02 11:01:00-05', 'Science B', 6),
15         ('2022-12-07 10:02:00-05', 'Science A', 6),
16         ('2022-12-17 16:00:00-05', 'Science B', 6);
```

Run the setup, then:

```
1 SELECT t.first_name, t.last_name, a.access_time, a.lab_name
2 FROM teachers t
3 LEFT JOIN LATERAL (
4     SELECT *
5     FROM teachers_lab_access
6     WHERE teacher_id = t.id
7     ORDER BY access_time DESC
8     LIMIT 2
9 ) a ON true
10 ORDER BY t.id;
```

For each teacher, this returns their 2 most recent lab accesses. Without LATERAL, you'd need window functions + CTEs. With LATERAL, it's clean and direct.

In data engineering: “for each customer, get their last 3 orders.” “For each sensor, get the most recent 5 readings.” LATERAL is the tool.

4.4 Try It: Top-1 Per Teacher

Modify the LATERAL query to return only each teacher's **most recent** lab access (not two). Teachers with no lab access should still appear (show NULLs).

```

1  -- Your query here
   ...
1  SELECT t.first_name, t.last_name, a.access_time, a.lab_name
2  FROM teachers t
3  LEFT JOIN LATERAL (
4      SELECT *
5      FROM teachers_lab_access
6      WHERE teacher_id = t.id
7      ORDER BY access_time DESC
8      LIMIT 1
9  ) a ON true
10 ORDER BY t.id;

```

Change LIMIT 2 to LIMIT 1. The LEFT JOIN ensures all teachers appear even with no matches.

5 Part 5: Common Table Expressions (CTEs)

Named temporary result sets that make complex queries readable.

5.1 CTEs with WITH

5.2 A Simple CTE

```

1  WITH large_counties (county_name, state_name, pop_est_2019) AS (
2      SELECT county_name, state_name, pop_est_2019
3      FROM us_counties_pop_est_2019
4      WHERE pop_est_2019 >= 100000
5  )
6  SELECT state_name, count(*)
7  FROM large_counties
8  GROUP BY state_name
9  ORDER BY count(*) DESC;

```

The CTE filters to counties with 100k+ population. The outer query counts them by state. It's the same as a subquery, but the logic flows top-to-bottom instead of inside-out.

5.3 CTEs for Joining Aggregations

Remember the derived table join from earlier? Here it is with CTEs:

```
1 WITH
2   counties (st, pop_est_2018) AS (
3     SELECT state_name, sum(pop_est_2018)
4     FROM us_counties_pop_est_2019
5     GROUP BY state_name
6   ),
7   establishments (st, establishment_count) AS (
8     SELECT st, sum(establishments) AS establishment_count
9     FROM cbp_naics_72_establishments
10    GROUP BY st
11  )
12 SELECT counties.st,
13        pop_est_2018,
14        establishment_count,
15        round((establishments.establishment_count /
16              counties.pop_est_2018::numeric(10,1)) * 1000, 1)
17              AS estabs_per_thousand
18 FROM counties JOIN establishments
19 ON counties.st = establishments.st
20 ORDER BY estabs_per_thousand DESC;
```

Compare this to the derived table version. Same result, dramatically more readable. Each CTE has a name and a clear purpose. In code review, this is the version that gets approved.

5.4 CTEs to Eliminate Redundancy

Remember the repeated subquery for median comparison? CTEs fix that:

```
1 WITH us_median AS (
2   SELECT percentile_cont(.5)
3     WITHIN GROUP (ORDER BY pop_est_2019) AS us_median_pop
4   FROM us_counties_pop_est_2019
5 )
6 SELECT county_name,
7        state_name AS st,
8        pop_est_2019,
9        us_median_pop,
10       pop_est_2019 - us_median_pop AS diff_from_median
```

```

11 FROM us_counties_pop_est_2019 CROSS JOIN us_median
12 WHERE (pop_est_2019 - us_median_pop) BETWEEN -1000 AND 1000;

```

The median is computed once in the CTE, then CROSS JOINed to every row. No repeated subqueries. Clean, efficient, maintainable.

5.5 Try It: State-Level Summary CTE

Using two CTEs, compute:

1. CTE 1: the total population per state
2. CTE 2: the number of counties per state

Join them and add a column for average population per county. Show state, total pop, county count, and avg pop per county. Order by average descending. Which state has the largest average county?

```

1  -- Your query here
. . .

1  WITH state_pop AS (
2      SELECT state_name, sum(pop_est_2019) AS total_pop
3      FROM us_counties_pop_est_2019
4      GROUP BY state_name
5  ),
6  state_counties AS (
7      SELECT state_name, count(*) AS num_counties
8      FROM us_counties_pop_est_2019
9      GROUP BY state_name
10 )
11 SELECT sp.state_name,
12        sp.total_pop,
13        sc.num_counties,
14        round(sp.total_pop::numeric / sc.num_counties, 0) AS avg_pop_per_county
15 FROM state_pop sp
16 JOIN state_counties sc ON sp.state_name = sc.state_name
17 ORDER BY avg_pop_per_county DESC;

```

6 Part 6: Cross Tabulations

Pivot tables in SQL. Every analyst's favorite party trick.

6.1 crosstab()

6.2 Enable the Extension

```
1 CREATE EXTENSION tablefunc;
```

This loads the `crosstab()` function from the `tablefunc` module. You only need to do this once per database.

6.3 The Ice Cream Survey

```
1 SELECT *
2 FROM ice_cream_survey
3 ORDER BY response_id
4 LIMIT 10;
```

200 employees across offices, each picking a flavor. We want a pivot: offices as rows, flavors as columns, counts as values.

6.4 Generating the Crosstab

```
1 SELECT *
2 FROM crosstab(
3     'SELECT office, flavor, count(*)
4     FROM ice_cream_survey
5     GROUP BY office, flavor
6     ORDER BY office',
7     'SELECT flavor
8     FROM ice_cream_survey
9     GROUP BY flavor
10    ORDER BY flavor'
11 )
12 AS (office text,
13     chocolate bigint,
14     strawberry bigint,
15     vanilla bigint);
```

Run this. The first argument is the source query (must have exactly 3 columns: row, category, value). The second argument defines the category labels. The `AS` clause names the output columns.

This transforms long data (one row per observation) into wide data (one row per group). Essential for reporting and dashboards.

6.5 Temperature Crosstab

A more complex example: median max temperature by station and month.

```
1 SELECT *
2 FROM crosstab(
3     'SELECT station_name,
4         date_part(''month'', observation_date),
5         percentile_cont(.5)
6         WITHIN GROUP (ORDER BY max_temp)
7     FROM temperature_readings
8     GROUP BY station_name,
9         date_part(''month'', observation_date)
10    ORDER BY station_name',
11    'SELECT month FROM generate_series(1,12) month'
12 )
13 AS (station text,
14     jan numeric(3,0), feb numeric(3,0), mar numeric(3,0),
15     apr numeric(3,0), may numeric(3,0), jun numeric(3,0),
16     jul numeric(3,0), aug numeric(3,0), sep numeric(3,0),
17     oct numeric(3,0), nov numeric(3,0), dec numeric(3,0));
```

Run this. Two stations, 12 months, median temperatures. Notice `generate_series(1,12)` creates the month column headers. The escaped single quotes (') inside the string are required.

6.6 Quick Quiz: Crosstab Requirements

The source query for `crosstab()` must return exactly how many columns?

- A. 2
- B. 3
- C. 4
- D. It depends on the output

...

B. Always three: (1) row identifier, (2) category, (3) value. The row identifier groups the pivot rows, the category determines which column gets the value.

7 Part 7: CASE Expressions

Conditional logic inside SQL. Your if/else for data transformation.

7.1 CASE for Classification

7.2 Reclassifying Data

```
1 SELECT max_temp,
2     CASE WHEN max_temp >= 90 THEN 'Hot'
3         WHEN max_temp >= 70 AND max_temp < 90 THEN 'Warm'
4         WHEN max_temp >= 50 AND max_temp < 70 THEN 'Pleasant'
5         WHEN max_temp >= 33 AND max_temp < 50 THEN 'Cold'
6         WHEN max_temp >= 20 AND max_temp < 33 THEN 'Frigid'
7         WHEN max_temp < 20 THEN 'Inhumane'
8         ELSE 'No reading'
9     END AS temperature_group
10 FROM temperature_readings
11 ORDER BY station_name, observation_date;
```

Each row gets classified into a group. CASE evaluates conditions top-to-bottom and returns the first match. ELSE catches anything that fell through.

7.3 CASE Inside a CTE: The Full Pattern

Combine CASE with a CTE to classify, then aggregate:

```
1 WITH temps_collapsed (station_name, max_temperature_group) AS (
2     SELECT station_name,
3         CASE WHEN max_temp >= 90 THEN 'Hot'
4             WHEN max_temp >= 70 AND max_temp < 90 THEN 'Warm'
5             WHEN max_temp >= 50 AND max_temp < 70 THEN 'Pleasant'
6             WHEN max_temp >= 33 AND max_temp < 50 THEN 'Cold'
7             WHEN max_temp >= 20 AND max_temp < 33 THEN 'Frigid'
8             WHEN max_temp < 20 THEN 'Inhumane'
9             ELSE 'No reading'
10        END
11     FROM temperature_readings
12 )
13 SELECT station_name, max_temperature_group, count(*)
14 FROM temps_collapsed
```

```

15 GROUP BY station_name, max_temperature_group
16 ORDER BY station_name, count(*) DESC;

```

Run this. Step 1 (CTE): classify every reading. Step 2 (outer query): count days per station per group. This two-step pattern – transform then aggregate – is the backbone of analytical queries.

7.4 Try It: Population Tiers

Using a CTE with CASE, classify each county as:

- ‘Metro’ (pop \geq 500,000)
- ‘Urban’ (pop 100,000 to 499,999)
- ‘Suburban’ (pop 50,000 to 99,999)
- ‘Rural’ (pop $<$ 50,000)

Then count the number of counties in each tier and compute the total population per tier. Which tier has the most counties? Which holds the most population?

```

1  -- Your query here
   ...
1  WITH county_tiers AS (
2     SELECT county_name, state_name, pop_est_2019,
3         CASE
4             WHEN pop_est_2019 >= 500000 THEN 'Metro'
5             WHEN pop_est_2019 >= 100000 THEN 'Urban'
6             WHEN pop_est_2019 >= 50000 THEN 'Suburban'
7             ELSE 'Rural'
8         END AS tier
9     FROM us_counties_pop_est_2019
10 )
11 SELECT tier,
12     count(*) AS num_counties,
13     sum(pop_est_2019) AS total_pop,
14     round(avg(pop_est_2019), 0) AS avg_county_pop
15 FROM county_tiers
16 GROUP BY tier
17 ORDER BY total_pop DESC;

```

Rural has the most counties by far. Metro has the most total population. The vast majority of U.S. counties are small.

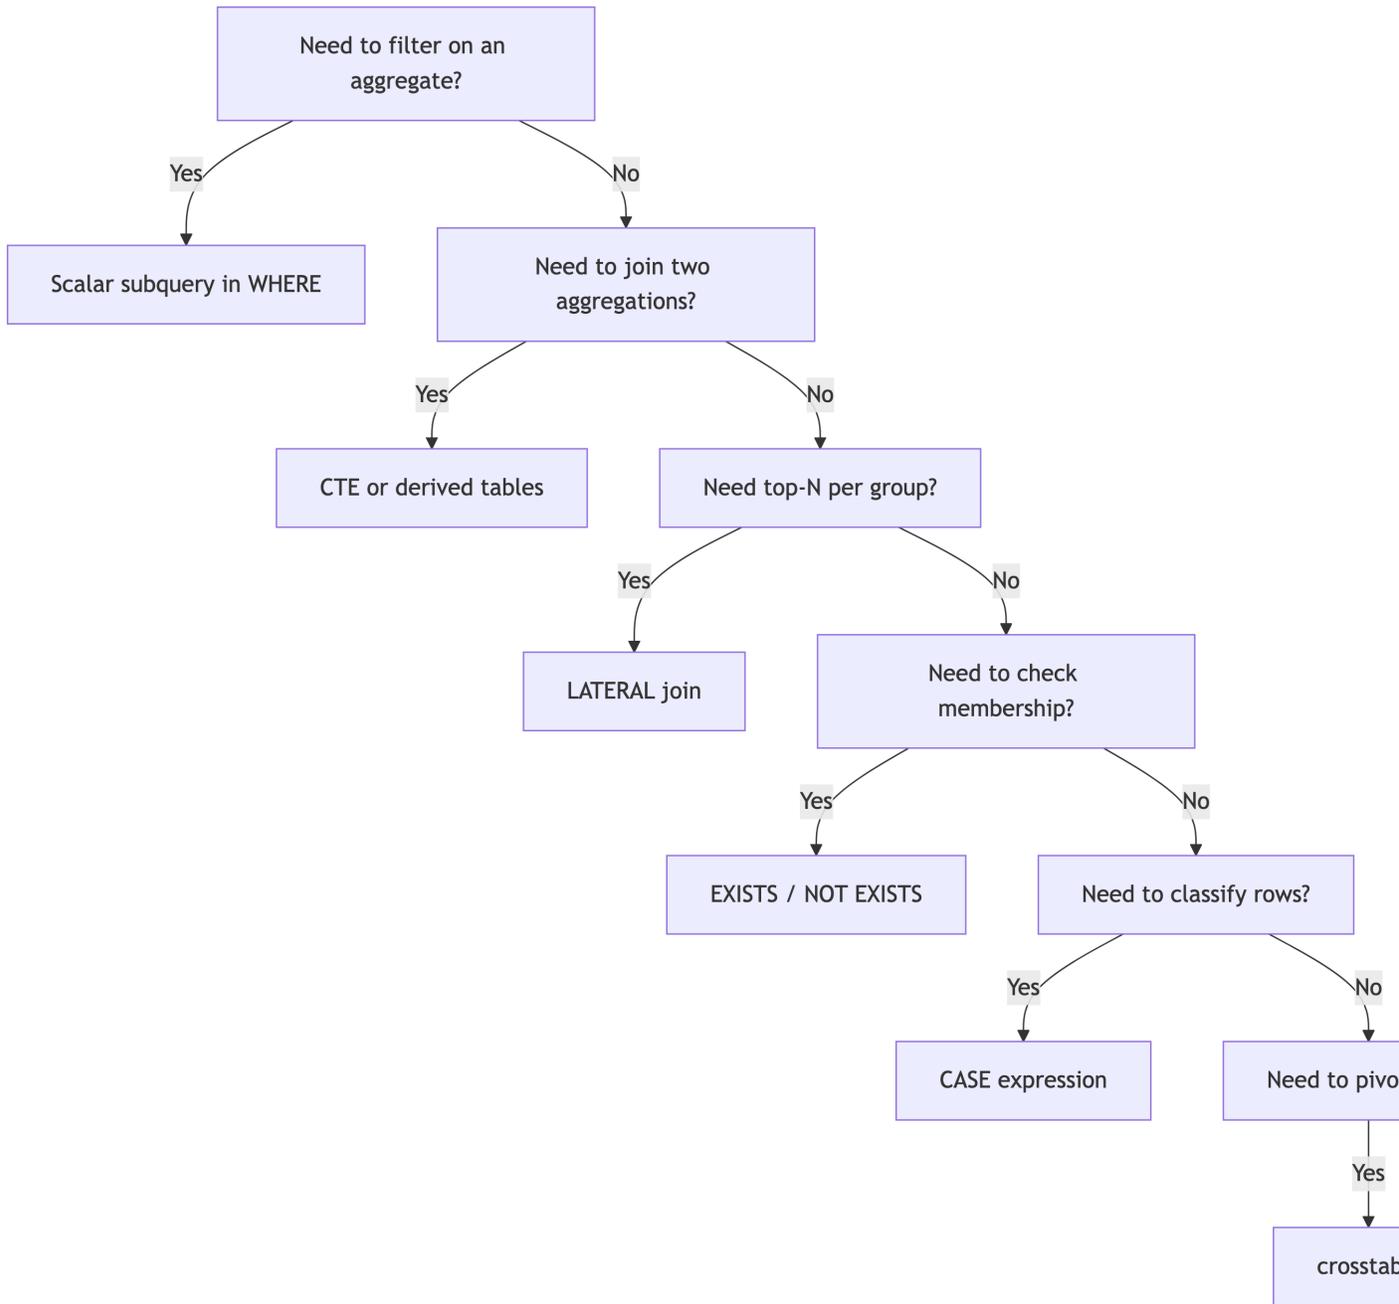
8 Part 8: Summary

8.1 What We Learned

8.2 Key Techniques

Technique	What It Does
Scalar subquery in WHERE	Compute a threshold, filter on it
Derived table (subquery in FROM)	Create a temporary table inline
Column subquery in SELECT	Add a computed column per row
Correlated subquery	Inner query references outer query (runs per row)
IN / EXISTS / NOT EXISTS	Check membership across tables
LATERAL join	Per-row subquery that can reference preceding tables
CTE (WITH)	Named temporary result sets for readability and reuse
crosstab()	Pivot long data to wide (requires <code>tablefunc</code> extension)
CASE	Conditional classification inside queries

8.3 When to Use What



8.4 The Big Ideas

1. **Subqueries go everywhere:** WHERE, FROM, SELECT, even inside other subqueries.
2. **CTEs > nested subqueries** for readability. Always. Fight me.
3. **LATERAL is the top-N-per-group tool.** Learn it. Love it.
4. **EXISTS short-circuits.** For large-table membership checks, prefer it over IN.
5. **CASE is your inline if/else.** Classify first, aggregate second.
6. **crosstab() pivots data.** Three columns in, pivot table out.

8.5 References

1. DeBarros, A. (2022). *Practical SQL* (2nd ed.). No Starch Press. Chapter 13.
2. PostgreSQL: [WITH Queries \(CTEs\)](#)
3. PostgreSQL: [Subquery Expressions](#)
4. PostgreSQL: [LATERAL Joins](#)
5. PostgreSQL: [tablefunc \(crosstab\)](#)