Lecture 6 SQL (Basic)

CS348 Spring 2025: Introduction to Database Management

> Instructor: Xiao Hu Sections: 001, 002, 003

Announcements

- Project Milestone o Due today
 - Not be graded
- Deadline extended for Assignment 1
 - From Thursday May 29 to Sunday Jun 1
- Online office hours by IAs for Assignment 1
 - See Piazza for Zoom information
 - Friday May 23 4pm 5pm
 - Wednesday May 28 4:30 pm 5:30pm

SQL features covered so far

Basic topics

- Data-definition language (DDL): define/modify schemas, drop relations
- Data-manipulation language (DML): query data
 - SELECT-FROM-WHERE statements
 - Set/Bag (DISTINCT, UNION/EXCEPT/INTERSECT (ALL))
 - Subqueries (table, scalar, IN, EXISTS, ALL, ANY)
 - Aggregation and grouping (GROUP BY, HAVING)
 - Ordering (ORDER)

 Using EXISTS, write a query to list the distinct IDs of groups that have at least two users

SELECT DISTINCT gid FROM Member m WHERE EXISTS (SELECT m1.gid FROM Member m1 WHERE m.gid = m1.gid and m.uid != m1.uid);

User (<u>uid</u> int, name string, age int, pop float) Group (<u>gid</u> string, name string) Member (<u>uid</u> int, <u>gid</u> string)

 Using WITH-AS and (NOT) IN, write a query to list the IDs of groups that Lisa belongs to but Ralph does not

WITH temp AS (SELECT gid FROM User u, Member m WHERE u.name = 'Ralph' and u.uid = m.uid)

SELECT DISTINCT gid FROM User u, Member m WHERE name = 'Lisa' AND u.uid = m.uid AND gid NOT IN (SELECT gid FROM temp);

User (<u>uid</u> int, name string, age int, pop float) Group (<u>gid</u> string, name string) Member (<u>uid</u> int, <u>gid</u> string)

• Write the same query but using EXCEPT (you may or may not use any other keywords)

SELECT gid FROM User u, Member m WHERE u.name = 'Lisa' AND u.uid = m.uid EXCEPT SELECT gid FROM User u, Member m WHERE u.name = 'Ralph' AND u.uid = m.uid;

User (<u>uid</u> int, name string, age int, pop float) Group (<u>gid</u> string, name string) Member (<u>uid</u> int, <u>gid</u> string)

What is next?

Basic topics

- Data-definition language (DDL): define/modify schemas, delete relations
- Data-manipulation language (DML): query data
 - SELECT-FROM-WHERE statements
 - Set/Bag (DISTINCT, UNION/EXCEPT/INTERSECT (ALL))
 - Subqueries (table, scalar, IN, EXISTS, ALL, ANY)
 - Aggregation and grouping (GROUP BY, HAVING)
 - Ordering (ORDER)

Next: NULL and JOIN

Incomplete information

- Example: User (<u>uid</u>, name, age, pop)
- Value unknown
 - We do not know Nelson's age
- Value not applicable
 - Suppose pop is based on interactions with others on our social networking site
 - Nelson is new to our site; what is their pop?

Possible Solution 1

Dedicate a value from each domain (type)

• pop cannot be -1, so use -1 as a special value to indicate a missing or invalid pop Incorrect answers

SELECT AVG(pop) FROM User WHERE pop != -1;

 Perhaps the value is not as special as you think!

SELECT AVG(pop) FROM User;

the Y2K bug



Complicated

http://www.90s411.com/images/y2k-cartoon.jpg

Possible Solution 2

- A valid-bit for every column
 - User (<u>uid</u>,

name_is_valid, age, age_is_valid, pop, pop_is_valid)

SELECT AVG(pop) FROM User WHERE pop_is_valid = 1;

- Complicates schema and queries
 - Need almost double the number of columns

Possible Solution 3

- Decompose the table; missing row = missing value
 - UserID (<u>uid</u>)
 - UserName (<u>uid</u>, name)
 - UserAge (uid, age)
 - UserPop (<u>uid</u>, pop)

- → Has a tuple for Nelson
- → Has a tuple for Nelson
 - \rightarrow No entry for Nelson
- → No entry for Nelson
- Conceptually the cleanest solution
- Still complicates schema and queries
 - How to get all information about users in a table?
 - Natural join doesn't work!

SQL's solution - NULL

- A special value NULL
 - For every domain (i.e., any datatype)
- Example: User (<u>uid</u>, name, age, pop)
 - <789, "Nelson", NULL, NULL>
- Special rules for dealing with NULL's

SELECT * FROM User WHERE name='Nelson' AND pop > 0.5 ??

Three-valued logic

TRUE = 1, FALSE = 0, UNKNOWN = 0.5 $x \text{ AND } y = \min(x, y)$ $x \text{ OR } y = \max(x, y)$ NOT x = 1 - x

X	у	x AND y	x OR y	NOT <i>x</i>
TRUE	TRUE	TRUE	TRUE	FALSE
TRUE	UNKNOWN	UNKNOWN	TRUE	FALSE
TRUE	FALSE	FALSE	TRUE	FALSE
UNKNOWN	TRUE	UNKNOWN	TRUE	UNKNOWN
UNKNOWN	UNKNOWN	UNKNOWN	UNKNOWN	UNKNOWN
UNKNOWN	FALSE	FALSE	UNKNOWN	UNKNOWN
FALSE	TRUE	FALSE	TRUE	TRUE
FALSE	UNKNOWN	FALSE	UNKNOWN	TRUE
FALSE	FALSE	FALSE	FALSE	TRUE

Other Rules

- Comparing a NULL with another value (including another NULL) using =, >, etc., the result is NULL
- WHERE and HAVING clauses only select rows for output if the condition evaluates to TRUE
 - FALSE and UNKNOWN are not sufficient
- Aggregate functions ignore NULL, except COUNT(*)
 - SUM, AVG, MIN, MAX all ignore NULLs
 - COUNT(age) also ignores NULL
 - If all inputs are NULL, SUM, AVG, MIN, MAX all return NULL

Will 789 be in the output?

(789, "Nelson", NULL, NULL)

SELECT uid FROM User WHERE name='Nelson' AND pop>0.5;

SELECT uid FROM User WHERE name='Nelson' OR pop>0.5;

SELECT uid FROM User WHERE NOT pop>0.5;

Unfortunate consequences

• Q1a = Q1b?

Q1a. SELECT AVG(pop) FROM User;

Q1b. SELECT SUM(pop)/COUNT(*) FROM User;

• Q2a = Q2b?

Q2a. SELECT * FROM User;	
Q2b. SELECT * FROM User WHERE pop = pop;	

• Be careful: NULL breaks many equivalences

Another problem

• Example: Who has NULL pop values?



 SQL introduced special, built-in predicates IS NULL and IS NOT NULL

SELECT * FROM User WHERE pop IS NULL;

GROUP BY with NULL

- In a GROUP BY clause, any NULL values in the grouping column(s) are treated all the same
 - they end up in one "NULL" group

SF	I FCT a	nge AVG(n	on), COUN	JT(*)		
FR		50; / 1 0(P	00,000		uid	name
GF	ROUP	BY age:			142	Bart
					123	Milhouse
	age	AVG(pop)	COUNT(*)		857	Lisa
	8	0.7	3		456	Nelson
	NULL	0.6	2		324	Alice

 If a group has only NULLs, then aggregate function (except COUNT(*)) returns NULL.

User

pop

0.9

0.7

0.3

NULL

NULL

age

8

8

8

NULL

NULL

18

User (<u>uid</u> int, name string, age int, pop float)¹⁹ Group (<u>gid</u> string, name string) Member (<u>uid</u> int, <u>gid</u> string)

• What is the output of these queries?

SELECTuid			l	Jser
FROM User	uid	name	age	рор
WHERE age > 5 OR pop < 0.5 ;	142	Bart	NULL	0.9
	123	Milhouse	8	NULL
	857	Lisa	8	0.7
SELECT uid	456	Nelson	8	NULL
FROM User	324	Alice	NULL	0.3
WHERE age > 5 AND pop < 0.5;				

SELECT AVG(pop), COUNT(*) FROM User WHERE age IS NOT NULL GROUP BY age;

User (<u>uid</u> int, name string, age int, pop float)²⁰ Group (<u>gid</u> string, name string) Member (<u>uid</u> int, <u>gid</u> string)

11----

• What is the output of these queries?

SELECT name FROM User WHERE age IN (SELECT age FROM User);

SELECT name FROM User WHERE age <= ANY(SELECT age FROM User) AND pop IS NOT NULL;

SELECT name FROM User WHERE age <= ALL(SELECT age FROM User);

		Ĺ	JSEI
uid	name	age	рор
142	Bart	NULL	0.9
123	Milhouse	8	NULL
857	Lisa	8	0.7
456	Nelson	8	NULL
324	Alice	NULL	0.3

User (<u>uid</u> int, name string, age int, pop float)²¹ Group (<u>gid</u> string, name string) Member (<u>uid</u> int, <u>gid</u> string)

• Write a query to find the ID of all users with non-null popularity who belong to at least one group.

SELECT uid FROM User WHERE pop IS NOT NULL AND uid IN (SELECT DISTINCT uid FROM Member);

JOIN

User (<u>uid</u> int, <u>uname</u> string, age int, pop float)²² Group (<u>gid</u> string, <u>gname</u> string) Member (<u>uid</u> int, <u>gid</u> string)

• Example: construct a group membership list with all groups and its members info (if non-empty)

SELECT g.gid, gname, u.uid, uname FROM Group g NATURAL JOIN Member m NATURAL JOIN User u;

SELECT g.gid, gname, u.uid, uname FROM Group g JOIN Member m ON g.gid = m.gid JOIN User u ON m.uid = u.uid;

SELECT g.gid, gname, u.uid, uname FROM Group g, Member m, User u WHERE g.gid = m.gid AND m.uid = u.uid;

Need a new JOIN operator

• Example: construct a master group membership list with all groups and its members info-(if non-empty)

SELECT g.gid, gname, u.uid, uname FROM Group g JOIN Member m ON g.gid = m.gid JOIN User u ON m.uid = u.uid;

- What if a group is empty?
- It may be reasonable for the master list to include empty groups as well
 - For these groups, user id and user name would be NULL

Outerjoin

- An extended relational algebra operator
- A full outerjoin ($R \bowtie S$) includes:
 - All rows in the natural join result of $R \bowtie S$, plus
 - "Dangling" *R* rows (those that do not join with any *S* rows) padded with NULL's for *S*'s columns
 - "Dangling" *S* rows (those that do not join with any *R* rows) padded with NULL's for *R*'s columns

Example of Outerjoin

Group

gid	gname
abc	Book Club
gov	Student Government
dps	Dead Putting Society
spr	Sports Club

mem		
uid	gid	
142	dps	
123	gov	
857	abc	
857	gov	
789	foo	

Member

Group ⋈ Member

gid	gname	uid
abc	Book Club	857
gov	Student Government	123
gov	Student Government	857
dps	Dead Putting Society	142
foo	NULL	789
spr	Sports Club	NULL

Left/Right Outerjoin

Group ⋈ Member

gid	gname	uid
abc	Book Club	857
gov	Student Government	123
gov	Student Government	857
dps	Dead Putting Society	142
spr	Sports Club	NULL

• A left outerjoin $(R \bowtie S)$ includes rows in $R \bowtie S$ plus dangling R rows padded with NULL's

Curry C Manakara	gid	gname	uid
Group 🔍 member	abc	Book Club	857
	gov	Student Government	123
	gov	Student Government	857
Z	dps	Dead Putting Society	142
'	foo	NULL	789

• A right outerjoin $(R \bowtie S)$ includes rows in $R \bowtie S$ plus dangling S rows padded with NULL's

Group

gid	gname
abc	Book Club
gov	Student Government
dps	Dead Putting Society
spr	Sports Club

Member				
uid	gid			
142	dps			
123	gov			
857	abc			
857	gov			
789	foo			

Outerjoin in SQL

SELECT * FROM Group JOIN Member ON Group.gid = Member.gid;

SELECT * FROM Group NATURAL JOIN Member;

SELECT * FROM Group LEFT OUTER JOIN Member ON Group.gid = Member.gid;

SELECT * FROM Group **RIGHT OUTER JOIN** Member **ON** Group.gid = Member.gid;

SELECT * FROM Group FULL OUTER JOIN Member ON Group.gid = Member.gid; $\approx Group \underset{Group.gid=Member.gid}{\bowtie} Member$

 $pprox Group \Join Member$

 $\approx Group \underset{Group.gid=Member.gid}{\rightarrowtail} Member$

 $\approx Group_{Group.gid=Member.gid} Member$

 $\approx Group \underset{Group.gid=Member.gid}{\rightarrowtail} Member$

User (<u>uid</u> int, <u>uname</u> string, age int, pop float)²⁸ Group (<u>gid</u> string, <u>gname</u> string) Member (<u>uid</u> int, <u>gid</u> string)

• What is the output of these queries?

Group		User				Member		
gid	gname	uid	uname	age	рор		uid	gid
abc	Book Club	142	Bart	10	0.9		857	dps
gov	Student Government	123	Milhouse	10	NULL		123	gov
dps	Dead Putting Society	857	Lisa	8	0.7		857	abc
spr	Sports Club	456	Alice	8	NULL		123	abc

SELECT uname, gname FROM User u NATURAL JOIN Member m NATURAL JOIN Group g;

SELECT uname, gid FROM User u LEFT OUTER JOIN Member m ON u.uid = m.uid; SELECT COUNT(uid), COUNT(gname) FROM Member m RIGHT OUTER JOIN Group g ON g.gid = m.gid;

SQL features covered so far

Basic topics

- Data-definition language (DDL): define/modify schemas, delete relations
- Data-manipulation language (DML): query data
 - SELECT-FROM-WHERE statements
 - Set/Bag (DISTINCT, UNION/EXCEPT/INTERSECT (ALL))
 - Subqueries (table, scalar, IN, EXISTS, ALL, ANY)
 - Aggregation and grouping (GROUP BY, HAVING)
 - Ordering (ORDER)
 - NULL and JOIN

Next: Modify data (INSERT/DELETE/UPDATE)

INSERT

User (<u>uid</u> int, name string, age int, pop float) Group (<u>gid</u> string, name string) Member (<u>uid</u> int, <u>gid</u> string)

- Insert one row
 - User 789 joins Dead Putting Society

INSERT INTO Member VALUES (789, 'dps');

INSERT INTO User (uid, name) VALUES (389, 'Marge');

- Insert the result of a query
 - Everybody joins Dead Putting Society!

INSERT INTO Member (SELECT uid, 'dps' FROM User WHERE uid NOT IN (SELECT uid FROM Member WHERE gid = 'dps'));

DELETE

User (<u>uid</u> int, name string, age int, pop float) Group (<u>gid</u> string, name string) Member (<u>uid</u> int, <u>gid</u> string)

• Delete everything from a table

DELETE FROM Member;

- Delete according to a WHERE condition
 - User 789 leaves Dead Putting Society

DELETE FROM Member WHERE uid=789 AND gid='dps';

 Users over age 18 must be removed from Book Club DELETE FROM Member WHERE uid IN (SELECT uid FROM User WHERE age > 18) AND gid = 'abc';

DELETE m, u FROM Member m NATURAL JOIN User u WHERE age > 18 AND gid = 'abc';

UPDATE

User (<u>uid</u> int, name string, age int, pop float) Group (<u>gid</u> string, name string) Member (<u>uid</u> int, <u>gid</u> string)

• Example: User 142 changes name to "Barney"

UPDATE User SET name = 'Barney' WHERE uid = 142;

• Example: We are all popular!

UPDATE User SET pop = (SELECT AVG(pop) FROM User);

• But won't update of every row causes average pop to change?

[©] Subquery is always computed over the old table

• What is the output of these queries?

Group		User				Member		
gid	gname	uid	uname	age	рор		uid	gid
abc	Book Club	142	Bart	10	0.9		857	dps
gov	Student Government	123	Milhouse	10	NULL		123	gov
dps	Dead Putting Society	857	Lisa	8	0.7		857	abc
spr	Sports Club	456	Alice	8	NULL		123	abc

DELETE m, g FROM Member m NATURAL JOIN Group g WHERE g.gid = 'dps';

INSERT INTO Member (SELECT uid, 'spr' FROM User WHERE age >= 10 AND pop IS NOT NULL);

UPDATE User u **NATURAL JOIN** Member m **SET** u.age = 11, u.pop = 0.4, m.gid = 'spr' WHERE u.uid = 123 and m.gid = 'gov';

What is next?

Basic topics

- Data-definition language (DDL): define/modify schemas, drop relations
- Data-manipulation language (DML): query data
 - SELECT-FROM-WHERE statements
 - Set/Bag (DISTINCT, UNION/EXCEPT/INTERSECT (ALL))
 - Subqueries (table, scalar, IN, EXISTS, ALL, ANY)
 - Aggregation and grouping (GROUP BY, HAVING)
 - Ordering (ORDER)
 - NULL and JOIN

and modify data (INSERT/DELETE/UPDATE)

Constraints