

Introduction to Data Management SQL Subqueries

Paul G. Allen School of Computer Science and Engineering
University of Washington, Seattle

Announcements

- HW1 is graded
- HW2 will probably be done early next week
- HW3 due next Friday
- Midterm on Friday, 4/26 in class
 - Closed books, no cheat sheet (you won't need it)

Predicates on Subqueries

- EXISTS (SELECT) checks if it is not empty
NOT EXISTS (SELECT ...) checks if it is empty

- X in (SELECT Y FROM ...) checks output has X
X not in (SELECT Y ...) checks if it doesn't have X

- X > ALL(SELECT ...)
X > ANY(SELECT ...)
checks if X is > than one or all values in output

Recap: EXISTS

Find people who **do** drive cars

```
SELECT P.UserID, P.Name
FROM Payroll P
WHERE exists
    (SELECT *
     FROM Regist R
     WHERE P.UserID = R.UserID);
```

Payroll

UserID	Name	Job	Salary
123	Jack	TA	50000
345	Allison	TA	60000
567	Magda	Prof	90000
789	Dan	Prof	100000

Regist

UserID	Car
123	Charger
567	Civic
567	Pinto

Recap: EXISTS

Find people who **do** drive cars

```
SELECT P.UserID, P.Name
FROM Payroll P
WHERE exists
  (SELECT *
   FROM Regist R
   WHERE P.UserID = R.UserID);
```

Payroll

UserID	Name	Job	Salary
123	Jack	TA	50000
345	Allison	TA	60000
567	Magda	Prof	90000
789	Dan	Prof	100000

Regist

UserID	Car
123	Charger
567	Civic
567	Pinto



UserID	Name
123	Jack
567	Magda

Recap: EXISTS

Find people who **do** drive cars

```
SELECT P.UserID, P.Name
FROM Payroll P
WHERE exists
  (SELECT *
   FROM Regist R
   WHERE P.UserID = R.UserID);
```


Same as
a semi-join

Payroll

UserID	Name	Job	Salary
123	Jack	TA	50000
345	Allison	TA	60000
567	Magda	Prof	90000
789	Dan	Prof	100000

Regist

UserID	Car
123	Charger
567	Civic
567	Pinto



UserID	Name
123	Jack
567	Magda

Recap: NOT EXISTS

Find people who **do not** drive cars

Payroll

UserID	Name	Job	Salary
123	Jack	TA	50000
345	Allison	TA	60000
567	Magda	Prof	90000
789	Dan	Prof	100000

Regist

UserID	Car
123	Charger
567	Civic
567	Pinto

Recap: NOT EXISTS

Find people who **do not** drive cars

```
SELECT P.UserID, P.Name
FROM Payroll P
WHERE exists
  (SELECT *
   FROM Regist R
   WHERE P.UserID != R.UserID);
```

Does this work?

Payroll

UserID	Name	Job	Salary
123	Jack	TA	50000
345	Allison	TA	60000
567	Magda	Prof	90000
789	Dan	Prof	100000

Regist

UserID	Car
123	Charger
567	Civic
567	Pinto

Recap: NOT EXISTS

Find people who **do not** drive cars

```
SELECT P.UserID, P.Name
FROM Payroll P
WHERE exists
  (SELECT *
   FROM Regist R
   WHERE P.UserID != R.UserID);
```

Does this work?

Payroll

UserID	Name	Job	Salary
123	Jack	TA	50000
345	Allison	TA	60000
567	Magda	Prof	90000
789	Dan	Prof	100000

Regist

UserID	Car
123	Charger
567	Civic
567	Pinto

NO!
It returns everybody

Recap: NOT EXISTS

Find people who **do not** drive cars

Payroll

UserID	Name	Job	Salary
123	Jack	TA	50000
345	Allison	TA	60000
567	Magda	Prof	90000
789	Dan	Prof	100000

Regist

UserID	Car
123	Charger
567	Civic
567	Pinto

Recap: NOT EXISTS

Find people who **do not** drive cars

```
SELECT P.UserID, P.Name
FROM Payroll P
WHERE not exists
    (SELECT *
     FROM Regist R
     WHERE P.UserID = R.UserID);
```

Payroll

UserID	Name	Job	Salary
123	Jack	TA	50000
345	Allison	TA	60000
567	Magda	Prof	90000
789	Dan	Prof	100000

Regist

UserID	Car
123	Charger
567	Civic
567	Pinto

Recap: NOT EXISTS

Find people who **do not** drive cars

```
SELECT P.UserID, P.Name
FROM Payroll P
WHERE not exists
  (SELECT *
   FROM Regist R
   WHERE P.UserID = R.UserID);
```

Payroll

UserID	Name	Job	Salary
123	Jack	TA	50000
345	Allison	TA	60000
567	Magda	Prof	90000
789	Dan	Prof	100000

Regist

UserID	Car
123	Charger
567	Civic
567	Pinto



UserID	Name
345	Allison
789	Dan

Unnesting EXISTS

Find people who **do** drive cars

```
SELECT P.UserID, P.Name
FROM Payroll P
WHERE exists
    (SELECT *
     FROM Regist R
     WHERE P.UserID = R.UserID);
```




```
SELECT DISTINCT P.UserID, P.Name
FROM Payroll P, Regist R
WHERE P.UserID = R.UserID;
```

How do we unnest NOT EXISTS?

Find people who **do not** drive cars

```
SELECT P.UserID, P.Name
FROM Payroll P
WHERE not exists
    (SELECT *
     FROM Regist R
     WHERE P.UserID = R.UserID);
```

This doesn't work...



```
SELECT DISTINCT P.UserID, P.Name
FROM Payroll P, Regist R
WHERE P.UserID != R.UserID;
```

How do we unnest NOT EXISTS?

Find people who **do not** drive cars

```
SELECT P.UserID, P.Name
FROM Payroll P
WHERE not exists
    (SELECT *
     FROM Regist R
     WHERE P.UserID = R.UserID);
```

This query cannot be unnested without aggregates.

Proof next

Monotone Functions

Definition

A function $f: R \rightarrow R$ is monotone if $x \leq y$ implies $f(x) \leq f(y)$


Monotone:

$$x^3 + x^2,$$

$$e^x,$$

$$\log(x),$$

...



Non-Monotone:

$$x^3 - x^2,$$

$$e^{-x},$$


$$\frac{1}{x},$$

...

Monotone Queries

Definition

A query Q is monotone if $I \subseteq J$ implies $q(I) \subseteq q(J)$




Monotone Queries

Definition

A query Q is monotone if $I \subseteq J$ implies $q(I) \subseteq q(J)$

Adding tuples to the input does
not remove tuples from the output



Monotone Queries

Find people who **do** drive cars

Is this query monotone?

Monotone Queries

Find people who **do** drive cars

Is this query monotone?

Payroll

UserID	Name	Job	Salary
123	Jack	TA	50000
345	Allison	TA	60000
567	Magda	Prof	90000
789	Dan	Prof	100000

Regist

UserID	Car
123	Charger
567	Civic
567	Pinto

I

Monotone Queries

Find people who **do** drive cars

Is this query monotone?

Payroll

UserID	Name	Job	Salary
123	Jack	TA	50000
345	Allison	TA	60000
567	Magda	Prof	90000
789	Dan	Prof	100000

Regist

UserID	Car
123	Charger
567	Civic
567	Pinto

I

UserID	Name
123	Jack
567	Magda

$q(I)$

Monotone Queries

Find people who **do** drive cars

Is this query monotone?

Payroll

UserID	Name	Job	Salary
123	Jack	TA	50000
345	Allison	TA	60000
567	Magda	Prof	90000
789	Dan	Prof	100000

Regist

UserID	Car
123	Charger
567	Civic
567	Pinto

I

UserID	Name
123	Jack
567	Magda

$q(I)$

Payroll

UserID	Name	Job	Salary
123	Jack	TA	50000
345	Allison	TA	60000
567	Magda	Prof	90000
789	Dan	Prof	100000

Regist

UserID	Car
123	Charger
567	Civic
567	Pinto
345	Tesla

J

Monotone Queries

Find people who **do** drive cars

Is this query monotone?

Payroll

UserID	Name	Job	Salary
123	Jack	TA	50000
345	Allison	TA	60000
567	Magda	Prof	90000
789	Dan	Prof	100000

Regist

UserID	Car
123	Charger
567	Civic
567	Pinto

I

UserID	Name
123	Jack
567	Magda

$q(I)$

Payroll

UserID	Name	Job	Salary
123	Jack	TA	50000
345	Allison	TA	60000
567	Magda	Prof	90000
789	Dan	Prof	100000

Regist

UserID	Car
123	Charger
567	Civic
567	Pinto
345	Tesla

J

UserID	Name
123	Jack
567	Magda
345	Allison

$q(J)$

Monotone Queries

Find people who **do** drive cars

Is this query monotone?

Yes, it is monotone

Monotone Queries

Find people who **do not** drive cars

Is this query monotone?

Payroll

UserID	Name	Job	Salary
123	Jack	TA	50000
345	Allison	TA	60000
567	Magda	Prof	90000
789	Dan	Prof	100000

Regist

UserID	Car
123	Charger
567	Civic
567	Pinto

I

UserID	Name
345	Allison
789	Dan

$q(I)$

Monotone Queries

Find people who **do not** drive cars

Is this query monotone?

Payroll

UserID	Name	Job	Salary
123	Jack	TA	50000
345	Allison	TA	60000
567	Magda	Prof	90000
789	Dan	Prof	100000

Regist

UserID	Car
123	Charger
567	Civic
567	Pinto

I

UserID	Name
345	Allison
789	Dan

$q(I)$

Payroll

UserID	Name	Job	Salary
123	Jack	TA	50000
345	Allison	TA	60000
567	Magda	Prof	90000
789	Dan	Prof	100000

Regist

UserID	Car
123	Charger
567	Civic
567	Pinto
345	Tesla

J

Monotone Queries

Find people who **do not** drive cars

Is this query monotone?

Payroll

UserID	Name	Job	Salary
123	Jack	TA	50000
345	Allison	TA	60000
567	Magda	Prof	90000
789	Dan	Prof	100000

Regist

UserID	Car
123	Charger
567	Civic
567	Pinto

I

UserID	Name
345	Allison
789	Dan

$q(I)$

Payroll

UserID	Name	Job	Salary
123	Jack	TA	50000
345	Allison	TA	60000
567	Magda	Prof	90000
789	Dan	Prof	100000

Regist

UserID	Car
123	Charger
567	Civic
567	Pinto
345	Tesla

J

UserID	Name
345	Allison
789	Dan

$q(J)$

Monotone Queries

Find people who **do not** drive cars

Is this query monotone?

No, this query is not monotone

Monotone Queries

Theorem

Every SELECT-FROM-WHERE query without subqueries and without aggregates is monotone

Monotone Queries

Theorem

Every SELECT-FROM-WHERE query without subqueries and without aggregates is monotone

Proof. Consider a SQL query:

```
SELECT attrs
FROM T1, T2, ...
WHERE condition
```

Monotone Queries

Theorem

Every SELECT-FROM-WHERE query without subqueries and without aggregates is monotone

Proof. Consider a SQL query:

```
SELECT attrs
FROM T1, T2, ...
WHERE condition
```

Its nested loop semantics is:

```
for each r1 in T1:
  for each t2 in T2:
    for each t3 in T3:
      ...
      if (condition):
        output (r1, r2, ...)
```

If we insert a tuple into one of the input relations T_i , we will not remove any tuples from the output.

Consequence

The query “Find people who **do not** drive cars” cannot be unnested without aggregates

- The property whether the query is monotone or not does not depend on its SQL writeup
- Instead, it depends on the meaning of the query, regardless of how we write it in SQL

Monotone Queries

Count the number of employees.

```
SELECT count (*)  
FROM Payroll
```

Monotone?

Monotone Queries

Count the number of employees.

```
SELECT count (*)  
FROM Payroll
```

Monotone?

Payroll

UserID	Name	Job	Salary
123	Jack	TA	50000
345	Allison	TA	60000
567	Magda	Prof	90000
789	Dan	Prof	100000

I

count

4

$q(I)$

Monotone Queries

Count the number of employees.

```
SELECT count (*)  
FROM Payroll
```

Monotone?

Payroll

UserID	Name	Job	Salary
123	Jack	TA	50000
345	Allison	TA	60000
567	Magda	Prof	90000
789	Dan	Prof	100000

I

count

4

$q(I)$

Payroll

UserID	Name	Job	Salary
123	Jack	TA	50000
345	Allison	TA	60000
567	Magda	Prof	90000
789	Dan	Prof	100000
555	Alice	TA	80000

J

count

5

$q(J)$

Monotone Queries

Count the number of employees.

```
SELECT count (*)  
FROM Payroll
```

Monotone?

Payroll

UserID	Name	Job	Salary
123	Jack	TA	50000
345	Allison	TA	60000
567	Magda	Prof	90000
789	Dan	Prof	100000

I

count

4

$q(I)$

Payroll

UserID	Name	Job	Salary
123	Jack	TA	50000
345	Allison	TA	60000
567	Magda	Prof	90000
789	Dan	Prof	100000
555	Alice	TA	80000

J

count

5

$q(J)$

$\{4\} \not\subseteq \{5\}$

Monotone Queries

Count the number of employees.

```
SELECT count (*)  
FROM Payroll
```

Monotone?

Payroll

UserID	Name	Job	Salary
123	Jack	TA	50000
345	Allison	TA	60000
567	Magda	Prof	90000
789	Dan	Prof	100000

I

count

4

$q(I)$

Payroll

UserID	Name	Job	Salary
123	Jack	TA	50000
345	Allison	TA	60000
567	Magda	Prof	90000
789	Dan	Prof	100000
555	Alice	TA	80000

J

count

5

$q(J)$

Not
monotone

$\{4\} \not\subseteq \{5\}$

IN and NOT IN

Subqueries in WHERE/HAVING

X in (SELECT Y FROM ...)


- Compute the subquery
- Check if $X \in Output$

Y
...
...

Subqueries in WHERE/HAVING

X in (SELECT Y FROM ...)

- Compute the subquery
- Check if $X \in Output$




X not in (SELECT Y ...) **not(X in (SELECT Y ...))**

- Compute the subquery
- Check if $X \notin Output$

EXISTS v.s. IN

Find people who **do** drive cars

```
SELECT P.UserID, P.Name
FROM Payroll P
WHERE exists
    (SELECT *
     FROM Regist R
     WHERE P.UserID = R.UserID);
```



Same output


```
SELECT P.UserID, P.Name
FROM Payroll P
WHERE P.UserID in
    (SELECT R.UserID
     FROM Regist R);
```

NOT EXISTS v.s. NOT IN

Find people who **do not** drive cars

```
SELECT P.UserID, P.Name
FROM Payroll P
WHERE not exists
    (SELECT *
     FROM Regist R
     WHERE P.UserID = R.UserID);
```

Same output



```
SELECT P.UserID, P.Name
FROM Payroll P
WHERE P.UserID not in
    (SELECT R.UserID
     FROM Regist R);
```

Computing NOT IN

Find people who **do not** drive cars

```
SELECT P.UserID, P.Name
FROM Payroll P
WHERE P.UserID not in
      (SELECT R.UserID
       FROM Regist R);
```

1. Compute subquery

Payroll

UserID	Name	Job	Salary
123	Jack	TA	50000
345	Allison	TA	60000
567	Magda	Prof	90000
789	Dan	Prof	100000

Regist

UserID	Car
123	Charger
567	Civic
567	Pinto

Computing NOT IN

Find people who **do not** drive cars

```
SELECT P.UserID, P.Name
FROM Payroll P
WHERE P.UserID not in
      (SELECT R.UserID
       FROM Regist R);
```

1. Compute subquery

UserID
123
567
567

Payroll

UserID	Name	Job	Salary
123	Jack	TA	50000
345	Allison	TA	60000
567	Magda	Prof	90000
789	Dan	Prof	100000

Regist

UserID	Car
123	Charger
567	Civic
567	Pinto

Computing NOT IN

Find people who **do not** drive cars

```
SELECT P.UserID, P.Name
FROM Payroll P
WHERE P.UserID not in
      (SELECT R.UserID
       FROM Regist R);
```

1. Compute subquery

UserID
123
567
567

2. For each Payroll,
check if UserID \notin subquery

Payroll

UserID	Name	Job	Salary
123	Jack	TA	50000
345	Allison	TA	60000
567	Magda	Prof	90000
789	Dan	Prof	100000

Regist

UserID	Car
123	Charger
567	Civic
567	Pinto



UserID	Name
345	Allison
789	Dan

NOT EXISTS v.s. NOT IN

Find people who **do not** drive cars

```
SELECT P.UserID, P.Name
FROM Payroll P
WHERE not exists
    (SELECT *
     FROM Regist R
     WHERE P.UserID = R.UserID);
```

Which query is
more efficient?

```
SELECT P.UserID, P.Name
FROM Payroll P
WHERE P.UserID not in
    (SELECT R.UserID
     FROM Regist R);
```

NOT EXISTS v.s. NOT IN

Find people who **do not** drive cars

```
SELECT P.UserID, P.Name  
FROM Payroll P  
WHERE not exists  
  (SELECT *  
   FROM Regist R  
   WHERE P.UserID = R.UserID);
```

Correlated subquery:
computed repeatedly,
once for each Payroll

Which query is
more efficient?

Computed only once

```
SELECT P.UserID, P.Name  
FROM Payroll P  
WHERE P.UserID not in  
  (SELECT R.UserID  
   FROM Regist R);
```


ANY and ALL

ANY and ALL

< or <= or > or ...

$X < ANY (SELECT Y FROM ...)$

- Compute the subquery
- Check if there exists $Y \in Output$ s.t. $X < Y$


Y
...
...

ANY and ALL

< or <= or > or ...

$X < ANY (SELECT Y FROM ...)$

- Compute the subquery
- Check if there exists $Y \in Output$ s.t. $X < Y$



$X < ALL (SELECT Y FROM ...)$

- Compute the subquery
- Check if for all $Y \in Output$, $X < Y$

ANY and ALL

Find people who drive **some** car made after 2017

Payroll

UserID	Name	Job	Salary
123	Jack	TA	50000
345	Allison	TA	60000
567	Magda	Prof	90000
789	Dan	Prof	100000

Regist

UserID	Car	Year
123	Charger	2016
123	Tesla	2018
567	Civic	2020
567	Pinto	2022

ANY and ALL

Find people who drive **some** car made after 2017

```
SELECT P.*
FROM Payroll P
WHERE 2017 <
  ANY (SELECT R.Year
        FROM Regist R
        WHERE P.UserID = R.UserID);
```

Payroll

UserID	Name	Job	Salary
123	Jack	TA	50000
345	Allison	TA	60000
567	Magda	Prof	90000
789	Dan	Prof	100000

Regist

UserID	Car	Year
123	Charger	2016
123	Tesla	2018
567	Civic	2020
567	Pinto	2022

ANY and ALL

Find people who drive **some** car made after 2017

```
SELECT P.*
FROM Payroll P
WHERE 2017 <
  ANY (SELECT R.Year
        FROM Regist R
        WHERE P.UserID = R.UserID);
```

Jack:


Year
2016
2018

Payroll

UserID	Name	Job	Salary
123	Jack	TA	50000
345	Allison	TA	60000
567	Magda	Prof	90000
789	Dan	Prof	100000

Regist

UserID	Car	Year
123	Charger	2016
123	Tesla	2018
567	Civic	2020
567	Pinto	2022



Name	...
Jack	

ANY and ALL

Find people who drive **some** car made after 2017

```
SELECT P.*
FROM Payroll P
WHERE 2017 <
  ANY (SELECT R.Year
        FROM Regist R
        WHERE P.UserID = R.UserID);
```

Allison:

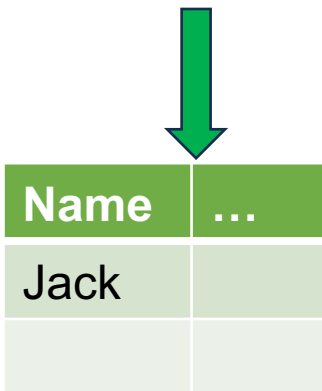
Year

Payroll

UserID	Name	Job	Salary
123	Jack	TA	50000
345	Allison	TA	60000
567	Magda	Prof	90000
789	Dan	Prof	100000

Regist

UserID	Car	Year
123	Charger	2016
123	Tesla	2018
567	Civic	2020
567	Pinto	2022



Name	...
Jack	

ANY and ALL

Find people who drive **some** car made after 2017

```
SELECT P.*
FROM Payroll P
WHERE 2017 <
  ANY (SELECT R.Year
        FROM Regist R
        WHERE P.UserID = R.UserID);
```

Magda:


Year
2020
2022

Payroll

UserID	Name	Job	Salary
123	Jack	TA	50000
345	Allison	TA	60000
567	Magda	Prof	90000
789	Dan	Prof	100000

Regist

UserID	Car	Year
123	Charger	2016
123	Tesla	2018
567	Civic	2020
567	Pinto	2022



Name	...
Jack	
Magda	

ANY and ALL

Find people who drive **some** car made after 2017

```
SELECT P.*  
FROM Payroll P  
WHERE 2017 <  
  ANY (SELECT R.Year  
        FROM Regist R  
        WHERE P.UserID = R.UserID);
```

Dan:


Year

Payroll

UserID	Name	Job	Salary
123	Jack	TA	50000
345	Allison	TA	60000
567	Magda	Prof	90000
789	Dan	Prof	100000

Regist

UserID	Car	Year
123	Charger	2016
123	Tesla	2018
567	Civic	2020
567	Pinto	2022



Name	...
Jack	
Magda	

ANY and ALL

Find people who drive **some** car made after 2017

```
SELECT P.*
FROM Payroll P
WHERE 2017 <
  ANY (SELECT R.Year
        FROM Regist R
        WHERE P.UserID = R.UserID);
```

Same as
a semi-join

```
SELECT DISTINCT P.*
FROM Payroll P, Regist R
WHERE P.UserID = R.UserID
and R.Year > 2017
```

Payroll

UserID	Name	Job	Salary
123	Jack	TA	50000
345	Allison	TA	60000
567	Magda	Prof	90000
789	Dan	Prof	100000

Regist

UserID	Car	Year
123	Charger	2016
123	Tesla	2018
567	Civic	2020
567	Pinto	2022

Name	...
Jack	
Magda	

ANY and ALL

Find people who drive **only** cars made after 2017

```
SELECT P.*
FROM Payroll P
WHERE 2017 <
  ALL (SELECT R.Year
        FROM Regist R
        WHERE P.UserID = R.UserID);
```

Payroll

UserID	Name	Job	Salary
123	Jack	TA	50000
345	Allison	TA	60000
567	Magda	Prof	90000
789	Dan	Prof	100000

Regist

UserID	Car	Year
123	Charger	2016
123	Tesla	2018
567	Civic	2020
567	Pinto	2022

ANY and ALL

Find people who drive **only** cars made after 2017

```
SELECT P.*
FROM Payroll P
WHERE 2017 <
  ALL (SELECT R.Year
        FROM Regist R
        WHERE P.UserID = R.UserID);
```

Jack:

Year
2016
2018

Do we output Jack?

Payroll

UserID	Name	Job	Salary
123	Jack	TA	50000
345	Allison	TA	60000
567	Magda	Prof	90000
789	Dan	Prof	100000

Regist

UserID	Car	Year
123	Charger	2016
123	Tesla	2018
567	Civic	2020
567	Pinto	2022

Name	...

ANY and ALL

Find people who drive **only** cars made after 2017

```
SELECT P.*
FROM Payroll P
WHERE 2017 <
  ALL (SELECT R.Year
        FROM Regist R
        WHERE P.UserID = R.UserID);
```

Jack:

Year
2016
2018

Do we output Jack?

The test $2017 < \text{ALL}(\dots)$
fails for 2016

Payroll

UserID	Name	Job	Salary
123	Jack	TA	50000
345	Allison	TA	60000
567	Magda	Prof	90000
789	Dan	Prof	100000

Regist

UserID	Car	Year
123	Charger	2016
123	Tesla	2018
567	Civic	2020
567	Pinto	2022

Name	...

ANY and ALL

Find people who drive **only** cars made after 2017

```
SELECT P.*
FROM Payroll P
WHERE 2017 <
  ALL (SELECT R.Year
        FROM Regist R
        WHERE P.UserID = R.UserID);
```

Allison:

Year


Do we output Allison?

Payroll

UserID	Name	Job	Salary
123	Jack	TA	50000
345	Allison	TA	60000
567	Magda	Prof	90000
789	Dan	Prof	100000

Regist

UserID	Car	Year
123	Charger	2016
123	Tesla	2018
567	Civic	2020
567	Pinto	2022



Name	...
Allison	

ANY and ALL

Find people who drive **only** cars made after 2017

```
SELECT P.*
FROM Payroll P
WHERE 2017 <
  ALL (SELECT R.Year
        FROM Regist R
        WHERE P.UserID = R.UserID);
```

Allison:

Year

Do we output Allison?


The test $2017 < ALL(...)$
does not fail anywhere.
Hence it holds!!

Payroll

UserID	Name	Job	Salary
123	Jack	TA	50000
345	Allison	TA	60000
567	Magda	Prof	90000
789	Dan	Prof	100000

Regist

UserID	Car	Year
123	Charger	2016
123	Tesla	2018
567	Civic	2020
567	Pinto	2022



Name	...
Allison	

ANY and ALL

Find people who drive **only** cars made after 2017

```
SELECT P.*
FROM Payroll P
WHERE 2017 <
  ALL (SELECT R.Year
        FROM Regist R
        WHERE P.UserID = R.UserID);
```

Magda:

Year
2020
2022


The test $2017 < ALL(\dots)$ does not fail anywhere.

Payroll

UserID	Name	Job	Salary
123	Jack	TA	50000
345	Allison	TA	60000
567	Magda	Prof	90000
789	Dan	Prof	100000

Regist

UserID	Car	Year
123	Charger	2016
123	Tesla	2018
567	Civic	2020
567	Pinto	2022



Name	...
Allison	
Magda	

ANY and ALL

Find people who drive **only** cars made after 2017

```
SELECT P.*
FROM Payroll P
WHERE 2017 <
  ALL (SELECT R.Year
        FROM Regist R
        WHERE P.UserID = R.UserID);
```

Dan:

Year


The test $2017 < ALL(...)$ does not fail anywhere. Hence it holds!!

Payroll

UserID	Name	Job	Salary
123	Jack	TA	50000
345	Allison	TA	60000
567	Magda	Prof	90000
789	Dan	Prof	100000

Regist

UserID	Car	Year
123	Charger	2016
123	Tesla	2018
567	Civic	2020
567	Pinto	2022



Name	...
Allison	
Magda	
Dan	

ANY and ALL

Find people who drive **only** cars made after 2017

```
SELECT P.*
FROM Payroll P
WHERE 2017 <
  ALL (SELECT R.Year
        FROM Regist R
        WHERE P.UserID = R.UserID);
```

Can we unnest this query?

Payroll

UserID	Name	Job	Salary
123	Jack	TA	50000
345	Allison	TA	60000
567	Magda	Prof	90000
789	Dan	Prof	100000

Regist

UserID	Car	Year
123	Charger	2016
123	Tesla	2018
567	Civic	2020
567	Pinto	2022

Name	...
Allison	
Magda	
Dan	

ANY and ALL

Find people who drive **only** cars made after 2017

```
SELECT P.*
FROM Payroll P
WHERE 2017 <
  ALL (SELECT R.Year
        FROM Regist R
        WHERE P.UserID = R.UserID);
```

Can we unnest this query?

NO!
Non-monotone

Payroll

UserID	Name	Job	Salary
123	Jack	TA	50000
345	Allison	TA	60000
567	Magda	Prof	90000
789	Dan	Prof	100000

Regist

UserID	Car	Year
123	Charger	2016
123	Tesla	2018
567	Civic	2020
567	Pinto	2022

Name	...
Allison	
Magda	
Dan	

Quantifiers

(We may need to leave for Friday)

Recap: Predicates on Subqueries

- EXISTS / NOT EXISTS
- IN / NOT IN
- ANY / ALL

The are “equivalent” meaning that a query that you can write using one, you can also write using the others

Quantifiers

Find people who drive **only** cars made after 2017

```
SELECT P.*  
FROM Payroll P  
WHERE 2017 <  
  ALL(SELECT R.Year  
       FROM Regist R  
       WHERE P.UserID = R.UserID);
```

Quantifiers

Find people who drive **only** cars made after 2017

```
SELECT P.*
FROM Payroll P
WHERE 2017 <
  ALL (SELECT R.Year
        FROM Regist R
        WHERE P.UserID = R.UserID);
```

```
SELECT P.*
FROM Payroll P
WHERE NOT EXISTS
  (SELECT R.Year
    FROM Regist R
    WHERE P.UserID = R.UserID
    and R.Year <= 2017);
```

Quantifiers

Find people who drive **only** cars made after 2017

```
SELECT P.*
FROM Payroll P
WHERE 2017 <
  ALL (SELECT R.Year
        FROM Regist R
        WHERE P.UserID = R.UserID);
```

```
SELECT P.*
FROM Payroll P
WHERE NOT EXISTS
  (SELECT R.Year
   FROM Regist R
   WHERE P.UserID = R.UserID
   and R.Year <= 2017);
```

```
SELECT P.*
FROM Payroll P
WHERE P.UserID NOT IN
  (SELECT R.UserID
   FROM Regist R
   WHERE R.Year <= 2017);
```


Quantifiers

Find people who drive **only** cars made after 2017

```
SELECT P.*
FROM Payroll P
WHERE 2017 <
  ALL (SELECT R.Year
        FROM Regist R
        WHERE P.UserID = R.UserID);
```

All these
compute the
same thing

```
SELECT P.*
FROM Payroll P
WHERE NOT EXISTS
  (SELECT R.Year
    FROM Regist R
    WHERE P.UserID = R.UserID
    and R.Year <= 2017);
```

```
SELECT P.*
FROM Payroll P
WHERE P.UserID NOT IN
  (SELECT R.UserID
    FROM Regist R
    WHERE R.Year <= 2017);
```

- SQL can express naturally queries that represent existential quantifiers

- To write a query that uses a universal quantifier, use DeMorgan's laws (next)

Quantifiers

There are two of quantifiers for a predicate

- **Exists** ($\exists x, \dots$) there is at least 1 that satisfies predicate
- **Forall**: ($\forall x, \dots$) all elements satisfy the predicate

SQL makes it easy to write **exists**

To write **forall**, use double negation

predicate holds **forall** elements
if and only if

not (**exists** element where not(predicate) holds)

Brief Review of Logic

- Implication: $A \rightarrow B$ is same as: $\text{not}(A) \text{ or } B$

- DeMorgan's Laws:

$$\text{not}(A \text{ and } B) = \text{not}(A) \text{ or } \text{not}(B)$$

$$\text{not}(A \text{ or } B) = \text{not}(A) \text{ and } \text{not}(B)$$

$$\text{not}(A \rightarrow B) = A \text{ and } \text{not}(B)$$

$$\text{not}(\exists x, P(x)) = \forall x, \text{not}(P(x))$$

$$\text{not}(\forall x, P(x)) = \exists x, \text{not}(P(x))$$

Brief Review of First Order Logic

$\exists x, \dots$ means “*there exists x such that...*”

$\forall x, \dots$ means “*for all x it holds that ...*”

Query1: Person P drives some car made on/before 2017

$\exists R \in \text{Regist}, (P.\text{UserID} = R.\text{UserID}) \text{ and } (R.\text{Year} \leq 2017)$

Brief Review of First Order Logic

$\exists x, \dots$ means “*there exists x such that...*”

$\forall x, \dots$ means “*for all x it holds that ...*”

Query1: Person P drives some car made on/before 2017:

$$\exists R \in \text{Regist}, (P.\text{UserID} = R.\text{UserID}) \text{ and } (R.\text{Year} \leq 2017)$$

Query2: Person P drives only cars made after 2017:

$$\forall R \in \text{Regist}, (P.\text{UserID} = R.\text{UserID}) \Rightarrow (R.\text{Year} > 2017)$$

Brief Review of First Order Logic

$\exists x, \dots$ means “*there exists x such that...*”

$\forall x, \dots$ means “*for all x it holds that ...*”

Query1: Person P drives some car made on/before 2017:

$$\exists R \in \text{Regist}, (P.\text{UserID} = R.\text{UserID}) \text{ and } (R.\text{Year} \leq 2017)$$

Query2: Person P drives only cars made after 2017:

$$\forall R \in \text{Regist}, (P.\text{UserID} = R.\text{UserID}) \Rightarrow (R.\text{Year} > 2017)$$

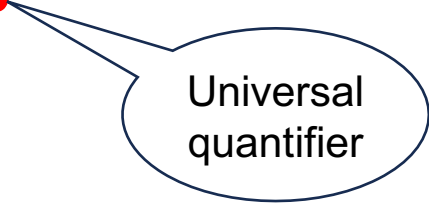
Check at home using DeMorgan: $\text{Not}(\text{Query1}) = \text{Query2}$

How to Write FORALL in SQL

Find person P drives **only** cars made after 2017

How to Write FORALL in SQL

Find person P drives **only** cars made after 2017



Universal
quantifier

How to Write FORALL in SQL

Find person P drives **only** cars made after 2017

Universal
quantifier

Existential
quantifier

Negate: find the other persons

Find person P drives **some** car made on or before 2017

```
SELECT P.*  
FROM Payroll P  
WHERE EXISTS  
  (SELECT R.Year  
   FROM Regist R  
   WHERE P.UserID = R.UserID  
    and R.Year <= 2017);
```

How to Write FORALL in SQL

Find person P drives **only** cars made after 2017

Universal
quantifier

Existential
quantifier

Negate: find the other persons

Find person P drives **some** car made on or before 2017

```
SELECT P.*  
FROM Payroll P  
WHERE EXISTS  
  (SELECT R.Year  
   FROM Regist R  
   WHERE P.UserID = R.UserID  
    and R.Year <= 2017);
```

Negate again:
find the other other persons

```
SELECT P.*  
FROM Payroll P  
WHERE NOT EXISTS  
  (SELECT R.Year  
   FROM Regist R  
   WHERE P.UserID = R.UserID  
    and R.Year <= 2017);
```

Discussion of SQL

- We have learned a lot of SQL!
 - Joins, outer-joins, 3VL, Group-by, Having, Nested queries, quantifiers

- It is a declarative language: write **what**, not **how**

- Next lecture: **relational algebra** expresses **how**