

A Survey of Deductive Databases

Raghu Ramakrishnan and Jeffrey D. Ullman

CS 848, Fall 2016

University of Waterloo

Presented by: Siddhartha Sahu

Overview

- Relational Databases
- Deductive Databases
- Datalog
- Example Queries
- Query Execution
- Conclusion and Discussion

Relational Databases



Relational Databases

Predominant model for data storage and processing



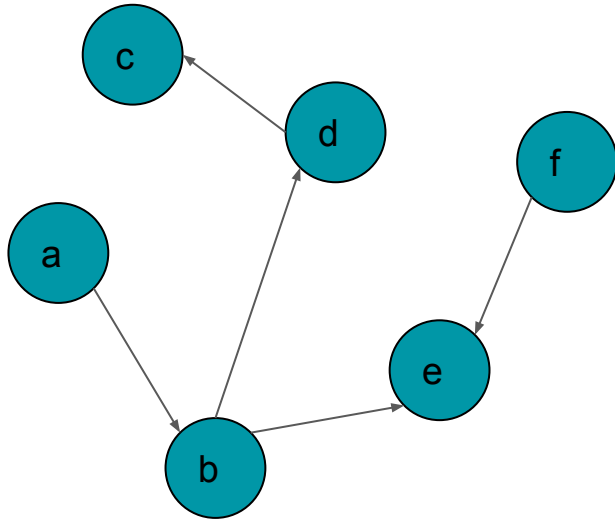
Relational Databases

Predominant model for data storage and processing

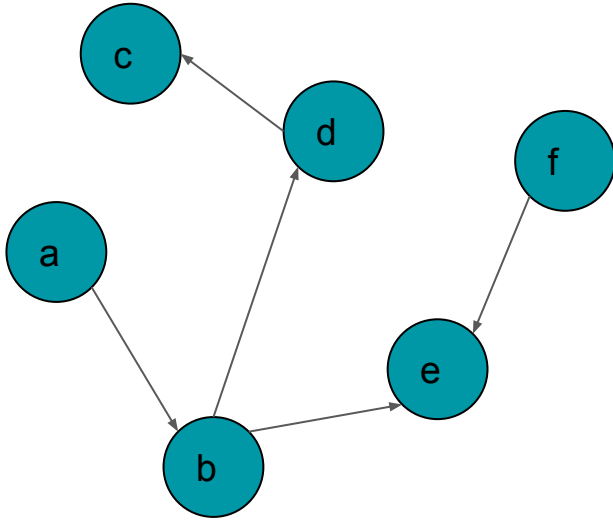
Declarative language: focus on **what** rather than how



Relational Databases



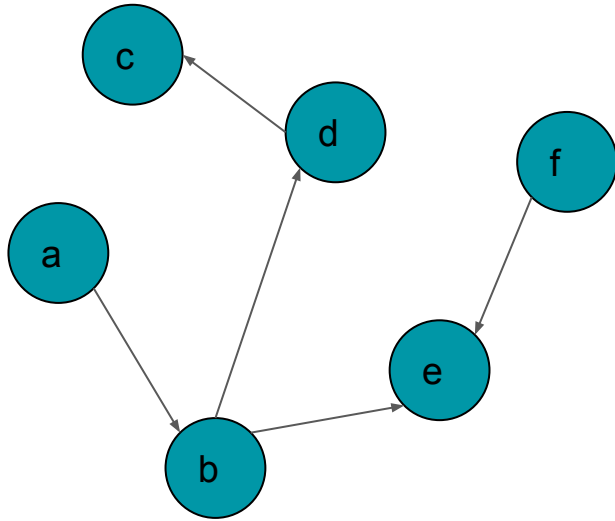
Relational Databases



INSERT INTO edges (...)

edges	
id_from	id_to
a	b
b	d
b	e
d	c
f	e

Relational Databases

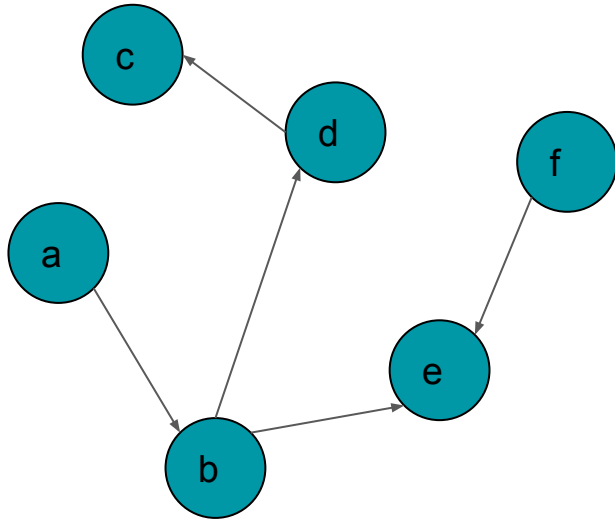


INSERT INTO edges (...)

edges	
id_from	id_to
a	b
b	d
b	e
d	c
f	e

Q: List vertices that vertex 'b' have an outgoing edge to.

Relational Databases



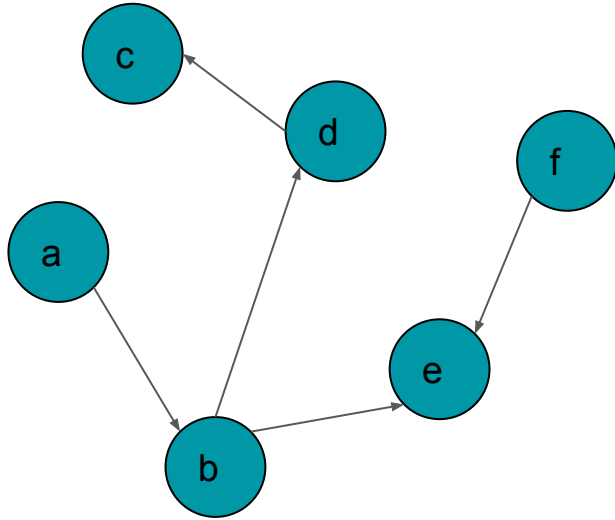
INSERT INTO edges (...)

edges	
id_from	id_to
a	b
b	d
b	e
d	c
f	e

Q: List vertices that vertex 'b' have an outgoing edge to.

A: `SELECT id_to from edges WHERE id_from = 'b'`

Relational Databases



INSERT INTO edges (...)

edges	
id_from	id_to
a	b
b	d
b	e
d	c
f	e

Q: List all vertex pairs (x,y) , such that y is reachable from x .

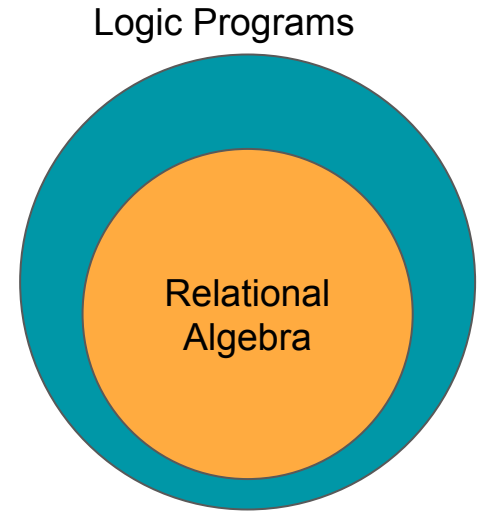
A: ?

Deductive Databases

Deductive Databases

Support a superset of relational algebra.

- Supports all queries from relational algebra.
- Supports recursions.



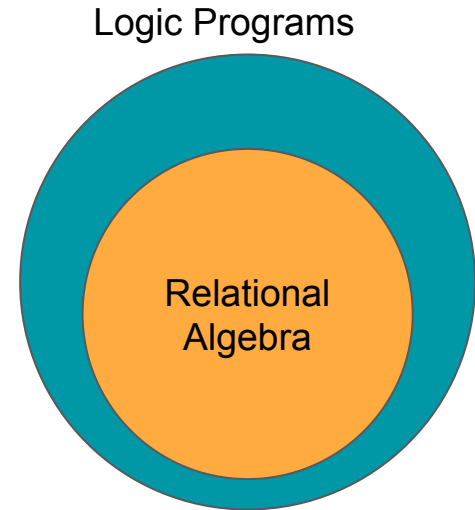
Deductive Databases

Support a superset of relational algebra.

- Supports all queries from relational algebra.
- Supports recursions.

Datalog: subset of Prolog, a logic programming language

- Database centric requirements
- Emphasis on completeness and termination
- Queries on data stored on secondary storage



Deductive Databases

Support a superset of relational algebra.

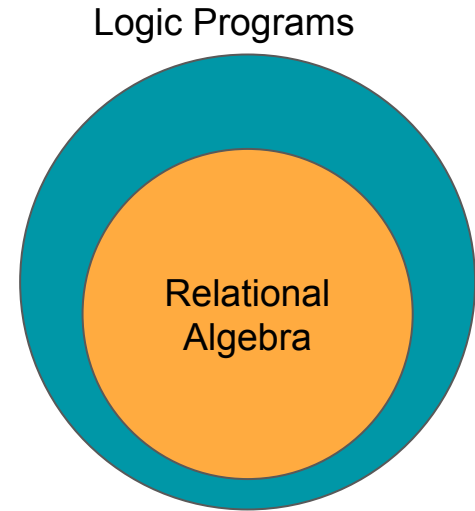
- Supports all queries from relational algebra.
- Supports recursions.

Datalog: subset of Prolog, a logic programming language

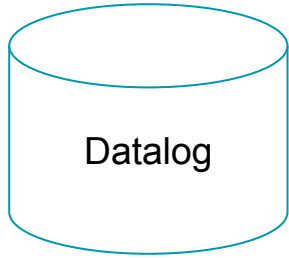
- Database centric requirements
- Emphasis on completeness and termination
- Queries on data stored on secondary storage

A database of facts.

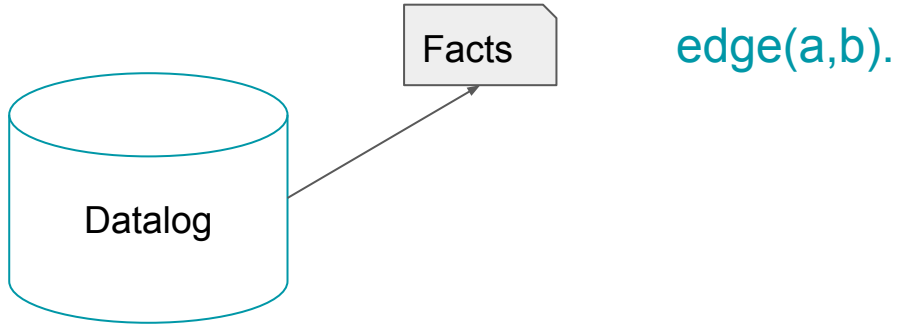
A set of rules for deriving new facts from existing facts.



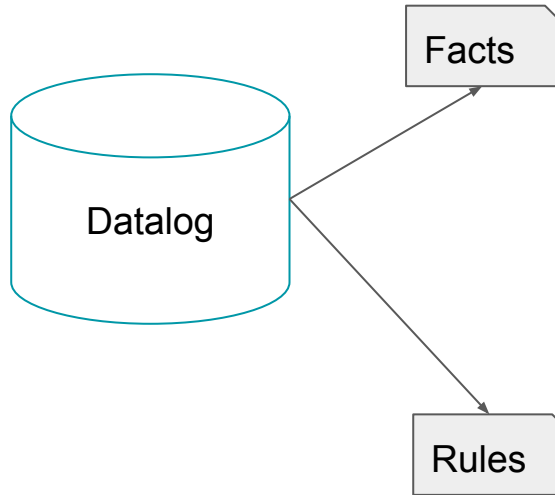
Datalog: Terminology



Datalog: Terminology



Datalog: Terminology

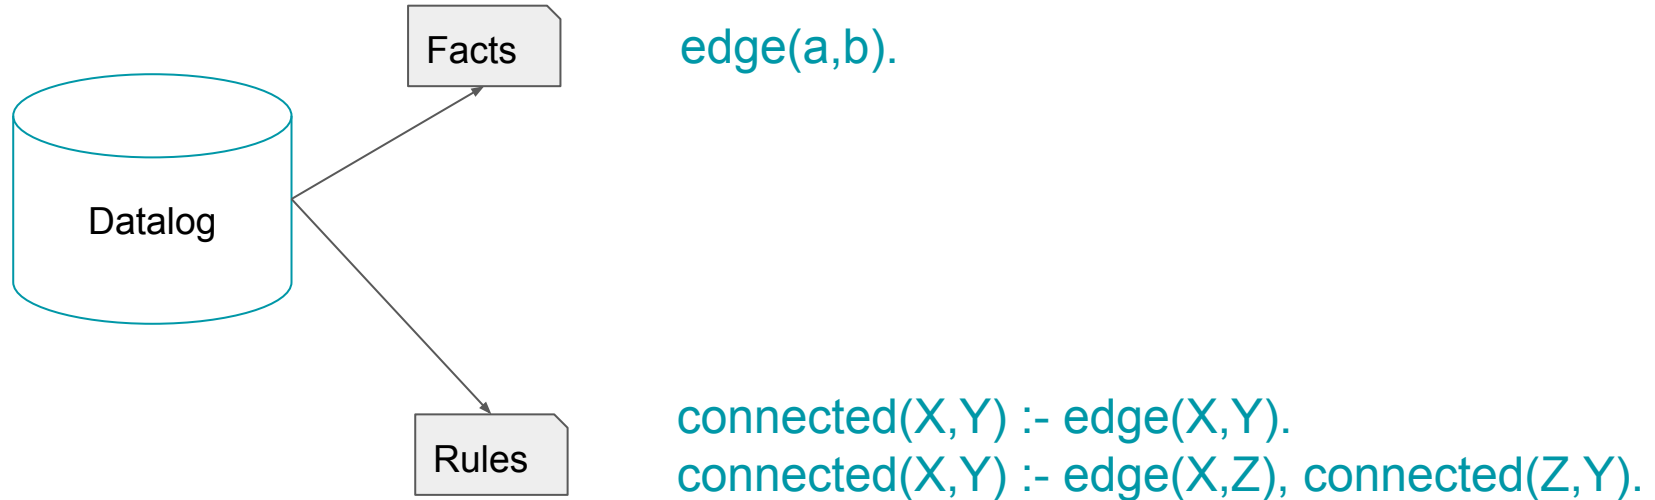


`edge(a,b).`

`connected(X,Y) :- edge(X,Y).`

`connected(X,Y) :- edge(X,Z), connected(Z,Y).`

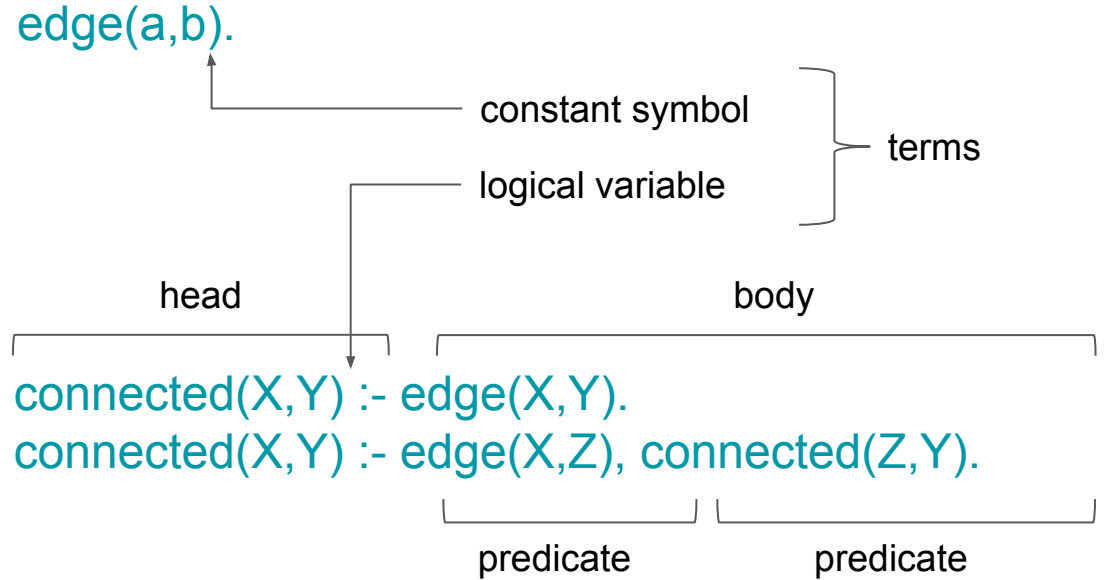
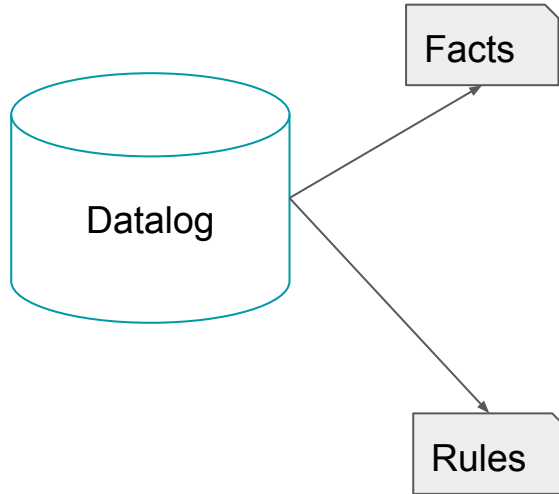
Datalog: Terminology



Implication/Clause: $A_0 :- A_1, A_2, \dots, A_k$ where A_0 is **true** if A_1 and A_2 ... and A_k are **true**.

$k = 0$: fact; $k > 0$: rule

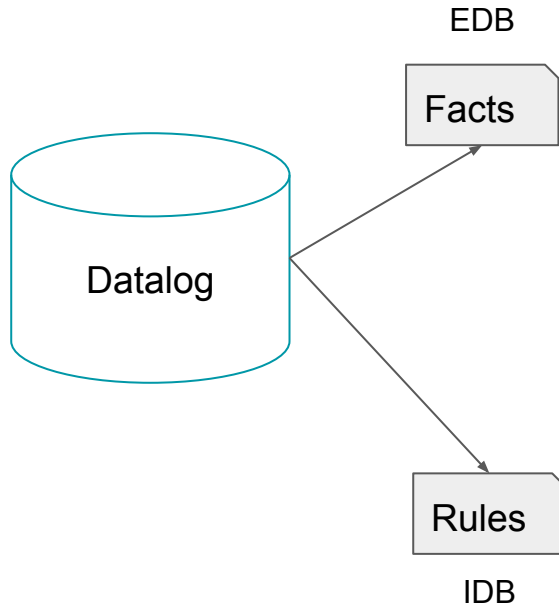
Datalog: Terminology



Implication/Clause: $A_0 :- A_1, A_2, \dots, A_k$ where A_0 is **true** if A_1 and $A_2 \dots$ and A_k are **true**.

$k = 0$: fact; $k > 0$: rule

Datalog: Terminology



`edge(a,b).`

constant symbol

logical variable

terms

head

body

`connected(X,Y) :- edge(X,Y).`

`connected(X,Y) :- edge(X,Z), connected(Z,Y).`

predicate

predicate

Implication/Clause: $A_0 :- A_1, A_2, \dots, A_k$ where A_0 is **true** if A_1 and A_2 ... and A_k are **true**.

$k = 0$: fact; $k > 0$: rule

Datalog: Examples

users
uid
name
age

accounts
uid
account_type
amount

Datalog: Examples

users
uid
name
age

accounts
uid
account_type
amount

`users(42, 'Jane Doe', 26).`

`accounts(42, 'savings', 5692.23)`

Datalog: Examples

Selection

Q: List all users with **age** > 23.

users
uid
name
age

accounts
uid
account_type
amount

users(42, 'Jane Doe', 26).

accounts(42, 'savings', 5692.23)

Datalog: Examples

Selection

Q: List all users with **age** > 23.

users
uid
name
age

accounts
uid
account_type
amount

users(42, 'Jane Doe', 26).

accounts(42, 'savings', 5692.23)

Relational Algebra: $\sigma_{\text{age} > 23}(\text{users})$

SQL: `SELECT * FROM users WHERE age > 23;`

Datalog: Examples

Selection

Q: List all users with **age** > 23.

users
uid
name
age

accounts
uid
account_type
amount

users(42, 'Jane Doe', 26).

accounts(42, 'savings', 5692.23)

Relational Algebra: $\sigma_{\text{age} > 23}(\text{users})$

SQL: SELECT * FROM users WHERE **age** > 23;

Datalog: S(Uid, Name, Age) :- users(Uid, Name, Age), **Age** > 23.

Datalog: Examples

Projection

Q: List **name** of users with age > 23.

users
uid
name
age

accounts
uid
account_type
amount

users(42, 'Jane Doe', 26).

accounts(42, 'savings', 5692.23)

Datalog: Examples

Projection

Q: List **name** of users with age > 23.

users
uid
name
age

accounts
uid
account_type
amount

users(42, 'Jane Doe', 26).

accounts(42, 'savings', 5692.23)

Relational Algebra:

$\pi_{\text{name}}(\sigma_{\text{age} > 23}(\text{users}))$

SQL:

SELECT **name** FROM users WHERE age > 23;

Datalog: Examples

Projection

Q: List **name** of users with age > 23.

users
uid
name
age

accounts
uid
account_type
amount

users(42, 'Jane Doe', 26).

accounts(42, 'savings', 5692.23)

Relational Algebra: $\pi_{\text{name}}(\sigma_{\text{age} > 23}(\text{users}))$

SQL: SELECT **name** FROM users WHERE age > 23;

Datalog: P(**Name**) :- users(Uid, Name, Age), Age > 23.

Datalog: Examples

Join

Q: List **name**, **amount** of users with age > 23.

users
uid
name
age

accounts
uid
account_type
amount

users(42, 'Jane Doe', 26).

accounts(42, 'savings', 5692.23)

Datalog: Examples

Join

Q: List **name**, **amount** of users with age > 23.

users
uid
name
age

accounts
uid
account_type
amount

users(42, 'Jane Doe', 26).

accounts(42, 'savings', 5692.23)

Relational Algebra:

$\pi_{\text{name,amount}}(\sigma_{\text{age} > 23}(\text{users} \bowtie_{\text{uid}} \text{accounts}))$

SQL:

SELECT **name,amount** FROM users,accounts

WHERE users.uid = accounts.uid AND age > 23;

Datalog: Examples

Join

Q: List **name**, **amount** of users with age > 23.

users
uid
name
age

accounts
uid
account_type
amount

users(42, 'Jane Doe', 26).

accounts(42, 'savings', 5692.23)

Relational Algebra:

$\pi_{\text{name, amount}}(\sigma_{\text{age} > 23}(\text{users} \bowtie_{\text{uid}} \text{accounts}))$

SQL:

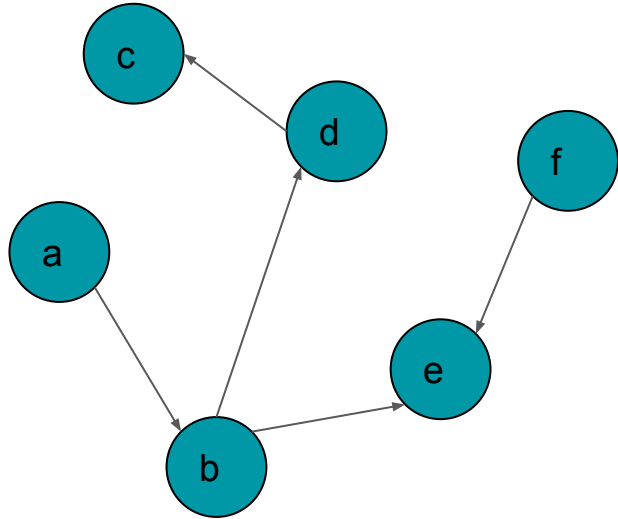
SELECT **name, amount** FROM users, accounts

WHERE users.uid = accounts.uid AND age > 23;

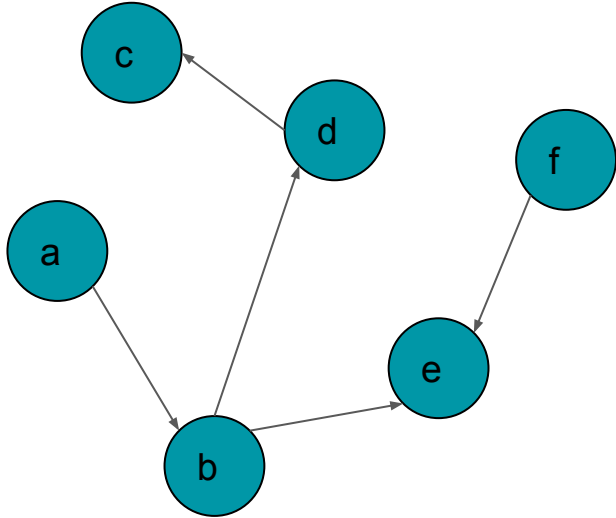
Datalog:

J(Name, Amount) :- users(Uid, Name, Age), accounts(Uid,
Account_type, Amount), Age > 23.

Datalog: Examples

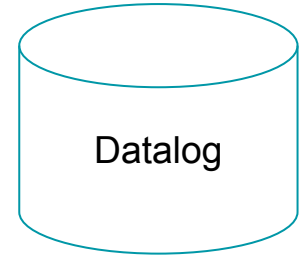


Datalog: Examples

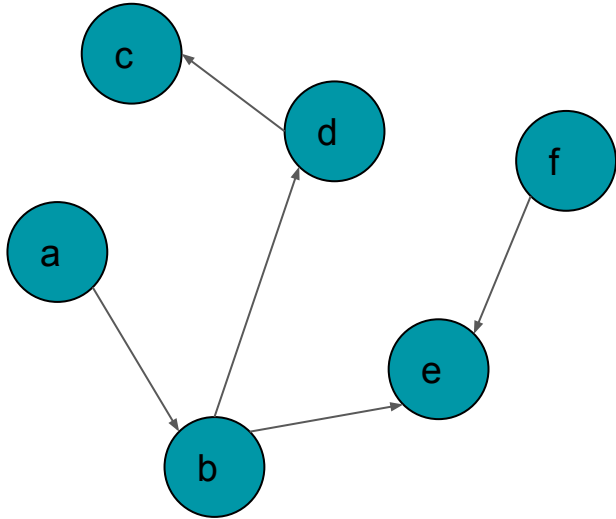


edge(a,b).
edge(b,d).
edge(b,e).
edge(d,c).
edge(f,e).

connected(X,Y) :- edge(X,Y).
connected(X,Y) :- edge(X,Z), connected(Z,Y).

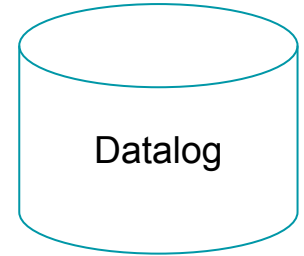


Datalog: Examples



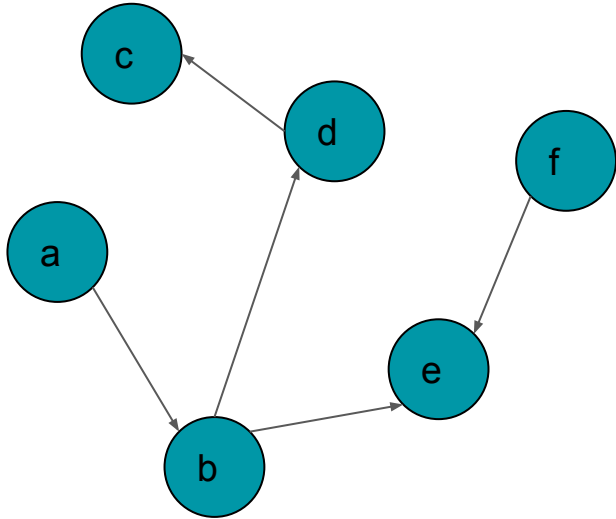
edge(a,b).
edge(b,d).
edge(b,e).
edge(d,c).
edge(f,e).

connected(X,Y) :- edge(X,Y).
connected(X,Y) :- edge(X,Z), connected(Z,Y).



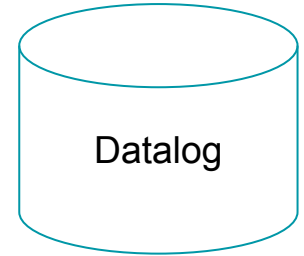
Q: List vertices that vertex 'b' have an outgoing edge to.

Datalog: Examples



edge(a,b).
edge(b,d).
edge(b,e).
edge(d,c).
edge(f,e).

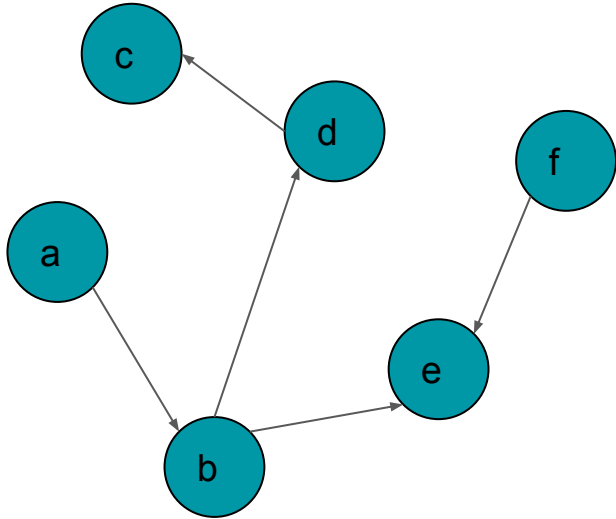
connected(X,Y) :- edge(X,Y).
connected(X,Y) :- edge(X,Z), connected(Z,Y).



Q: List vertices that vertex 'b' have an outgoing edge to.

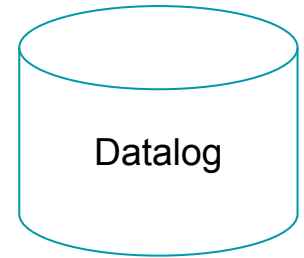
A: query(X) :- edge(b,X).

Datalog: Examples



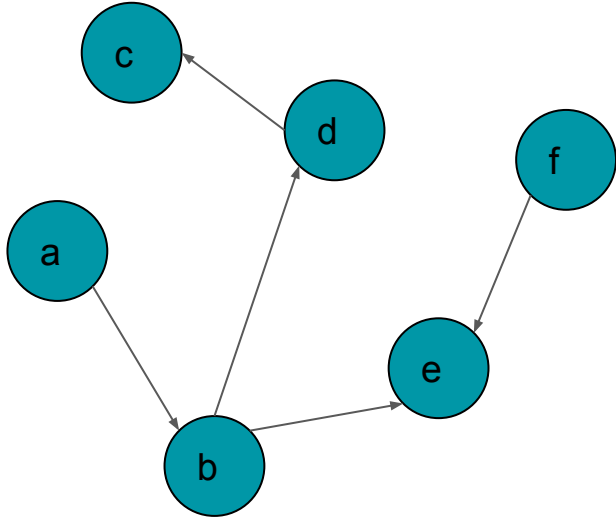
edge(a,b).
edge(b,d).
edge(b,e).
edge(d,c).
edge(f,e).

connected(X,Y) :- edge(X,Y).
connected(X,Y) :- edge(X,Z), connected(Z,Y).



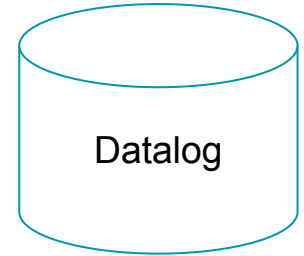
Q: List all vertex pairs (x,y), such that y is reachable from x.

Datalog: Examples



edge(a,b).
edge(b,d).
edge(b,e).
edge(d,c).
edge(f,e).

connected(X,Y) :- edge(X,Y).
connected(X,Y) :- edge(X,Z), connected(Z,Y).



Q: List all vertex pairs (x,y), such that y is reachable from x.

A: query(X,Y) :- connected(X,Y).

Query Evaluation: Naïve algorithm

Query Evaluation: Naïve algorithm

```
P0 = InitialValue  
Repeat  
    Pk = f(Pk-1)  
Until no-more-change
```

Query Evaluation: Naïve algorithm

1. Begin by assuming all IDB relations are empty.

```
P0 = InitialValue  
Repeat  
    Pk = f(Pk-1)  
Until no-more-change
```


Query Evaluation: Naïve algorithm

1. Begin by assuming all IDB relations are empty.
2. Repeatedly evaluate the rules using the EDB and the previous IDB to get a new IDB.

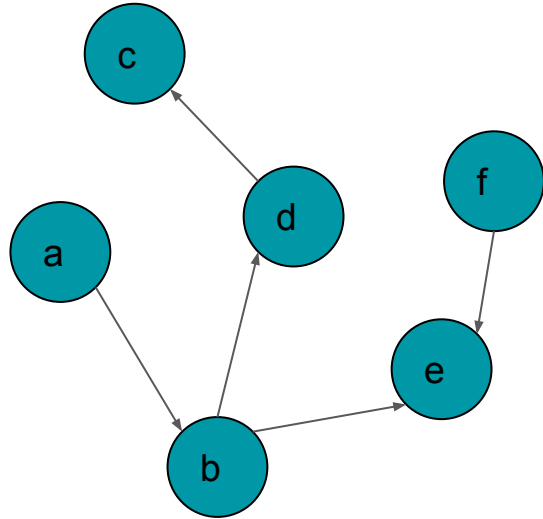
```
P0 = InitialValue  
Repeat  
    Pk = f(Pk-1)  
Until no-more-change
```

Query Evaluation: Naïve algorithm

1. Begin by assuming all IDB relations are empty.
2. Repeatedly evaluate the rules using the EDB and the previous IDB to get a new IDB.
3. End when there is no change to the IDB.

```
P0 = InitialValue  
Repeat  
    Pk = f(Pk-1)  
Until no-more-change
```

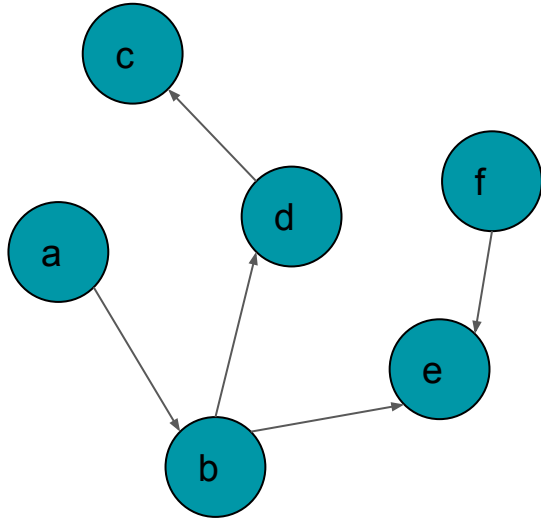
Query Evaluation: Naïve algorithm



`connected(X,Y) :- edge(X,Y).`
`connected(X,Y) :- edge(X,Z), connected(Z,Y).`

`connected(X,Y) .`

Query Evaluation: Naïve algorithm

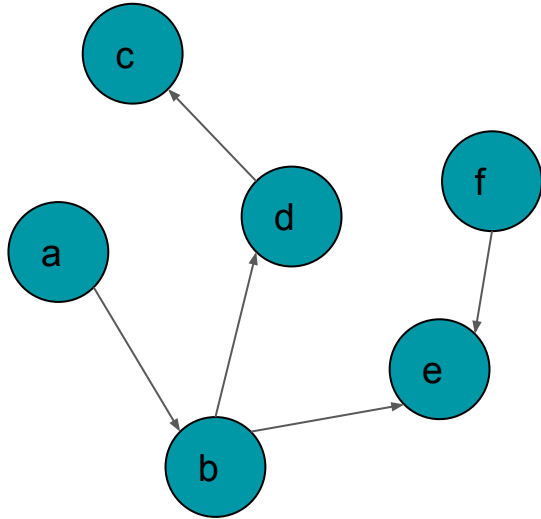


edges	
a	b
b	d
d	c
b	e
f	e

`connected(X,Y) :- edge(X,Y).`
`connected(X,Y) :- edge(X,Z), connected(Z,Y).`

`connected(X,Y) .`

Query Evaluation: Naïve algorithm



edges	
a	b
b	d
d	c
b	e
f	e

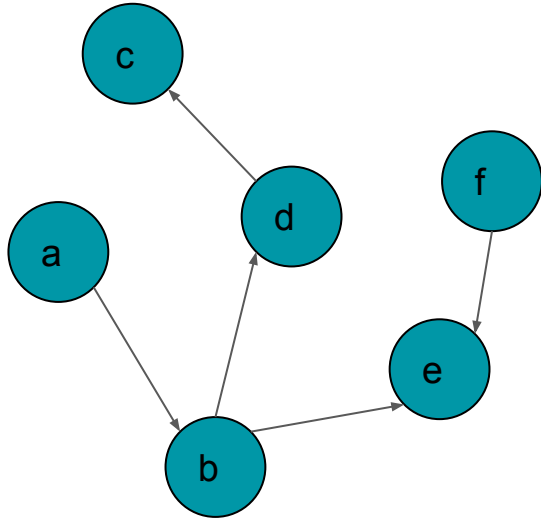
`connected(X,Y) :- edge(X,Y).`
`connected(X,Y) :- edge(X,Z), connected(Z,Y).`

`connected(X,Y) .`

\emptyset

`I = 0`

Query Evaluation: Naïve algorithm



edges	
a	b
b	d
d	c
b	e
f	e

connected(X,Y) :- edge(X,Y).
connected(X,Y) :- edge(X,Z), connected(Z,Y).

connected(X,Y) .

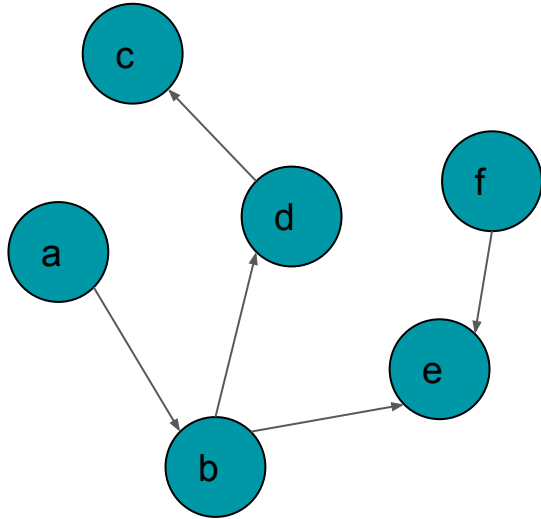
∅

l = 0

a	b
b	d
d	c
b	e
f	e

l = 1

Query Evaluation: Naïve algorithm



edges	
a	b
b	d
d	c
b	e
f	e

`connected(X,Y) :- edge(X,Y).`
`connected(X,Y) :- edge(X,Z), connected(Z,Y).`

`connected(X,Y) .`

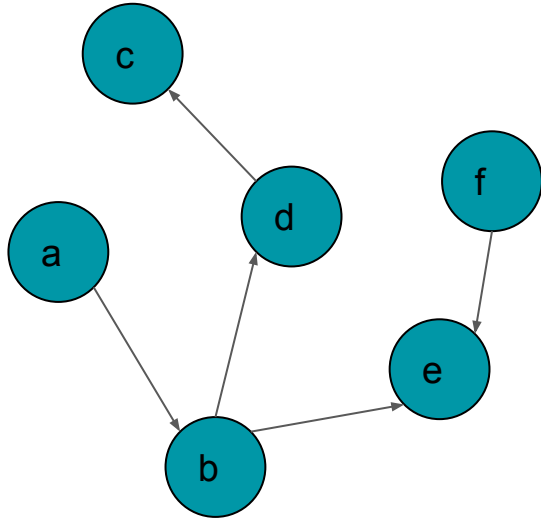
\emptyset

$l = 0$

a	b
b	d
d	c
b	e
f	e

$l = 1$

Query Evaluation: Naïve algorithm



edges	
a	b
b	d
d	c
b	e
f	e

`connected(X,Y) :- edge(X,Y).`
`connected(X,Y) :- edge(X,Z), connected(Z,Y).`

`connected(X,Y) .`

\emptyset

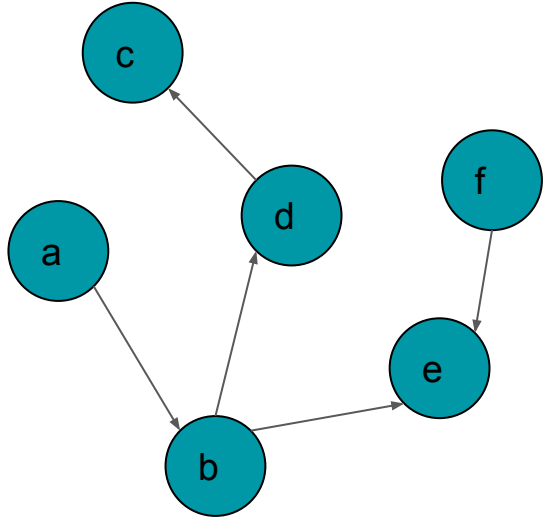
$l = 0$

a	b
b	d
d	c
b	e
f	e

$l = 1$

$l = 2$

Query Evaluation: Naïve algorithm



edges	
a	b
b	d
d	c
b	e
f	e

connected(X,Y) :- edge(X,Y).
connected(X,Y) :- edge(X,Z), connected(Z,Y).

connected(X,Y) .

∅

l = 0

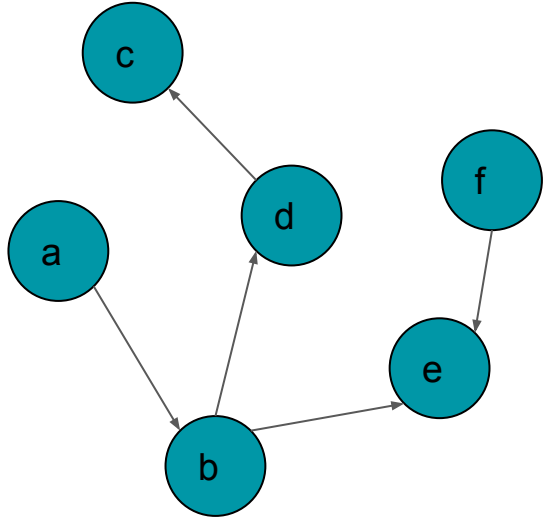
a	b
b	d
d	c
b	e
f	e

l = 1

a	b
b	d
d	c
b	e
f	e

l = 2

Query Evaluation: Naïve algorithm



edges	
a	b
b	d
d	c
b	e
f	e

`connected(X,Y) :- edge(X,Y).`
`connected(X,Y) :- edge(X,Z), connected(Z,Y).`

`connected(X,Y) .`

\emptyset

$l = 0$

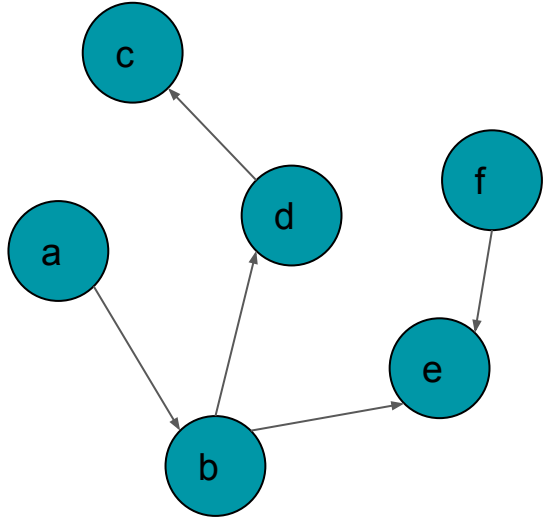
a	b
b	d
d	c
b	e
f	e

$l = 1$

a	d
a	e
a	b
b	d
d	c
b	e
f	e

$l = 2$

Query Evaluation: Naïve algorithm



edges	
a	b
b	d
d	c
b	e
f	e

`connected(X,Y) :- edge(X,Y).`
`connected(X,Y) :- edge(X,Z), connected(Z,Y).`

`connected(X,Y) .`

\emptyset

$l = 0$

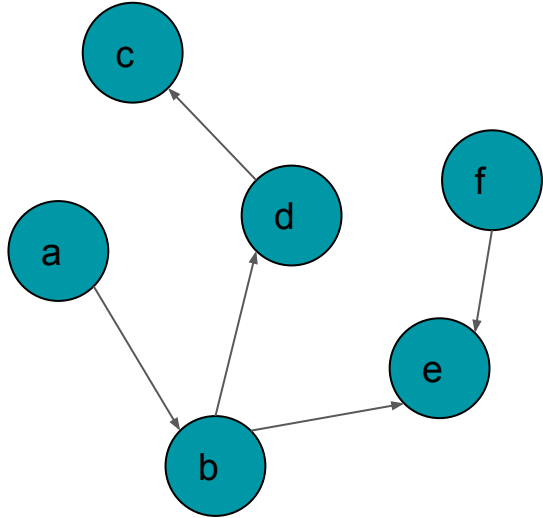
a	b
b	d
d	c
b	e
f	e

$l = 1$

b	c
a	d
a	e
a	b
b	d
d	c
b	e
f	e

$l = 2$

Query Evaluation: Naïve algorithm



edges	
a	b
b	d
d	c
b	e
f	e

connected(X,Y) :- edge(X,Y).
 connected(X,Y) :- edge(X,Z), connected(Z,Y).

connected(X,Y) .

∅

l = 0

a	b
b	d
d	c
b	e
f	e

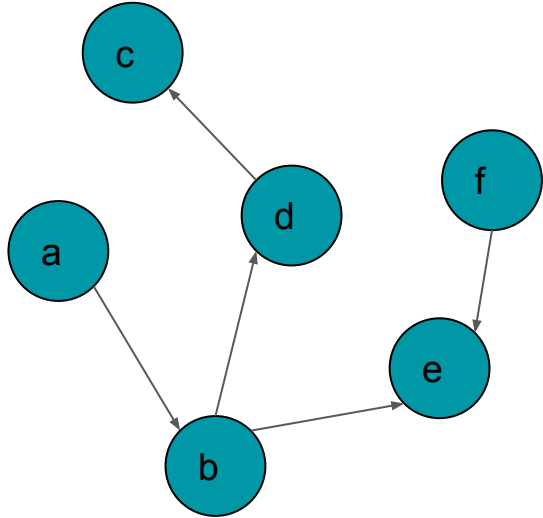
l = 1

b	c
a	d
a	e
a	b
b	d
d	c
b	e
f	e

l = 2

l = 3

Query Evaluation: Naïve algorithm



edges	
a	b
b	d
d	c
b	e
f	e

connected(X,Y) :- edge(X,Y).
connected(X,Y) :- edge(X,Z), connected(Z,Y).

connected(X,Y) .

∅

l = 0

a	b
b	d
d	c
b	e
f	e

l = 1

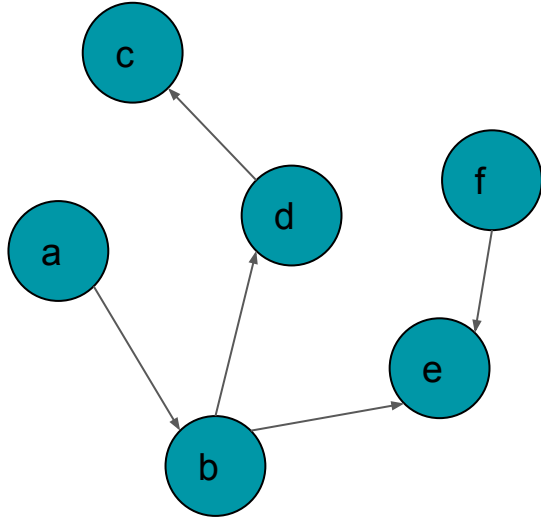
b	c
a	d
a	e
a	b
b	d
d	c
b	e
f	e

l = 2

b	c
a	d
a	e
a	b
a	b
b	d
d	c
b	e
f	e

l = 3

Query Evaluation: Naïve algorithm



edges	
a	b
b	d
d	c
b	e
f	e

connected(X,Y) :- edge(X,Y).
connected(X,Y) :- edge(X,Z), connected(Z,Y).

connected(X,Y) .

∅

l = 0

a	b
b	d
d	c
b	e
f	e

l = 1

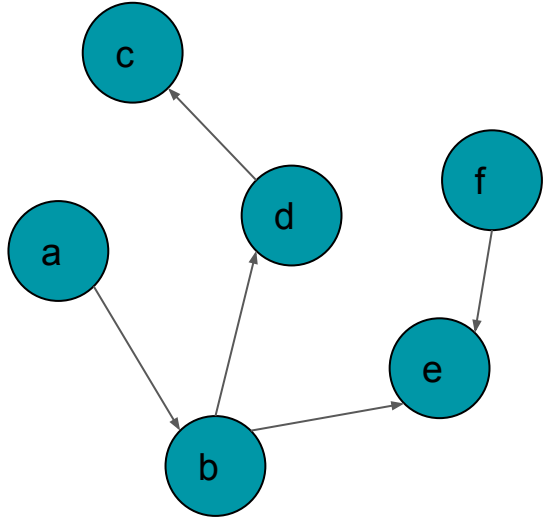
b	c
a	d
a	e
a	b
b	d
d	c
b	e
f	e

l = 2

a	c
b	c
a	d
a	e
a	b
b	d
d	c
b	e
f	e

l = 3

Query Evaluation: Naïve algorithm



edges	
a	b
b	d
d	c
b	e
f	e

connected(X,Y) :- edge(X,Y).
connected(X,Y) :- edge(X,Z), connected(Z,Y).

connected(X,Y) .

∅

l = 0

a	b
b	d
d	c
b	e
f	e

l = 1

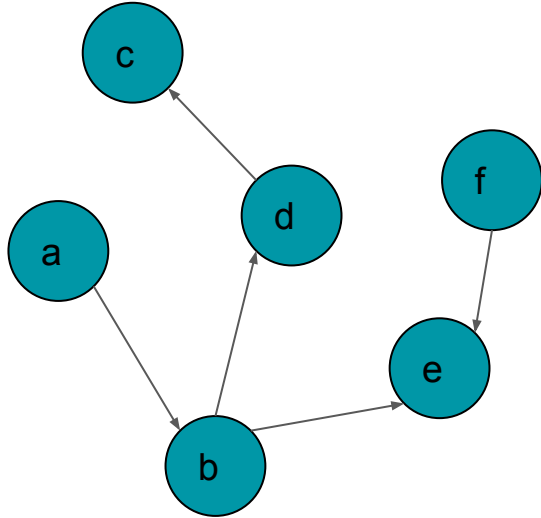
b	c
a	d
a	e
a	b
b	d
d	c
b	e
f	e

l = 2

a	c
b	c
a	d
a	e
a	b
b	d
d	c
b	e
f	e

l = 3

Query Evaluation: Naïve algorithm



edges	
a	b
b	d
d	c
b	e
f	e

connected(X,Y) :- edge(X,Y).
 connected(X,Y) :- edge(X,Z), connected(Z,Y).

connected(X,Y).

∅

l = 0

a	b
b	d
d	c
b	e
f	e

l = 1

b	c
a	d
a	e
a	b
b	d
d	c
b	e
f	e

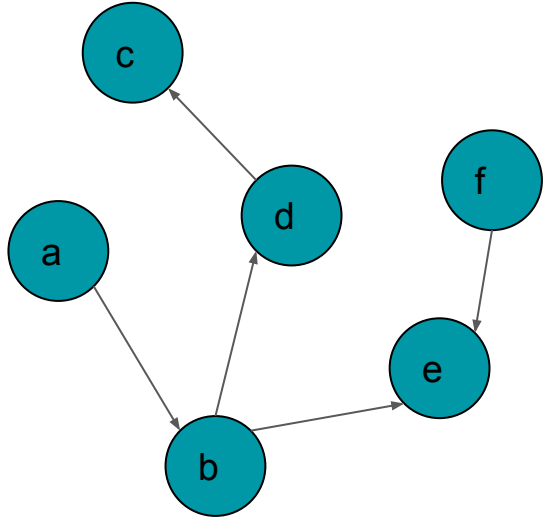
l = 2

a	c
b	c
a	d
a	e
a	b
b	d
d	c
b	e
f	e

l = 3

l = 4

Query Evaluation: Naïve algorithm



edges	
a	b
b	d
d	c
b	e
f	e

`connected(X,Y) :- edge(X,Y).`
`connected(X,Y) :- edge(X,Z), connected(Z,Y).`

`connected(X,Y).`

∅

l = 0

a	b
b	d
d	c
b	e
f	e

l = 1

b	c
a	d
a	e
a	b
b	d
d	c
b	e
f	e

l = 2

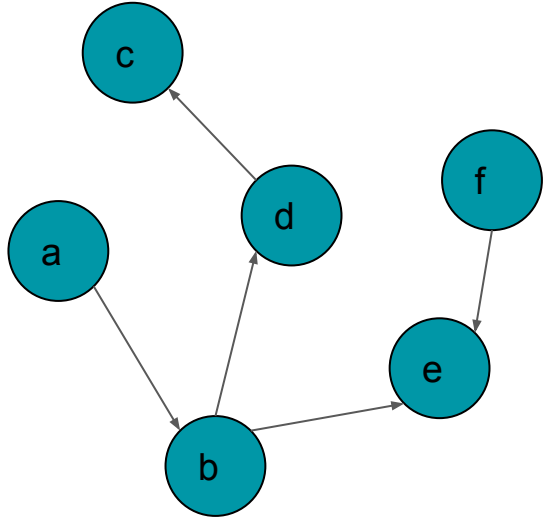
a	c
b	c
a	d
a	e
a	b
b	d
d	c
b	e
f	e

l = 3

a	c
b	c
a	d
a	e
a	b
a	b
b	d
d	c
b	e
f	e

l = 4

Query Evaluation: Naïve algorithm



edges	
a	b
b	d
d	c
b	e
f	e

connected(X,Y) :- edge(X,Y).
connected(X,Y) :- edge(X,Z), connected(Z,Y).

connected(X,Y) .

∅

l = 0

a	b
b	d
d	c
b	e
f	e

l = 1

b	c
a	d
a	e
a	b
b	d
d	c
b	e
f	e

l = 2

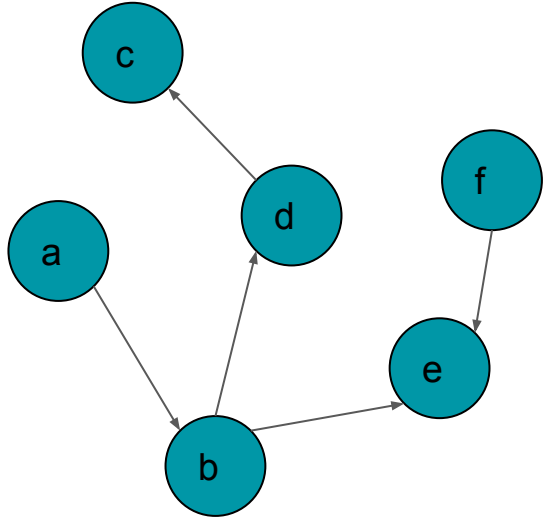
a	c
b	c
a	d
a	e
a	b
b	d
d	c
b	e
f	e

l = 3

a	c
b	c
a	d
a	e
a	b
b	d
d	c
b	e
f	e

l = 4

Query Evaluation: Naïve algorithm



edges	
a	b
b	d
d	c
b	e
f	e

connected(X,Y) :- edge(X,Y).
 connected(X,Y) :- edge(X,Z), connected(Z,Y).

connected(X,Y) .

∅

l = 0

a	b
b	d
d	c
b	e
f	e

l = 1

b	c
a	d
a	e
a	b
b	d
d	c
b	e
f	e

l = 2

a	c
b	c
a	d
a	e
a	b
b	d
d	c
b	e
f	e

l = 3

a	c
b	c
a	d
a	e
a	b
b	d
d	c
b	e
f	e

l = 4

Query Evaluation: Semi-Naïve algorithm

- * Avoid repeating computations already done in previous iterations.
- * Focus on only the newly derived tuples (deltas) from previous iterations.

Query Evaluation: Semi-Naïve algorithm

* Avoid repeating computations already done in previous iterations.

* Focus on only the newly derived tuples (deltas) from previous iterations.

```
for each IDB predicate  $p$ 
  do  $\begin{cases} p^{[0]} := \emptyset \\ \delta(p)^{[0]} := \text{tuples produced by rules using only EDB's} \end{cases}$ 
   $i := 1$ 
  repeat
     $p^{[i]} := p^{[i-1]} \cup \delta(p)^{[i-1]}$ 
    evaluate  $\Delta(p)^{[i]}$ 
     $\delta(p)^{[i]} := \Delta(p)^{[i]} - p^{[i]}$ 
     $i := i + 1$ 
  until  $\delta(p)^{[i]} = \emptyset$  for each IDB predicate  $p$ 
```

Query Evaluation: Semi-Naïve algorithm

* Avoid repeating computations already done in previous iterations.

* Focus on only the newly derived tuples (deltas) from previous iterations.

```
for each IDB predicate p
  do  $\begin{cases} p^{[0]} := \emptyset \\ \delta(p)^{[0]} := \text{tuples produced by rules using only EDB's} \end{cases}$ 
  i := 1
  repeat
     $p^{[i]} := p^{[i-1]} \cup \delta(p)^{[i-1]}$ 
    evaluate  $\Delta(p)^{[i]}$ 
     $\delta(p)^{[i]} := \Delta(p)^{[i]} - p^{[i]}$ 
    i := i + 1
  until  $\delta(p)^{[i]} = \emptyset$  for each IDB predicate p
```

$$\begin{aligned} \Delta(p)^{[i]} & :- \delta(p_1)^{[i-1]}, p_2^{[i-1]}, \dots, p_n^{[i-1]}, q_1, \dots, q_m. \\ \Delta(p)^{[i]} & :- p_1^{[i]}, \delta(p_2)^{[i-1]}, p_3^{[i-1]}, \dots, p_n^{[i-1]}, q_1, \dots, q_m. \\ & \dots \\ \Delta(p)^{[i]} & :- p_1^{[i]}, \dots, p_{n-1}^{[i]}, \delta(p_n)^{[i-1]}, q_1, \dots, q_m. \end{aligned}$$

Query Evaluation: Semi-Naïve algorithm

edges	
a	b
b	d
d	c
b	e
f	e

connected(X,Y) :- edge(X,Y).
connected(X,Y) :- edge(X,Z), connected(Z,Y).

\emptyset

$P_{I=0}$

Query Evaluation: Semi-Naïve algorithm

edges	
a	b
b	d
d	c
b	e
f	e

a	b
b	d
d	c
b	e
f	e

\emptyset

$P_{l=0}$

$\delta_{l=0}$

`connected(X,Y) :- edge(X,Y).`

`connected(X,Y) :- edge(X,Z), connected(Z,Y).`

Query Evaluation: Semi-Naïve algorithm

edges	
a	b
b	d
d	c
b	e
f	e

`connected(X,Y) :- edge(X,Y).`
`connected(X,Y) :- edge(X,Z), connected(Z,Y).`

\emptyset

$P_{l=0}$

a	b
b	d
d	c
b	e
f	e

$\delta_{l=0}$

a	b
b	d
d	c
b	e
f	e

$P_{l=1}$

Query Evaluation: Semi-Naïve algorithm

`connected(X,Y) :- edge(X,Y).`
`connected(X,Y) :- edge(X,Z), connected(Z,Y).`

edges	
a	b
b	d
d	c
b	e
f	e

a	b
b	d
d	c
b	e
f	e

\emptyset

$P_{l=0}$

$\delta_{l=0}$

a	b
b	d
d	c
b	e
f	e

$P_{l=1}$

a	d
a	e
b	c

$\Delta_{l=1}$

Query Evaluation: Semi-Naïve algorithm

`connected(X,Y) :- edge(X,Y).`
`connected(X,Y) :- edge(X,Z), connected(Z,Y).`

edges	
a	b
b	d
d	c
b	e
f	e

\emptyset

$P_{l=0}$

a	b
b	d
d	c
b	e
f	e

$\bar{\delta}_{l=0}$

a	b
b	d
d	c
b	e
f	e

$P_{l=1}$

a	d
a	e
b	c

$\Delta_{l=1}$

a	d
a	e
b	c

$\bar{\delta}_{l=1}$

Query Evaluation: Semi-Naïve algorithm

edges	
a	b
b	d
d	c
b	e
f	e

connected(X,Y) :- edge(X,Y).
 connected(X,Y) :- edge(X,Z), connected(Z,Y).

∅

$P_{l=0}$

a	b
b	d
d	c
b	e
f	e

$\delta_{l=0}$

a	b
b	d
d	c
b	e
f	e

$P_{l=1}$

a	d
a	e
b	c

$\Delta_{l=1}$

a	d
a	e
b	c

$\delta_{l=1}$

a	d
a	e
b	c
a	b
b	d
d	c
b	e
f	e

$P_{l=2}$

Query Evaluation: Semi-Naïve algorithm

edges	
a	b
b	d
d	c
b	e
f	e

connected(X,Y) :- edge(X,Y).
 connected(X,Y) :- edge(X,Z), connected(Z,Y).

\emptyset
 $P_{l=0}$ $\delta_{l=0}$

a	b
b	d
d	c
b	e
f	e

$P_{l=1}$ $\delta_{l=1}$

a	b
b	d
d	c
b	e
f	e

$\Delta_{l=1}$

a	d
a	e
b	c

$P_{l=2}$

a	d
a	e
b	c
a	b
b	d
d	c
b	e
f	e

$\Delta_{l=2}$

a	c
---	---

Query Evaluation: Semi-Naïve algorithm

edges	
a	b
b	d
d	c
b	e
f	e

connected(X,Y) :- edge(X,Y).
 connected(X,Y) :- edge(X,Z), connected(Z,Y).

\emptyset
 $P_{l=0}$ $\delta_{l=0}$

a	b
b	d
d	c
b	e
f	e

$P_{l=1}$ $\delta_{l=1}$

a	b
b	d
d	c
b	e
f	e

a	d
a	e
b	c

$\Delta_{l=1}$

a	d
a	e
b	c

$P_{l=2}$ $\delta_{l=2}$

a	d
a	e
b	c
a	b
b	d
d	c
b	e
f	e

a	c
---	---

$\Delta_{l=2}$

a	c
---	---

Query Evaluation: Semi-Naïve algorithm

edges	
a	b
b	d
d	c
b	e
f	e

connected(X,Y) :- edge(X,Y).
 connected(X,Y) :- edge(X,Z), connected(Z,Y).

\emptyset
 $P_{l=0}$ $\delta_{l=0}$

a	b
b	d
d	c
b	e
f	e

a	b
b	d
d	c
b	e
f	e

a	d
a	e
b	c

$\Delta_{l=1}$

a	d
a	e
b	c

$P_{l=1}$ $\delta_{l=1}$

a	d
a	e
b	c
a	b
b	d
d	c
b	e
f	e

a	c
---	---

$\Delta_{l=2}$

a	c
---	---

$P_{l=2}$ $\delta_{l=2}$

a	c
a	d
a	e
b	c
a	b
b	d
d	c
b	e
f	e

Query Evaluation: Semi-Naïve algorithm

edges	
a	b
b	d
d	c
b	e
f	e

connected(X,Y) :- edge(X,Y).
 connected(X,Y) :- edge(X,Z), connected(Z,Y).

\emptyset
 $P_{l=0}$ $\delta_{l=0}$

a	b
b	d
d	c
b	e
f	e

$P_{l=1}$ $\delta_{l=1}$

a	b
b	d
d	c
b	e
f	e

a	d
a	e
b	c

$\Delta_{l=1}$

a	d
a	e
b	c

$P_{l=2}$ $\delta_{l=2}$

a	d
a	e
b	c
a	b
b	d
d	c
b	e
f	e

a	c
---	---

$\Delta_{l=2}$

a	c
---	---

$\delta_{l=3}$

a	c
a	d
a	e
b	c
a	b
b	d
d	c
b	e
f	e

\emptyset
 $\Delta_{l=3}$

Deductive Databases: Additional concepts

Deductive Databases: Additional concepts

Negation predicates

Deductive Databases: Additional concepts

Negation predicates

Safe rules

Deductive Databases: Additional concepts

Negation predicates

Safe rules

Query optimization

- Magic Sets

- Rule-Rewriting Techniques

- Iterative Fixpoint Evaluation

Deductive Databases: Additional concepts

Negation predicates

Safe rules

Query optimization

- Magic Sets

- Rule-Rewriting Techniques

- Iterative Fixpoint Evaluation

Aggregations

Conclusion

Conclusion

Deductive databases are **more expressive** than relational databases.

Support for **recursive** queries

Conclusion

Deductive databases are **more expressive** than relational databases.

Support for **recursive** queries

Datalog: query language adapted from Prolog

Conclusion

Deductive databases are **more expressive** than relational databases.

Support for **recursive** queries

Datalog: query language adapted from Prolog

Focus on **query optimization**

Conclusion

Deductive databases are **more expressive** than relational databases.

Support for **recursive** queries

Datalog: query language adapted from Prolog

Focus on **query optimization**

Naive vs Semi Naive query execution algorithms

Avoid **repeated computations**

Discussion

SQL has recursion techniques like CTE

How does that compare to Datalog in terms of **expressiveness**?

Application domains best suited for Datalog?

Program analysis (recursion)

Declarative networking (NDlog)

Security (SeNDlog)

Applicability to general **processing frameworks**?

Hive

Spark SQL