Oregon State University

Learning Strategies in Game-Theoretic Data Interaction

Ben McCamish, Arash Termehchy, Behrouz Touri, Liang Huang

Information & **D**ata Manag**e**ment and **A**nalytics Laboratory (IDEA)

IDEA Lab.

Querying a database of student grades

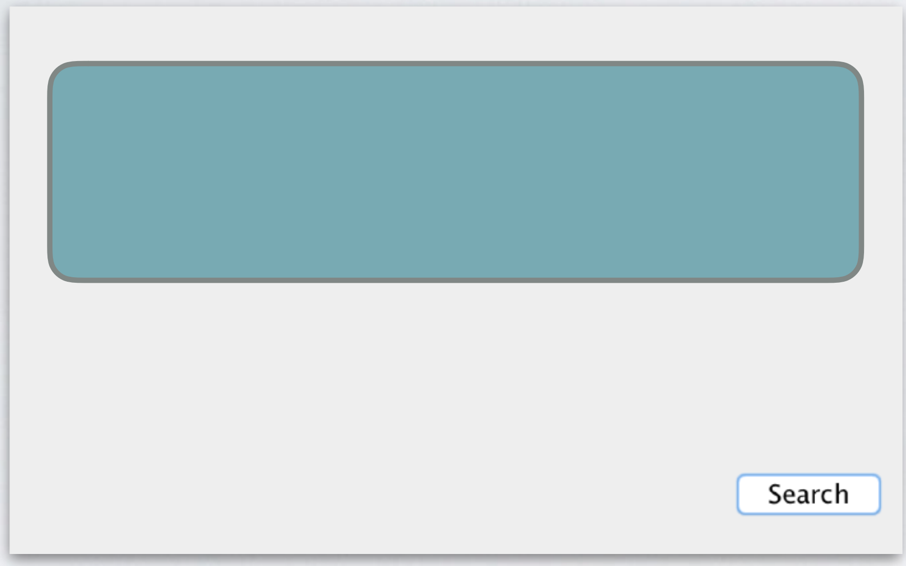

Grades

First_Name	Last_Name	Dept.	Grade
...	...	...	...
Sarah	Smith	CE	A
John	Smith	EE	B
Kerry	Smith	CS	D
...	...	...	...

Results

First_Name	Last_Name	Dept.	Grade

- A user's **intent** is the content they wish to find in the database
- They use **queries** attempting to communicate their intent

Most users cannot precisely express their intents

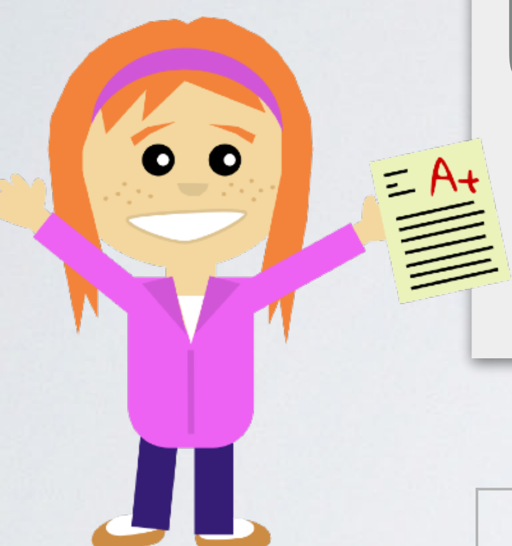

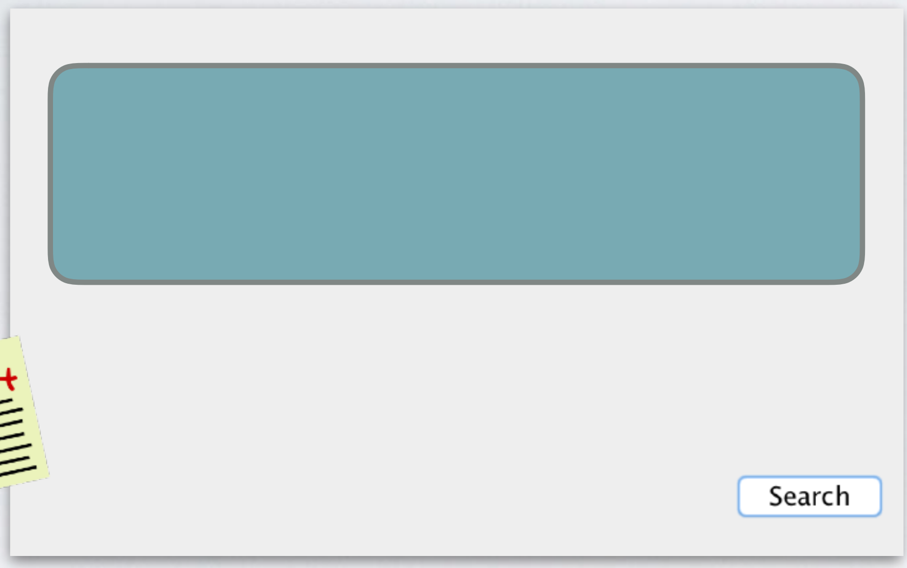

Grades			
First_Name	Last_Name	Dept.	Grade
...	...	...	...
Sarah	Smith	CE	A
John	Smith	EE	B
Kerry	Smith	CS	D
...	...	...	...

Results			
First_Name	Last_Name	Dept.	Grade

- **Intent:** user looking for grade of student *Kerry Smith*
- Not sufficiently familiar with the database content and structure

Most users cannot precisely express their intents

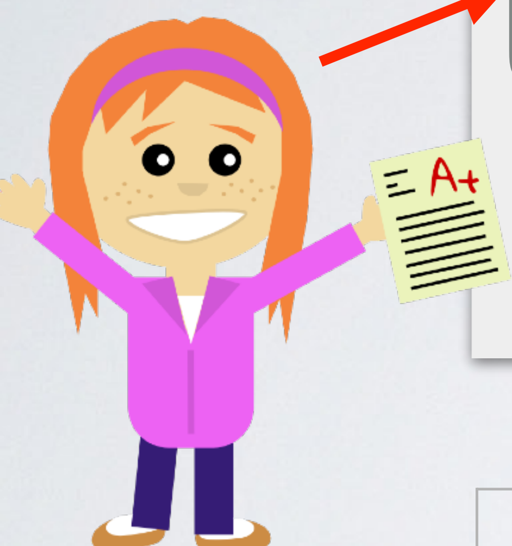

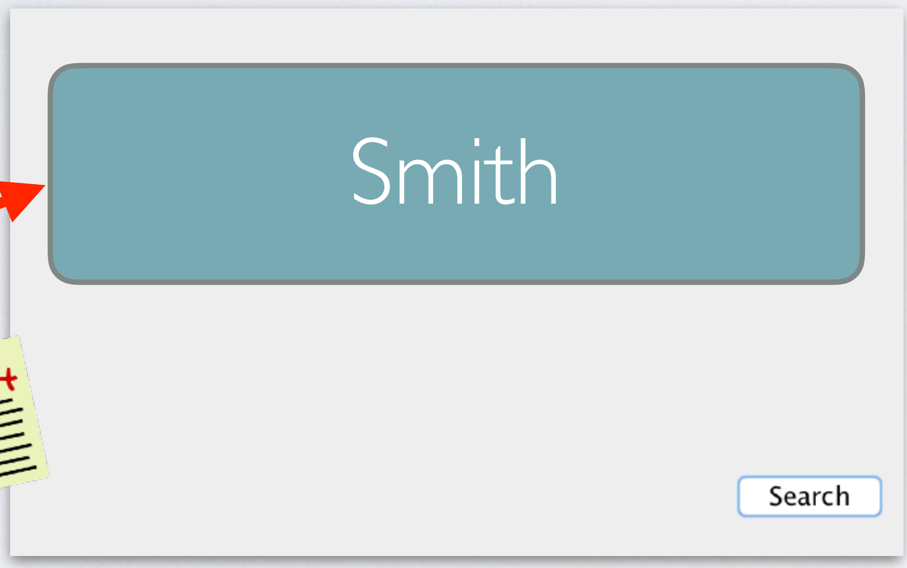

Grades			
First_Name	Last_Name	Dept.	Grade
...	...	...	...
Sarah	Smith	CE	A
John	Smith	EE	B
Kerry	Smith	CS	D
...	...	...	...

Results			
First_Name	Last_Name	Dept.	Grade

- **Query:** Has last name "Smith"
- Does not precisely express intent

Most users cannot precisely express their intents

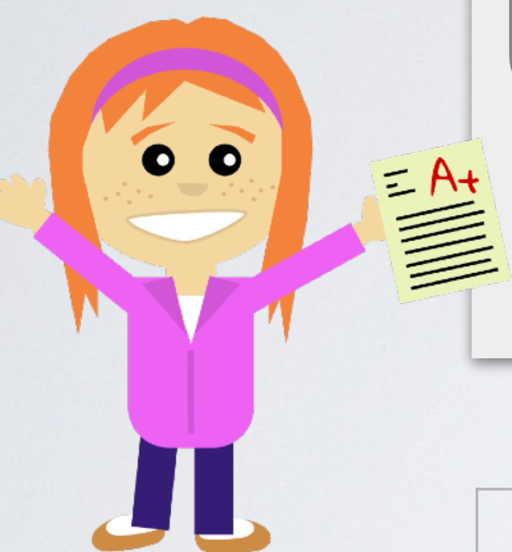

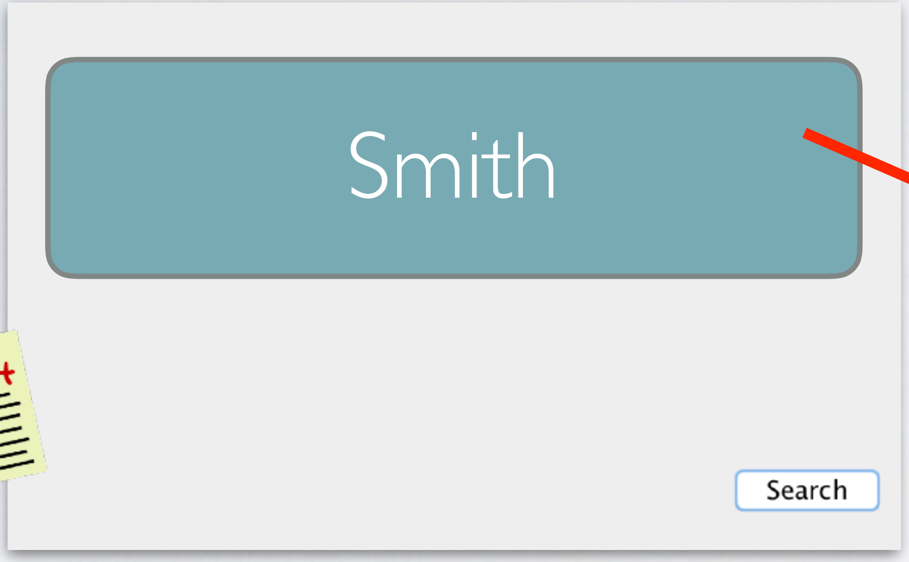

Grades

First_Name	Last_Name	Dept.	Grade
...	...	...	...
Sarah	Smith	CE	A
John	Smith	EE	B
Kerry	Smith	CS	D
...	...	...	...

Results

First_Name	Last_Name	Dept.	Grade

- Database has too many tuples matching query, mostly non-relevant.

Most users cannot precisely express their intents

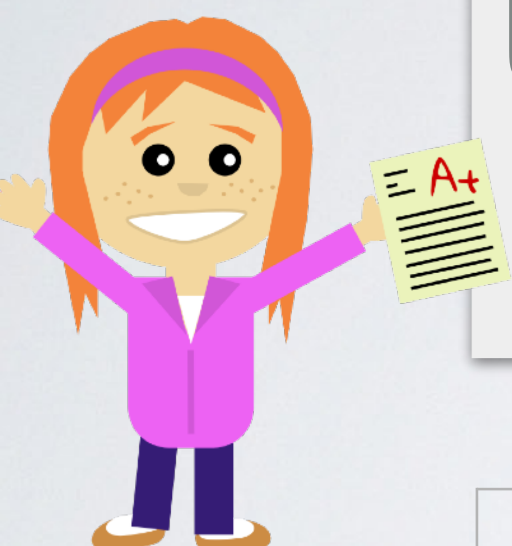

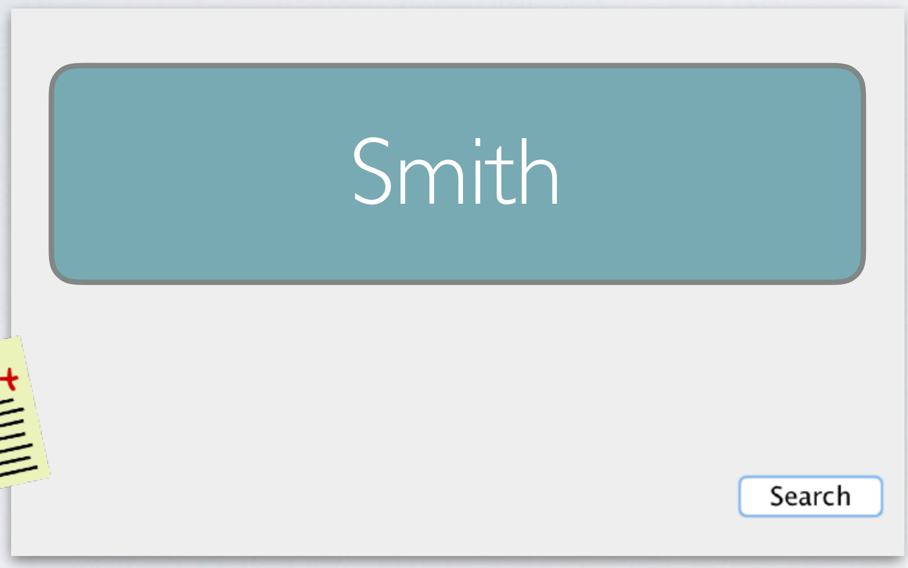

Grades

First_Name	Last_Name	Dept.	Grade
...	...	...	...
Sarah	Smith	CE	A
John	Smith	EE	B
Kerry	Smith	CS	D
...	...	...	...

Results

First_Name	Last_Name	Dept.	Grade
Sarah	Smith	CE	A
John	Smith	EE	B

- Database system returns only a subset of matching tuples

Most users cannot precisely express their intents

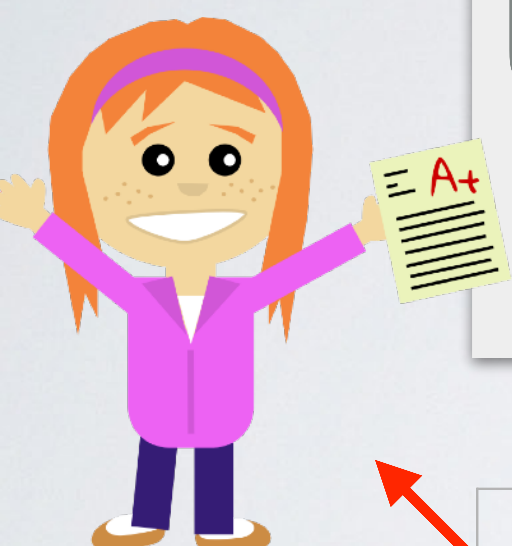

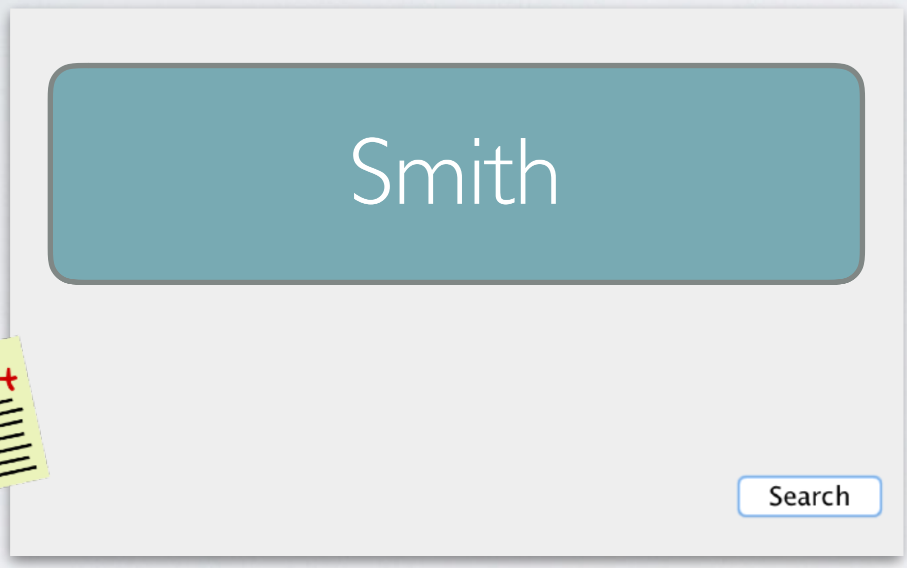

Grades			
First_Name	Last_Name	Dept.	Grade
…	…	…	…
Sarah	Smith	CE	A
John	Smith	EE	B
Kerry	Smith	CS	D
…	…	…	…

Results			
First_Name	Last_Name	Dept.	Grade
Sarah	Smith	CE	A
John	Smith	EE	B

- The user doesn't find the student she is looking for

Users learn by interacting with database systems

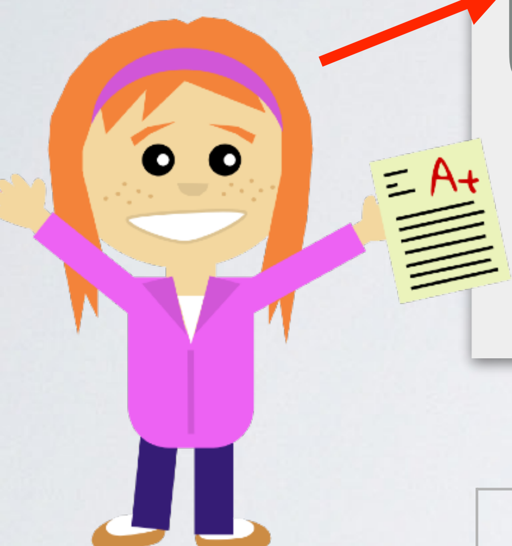

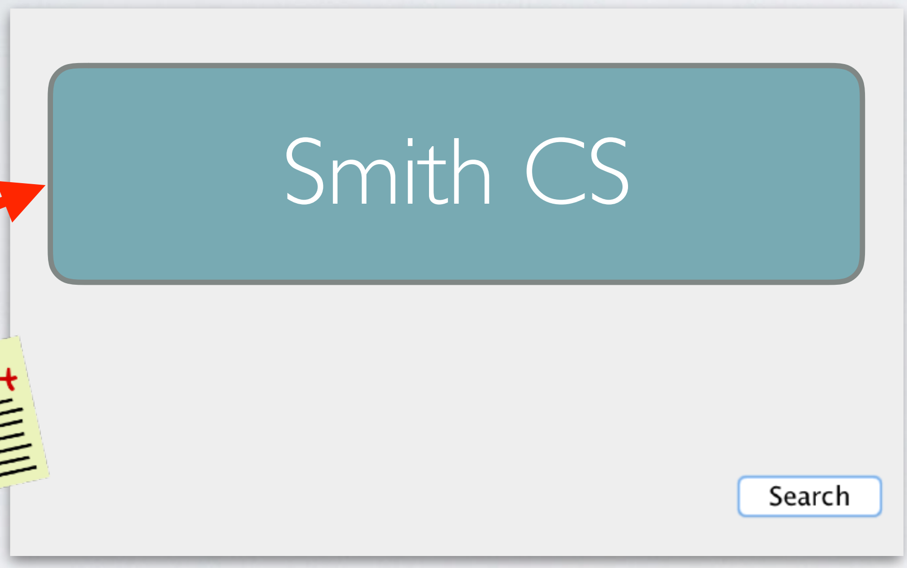

Grades			
First_Name	Last_Name	Dept.	Grade
...	...	...	...
Sarah	Smith	CE	A
John	Smith	EE	B
Kerry	Smith	CS	D
...	...	...	...

Results			
First_Name	Last_Name	Dept.	Grade

- Reformulates query after learning about the database and it's content.
- **Reformulated Query:** Has last name "Smith" and is in the Department "CS"
- New query expresses user's intent much more accurately

But they learn by interacting with database systems

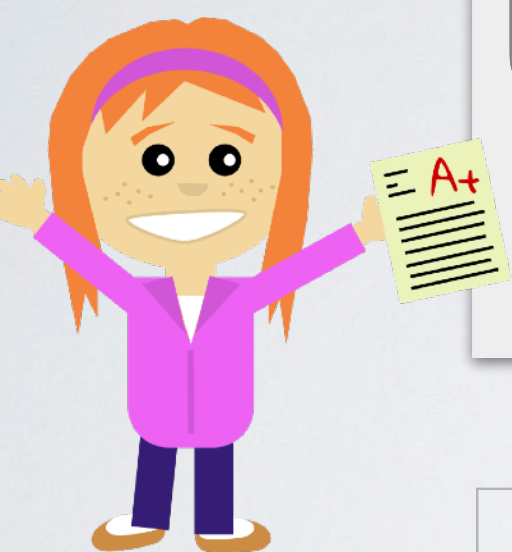

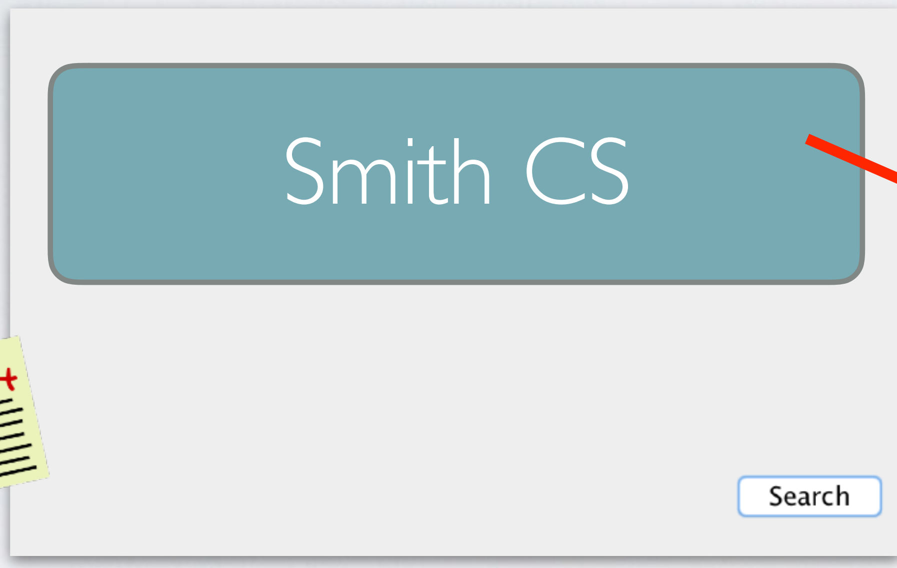

Grades

First_Name	Last_Name	Dept.	Grade
...	...	...	...
Sarah	Smith	CE	A
John	Smith	EE	B
Kerry	Smith	CS	D
...	...	...	...

Results

First_Name	Last_Name	Dept.	Grade

- Database system finds the desired tuple

But they learn by interacting with database systems

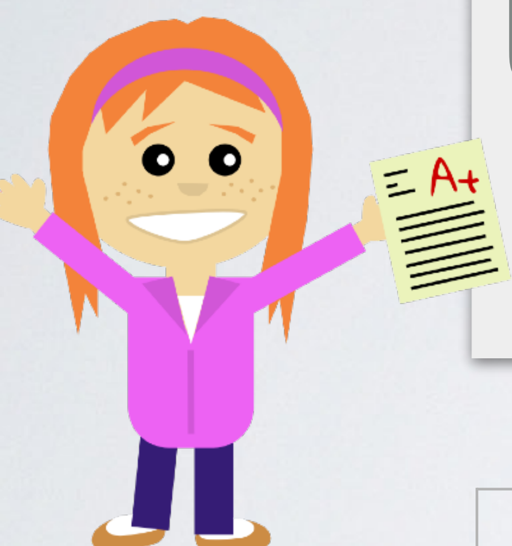

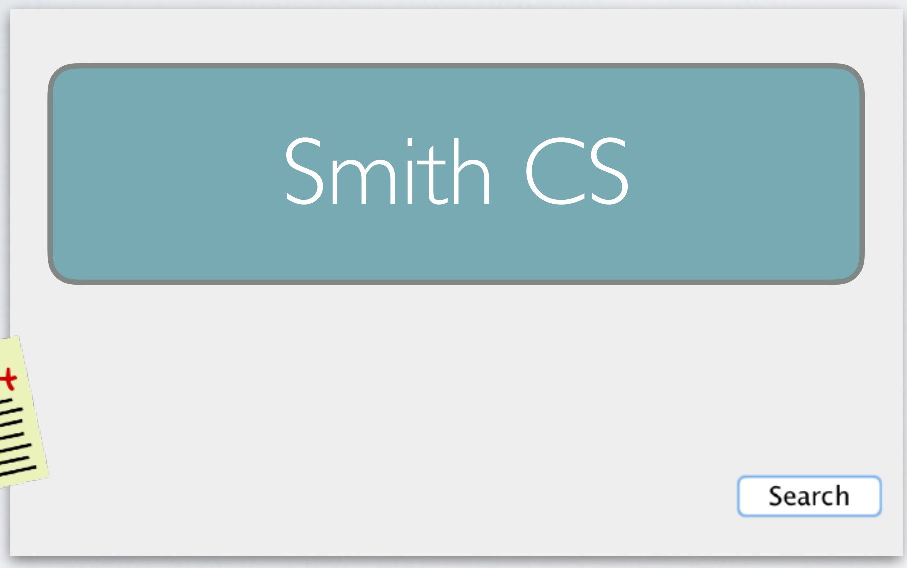

Grades

First_Name	Last_Name	Dept.	Grade
...	...	...	...
Sarah	Smith	CE	A
John	Smith	EE	B
Kerry	Smith	CS	D
...	...	...	...

Results

First_Name	Last_Name	Dept.	Grade
Kerry	Smith	CS	D

- Database system returns the desired tuple

But they learn by interacting with database systems

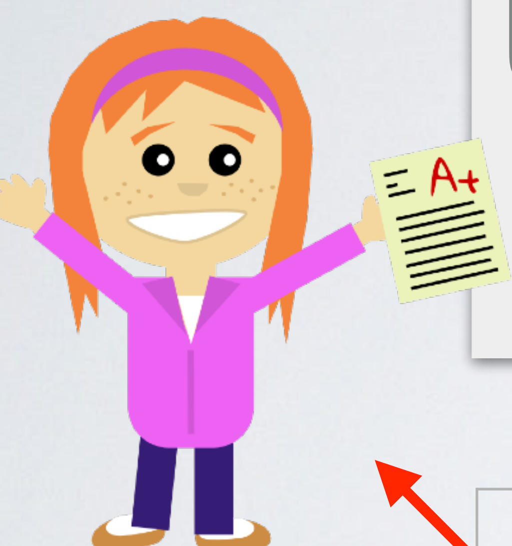

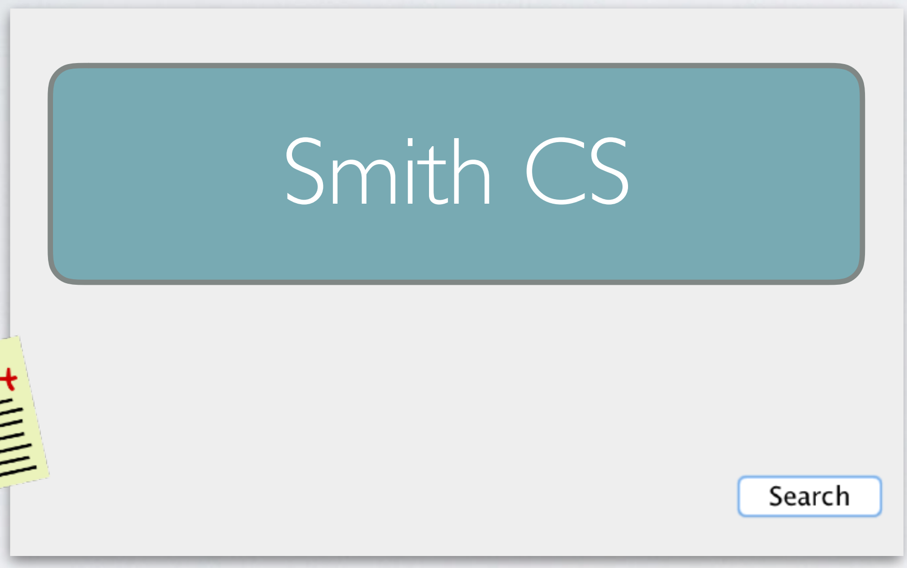

Grades			
First_Name	Last_Name	Dept.	Grade
...	...	...	...
Sarah	Smith	CE	A
John	Smith	EE	B
Kerry	Smith	CS	D
...	...	...	...

Results			
First_Name	Last_Name	Dept.	Grade
Kerry	Smith	CS	D

- User selects the returned tuple
- Learning and reformulating query allowed the user to find the desired student

Database system can learn as well

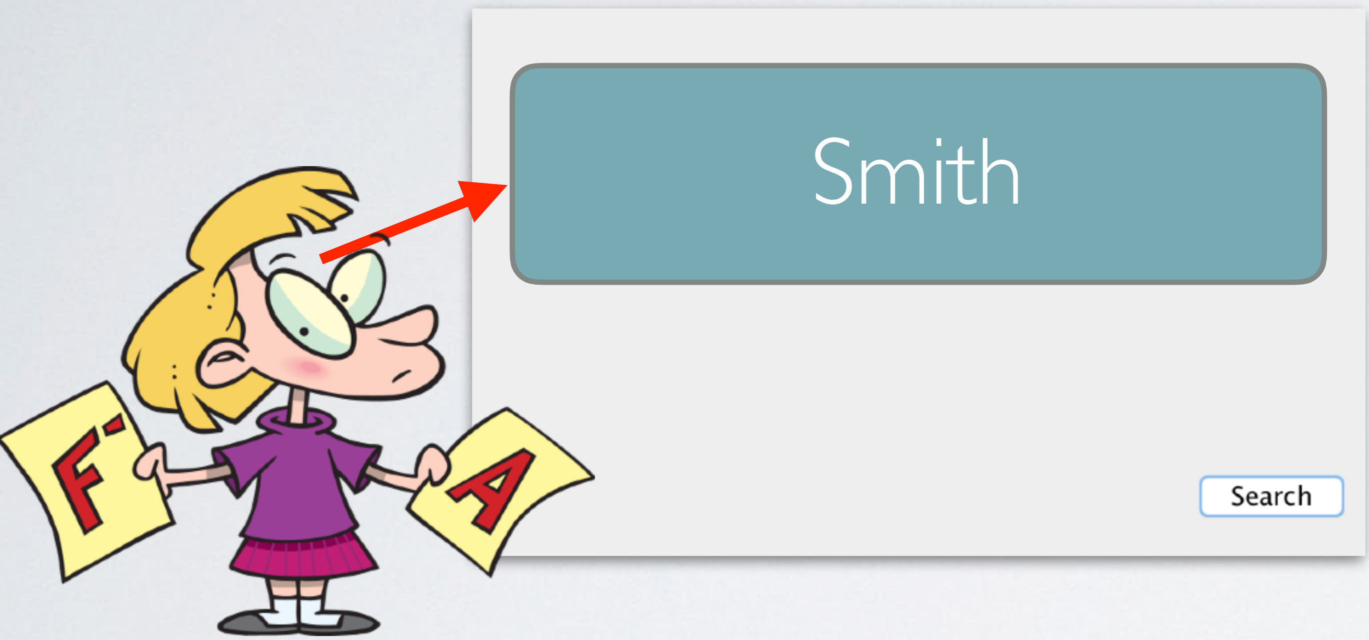

Grades			
First_Name	Last_Name	Dept.	Grade
…	…	…	…
Sarah	Smith	CE	A
John	Smith	EE	B
Kerry	Smith	CS	D
…	…	…	…

Results			
First_Name	Last_Name	Dept.	Grade

- **Intent:** User looking for grade of student *Kerry Smith*
- **Query:** Has Last Name "Smith"
- Does not precisely express intent

Database system can learn as well

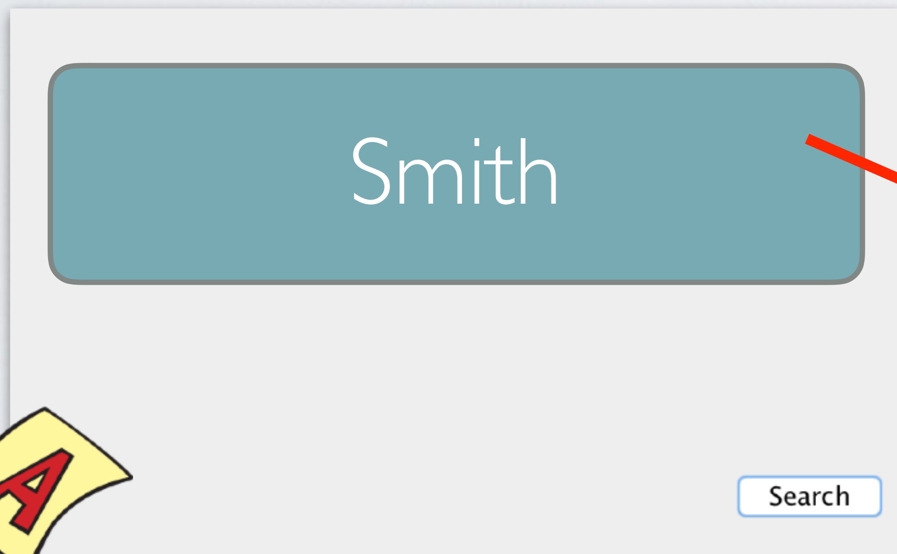

Grades

First_Name	Last_Name	Dept.	Grade
...	...	...	...
Sarah	Smith	CE	A
John	Smith	EE	B
Kerry	Smith	CS	D
...	...	...	...

Results

First_Name	Last_Name	Dept.	Grade

- Database has too many tuples matching query

Database system can learn as well

Smith

Search

Grades			
First_Name	Last_Name	Dept.	Grade
…	…	…	…
Sarah	Smith	CE	A
John	Smith	EE	B
Kerry	Smith	CS	D
…	…	…	…

Results			
First_Name	Last_Name	Dept.	Grade
Kerry	Smith	CS	D

- Database system has learned to return *Kerry Smith* in CS department

Database system can learn as well

Smith

Search

Grades			
First_Name	Last_Name	Dept.	Grade
…	…	…	…
Sarah	Smith	CE	A
John	Smith	EE	B
Kerry	Smith	CS	D
…	…	…	…

Results			
First_Name	Last_Name	Dept.	Grade
Kerry	Smith	CS	D

- The user finds and selects the tuple

Interaction is a game between two potentially rational agents

- Two Players: **user** and **database system**
- They have common interests and work together
  - Want to reach a mutual understanding such that user gets desired information
- Strategy of the user is how intents are expressed using queries
- Strategy of the database system is how to decode queries

User strategy

Intent #	Intent
e_1	*John Smith in EE*
e_2	*Sarah Smith in CE*
e_3	*Kerry Smith in CS*

Query #	Query
q_1	*"Smith CE"*
q_2	*"Smith"*

- Row-stochastic mapping from intents to queries.

User Strategy (U)

	q_1	q_2
e_1	0	1
e_2	1	0
e_3	0	1

Grades

First_Name	Last_Name	Dept.	Grade
...	...	...	...
Sarah	Smith	CE	A
John	Smith	EE	B
Kerry	Smith	CS	D
...	...	...	...

User may use a single query for multiple intents

Intent #	Intent
e_1	*John Smith in EE*
e_2	*Sarah Smith in CE*
e_3	*Kerry Smith in CS*

Query #	Query
q_1	*"Smith CE"*
q_2	*"Smith"*

- Due to the lack of knowledge, saving time, …
- Makes it hard to interpret the exact intent behind the query.

User Strategy (U)

	q_1	q_2
e_1	0	1
e_2	1	0
e_3	0	1

Grades

First_Name	Last_Name	Dept.	Grade
…	…	…	…
Sarah	Smith	CE	A
John	Smith	EE	B
Kerry	Smith	CS	D
…	…	…	…

Oregon State University

Database system strategy

Intent #	Intent
e_1	*ans(y) ← Grades(x,'Smith', 'EE', y)*
e_2	*ans(y) ← Grades(x,'Smith', 'CE', y)*
e_3	*ans(y) ← Grades(x,'Smith', 'CS', y)*

Sarah Smith in CE

Database Strategy (D)

	e_1	e_2	e_3
q_1	0	1	0
q_2	0.5	0	0.5

Query #	Query
q_1	*"Smith CE"*
q_2	*"Smith"*

Grades

First_Name	Last_Name	Dept.	Grade
...	...	...	...
Sarah	Smith	CE	A
John	Smith	EE	B
Kerry	Smith	CS	D
...	...	...	...

- Row-stochastic mapping from queries to intents

Payoff: expected effectiveness of communicating every intent

Intent #	Intent
e_1	*John Smith in EE*
e_2	*Sarah Smith in CE*
e_3	*Kerry Smith in CS*

Query #	Query
q_1	*"Smith CE"*
q_2	*"Smith"*

User Strategy (U)

	q_1	q_2
e_1	0	1
e_2	1	0
e_3	0	1

Database Strategy (D)

	e_1	e_2	e_3
q_1	0	1	0
q_2	0.5	0	0.5

- Prior probability of intent

$$r(U, D) = \sum_{i=1}^{m} \pi_i \sum_{j=1}^{n} U_{ij} \sum_{\ell=1}^{o} D_{j\ell} \; prec(e_i, e_\ell)$$

Oregon State University

Payoff: expected effectiveness of communicating every intent

Intent #	Intent
e_1	*John Smith in EE*
e_2	*Sarah Smith in CE*
e_3	*Kerry Smith in CS*

Query #	Query
q_1	*"Smith CE"*
q_2	*"Smith"*

User Strategy (U)

	q_1	q_2
e_1	0	1
e_2	1	0
e_3	0	1

Database Strategy (D)

	e_1	e_2	e_3
q_1	0	1	0
q_2	0.5	0	0.5

$$r(U,D) = \sum_{i=1}^{m} \pi_i \sum_{j=1}^{n} U_{ij} \sum_{\ell=1}^{o} D_{j\ell}\ prec(e_i, e_\ell)$$

Payoff: expected effectiveness of communicating every intent

Intent #	Intent
e_1	*John Smith in EE*
e_2	*Sarah Smith in CE*
e_3	*Kerry Smith in CS*

Query #	Query
q_1	*"Smith CE"*
q_2	*"Smith"*

User Strategy (U)

	q_1	q_2
e_1	0	1
e_2	1	0
e_3	0	1

Database Strategy (D)

	e_1	e_2	e_3
q_1	0	1	0
q_2	0.5	0	0.5

$$r(U,D) = \sum_{i=1}^{m} \pi_i \sum_{j=1}^{n} U_{ij} \sum_{\ell=1}^{o} D_{j\ell} \; prec(e_i, e_\ell)$$

Payoff: expected effectiveness of communicating every intent

Intent #	Intent
e_1	*John Smith in EE*
e_2	*Sarah Smith in CE*
e_3	*Kerry Smith in CS*

Query #	Query
q_1	*"Smith CE"*
q_2	*"Smith"*

User Strategy (U)

	q_1	q_2
e_1	0	1
e_2	1	0
e_3	0	1

Database Strategy (D)

	e_1	e_2	e_3
q_1	0	1	0
q_2	0.5	0	0.5

$$r(U,D) = \sum_{i=1}^{m} \pi_i \sum_{j=1}^{n} U_{ij} \sum_{\ell=1}^{o} D_{j\ell}\, prec(e_i, e_\ell)$$

- Precision is the fraction of the returned tuples that are desired
- Computed using user feedback

Interesting problems

1. What are the stable states (equilibria) of the game? Is there any undesirable (sub-optimal) equilibria?

2. What are the user's learning mechanisms?

3. What learning algorithms should the database system adopt so the collaboration converges to desirable equilibria?

    1. Learning may not converge or converge to a desired equilibrium in games, e.g., *Shapely game*.

Equilibria of the game

- **Nash Equilibrium:** A strategy profile in which no player can increase its payoff by unilaterally deviating from the current strategy

Intent #	Intent
e_1	*John Smith in EE*
e_2	*Sarah Smith in CE*
e_3	*Kerry Smith in CS*

Query #	Query
q_1	*"Smith CE"*
q_2	*"Smith"*

User Strategy (U)

	q_1	q_2
e_1	0	1
e_2	1	0
e_3	0	1

Database Strategy (D)

	e_1	e_2	e_3
q_1	0	1	0
q_2	0.5	0	0.5

$$r(U,D) = 2$$

The game has Nash equilibria with sub-optimal payoff

- User express all intents using q_2: "Smith"
- Database system aways returns e_2: *Sara Smith in CE*

Intent #	Intent
e_1	*John Smith in EE*
e_2	*Sarah Smith in CE*
e_3	*Kerry Smith in CS*

Query #	Query
q_1	*"Smith CE"*
q_2	*"Smith"*

User Strategy (U)

	q_1	q_2
e_1	0	1
e_2	0	1
e_3	0	1

Database Strategy (D)

	e_1	e_2	e_3
q_1	0	1	0
q_2	0	1	0

$$r(U,D) = 1$$

- Detailed analyses are at http://tinyurl.com/charmarxiv

The game has Nash equilibria with sub-optimal payoff

- If user learns query q_1 to represent e_2, payoff will not increase

Intent #	Intent
e_1	*John Smith in EE*
e_2	*Sarah Smith in CE*
e_3	*Kerry Smith in CS*

Query #	Query
q_1	*"Smith CE"*
q_2	*"Smith"*

User Strategy (U)

	q_1	q_2
e_1	0	1
e_2	0	1
e_3	0	1

Database Strategy (D)

	e_1	e_2	e_3
q_1	0	1	0
q_2	0	1	0

$$r(U,D) = 1$$

.User Strategy (U)

	q_1	q_2
e_1	0	1
e_2	1	0
e_3	0	1

Database Strategy

	e_1	e_2	e_3
q_1	0	1	0
q_2	0	1	0

$$r(U,D) = 1$$

- Details at http://tinyurl.com/charmarxiv

How users may learn?

- Research in psychology shows that humans exhibit reinforcement learning behavior
- Select a query based on its past payoff, i.e., exploitation.
- Explore and try new/ less successful queries to gain new knowledge, i.e., exploration.
    - Sacrifice payoff in the short-term in the hope of more payoff over the long run.

User learning mechanism

Short-term memory

- Win-Stay/Lose-Randomize: keeps using a query with non-zero payoff, randomly picks a query otherwise.
- Latest-Reward: uses a query with probability proportional to its latest payoff

User learning mechanism

Long-term memory

- Bush and Mosteller's: Reinforces probability of using a query with non-zero payoff by an amount independent of payoff
- Roth and Erev's: Reinforces probability of using a query proportional to its accumulated payoff
- Roth and Erev's Modified: Adds the ability to forget to Roth and Erev
- Cross's: Reinforces probability of using a query proportional to a linear adjustment of its accumulated payoff

Oregon State University

Empirical evaluation

- Yahoo query log over 300,00 interactions
  - https://webscope.sandbox.yahoo.com

Method	Mean Squared Distance
Bush and Mosteller's	0.0112
Cross's	0.01131
Roth and Erev	0.00993
Roth and Erev Modified	0.00994
Win-Stay/Lose-Randomize	0.01752
Latest-Reward	0.15167

What learning strategies should database system use?

- Current database systems assume that user strategy is fixed.
  - They model the problem as stochastic multi-armed bandit and use online learning algorithms, such as UCB-1
- We have used Roth and Erev reinforcement algorithms for the database system learning.
  - Uses randomization to explore available options
- **Theorem:** If players use the Roth and Erev method, the sequence of payoffs is a submartingale (statistically non-decreasing) and converges almost surely.

Roth and Erev outperforms UCB-1 in the long run

- Yahoo! interaction log

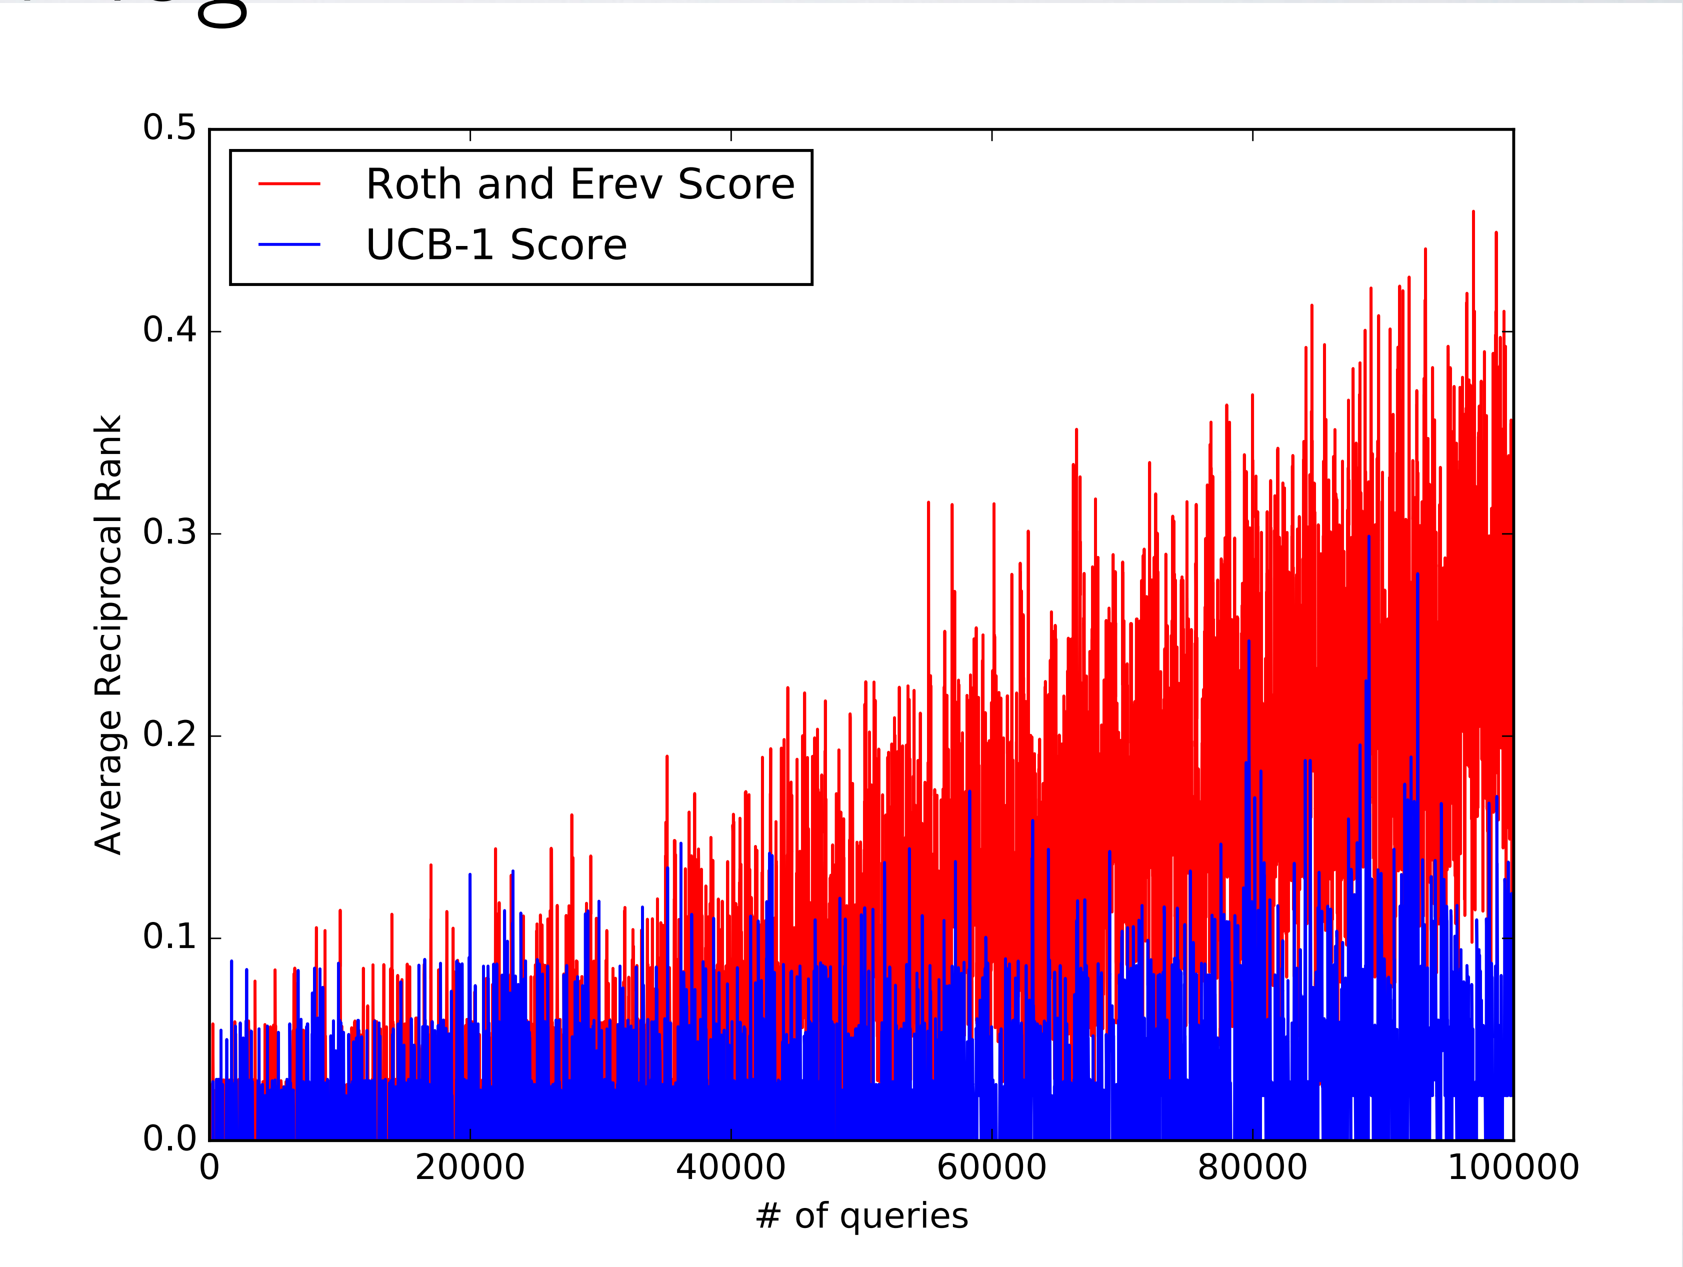

Conclusion

- The interaction between user and database systems is better modeled as a collaborative game
- The game has both desirable and undesirable equilibria
- Users are rather surprisingly intelligent learners
- Database system should use randomized learning strategies.
    - More information at our technical report: http://tinyurl.com/charmarxiv