Teaching Machines like we Teach People

Igor Labutov

Joint work with: Bishan Yang, Shashank Srivastava, Amos Azaria, Anusha Prakash, Tom Mitchell

LAER AI, Inc.

Cornell Tech
Jacobs-Technion Institute

Machine Learning Department
Carnegie Mellon University

Machine Learning *for* Machine Teaching

Machine Learning *for* Machine Teaching

There are six homeworks this semester, each due on …

In addition to homeworks, there are projects …

All assignments have a deadline

If a student emails a late homework, deduct 1 point per day …

A Personal Ontology

There are six homeworks this semester, each due on …

In addition to homeworks, there are projects …

All assignments have a deadline

If a student emails a late homework, deduct 1 point per day …

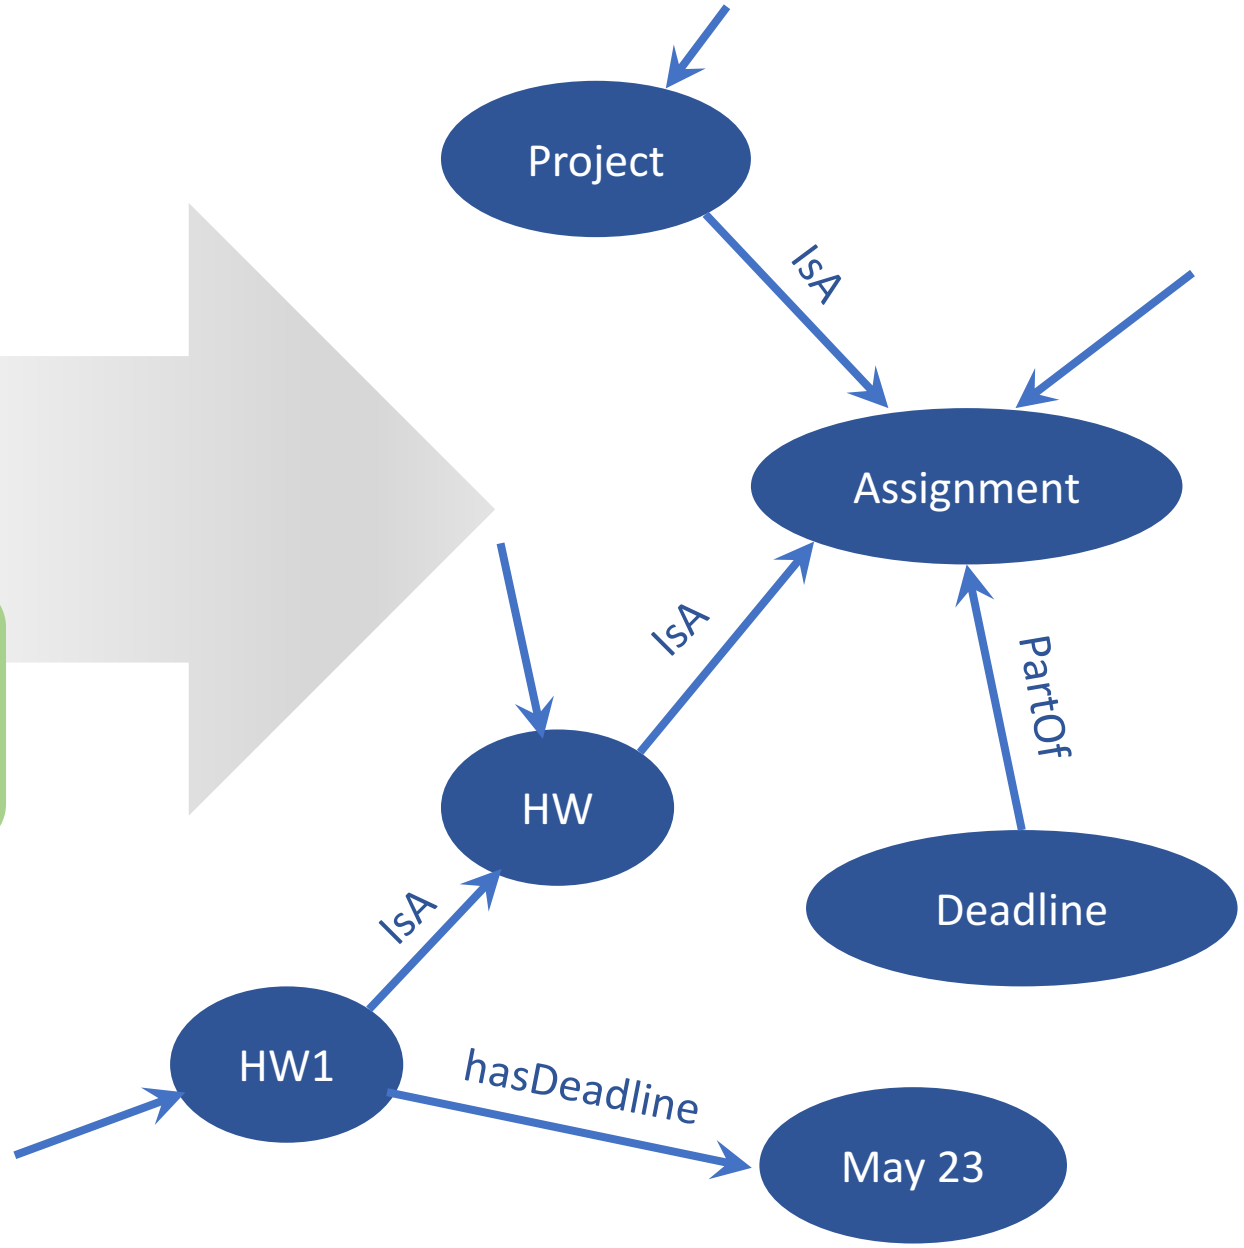

A Personal Ontology

There are six homeworks this semester, each due on …

In addition to homeworks, there are projects …

All assignments have a deadline

If a student emails a late homework, deduct 1 point per day …

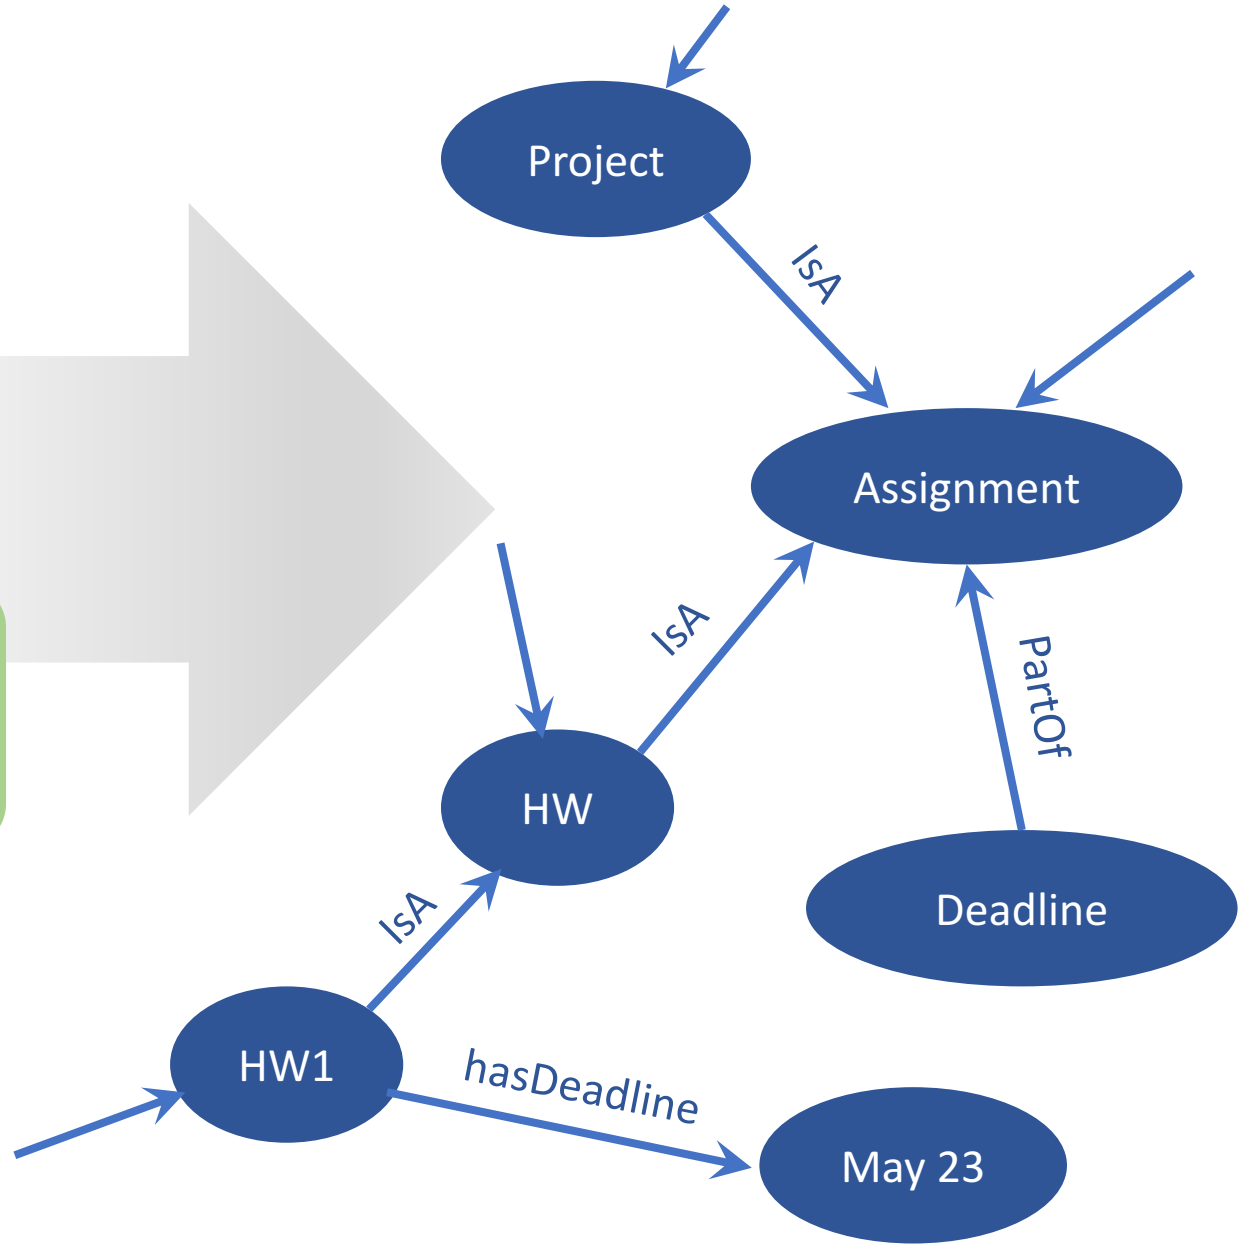

DEMO

for example his affiliation is CMU

Setting university affiliation field of oscar to cmu.

most of my colleagues have a university affiliation

create a field called "university affiliation" for concept colleague, right?

for example his affiliation is CMU

Setting university affiliation field of oscar to cmu.

most of my colleagues have a university affiliation

create a field called "university affiliation" for concept colleague, right?

also make sure that the sender has an email address of a colleague of mine

check that some email address fields of colleagues instances appears in sender field of email that arrived, right?

if there is an important meeting request from a colleague then put it on my calendar

I don't understand what you mean by "there is an important meeting request from a colleague", you can explain or demonstrate if you want to teach me, or just rephrase

for example his affiliation is CMU

Setting university affiliation field of oscar to cmu.

most of my colleagues have a university affiliation

create a field called "university affiliation" for concept colleague, right?

also make sure that the sender has an email address of a colleague of mine

check that some email address fields of colleagues instances appears in sender field of email that arrived, right?

if there is an important meeting request from a colleague then put it on my calendar

I don't understand what you mean by "there is an important meeting request from a colleague", you can explain or demonstrate if you want to teach me, or just rephrase

check if the subject mentions words like meet or meeting

I will check that subject field of email that arrived is "like meet or meeting", right?

for example his affiliation is CMU

Setting university affiliation field of oscar to cmu.

most of my colleagues have a university affiliation

create a field called "university affiliation" for concept colleague, right?

also make sure that the sender has an email address of a colleague of mine

check that some email address fields of colleagues instances appears in sender field of email that arrived, right?

if there is an important meeting request from a colleague then put it on my calendar

I don't understand what you mean by "there is an important meeting request from a colleague", you can explain or demonstrate if you want to teach me, or just rephrase

check if the subject mentions words like meet or meeting

I will check that subject field of email that arrived is "like meet or meeting", right?

no

Sorry, did you actually mean I will check that subject field of email that arrived is either "meet", "meeting"?

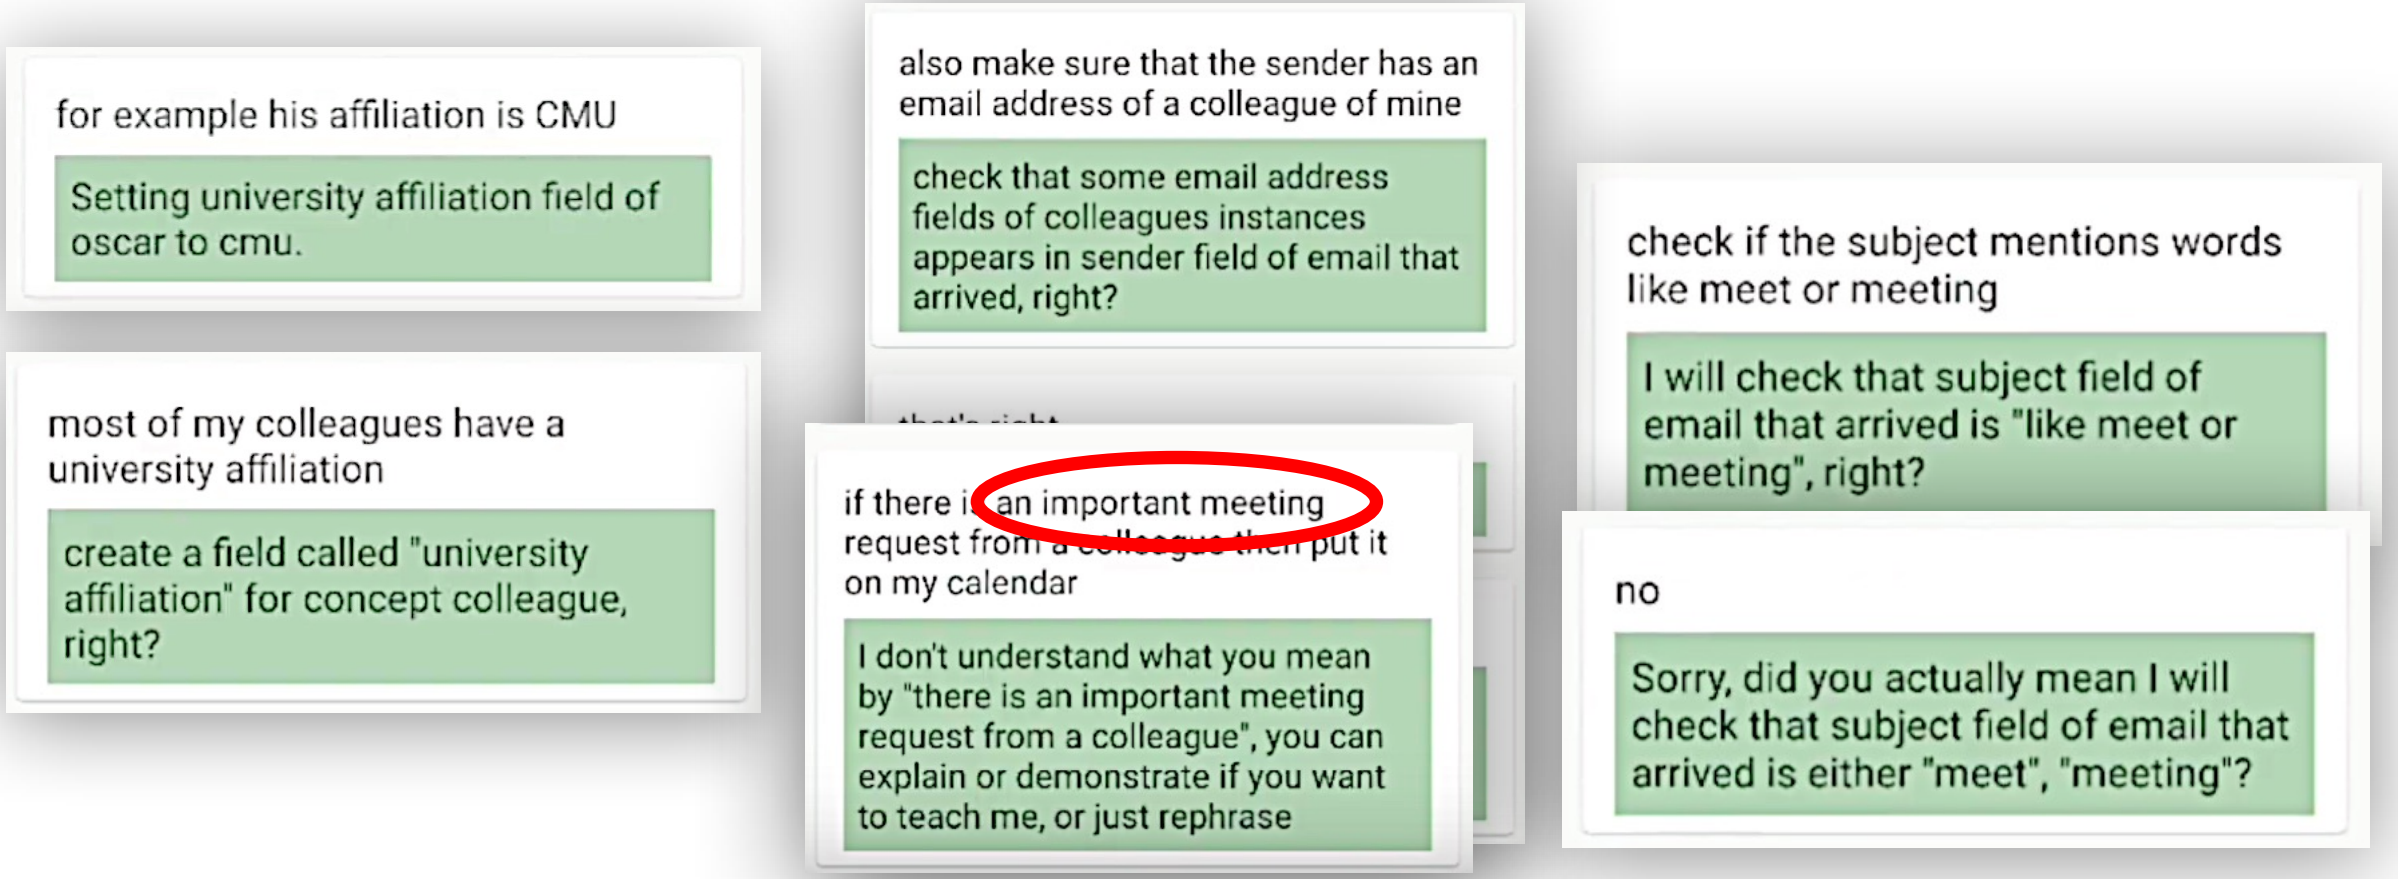

Knowledge-Base Construction

Concept Learning

Semantic Parsing

Learning To Interpret Language (Semantic Parsing)

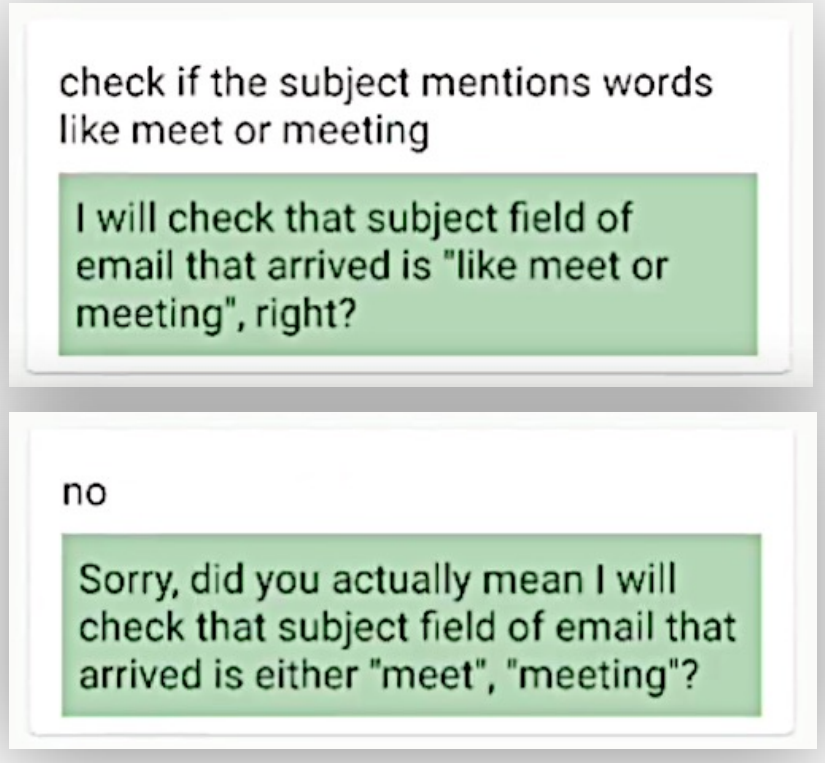

Learning To Interpret Language (Semantic Parsing)

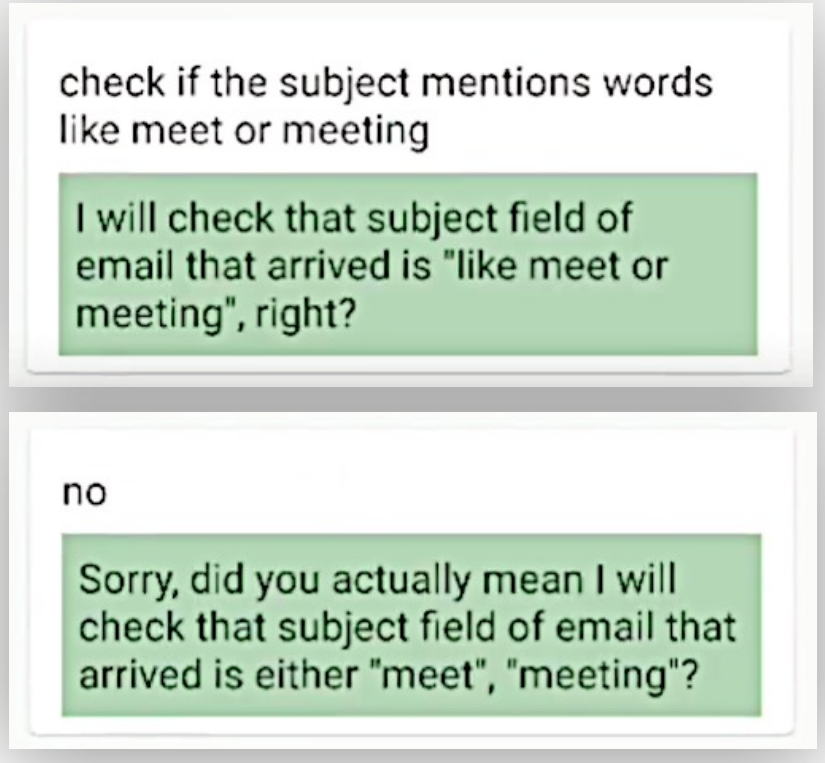

Negative Feedback

Learning To Interpret Language (Semantic Parsing)

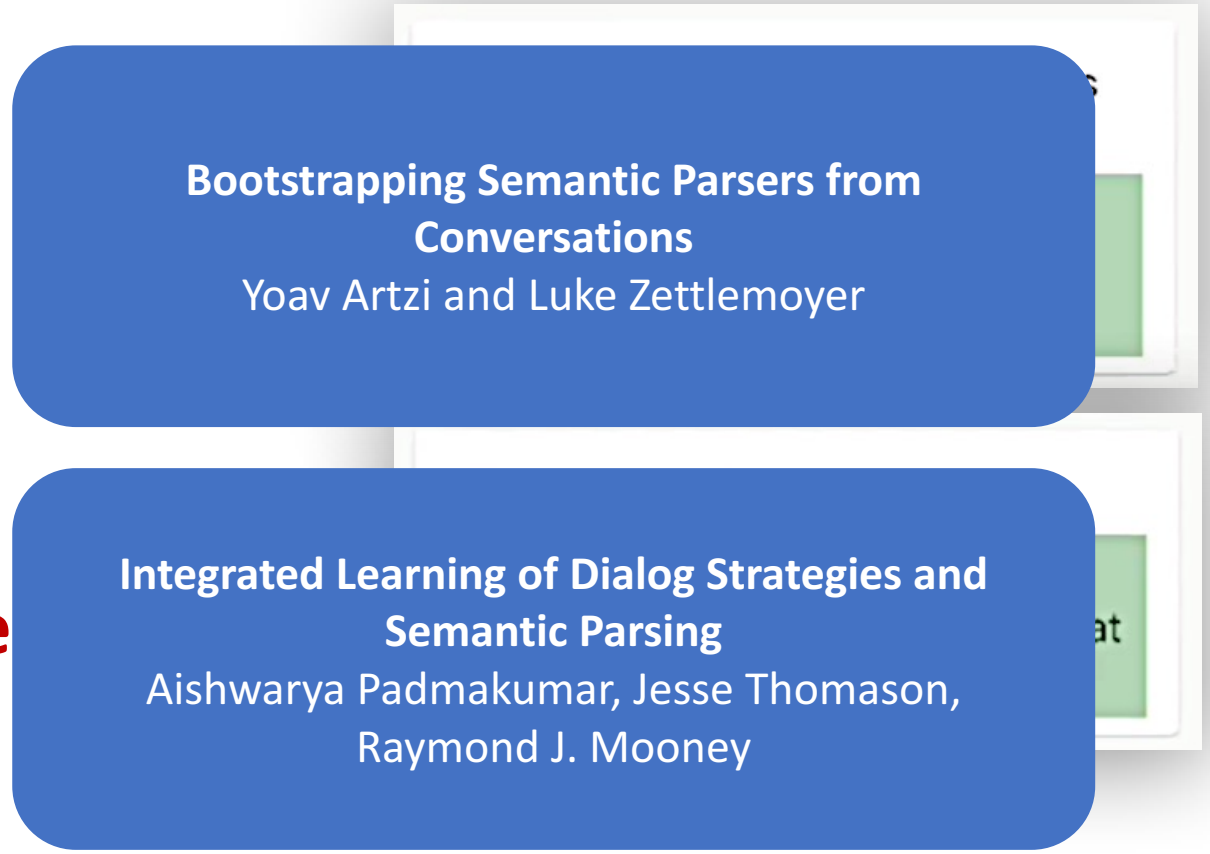

Negative

Learning To Interpret Language (Semantic Parsing)

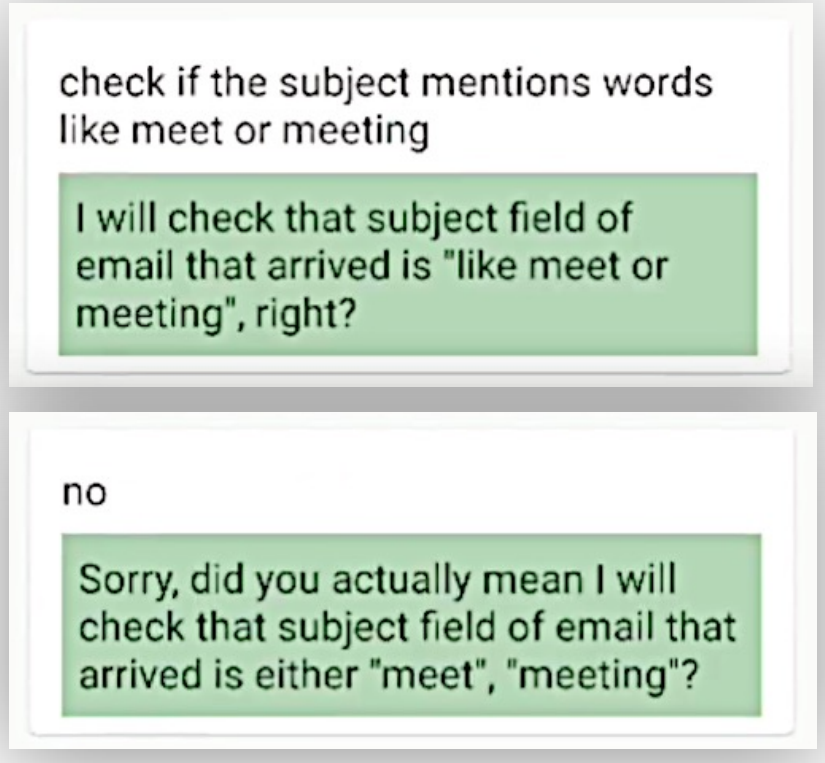

Negative Feedback

Learning To Interpret Language (Semantic Parsing)

Implicit Instruction vs Explicit Instruction

Annotated

Execution
(denotatio

ions words

field of
meet or

no , I meant the word meeting, not like meeting

Ah, got it!

Negative Feedback

Learning To Interpret Language (Semantic Parsing)

Implicit Instruction vs Explicit Instruction

Annotated

Execution
(denotatic

ions words

field of
meet or

no , I meant the word meeting, not like meeting

Ah, got it!

ative Feedback

EMNLP 2018
Learning to Learn Semantic Parsers from Natural Language Supervision
Igor Labutov, Bishan Yang,
Tom Mitchell

User Utterance	In [2014] how many jobs did he hold?
System NLG	find number of employments [Bob] had, prior to [2014]?
User Feedback	*During 2014, not prior to.*

User Utterance	Before [September 10th], how many places was [Brad] employed?
System NLG	find employment [Brad] had, prior to [September 10th]
User Feedback	*Yes, but I also want you to count how many placed Brad was employed at.*

User Utterance	anytime after [1995], did [Matt] talk to anybody about this at [JP Morgan]?
System NLG	find communications between [Matt] and people employed at [JP Morgan] about anything, that occured after [1995]
User Feedback	*No, I mean about the stock market, not just anything.*

User Utterance	Do we know if [John] and [Neal] spoke?
System NLG	find communications between [John] and [Neal] about [the financial reports] at any time
User Feedback	*I mean about anything, not about the financial reports.*

User Utterance	what was his last day at [Samsung]?
System NLG	find when employment of [Andy] at [Samsung] started
User Feedback	*No, when it ended.*

User-Agent Dialog

U

User's task utterance

NLG - Did you mean…

System's confirmation

F – No, what I meant was

User's feedback

User-Agent Dialog

U

User's task utterance

NLG - Did you mean…

=

Old parse

Original parse (possibly incorrect)

F – No, what I meant was

User's feedback

User-Agent Dialog

Hidden

Logical Form parse of *u*

True parse

U

User's task utterance

NLG - Did you mean…

=

Old parse

Original parse (possibly incorrect)

F – No, what I meant was

User's feedback

User-Agent Dialog

Hidden

Logical Form parse of *u*

True parse

U

User's task utterance

NLG - Did you mean…

=

Old parse

Original parse (possibly incorrect)

f – No, what I meant was

User's feedback

Task Parser: $P(y \mid u; \theta_t)$

Feedback Parser: $P(y \mid u, f, \hat{y}; \theta_f)$

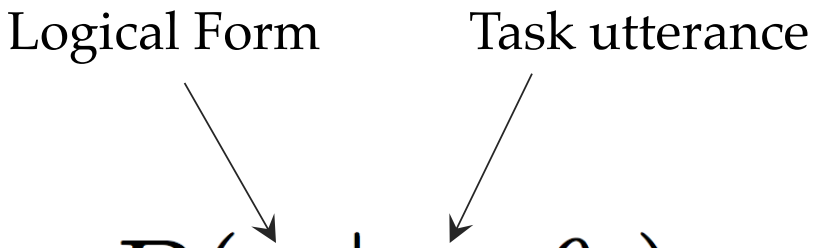

Task Parser: $P(y \mid u; \theta_t)$

Feedback Parser: $P(y \mid u, f, \hat{y}; \theta_f)$

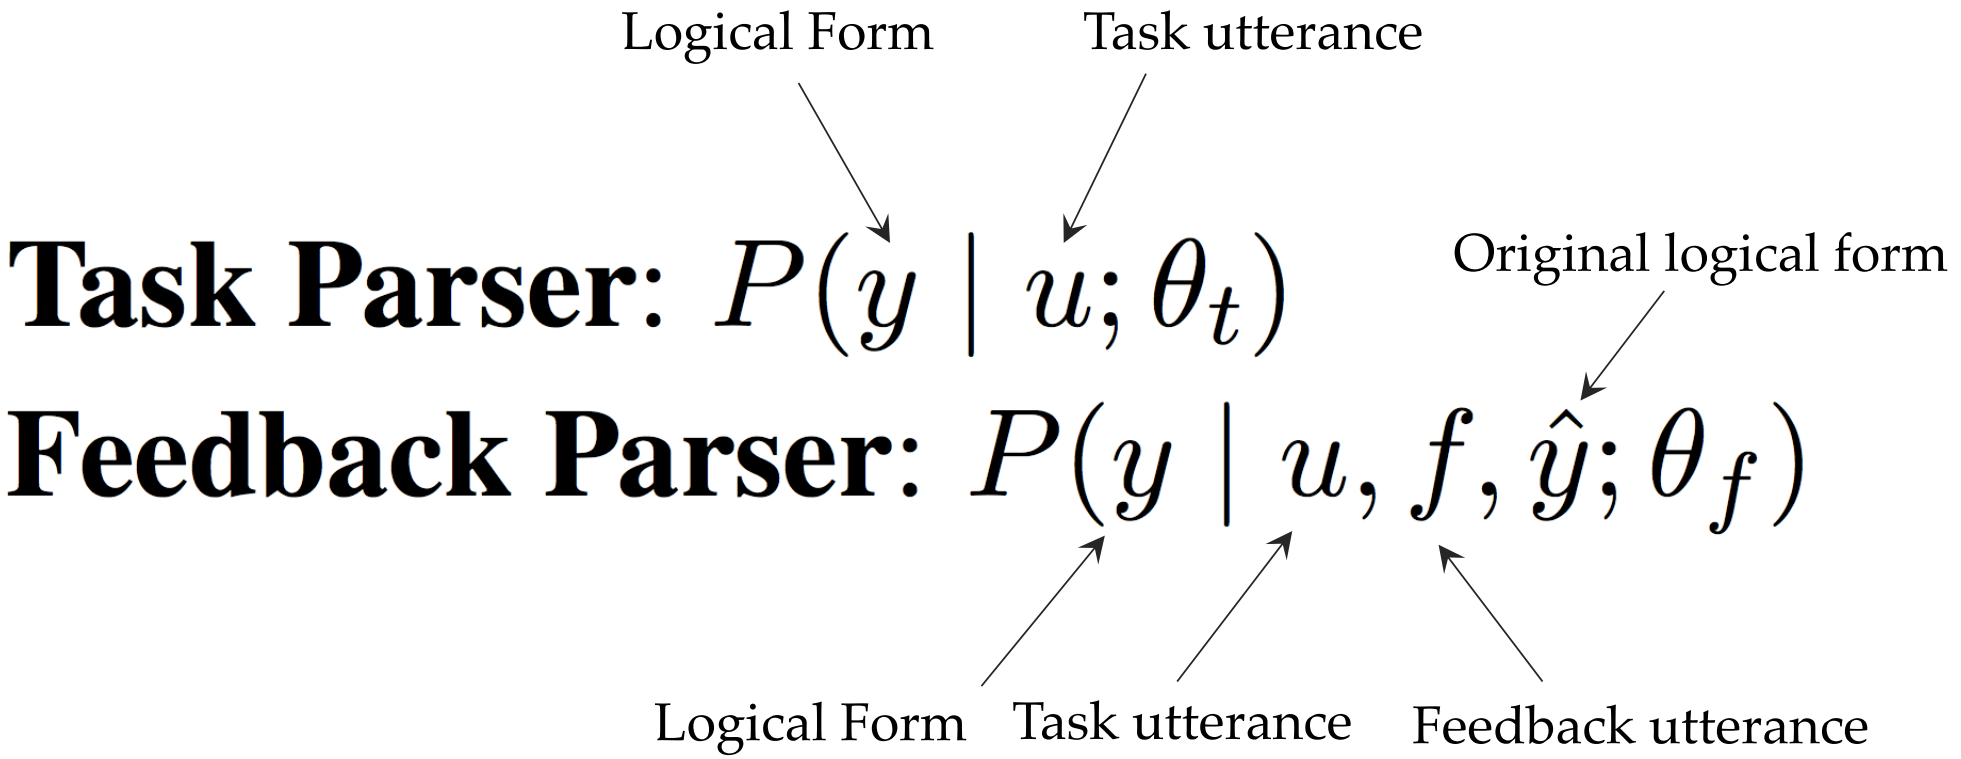
Logical Form
Task utterance
Task Parser: $P(y \mid u; \theta_t)$
Original logical form
Feedback Parser: $P(y \mid u, f, \hat{y}; \theta_f)$
Logical Form
Task utterance
Feedback utterance

Logical Form Task utterance

Task Parser: $P(y \mid u; \theta_t)$

Original logical form

Feedback Parser: $P(y \mid u, f, \hat{y}; \theta_f)$

Logical Form Task utterance Feedback utterance

Joint Model

$$P(y \mid u, f, \hat{y}; \theta_t, \theta_f) = \frac{1}{Z} \underbrace{P(y \mid u; \theta_t)}_{\text{task parser}} \underbrace{P(y \mid u, f, \hat{y}; \theta_f)}_{\text{feedback parser}}$$

Parsing Accuracy on held-out data

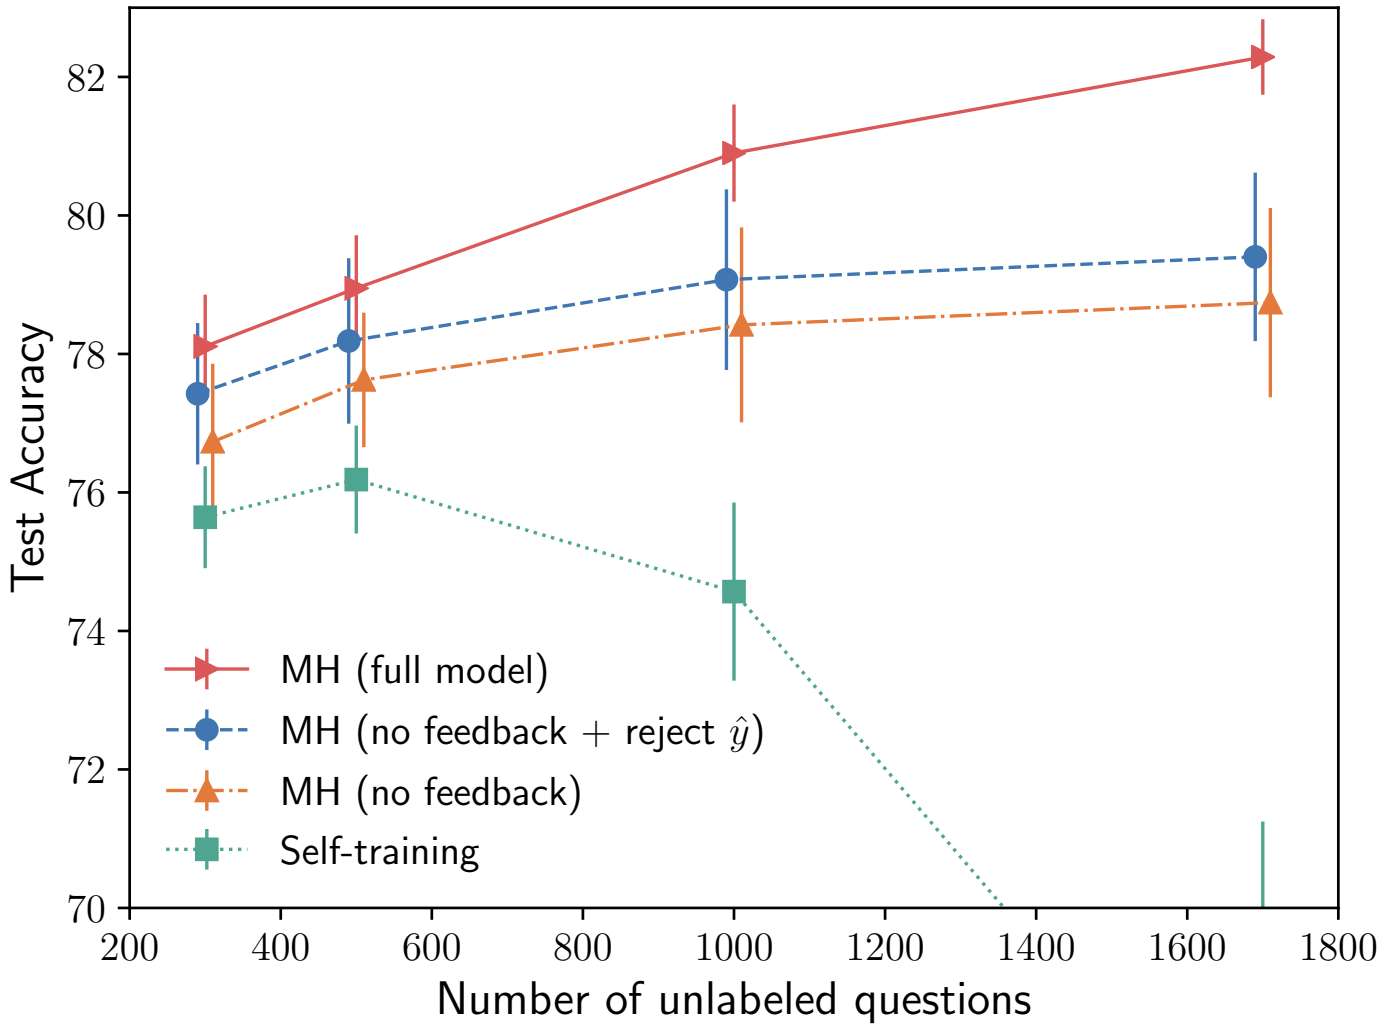

Inferring correct logical form from feedback of different "complexity"

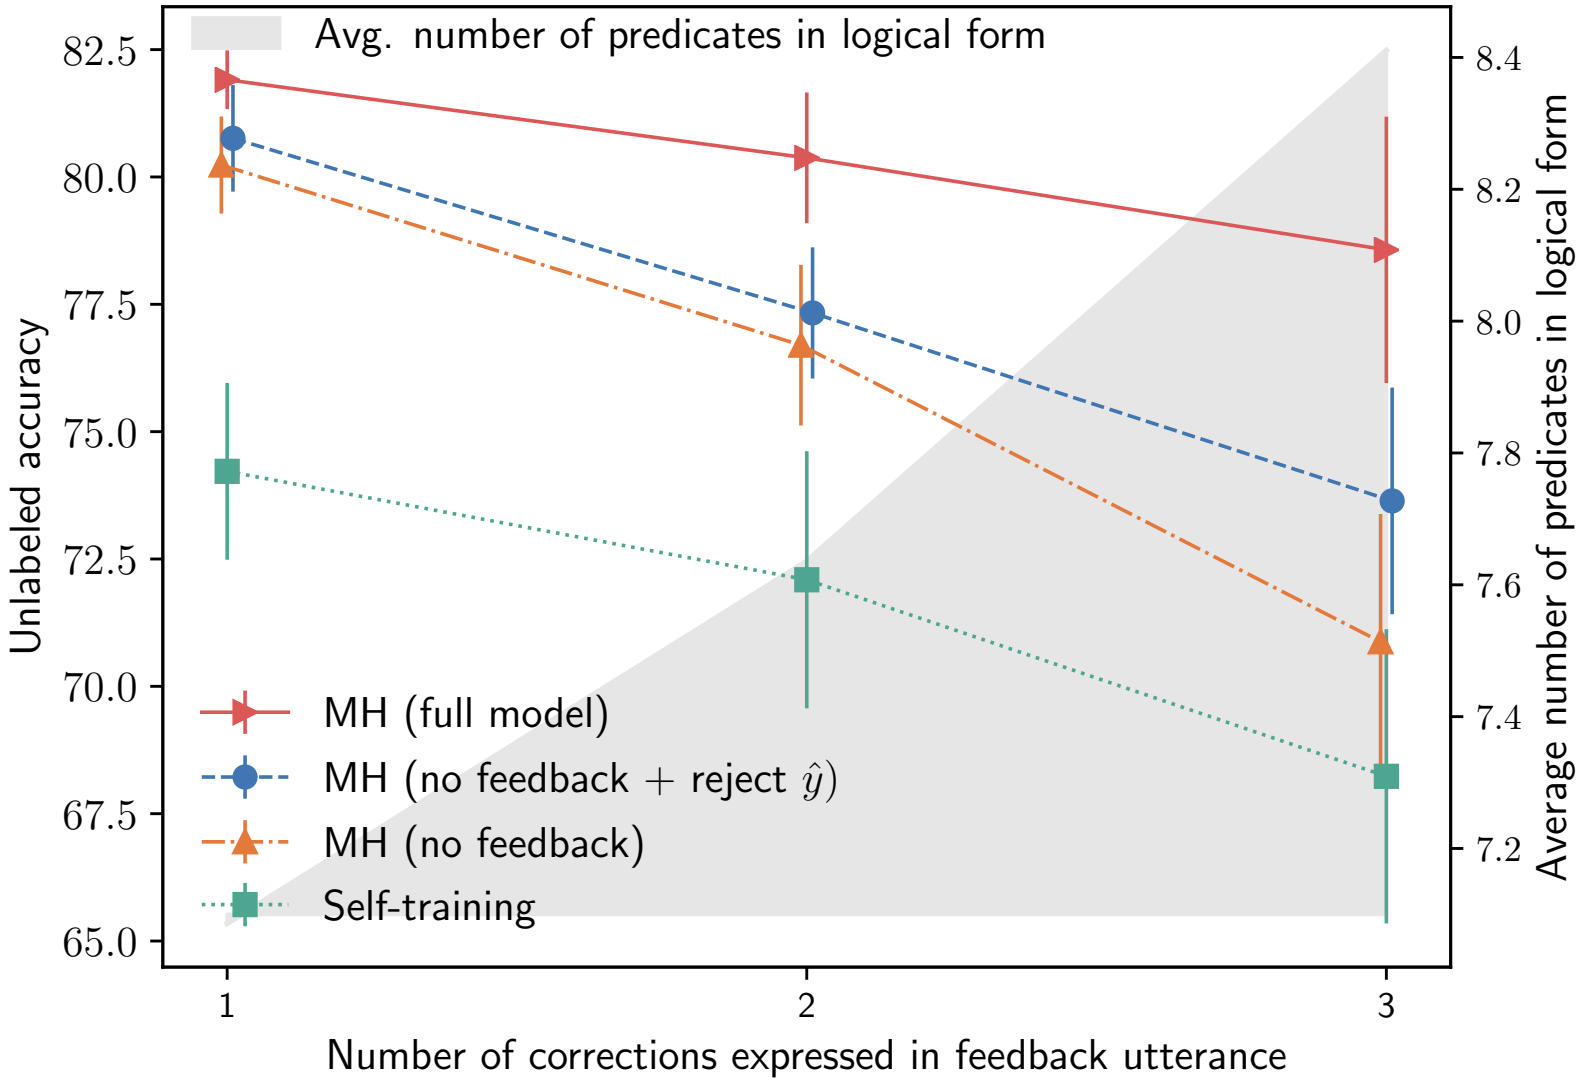

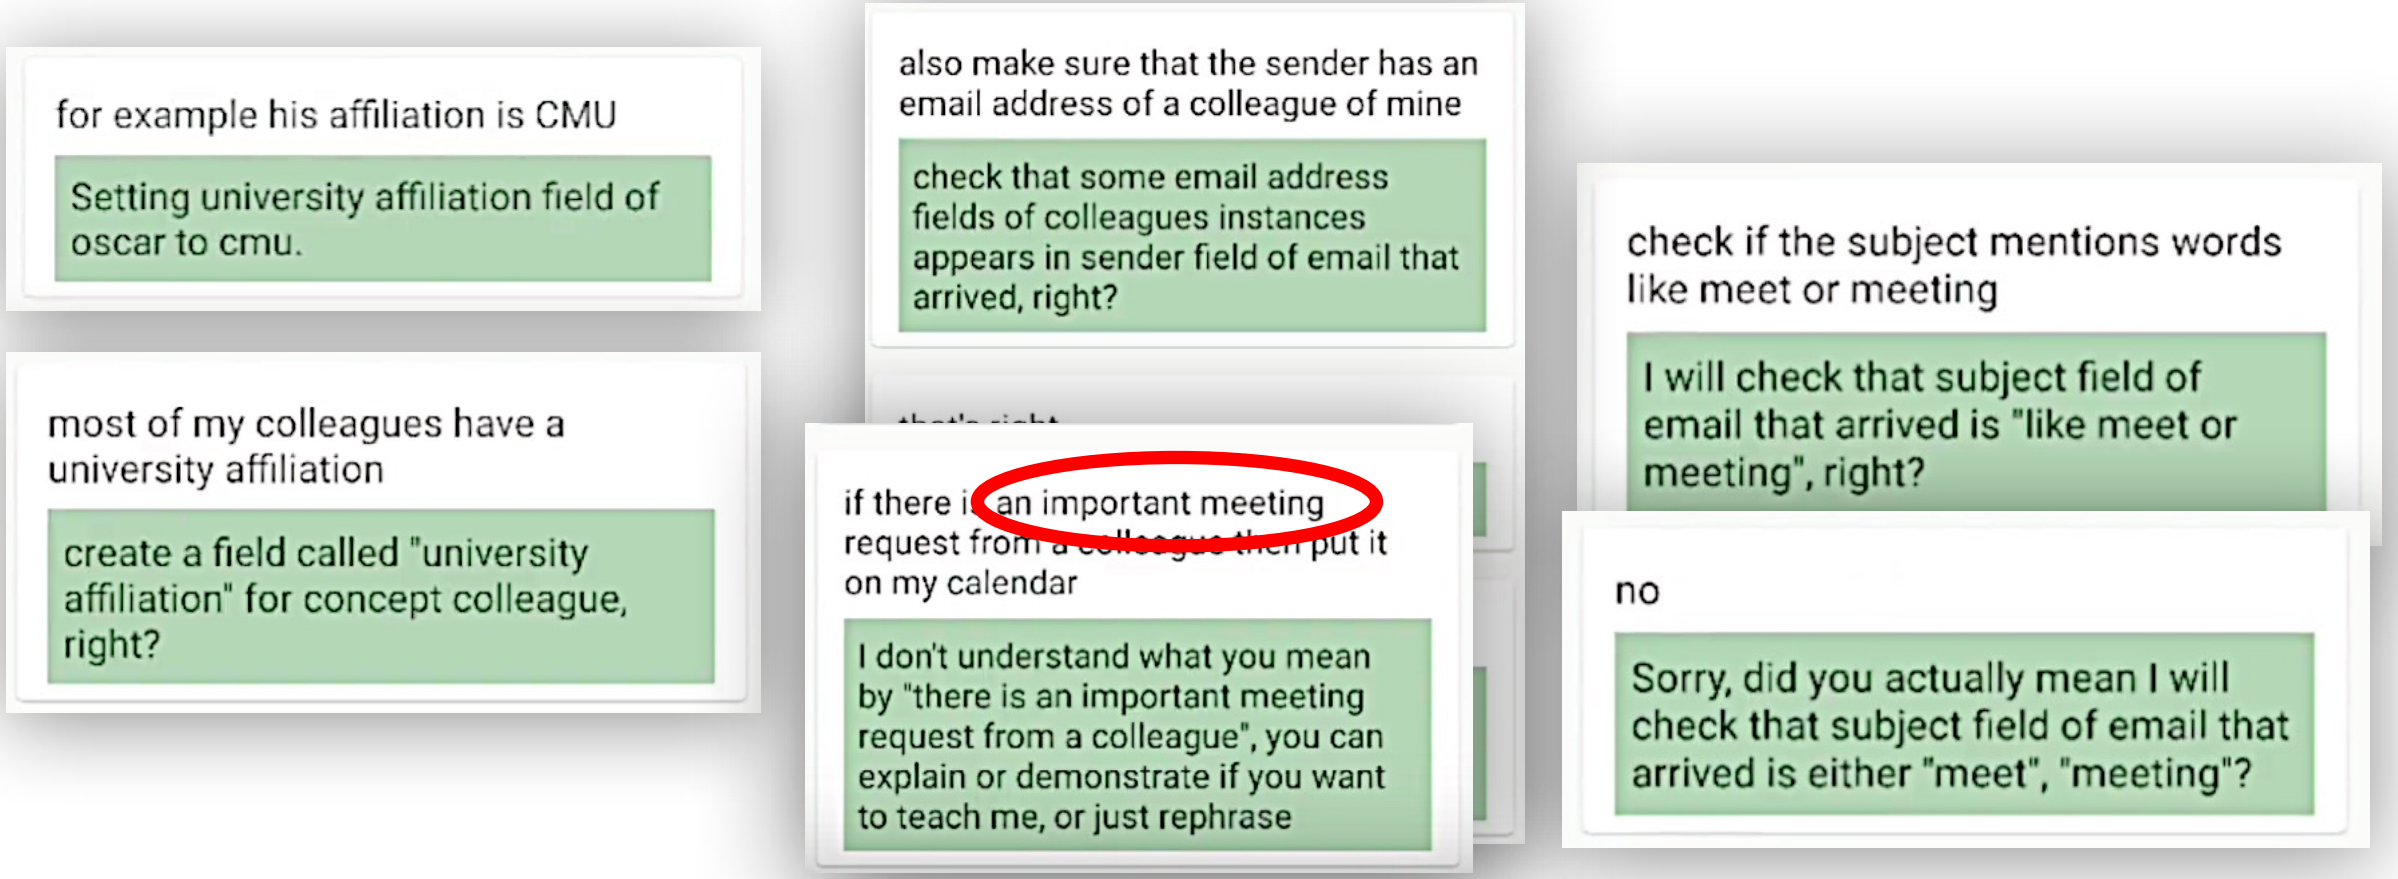

Knowledge-Base Construction

Concept Learning

Semantic Parsing

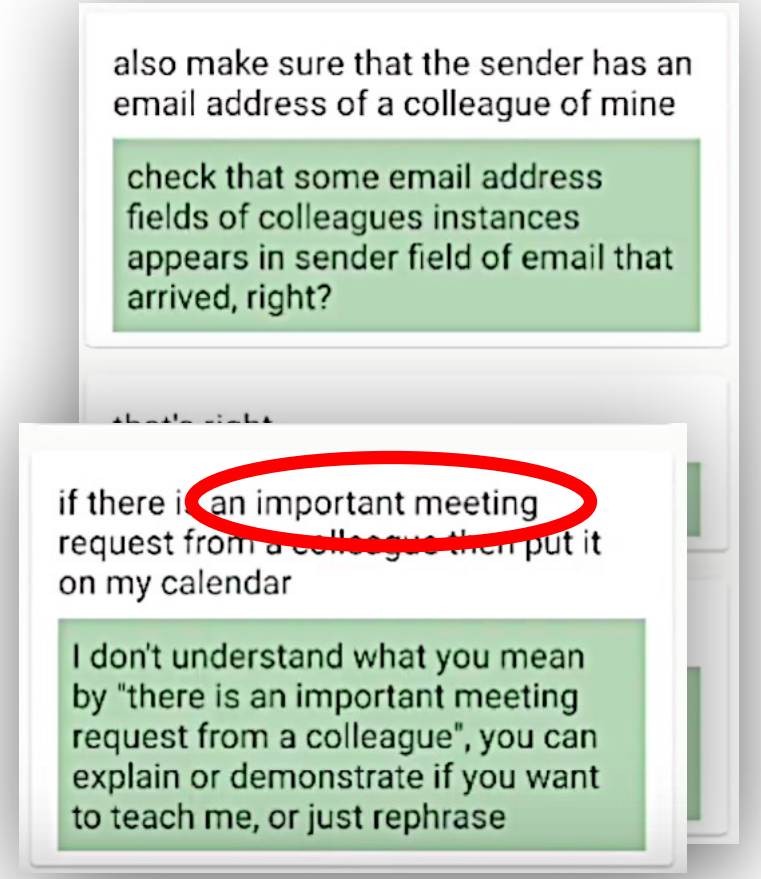

Concept Learning

Grounding Concepts

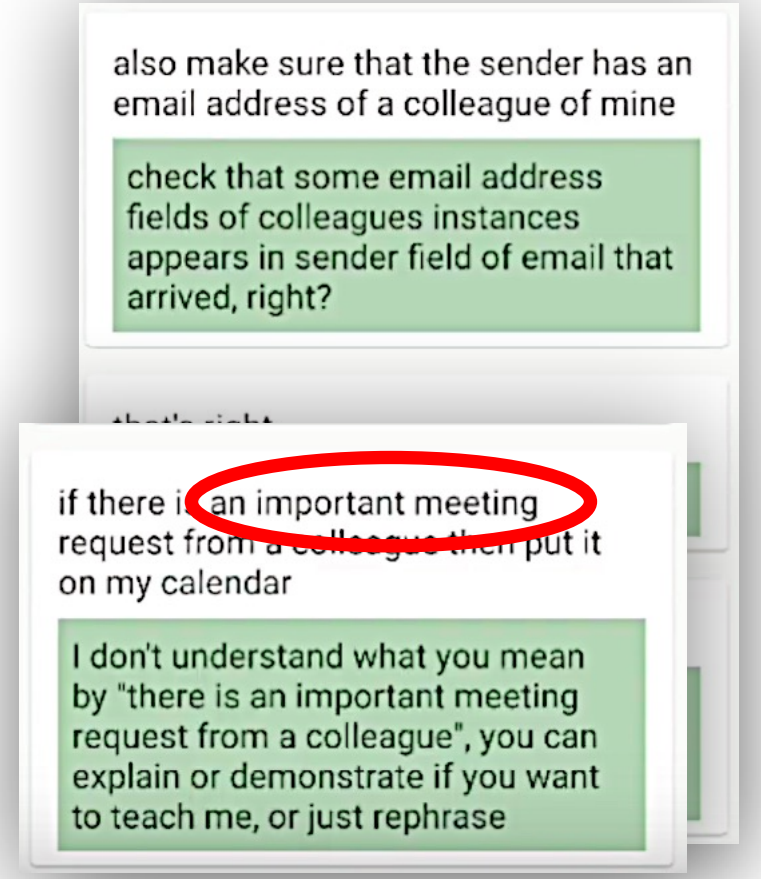

Concept Learning

EMNLP 2017

Joint concept learning and semantic parsing from natural language explanations

Shashank Srivastava, Igor Labutov, Tom Mitchell

Teaching fuzzy concepts

“Important Email”
“Meeting Request”
“Job opportunity”

Teaching fuzzy concepts

"These emails usually closes with a name or title"

"Emails from a public domain are most likely not office requests"

"Talk announcements will usually be from a colleague"

Learning a classifier from Natural Language Descriptions

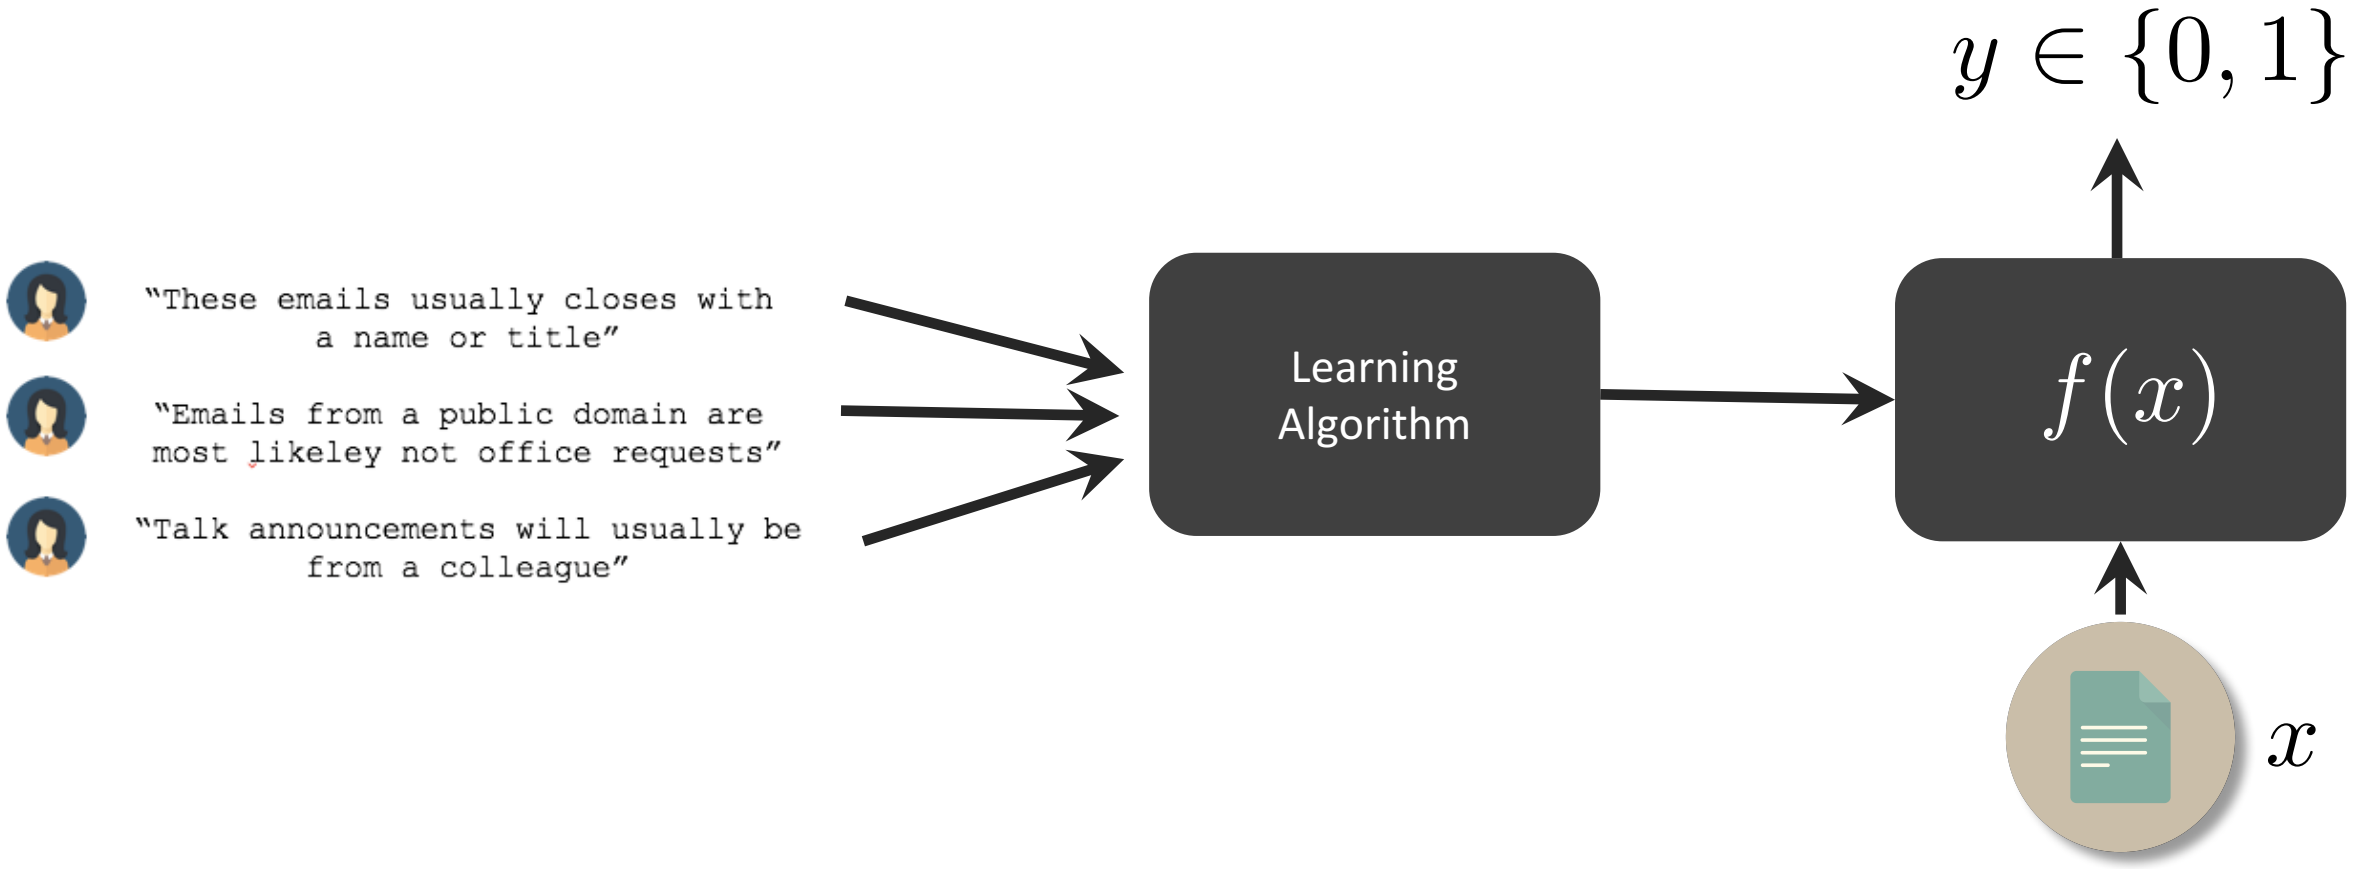

Learning a classifier from Natural Language Descriptions

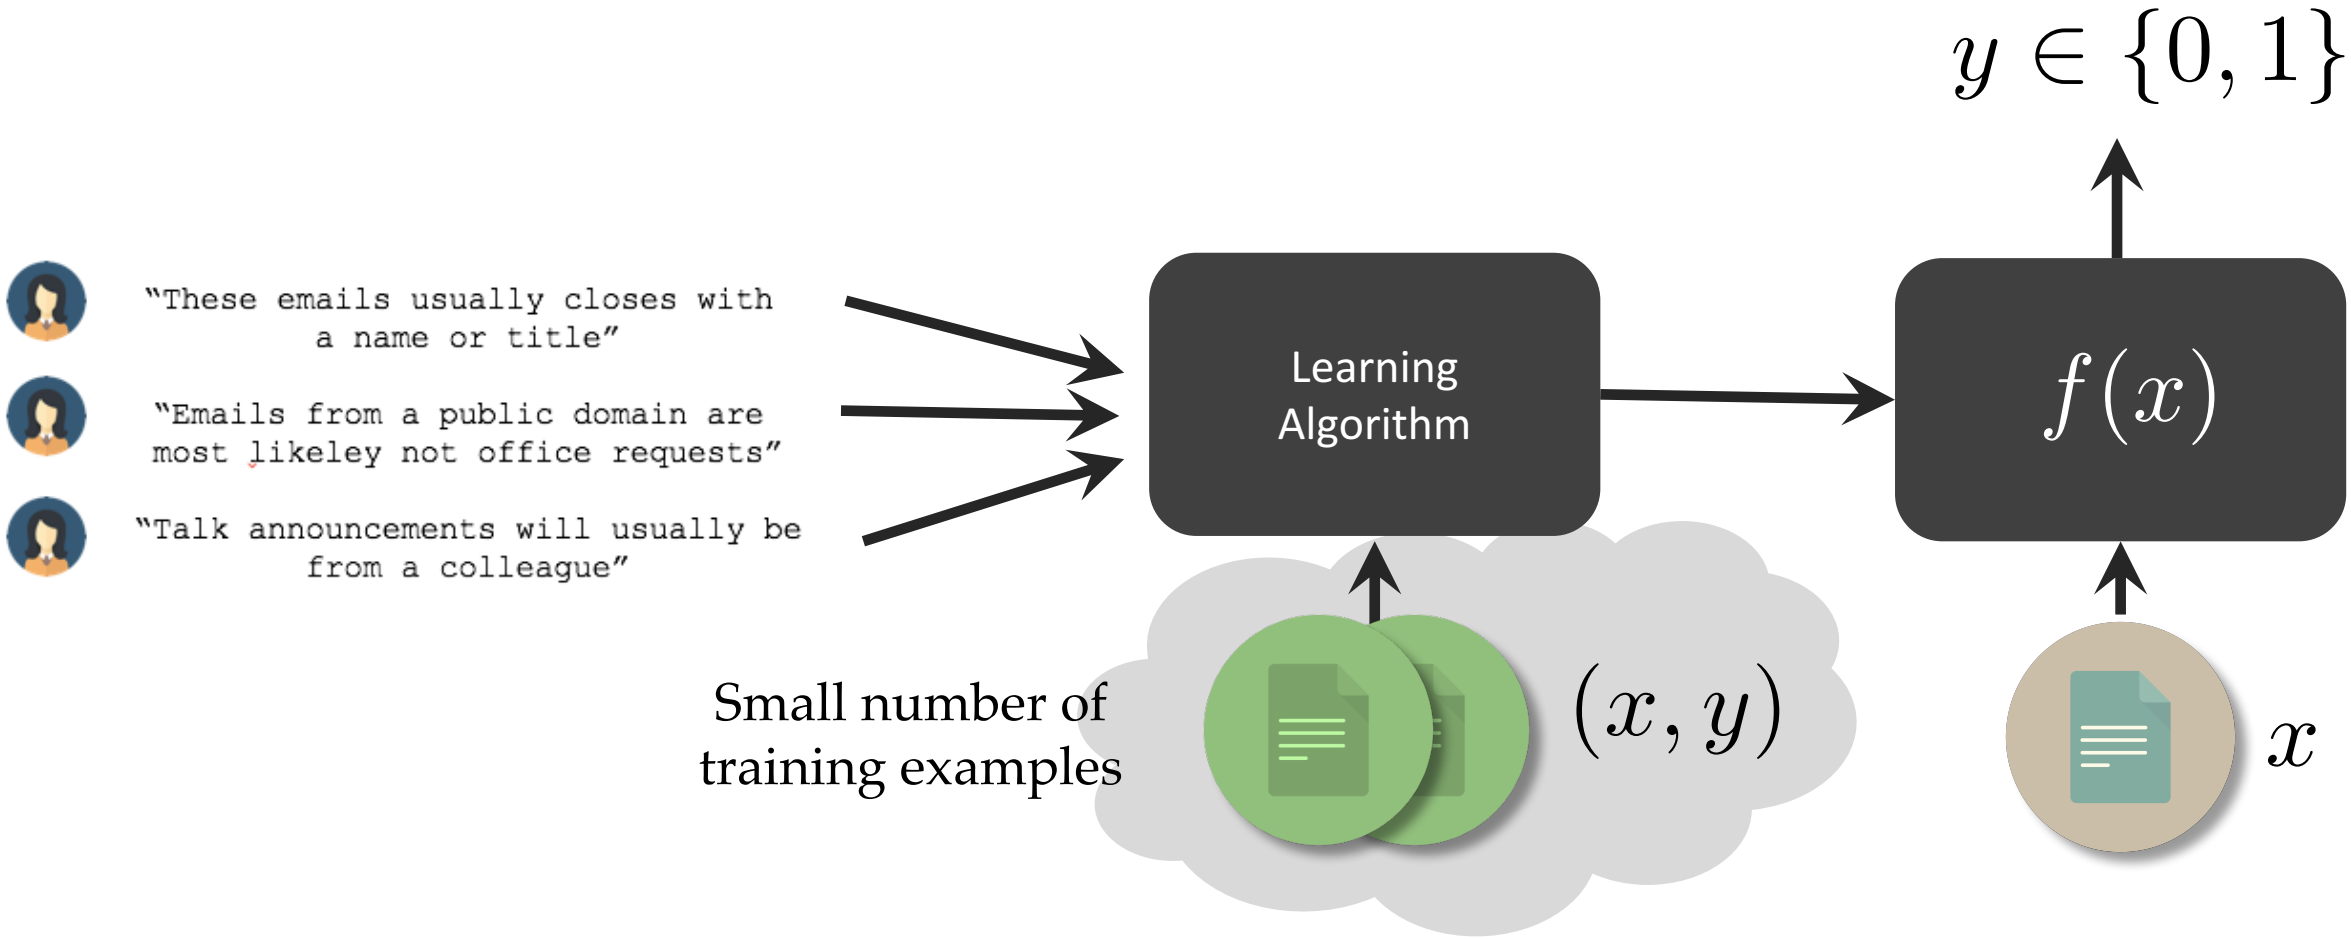

Learning a classifier from Natural Language Descriptions *(under the hood)*

1. Parse NL concept descriptions to feature functions

"Emails from my boss are usually important" → equals(email.sender, getContactEmail("boss"))

Learning a classifier from Natural Language Descriptions *(under the hood)*

1. Parse NL concept descriptions to feature functions

"Emails from my boss are usually important" → equals(email.sender, getContactEmail("boss"))

2. Learn weights over those features

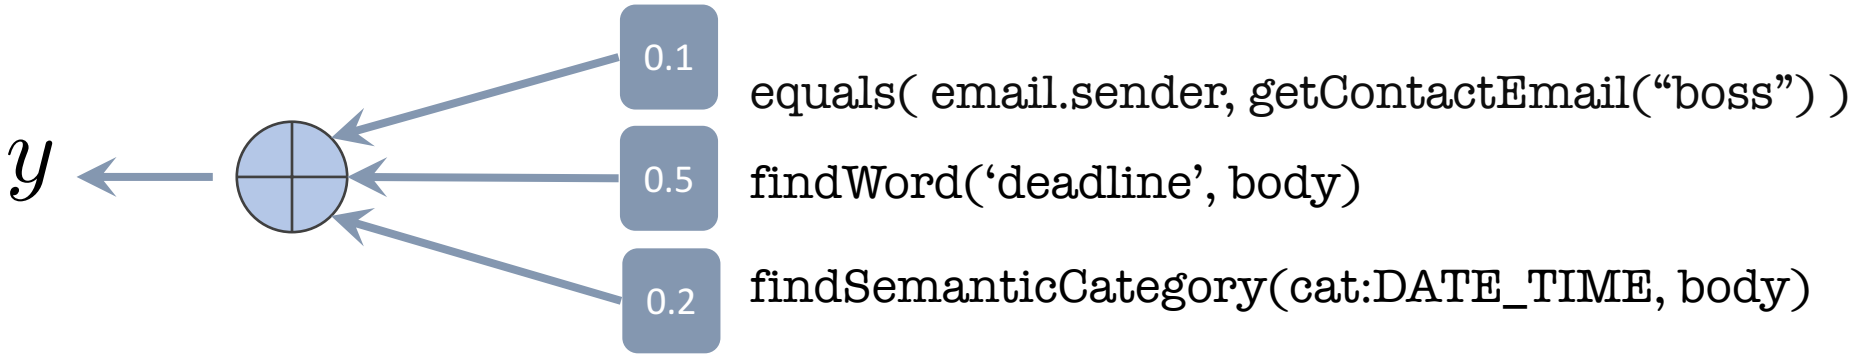

Learning a classifier from Natural Language Descriptions *(under the hood)*

1. Parse NL con… …ctions

"Emails from … usually impor…

…(email.sender, getContactEmail("boss"))

Where to get training data for **Semantic Parsing?**

2. Learn weights over those features

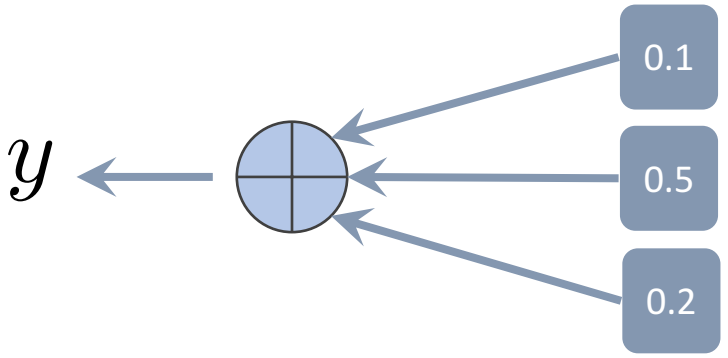

equals(email.sender, getContactEmail("boss"))

findWord('deadline', body)

findSemanticCategory(cat:DATE_TIME, body)

Learning a classifier from Natural Language Descriptions *(under the hood)*

1. Parse NL con… …ctions

"Emails from … usually impor…

…email.sender, getContactEmail("boss"))

Where to get training data for **Semantic Parsing?**

2. Learn weights over t…

…sender, getContactEmail("boss"))

findWord('deadline', body)

findSemanticCategory(cat:DATE_TIME, body)

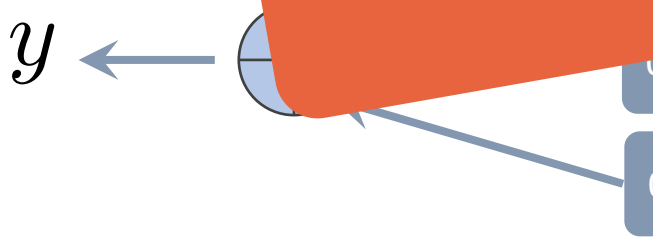

Where to get training data for **Binary Classification?**

Learning a classifier from Natural Language Descriptions *(under the hood)*

1. Parse NL concept descriptions to feature functions

"Emails from my boss are usually important" → equals(email.sender, getContactEmail("boss"))

2. Learn weights over those features

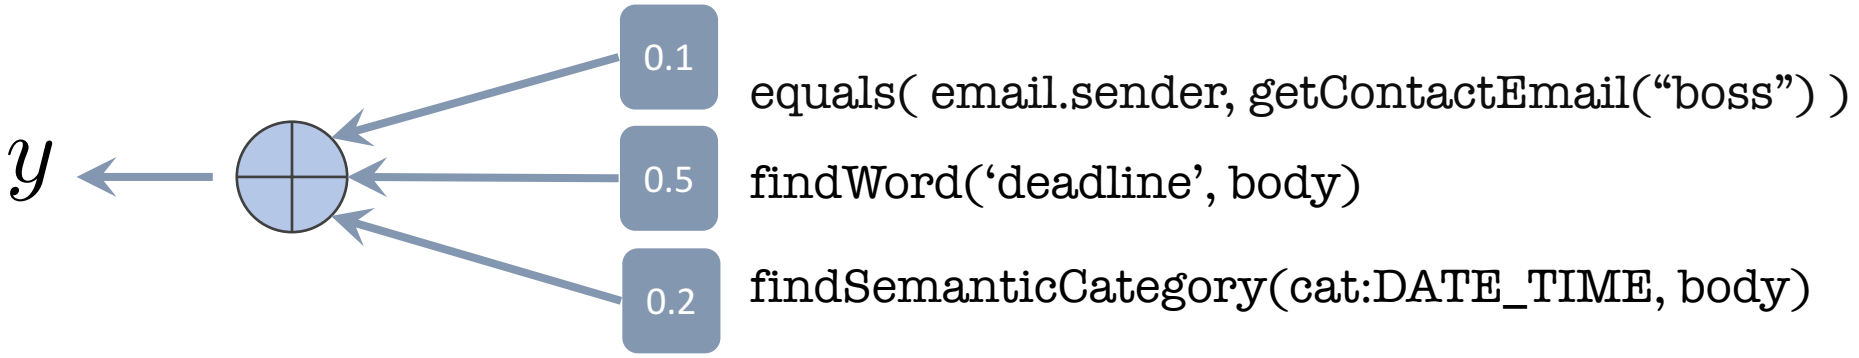

Learning a classifier from Natural Language Descriptions *(under the hood)*

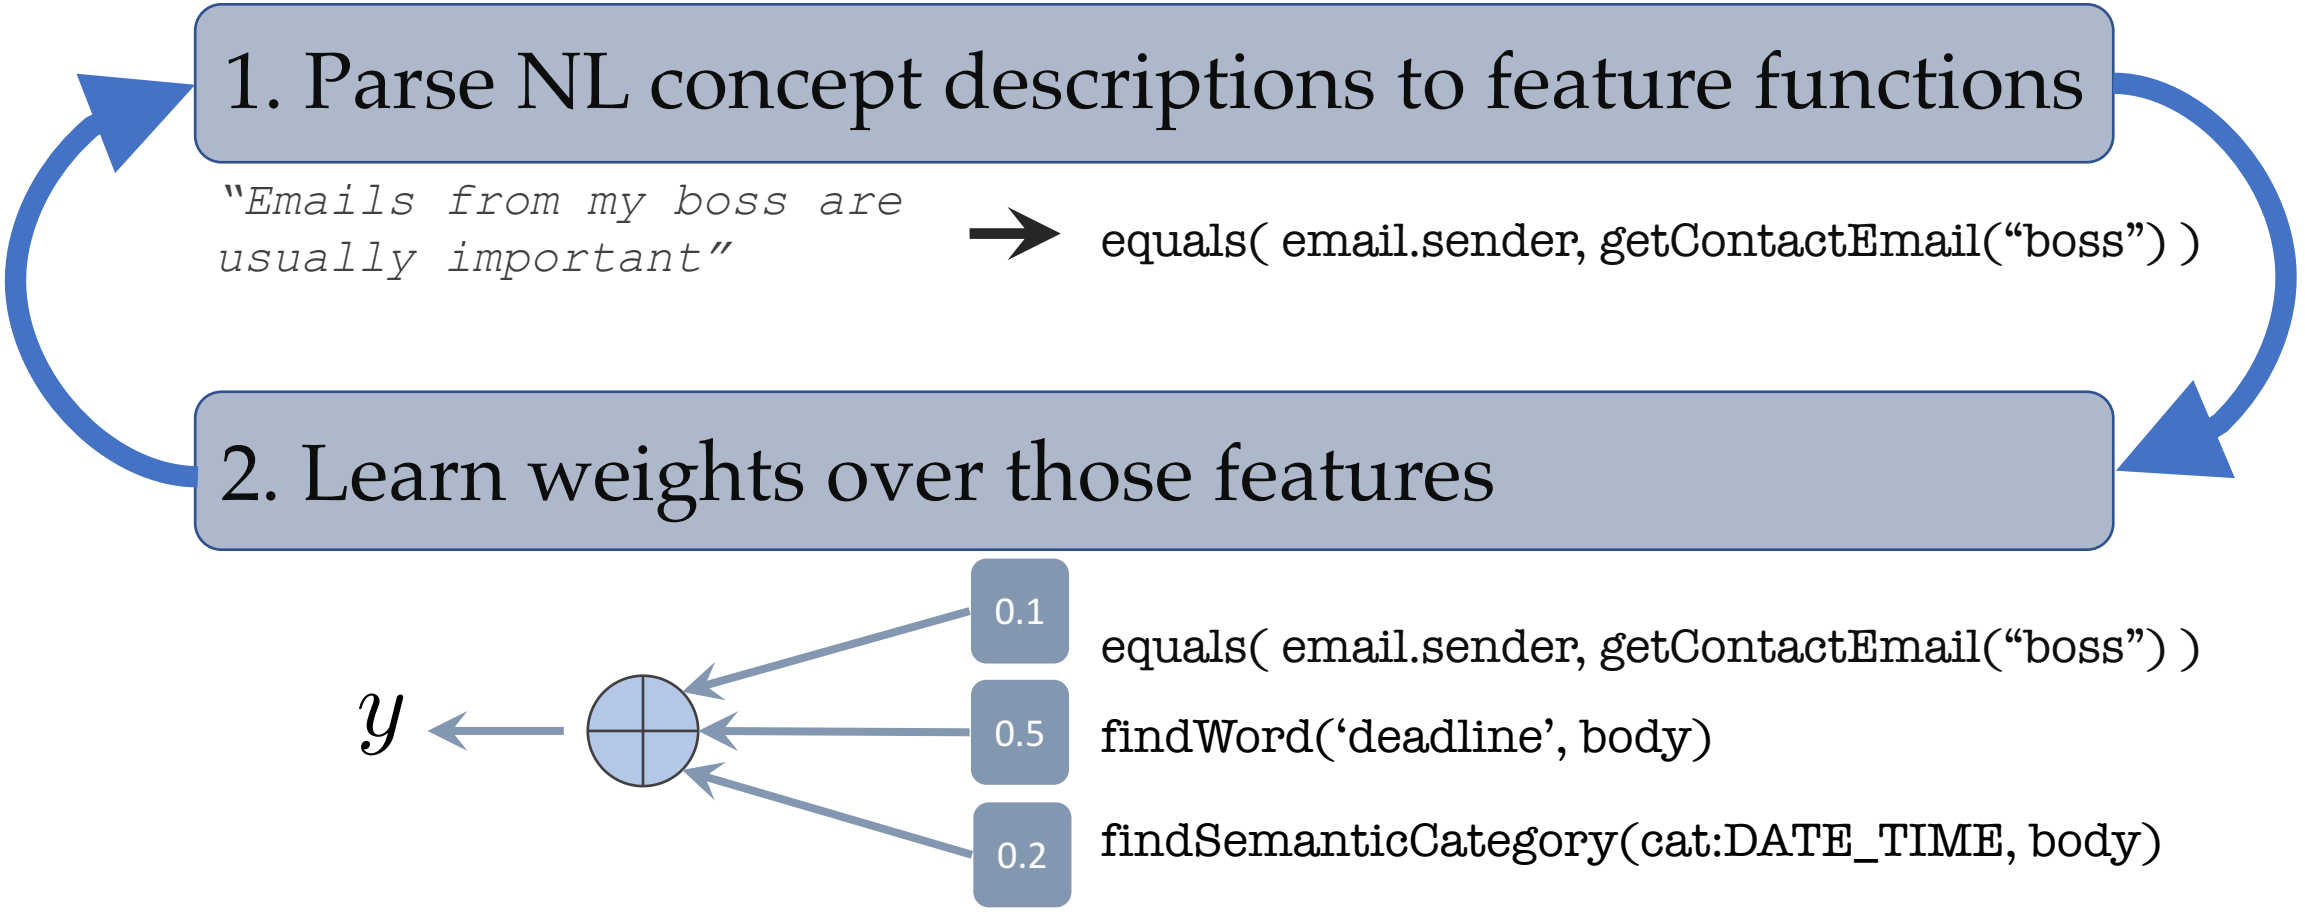

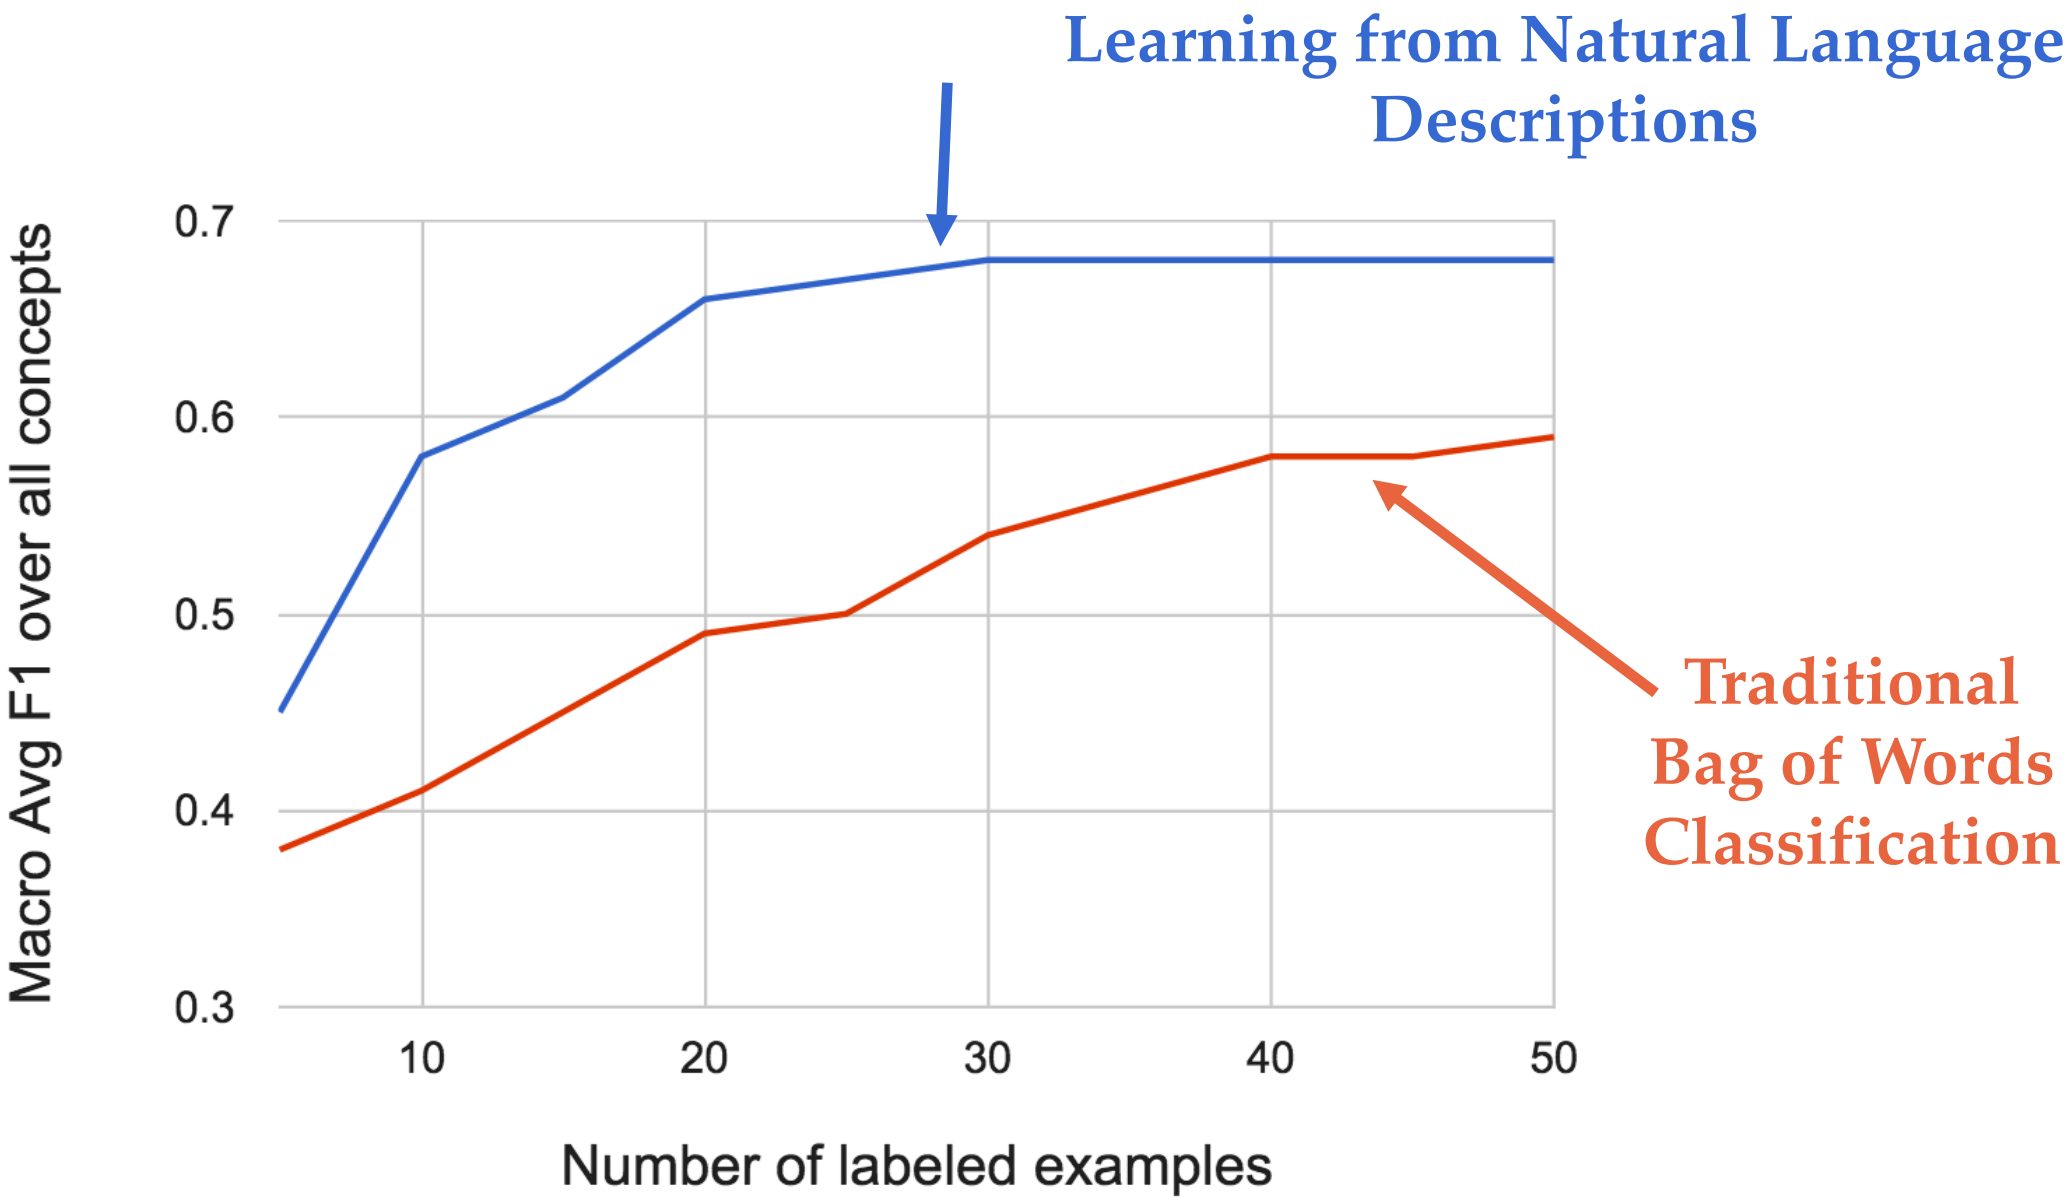
Learning from Natural Language Descriptions
Traditional Bag of Words Classification
Macro Avg F1 over all concepts
0.7
0.6
0.5
0.4
0.3
10
20
30
40
50
Number of labeled examples

Results: Semantic Parsing

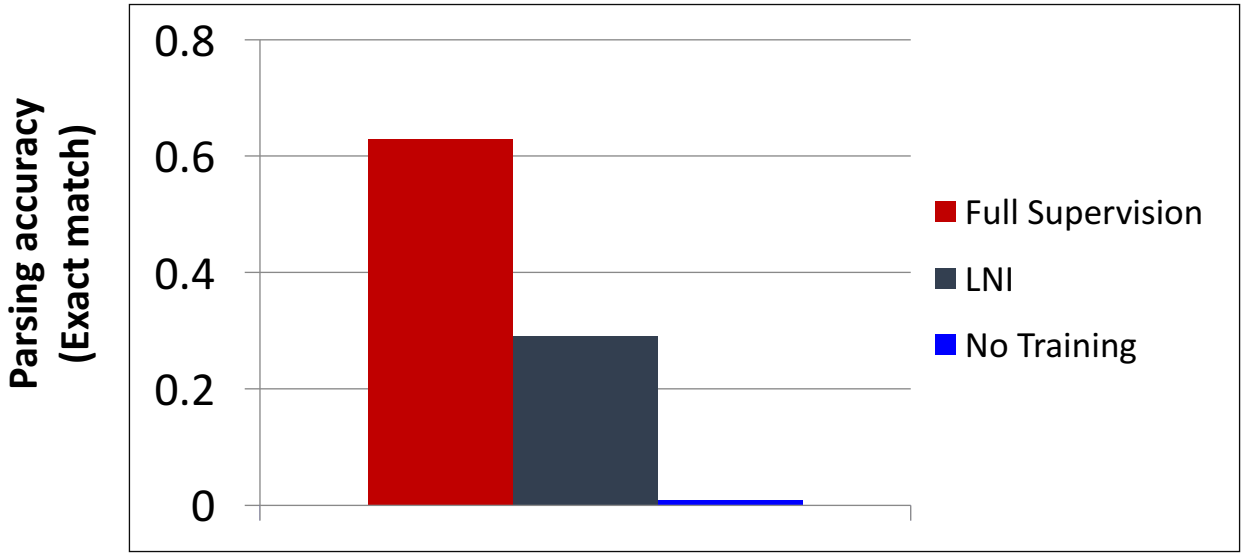

- Baseline (red): traditional supervised model trained on statements paired with logical forms

Predicted logical forms are often highly correlated

getPhraseMention(email, stringVal('meeting'))

getPhraseMention(body, stringVal('meeting'))

Hancock Braden, Varma, P., Wamg, S., Bringmann, M., Liang, P. & Re, C. (2018) Training classifiers with natural language explanations, ACL 2018

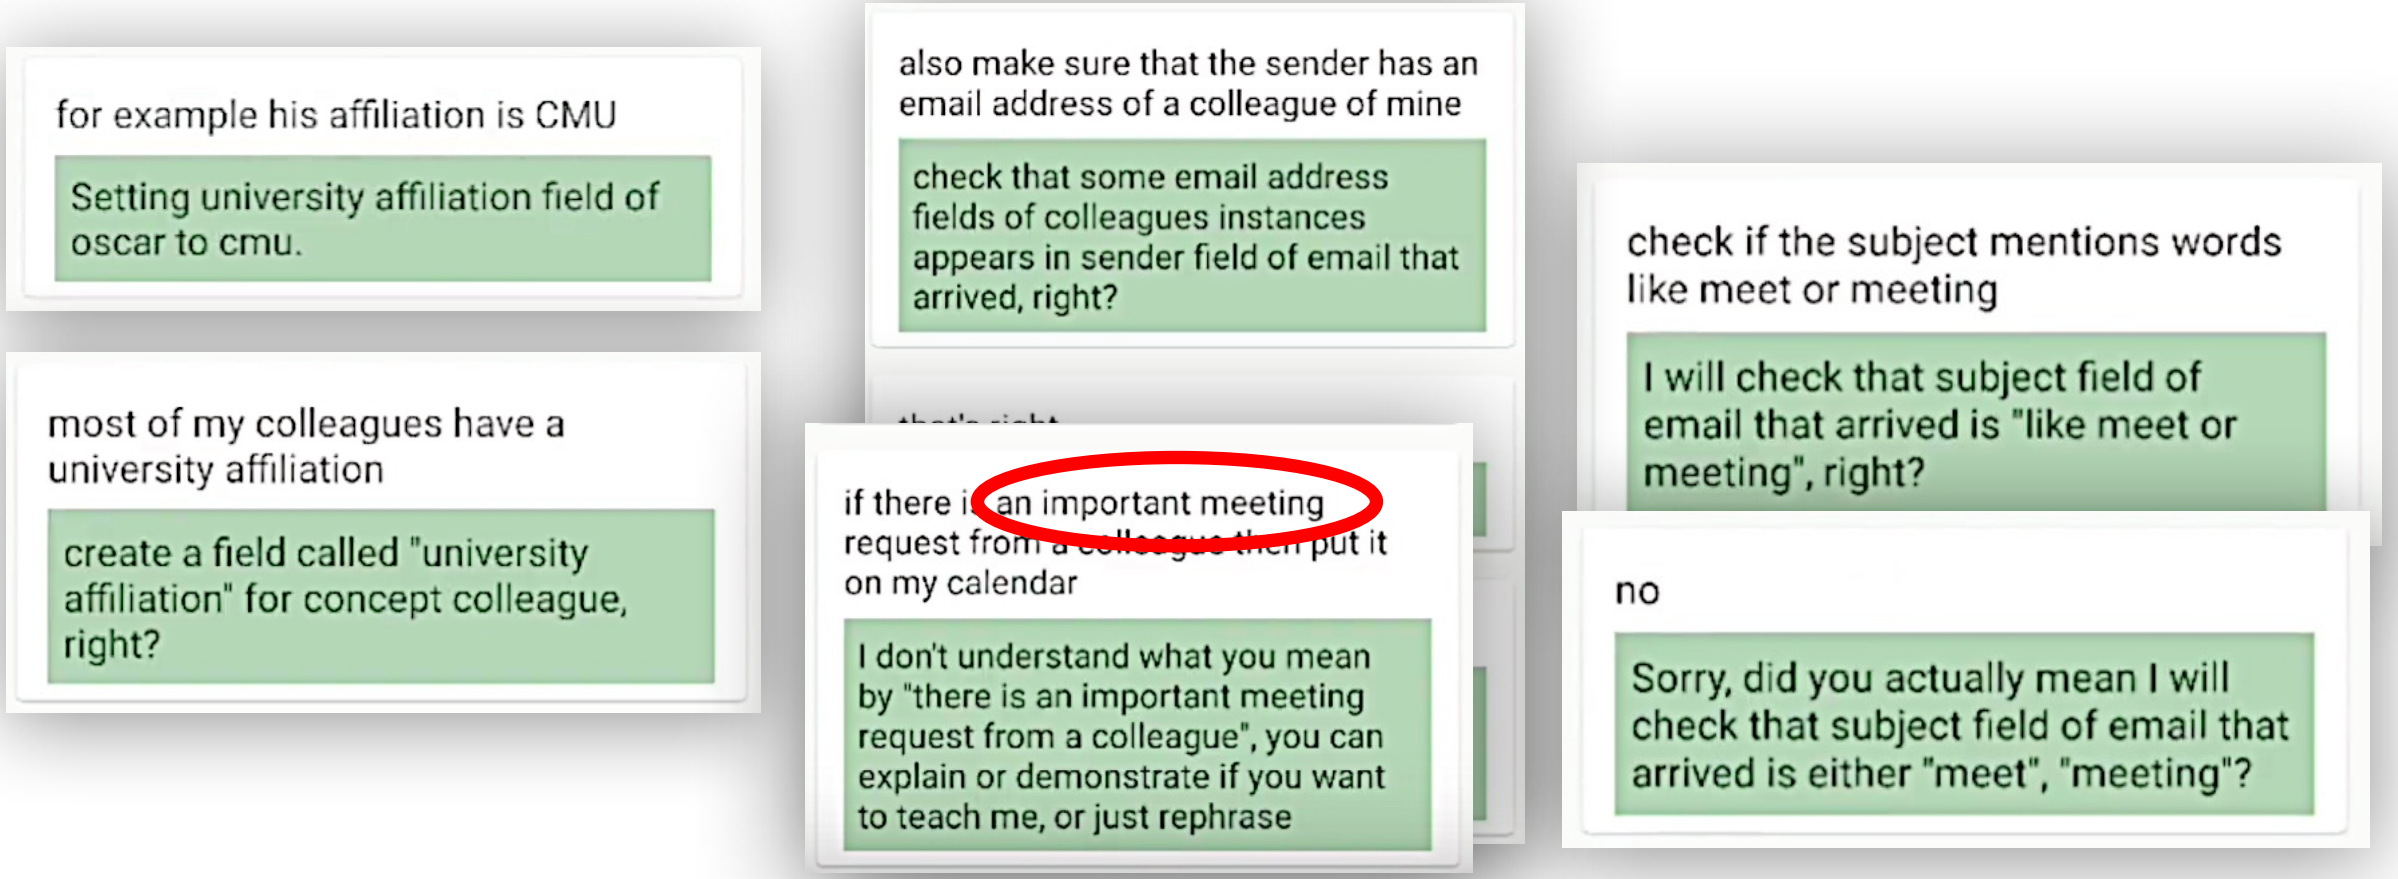

Knowledge-Base Construction

Concept Learning

Semantic Parsing

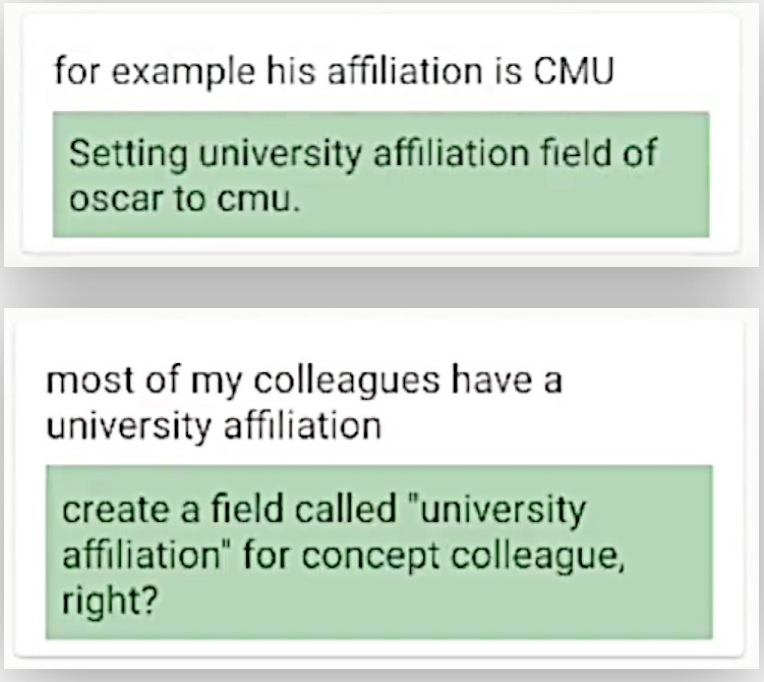

Knowledge-Base Construction

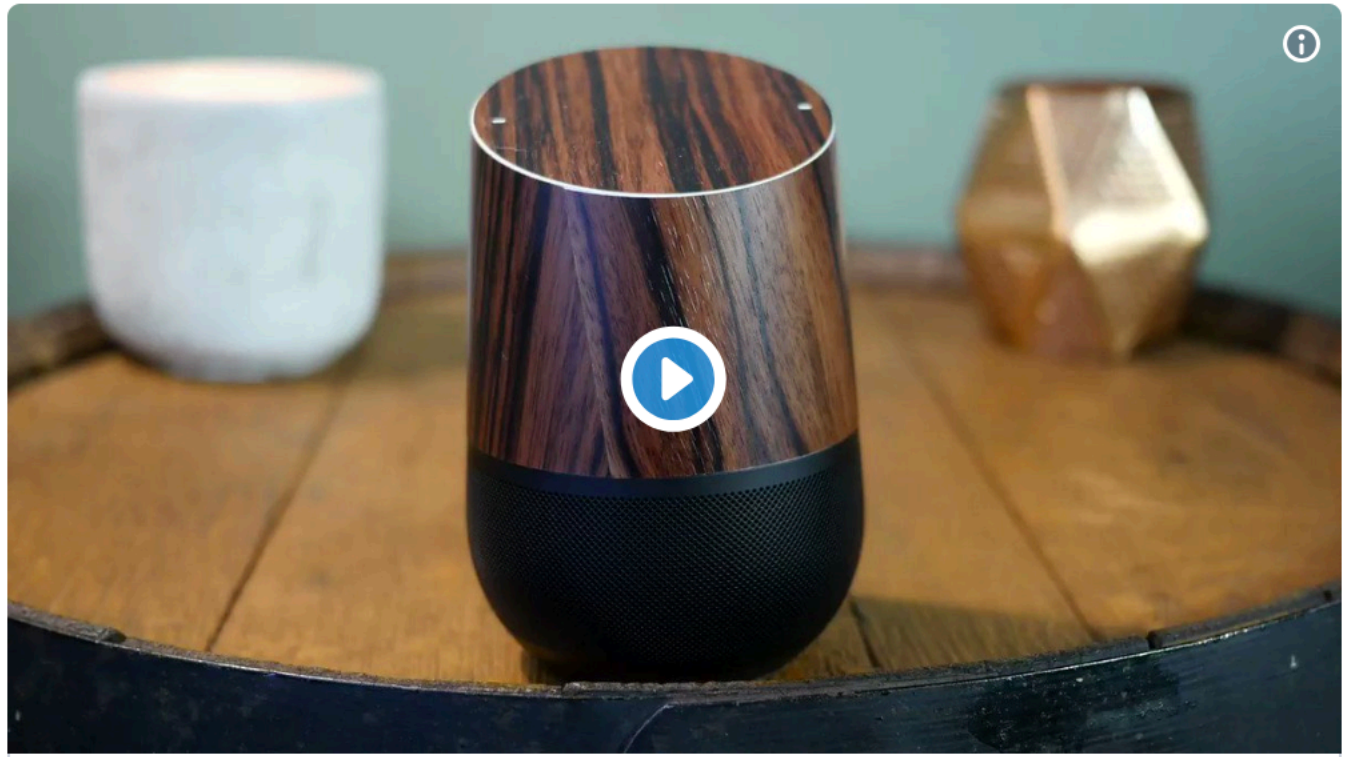

Taylor Martin
@caspertek

This is probably my new favorite Google Home feature.

5:05 PM - Jan 25, 2017

68 32 people are talking about this

"OK Google, remember that I put my passport in the bottom drawer of my filing cabinet."

"OK Google, remember that my spare key to the truck is hidden outside."

"OK Google, remember that Alex has my wallet."

"OK Google, remember that my password is money."

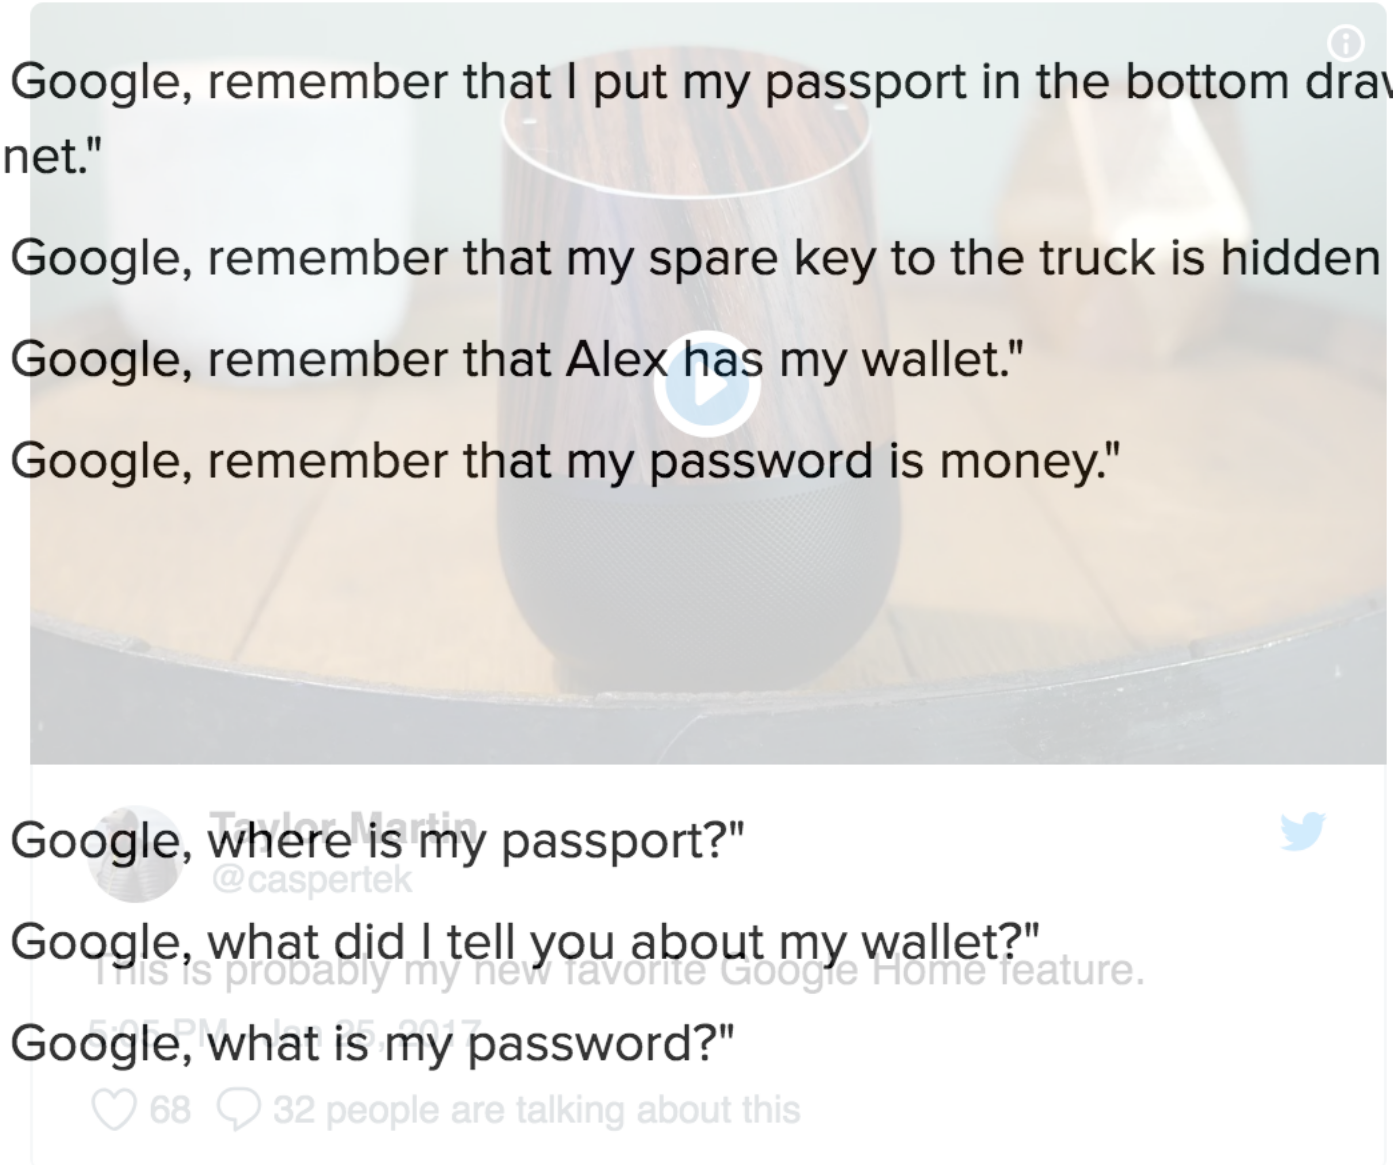

"OK Google, where is my passport?"

"OK Google, what did I tell you about my wallet?"

"OK Google, what is my password?"

Software Engineering World

1. There is a new important mobile project
2. That project is in the implementation stage
3. Hiram is a tester on mobile project
4. Mobile project has moved to the deployment stage
5. Andrew created a new issue for mobile project: fails with apache stack
6. Andrew is no longer assigned to that project
7. That developer resolved the changelog needs to be added issue

...

Software Engineering World

1. There is a new important mobile project
2. That project is in the implementation stage
3. Hiram is a tester on mobile project
4. Mobile project has moved to the deployment stage
5. Andrew created a new issue for mobile project: fails with apache stack
6. Andrew is no longer assigned to that project
7. That developer resolved the changelog needs to be added issue

...

ACL 2018

Multi-relational Question Answering from Narratives: Machine Reading and Reasoning in Simulated Worlds

Igor Labutov, Bishan Yang, Anusha Prakash, Amos Azaria

Software Engineering World

1. There is a new important mobile project
2. That project is in the implementation stage
3. Hiram is a tester on mobile project
4. Mobile project has moved to the deployment stage
5. Andrew created a new issue for mobile project: fails with apache stack
6. Andrew is no longer assigned to that project
7. That developer resolved the changelog needs to be added issue

...

Are there any developers assigned to projects in the evaluation stage?

Software Engineering World

1. There is a new important mobile project
2. That project is in the implementation stage
3. Hiram is a tester on mobile project
4. Mobile project has moved to the deployment stage
5. Andrew created a new issue for mobile project: fails with apache stack
6. Andrew is no longer assigned to that project
7. That developer resolved the changelog needs to be added issue

...

Are there any developers assigned to projects in the evaluation stage?

[Tawnya, Charlott, Hiram]

Software Engineering World

1. There is a new important mobile project
2. That project is in the implementation stage
3. Hiram is a tester on mobile project
4. Mobile project has moved to the deployment stage
5. Andrew created a new issue for mobile project: fails with apache stack
6. Andrew is no longer assigned to that project
7. That developer resolved the changelog needs to be added issue

...

Are there any developers assigned to projects in the evaluation stage?
[Tawnya, Charlott, Hiram]

Who is the null pointer exception during parsing issue assigned to?
Hiram

Software Engineering World

1. There is a new important mobile project
2. That project is in the implementation stage
3. Hiram is a tester on mobile project
4. Mobile project has moved to the deployment stage
5. Andrew created a new issue for mobile project: fails with apache stack
6. Andrew is no longer assigned to that project
7. That developer resolved the changelog needs to be added issue

...

Are there any developers assigned to projects in the evaluation stage?
[Tawnya, Charlott, Hiram]

Who is the null pointer exception during parsing issue assigned to?
Hiram

Are there any issues that are resolved for experimental projects?
[saving data throws exception,
wrong pos tag on consecutive words]

User teaches
new knowledge

She and Amy are both TAs for this course

John is now the department head

Which phd students are advised by the department head?

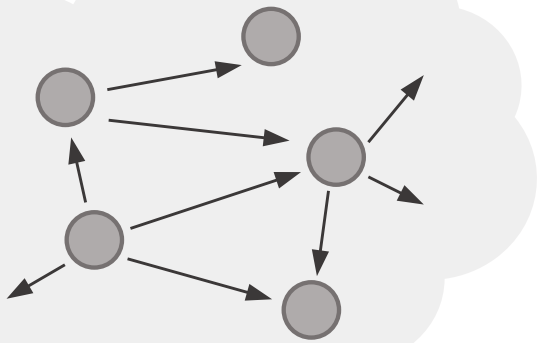

User queries
taught
knowledge

Simulated Worlds

Academic Faculty World

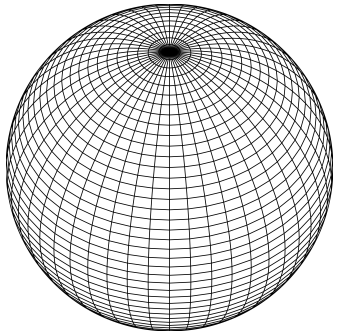

Software Project World

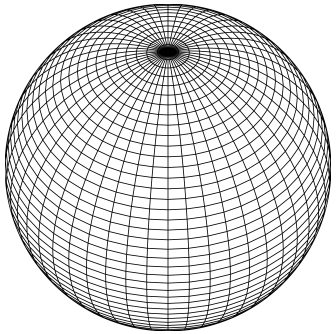

Student World

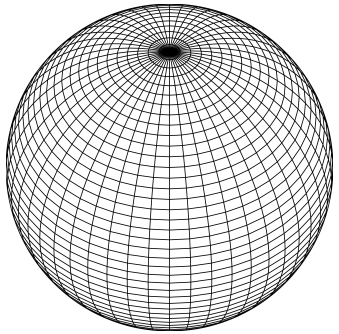

Shopping World

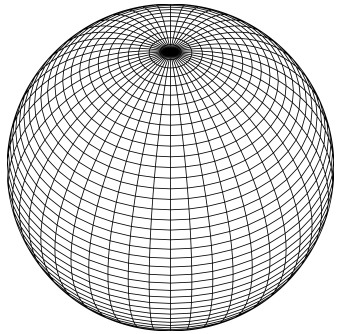

Meeting Scheduling World

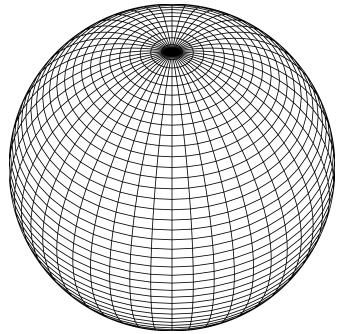

Software Engineering World

1. There is a new important mobile project
2. That project is in the implementation stage
3. Hiram is a tester on mobile project
4. Mobile project has moved to the deployment stage
5. Andrew created a new issue for mobile project: fails with apache stack
6. Andrew is no longer assigned to that project
7. That developer resolved the changelog needs to be added issue

...

Are there any developers assigned to projects in the evaluation stage?
[Tawnya, Charlott, Hiram]

Who is the null pointer exception during parsing issue assigned to?
Hiram

Are there any issues that are resolved for experimental projects?
[saving data throws exception,
wrong pos tag on consecutive words]

Software Engineering World

1. There is a new important mobile project
2. That project is in the implementation stage
3. Hiram is a tester on mobile project
4. Mobile project has moved to the deployment stage
5. Andrew created a new issue for mobile project: fails with apache stack
6. Andrew is no longer assigned to that project

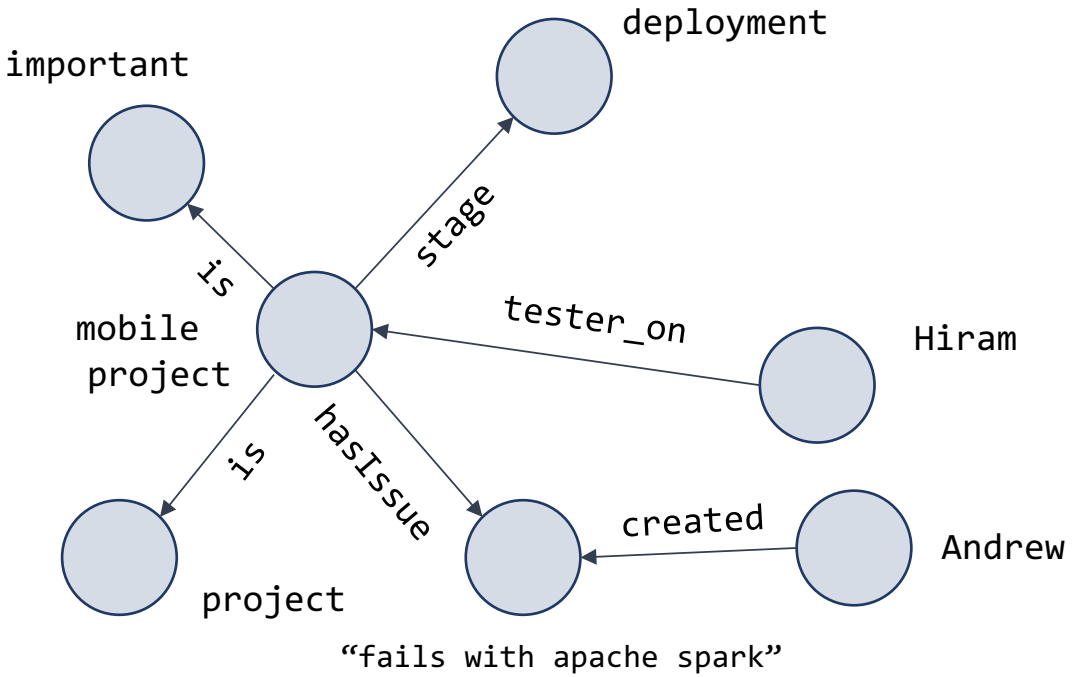

Are there any developers assigned to projects in the evaluation stage?
[Tawnya, Charlott, Hiram]

Who is the null pointer exception during parsing issue assigned to?
Hiram

Are there any issues that are resolved for experimental projects?
[saving data throws exception,
wrong pos tag on consecutive words]

Types of Questions

```
Set John's email address to john.adams@gmail.com
```

Questions

What is John's email?

Types of Questions

John and Mike just enrolled in my ML class

Andy just dropped this course

Questions

Who is taking the ML class?

Types of Questions

John and Mike just enrolled in my ML class

Andy just dropped this course

Questions

Who is taking the ML class?

$$\lambda x.\texttt{TakingClass}(x, \texttt{ML})$$

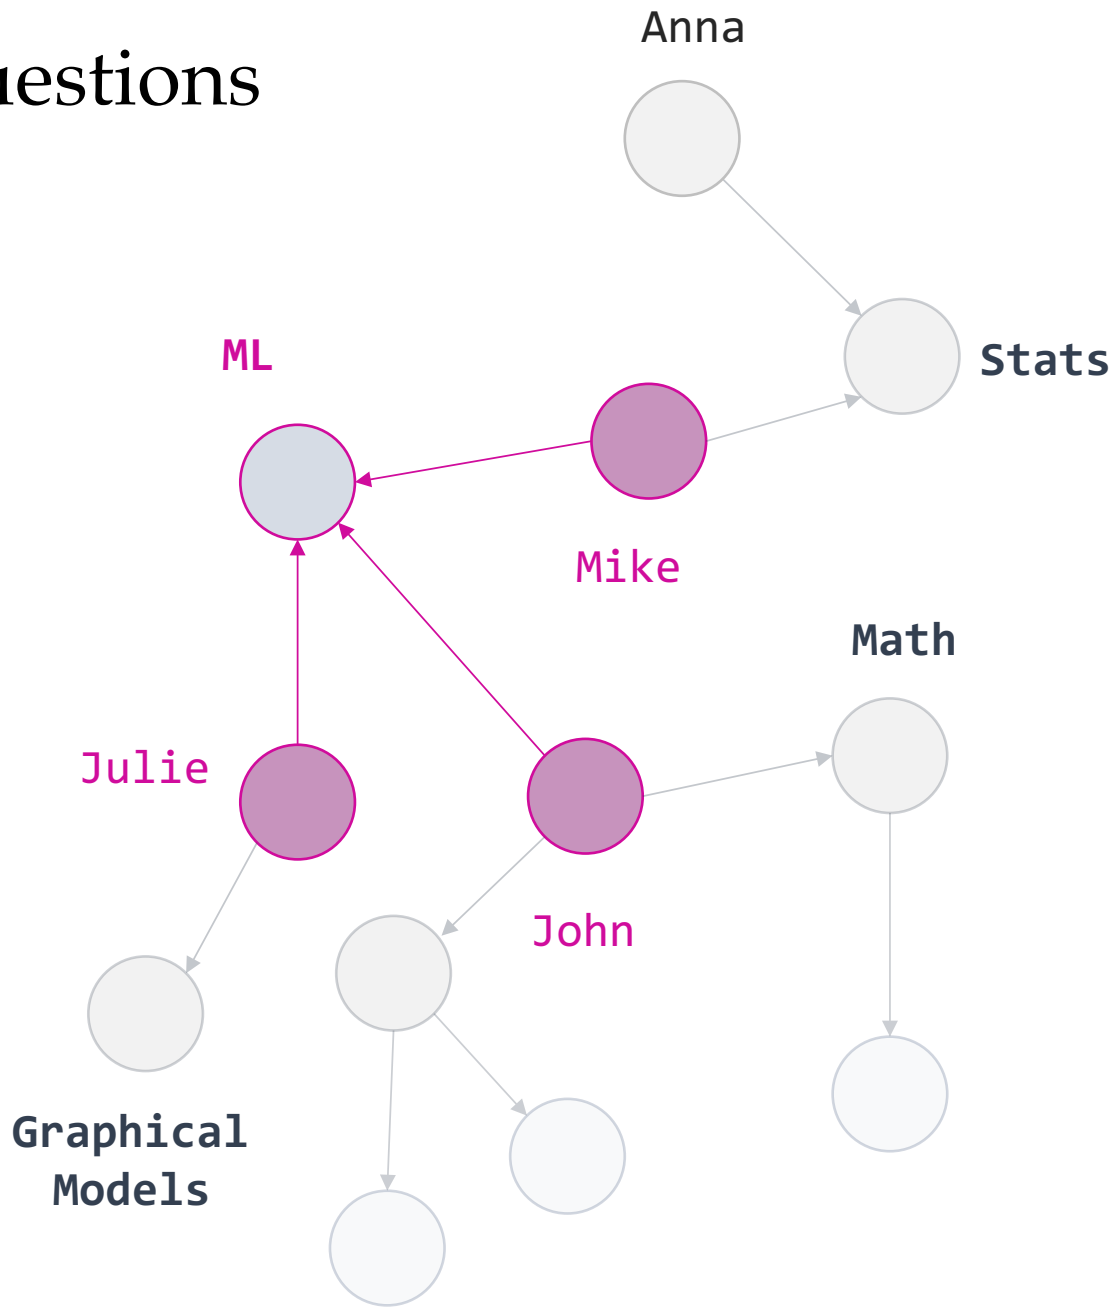

Types of Questions

I am teaching both ML and stats this semester

John and Mike just enrolled in my ML class

Andy just dropped this course

Questions

Which students are enrolled in one of the classes I am teaching?

$$\lambda x.\exists y.\texttt{TakingClass}(x, y) \wedge \texttt{CourseTaughtByMe}(y)$$

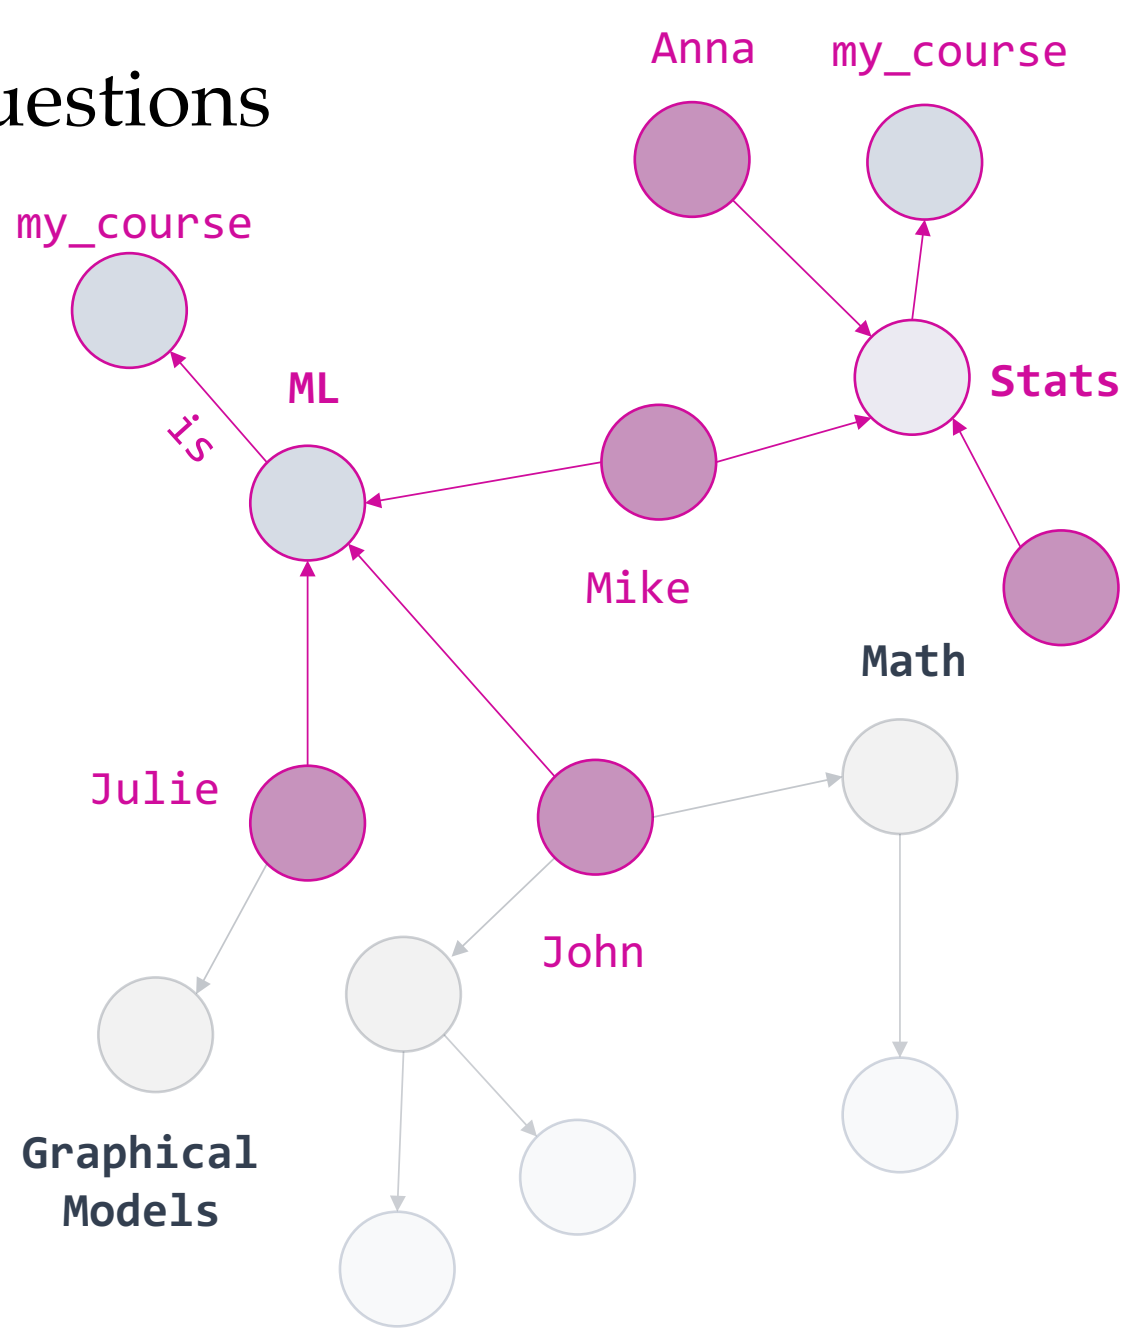

Types of Questions

I am teaching both ML and stats this semester

John and Mike just enrolled in my ML class

They are both undergrads in the CS department

Andy just dropped this course

Questions

Which undergrads are enrolled in one of the classes I am teaching?

$$\lambda x.\exists y.\texttt{TakingClass}(x, y) \\ \wedge \texttt{CourseTaughtByMe}(y) \\ \wedge \texttt{Undergraduate}(x)$$

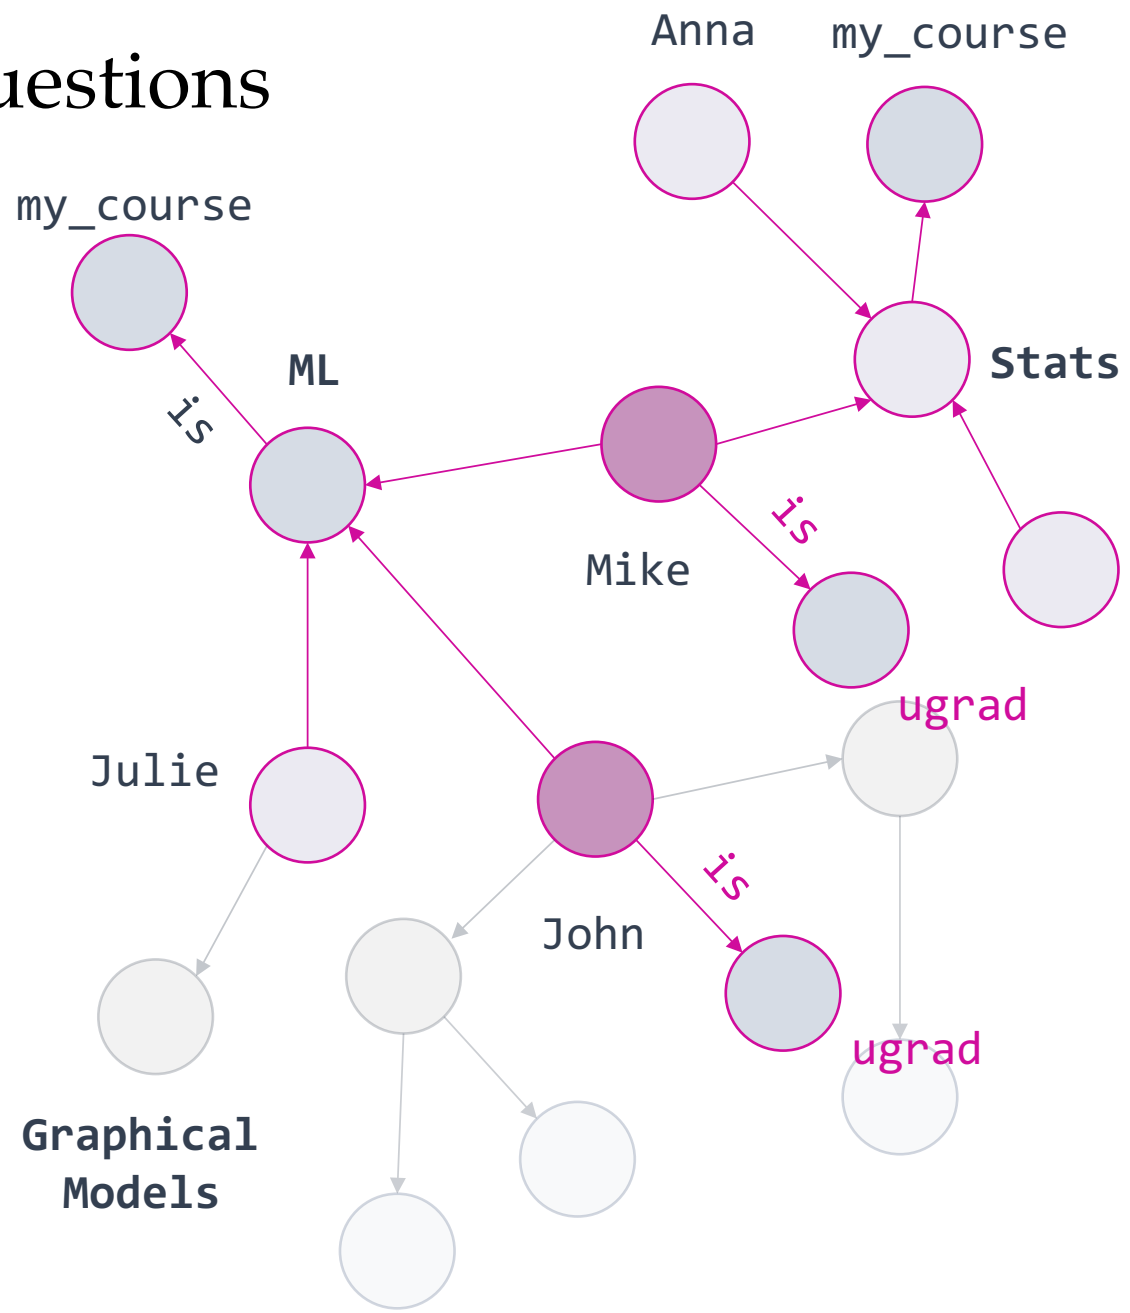

$\lambda x.\texttt{EmailAddress}(\texttt{John}, x)$	Single Entity / Single Relation
$\lambda x.\texttt{TakingClass}(x, \texttt{ML})$	Multi-Entity / Single Relation
$\lambda x.\exists y.\texttt{TakingClass}(x, y)$ $\wedge\ \texttt{CourseTaughtByMe}(y)$	Multi-Entity / Two Relations
$\lambda x.\exists y.\texttt{TakingClass}(x, y)$ $\wedge\ \texttt{CourseTaughtByMe}(y)$ $\wedge\ \texttt{Undergraduate}(x)$	Multi-Entity / Three Relations

Narrative

1. There is an associate professor named Andy
2. He returned from a sabbatical
3. This professor currently has funding
4. There is a masters level course called G301
5. That course is taught by him
6. That class is part of the mechanical engineering department
7. Roslyn is a student in this course
8. U203 is a undergraduate level course
9. Peggy and that student are TAs for this course

Question

"Who are the undergrads enrolled in one of the classes I am teaching?"

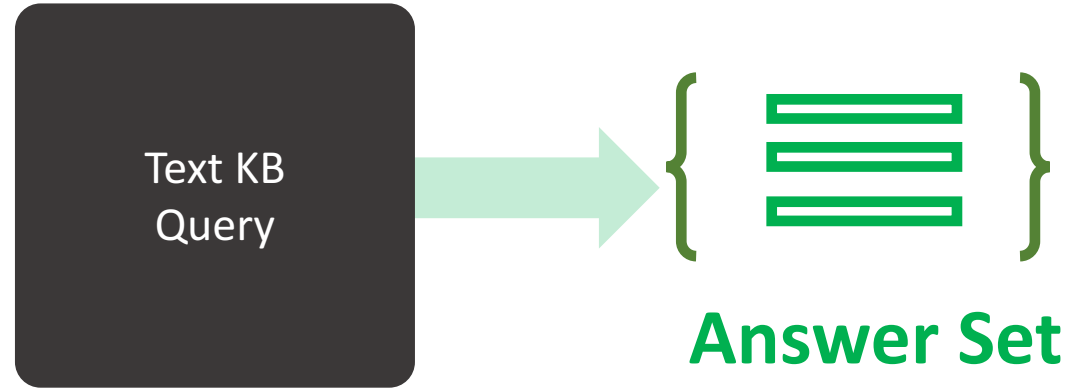

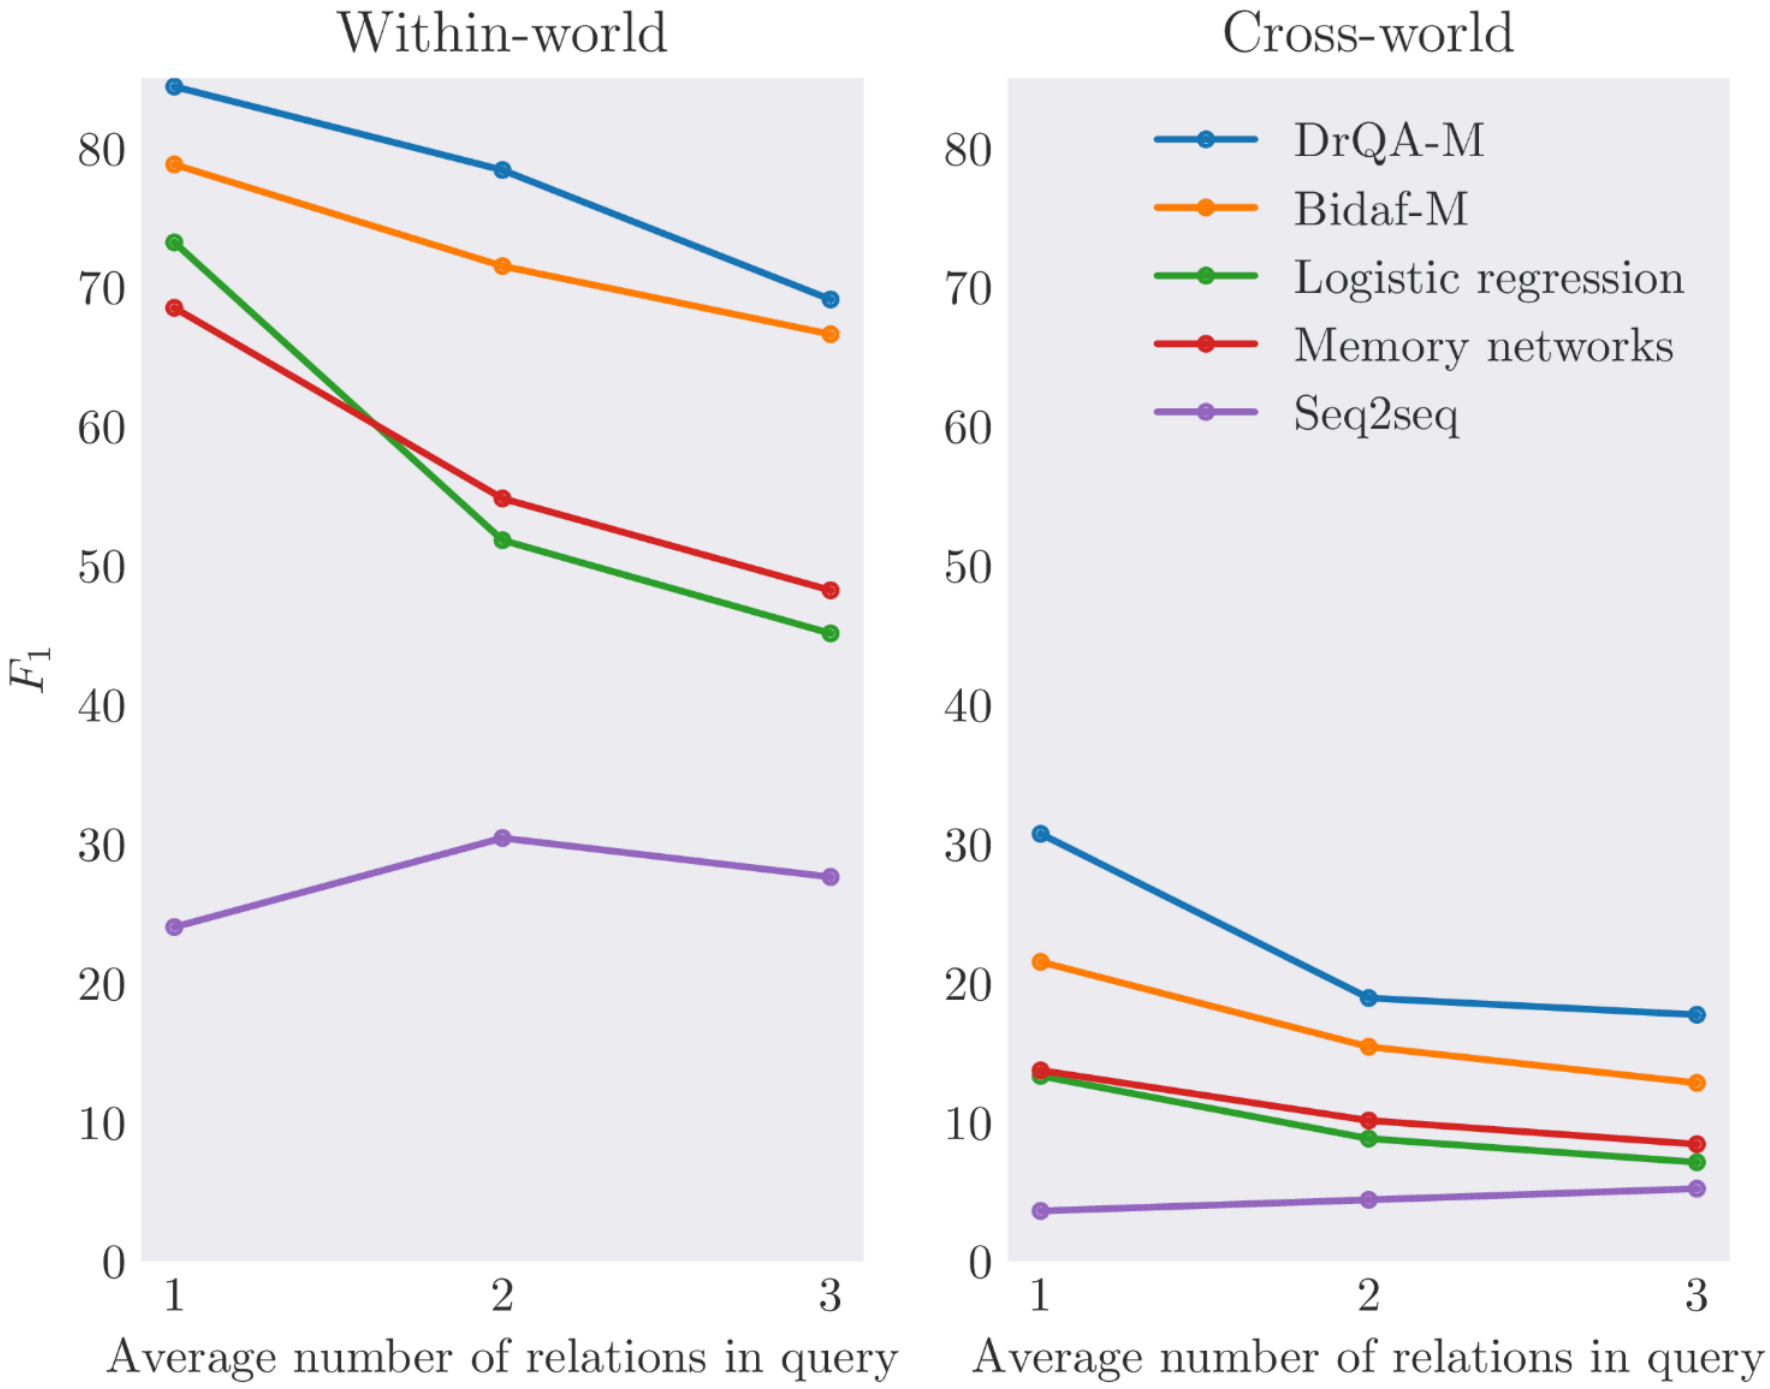
Within-world
Cross-world
DrQA-M
Bidaf-M
Logistic regression
Memory networks
Seq2seq
F_1
0
10
20
30
40
50
60
70
80
1
2
3
Average number of relations in query
Average number of relations in query

Supervised Learning

Learning from Labeled Examples

Supervised Learning

Learning from Labeled Examples

Supervised Learning

Active Learning

Learning from Labeled Examples

Supervised Learning

Labeling Features

Active Learning

Learning from Labeled Examples

Supervised Learning

Learning with constraints

Labeling Features

Active Learning

Learning from Labeled Examples

Supervised Learning

Active learning with features

Learning with constraints

Labeling Features

Active Learning

Learning from Labeled Examples

Supervised Learning

Problem structure

Active learning with features

Learning with constraints

Labeling Features

Active Learning

Learning from Labeled Examples

Supervised Learning

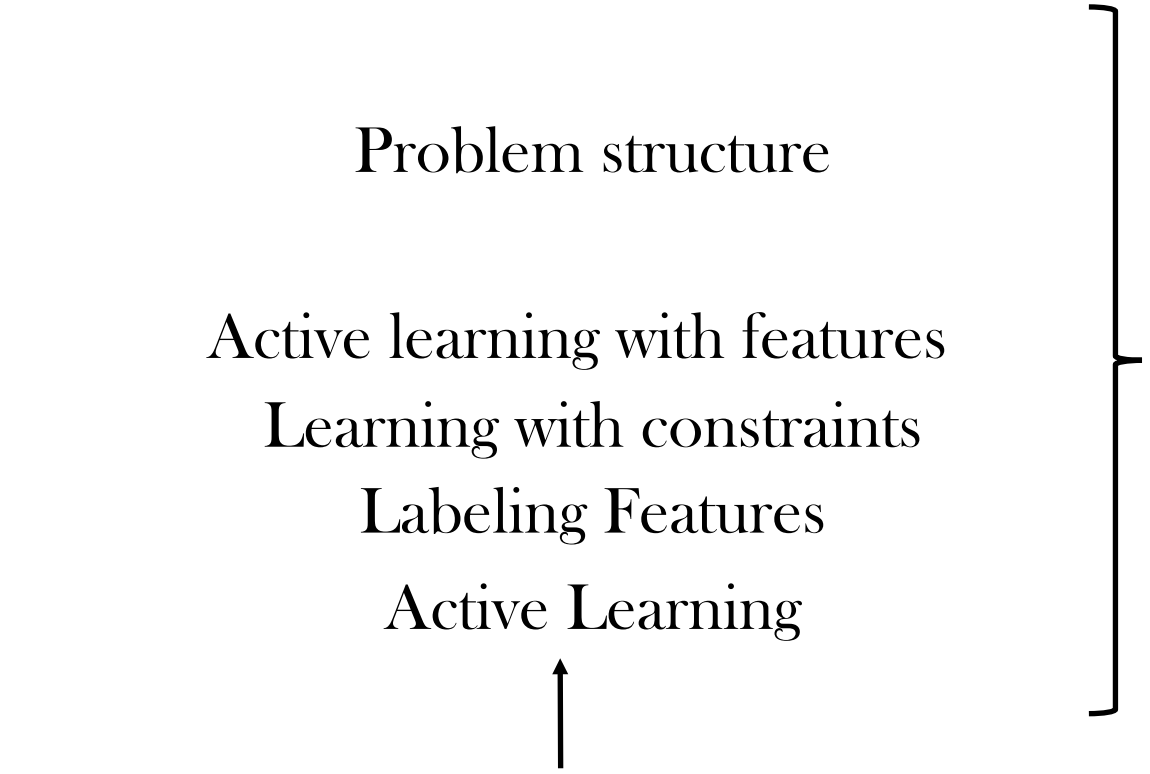
Problem structure
Active learning with features
Learning with constraints
Labeling Features
Active Learning
Learning from Labeled Examples
Supervised Learning
Language
as
interface

/lɛr/ *(imperative)*

1. To **Teach**
2. To **Learn**

LAER.AI

/lɛr/ *(imperative)*
1. To **Teach**
2. To **Learn**

Conversational assistant

for the workplace that can:

- Answer questions
- Solve problems
- **Learn new things**

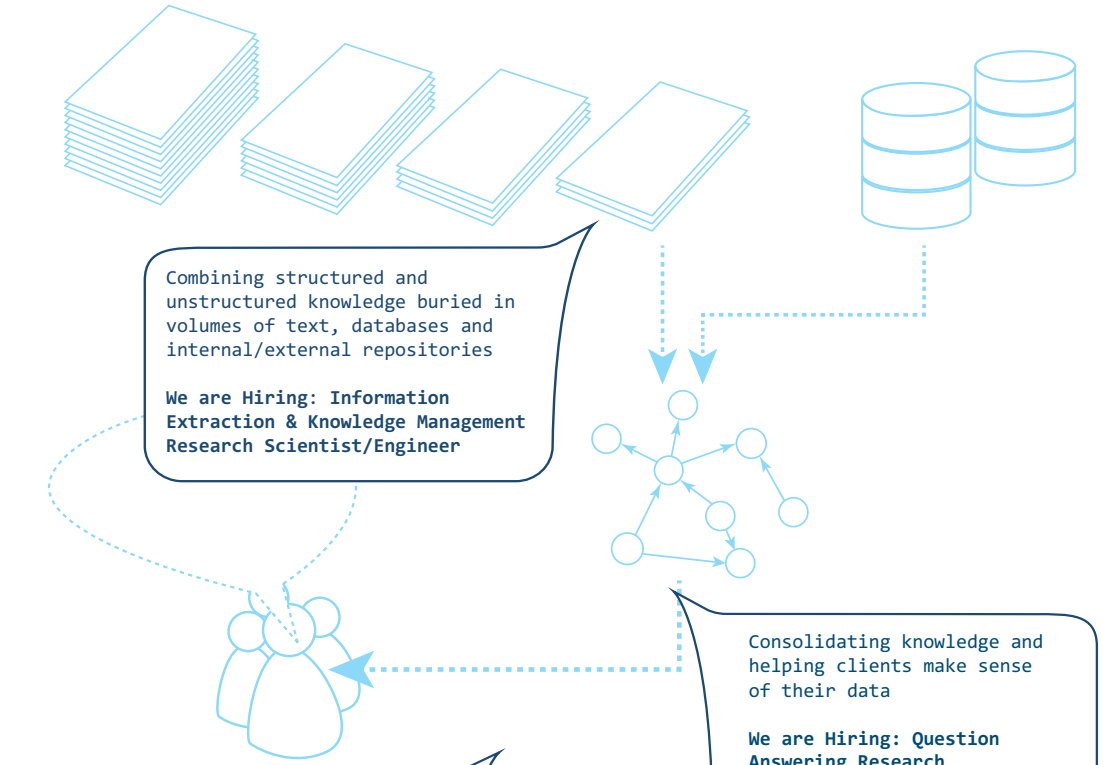

Join us in building the future of search

- jobs@laer.ai

Conversational assistant for the workplace that can:

- Answer questions
- Solve problems
- **Learn new things**

LAER.AI

/lɛr/ *(imperative)*
1. To **Teach**
2. To **Learn**

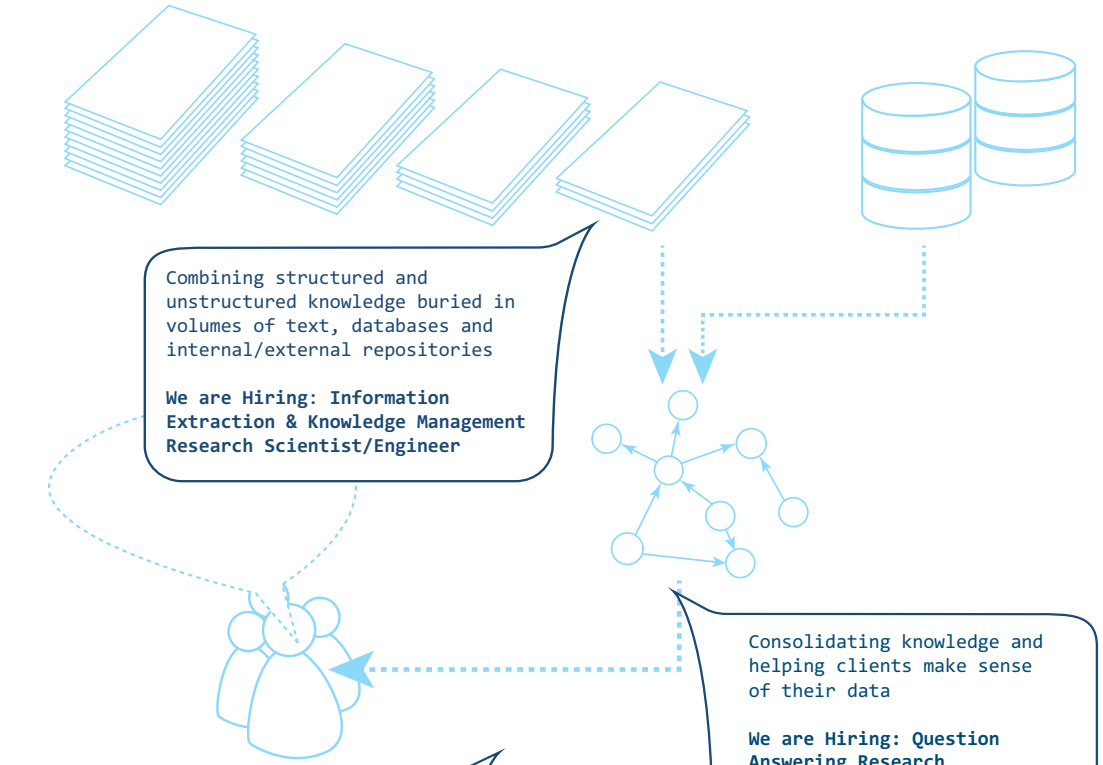

Join us in building the future of search

- jobs@laer.ai

Igor.labutov@laer.ai
Bishan.yang@laer.ai