

# Completeness, Recall and Negation in Open-World Knowledge Bases

Simon Razniewski, Hiba Arnaout, Shrestha Ghosh, Fabian Suchanek

1. Introduction and Foundations (Simon)
2. Predictive recall assessment (Fabian)
3. Counts from text and KB (Shrestha)
4. Negation (Hiba)
5. Relative completeness & Wrap-up (Simon)

# What is count information?

Relation between an entity and a set of entities



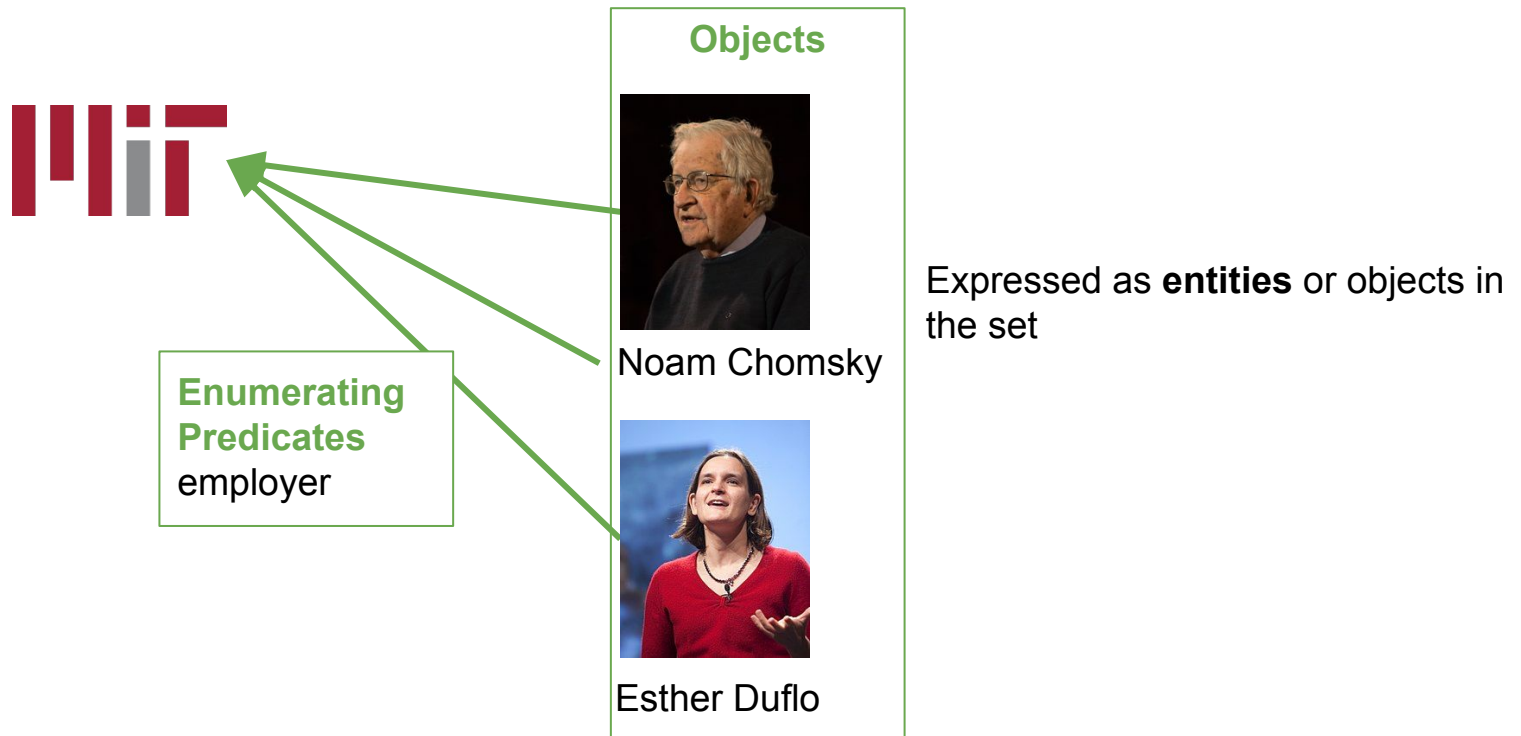
Noam Chomsky



Esther Duflo

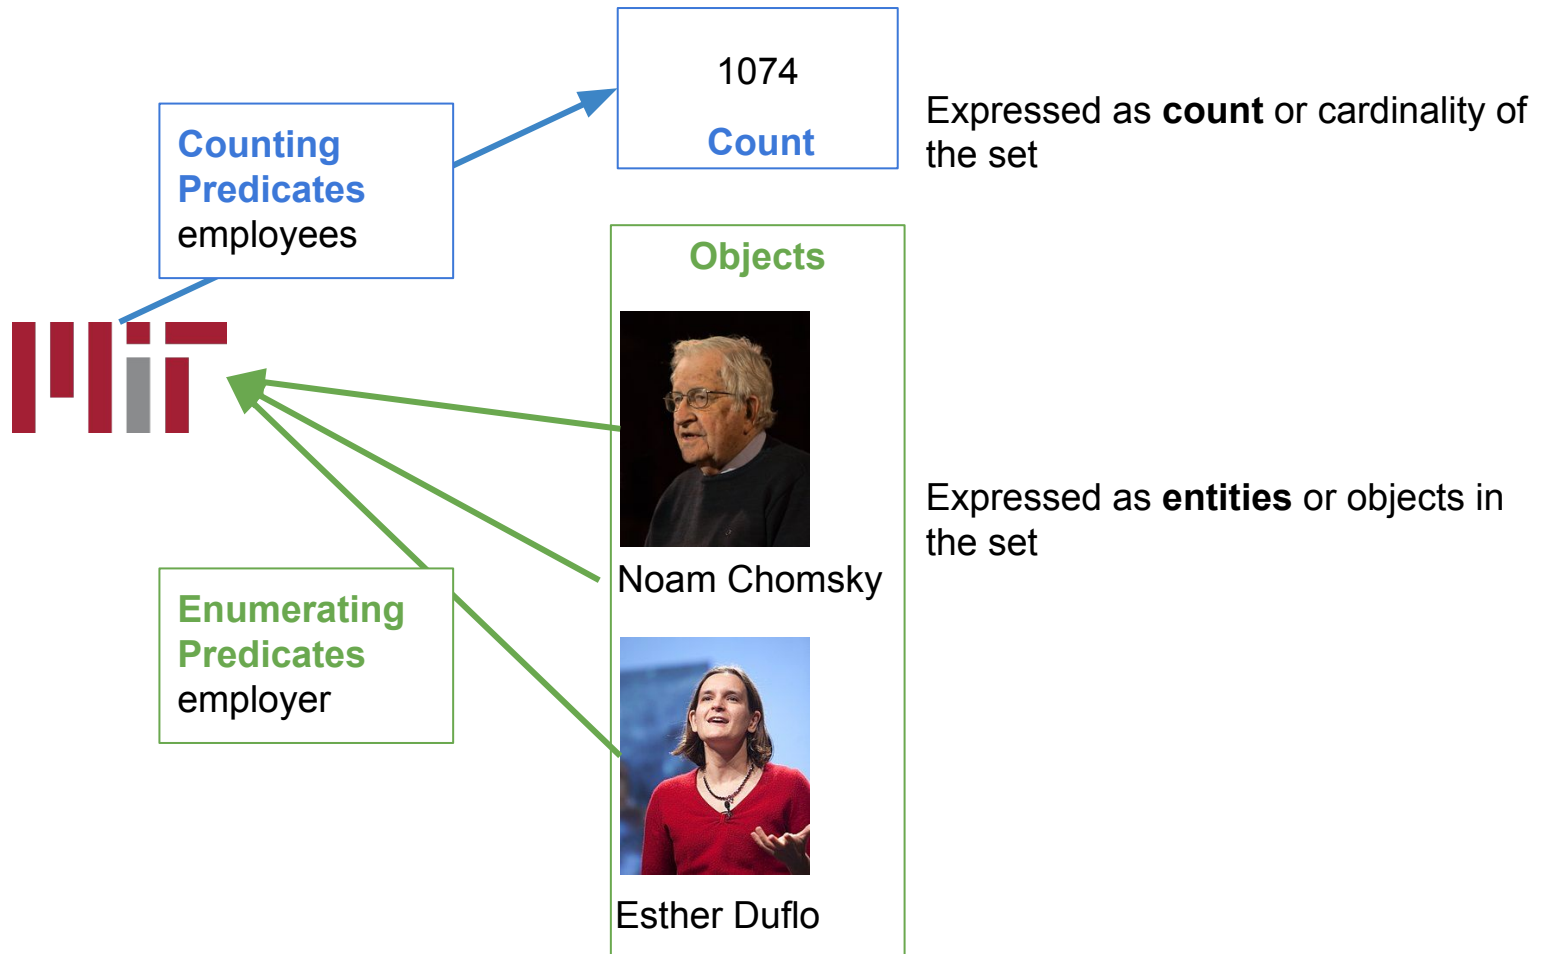
# What is count information?

Relation between an entity and a set of entities



# What is count information?

Relation between an entity and a set of entities



1. Count information for recall assessment
2. How can we extract count information from text?
3. Variants of count information in KB
4. How much count information is accounted for?
5. Counts for KB curation

# Count information for recall assessment

Counts and entities benefit from each other

Only entities

(?x, employer, MIT)

returns a handful of names from KB



Enumerating  
Predicates  
employer

Objects



Noam Chomsky



Esther Duflo

# Count information for recall assessment

Counts and entities benefit from each other

## Only entities

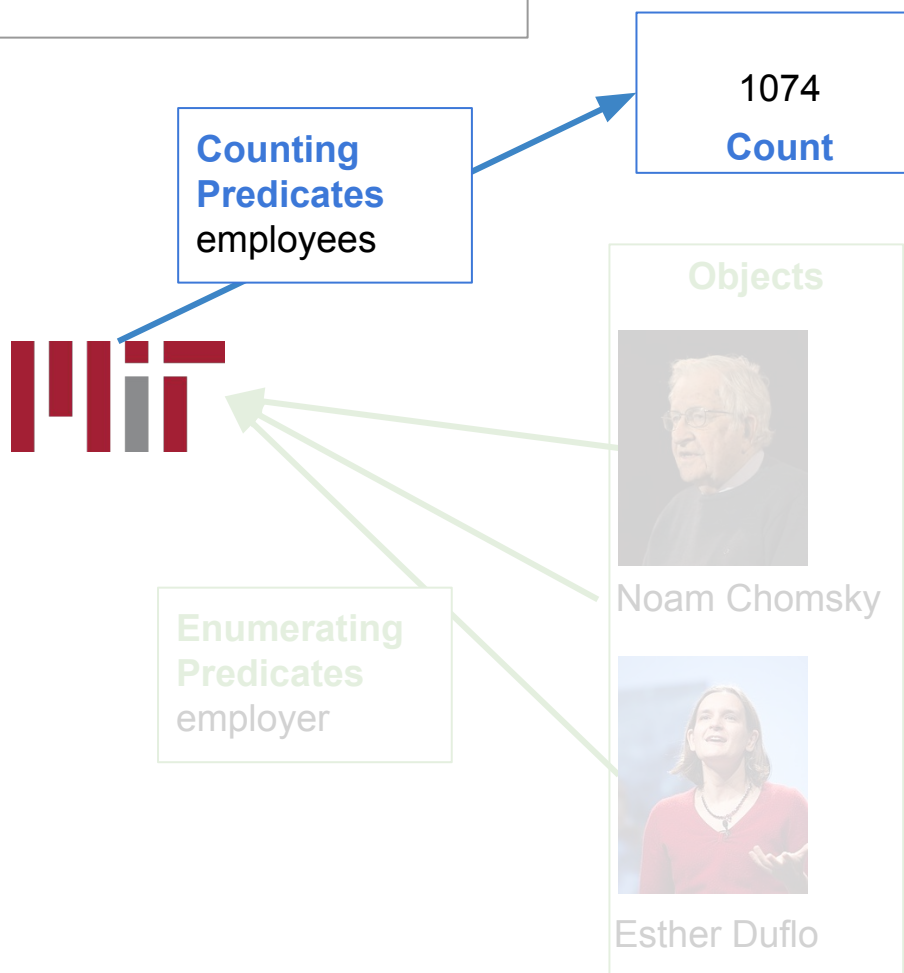
(?x, **employer**, MIT)

returns a handful of names from KB

## Only counts

(MIT, **employees**, ?y)

gives no insight about the entities



# Count information for recall assessment

## Only entities

(?x, **employer**, MIT)

returns a handful of names from KB

## Only counts

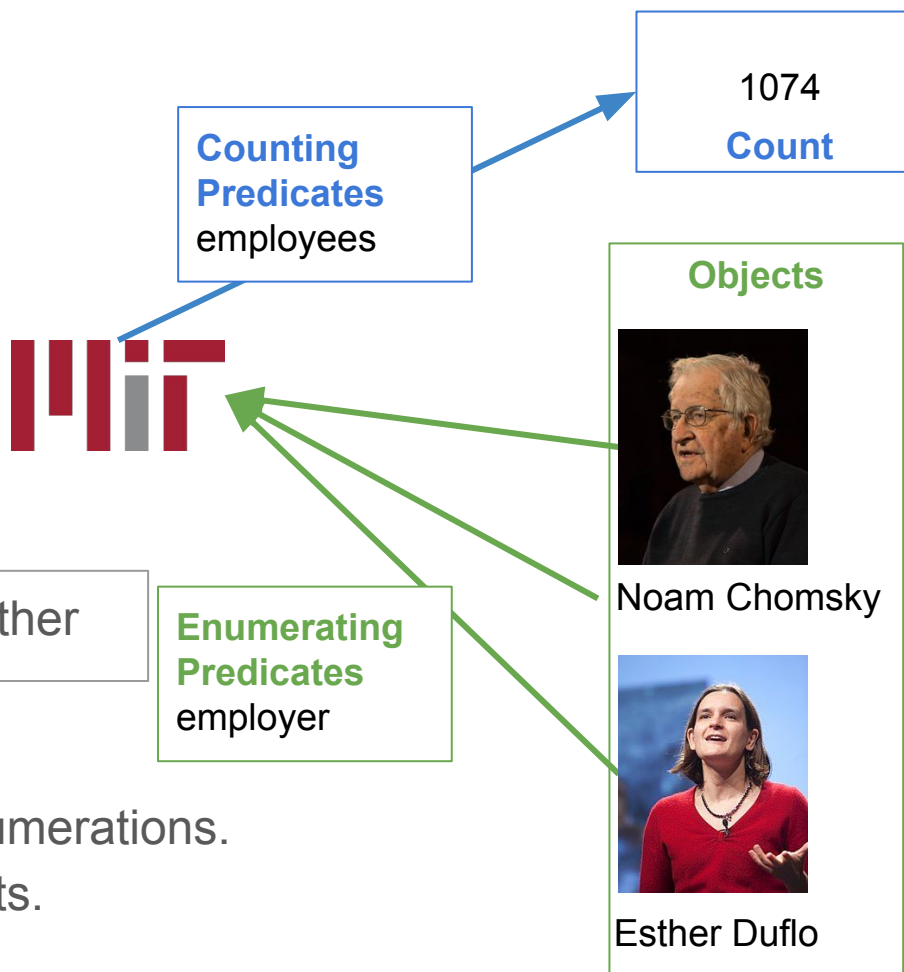
(MIT, **employees**, ?y)

gives no insight about the entities

Counts and entities benefit from each other

## Count and Entities

- Counts enhance incomplete entity enumerations.
- Representative entities enhance counts.





# Count information for recall assessment

KB mixes counts with standard facts



number of children

2

Tim Berners-Lee

How many children does Tim Berners-Lee have?

2 (KB fact)



child

Anne Blunt

Ralph King-Milbanke

Byron King-Noel

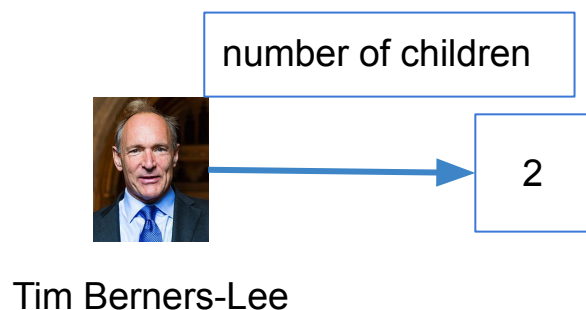
Ada Lovelace

How many children did Ada Lovelace have?

3 (Maybe?)

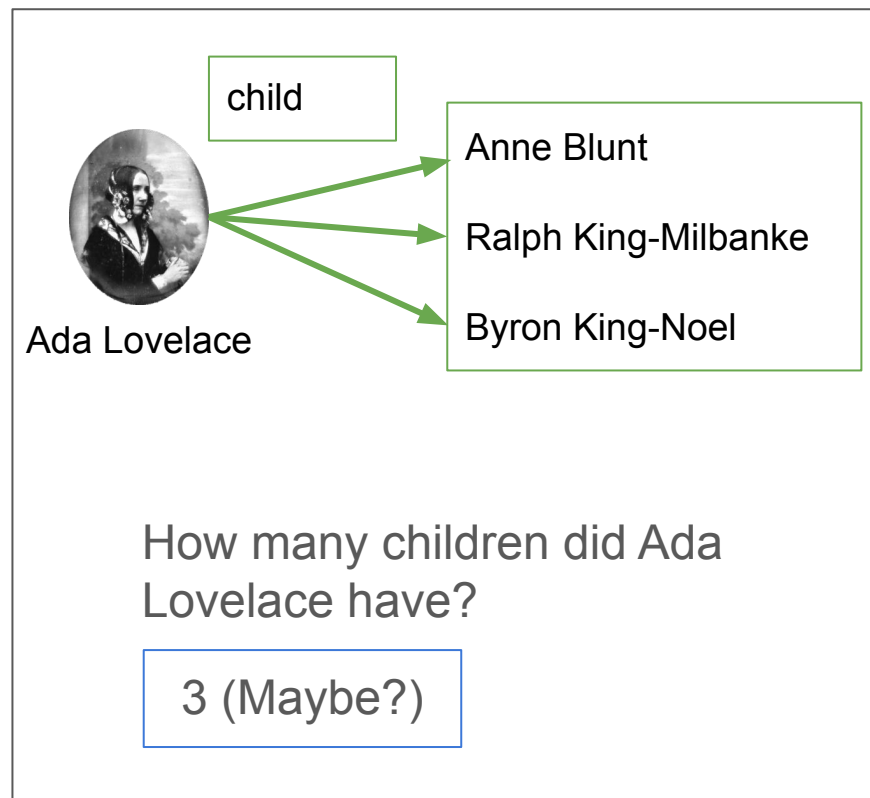
# Count information for recall assessment

KB mixes counts with standard facts



How many children does Tim Berners-Lee have?

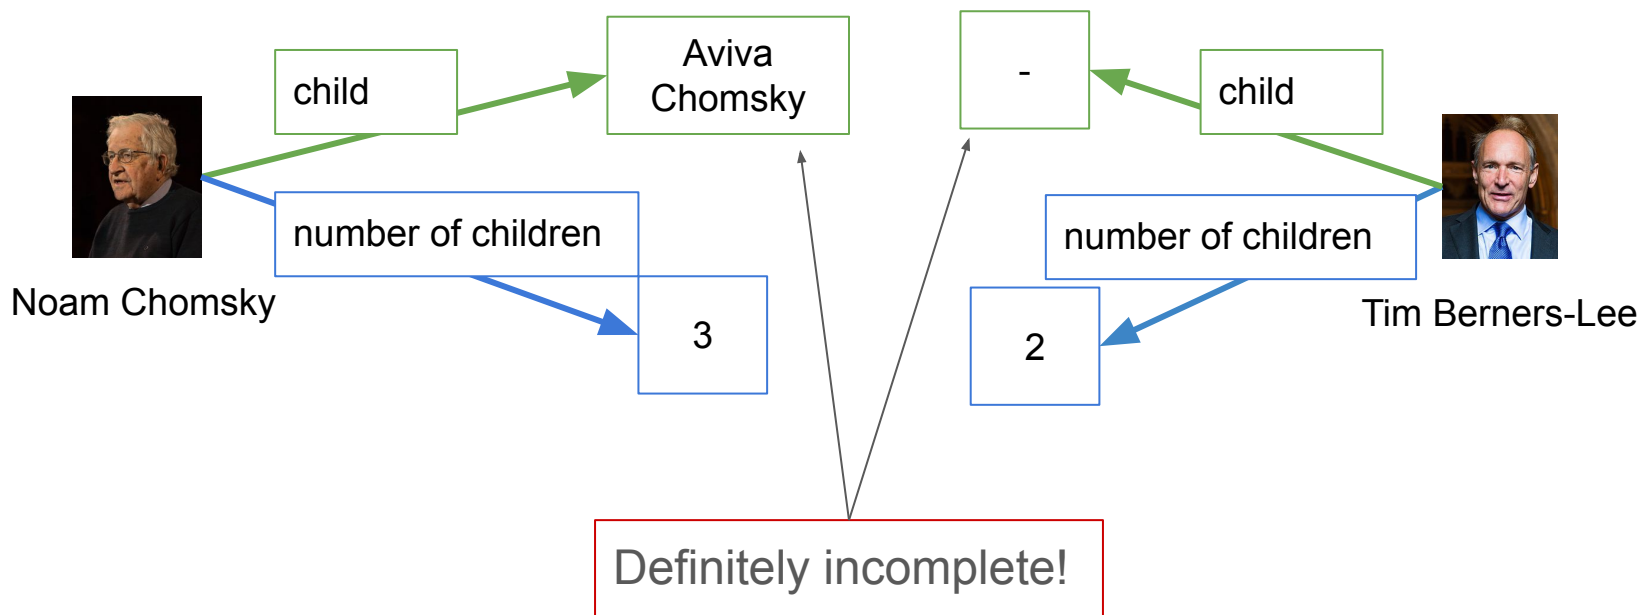
2 (KB fact)



Enumeration is often of known entities

# Count information for recall assessment

Count information can highlight KB inconsistencies



1. Count information for recall assessment
2. How can we extract count information from text?
3. Variants of count information in KB
4. How much count information is accounted for?
5. Counts for KB curation

# Count information from text

**Problem:** Counting Quantifier Extraction

**Input:**

- a text about a subject S
- a predicate P

**Task:** Determine the number of objects in which S stands in relation with P

**Subject:** Noam Chomsky

**Predicate:** number\_of\_children



Chomsky was married to Carol. They had **three children** together

**3**



# Count information from text

**Task 1:** Identify the **count tokens** and the **compositional cues**.

**Sequence Labelling of tokens** in a sentence on subject S and predicate P with:

- COUNT - for counts
- COMP - for compositional cues
- O - all other tokens

**Subject:** Noam Chomsky  
**Predicate:** number\_of\_children



Chomsky was married to Carol. They had **three** children together

O O O O O O **COUNT** O O



# Count information from text

**Task 1:** Identify the **count tokens** and the **compositional cues**.

**Sequence Labelling of tokens** in a sentence on subject S and predicate P with:

- COUNT - for counts
- COMP - for compositional cues
- O - all other tokens



**Subject:** Angelina Jolie

**Predicate:** number\_of\_children

Jolie has **three** sons **and** **three** daughters.  
O O **COUNT** O **COMP** **COUNT** O



# Count information from text

**Task 1:** Identify the **count tokens** and the **compositional cues**.

**COUNT** tokens are **linguistically diverse**

Cardinals

two sons,  
three books



# Count information from text

**Task 1:** Identify the **count tokens** and the **compositional cues**.

**COUNT** tokens are **linguistically diverse**

## Cardinals

two sons,  
three books

## Ordinals

second son,  
third book

# Count information from text

**Task 1:** Identify the **count tokens** and the **compositional cues**.

**COUNT** tokens are **linguistically diverse**

## Cardinals

two sons,  
three books

## Ordinals

second son,  
third book

## Number-related terms

twins, trilogy

# Count information from text

**Task 1:** Identify the **count tokens** and the **compositional cues**.

**COUNT** tokens are **linguistically diverse**

## Cardinals

two sons,  
three books

## Ordinals

second son,  
third book

## Number-related terms

twins, trilogy

## Indefinite Articles

a son,  
the book

# Count information from text

**Task 1:** Identify the **count tokens** and the **compositional cues**.

**COMP** cues for counts occur

- between consecutive count tokens, as
- **comma-separated**, **and-separated** counts

**Subject:** Angelina Jolie

**Predicate:** number\_of\_children

Jolie brought her **twins** , **one** daughter **and** **three** adopted children to the gala.

COMP

COMP

# Count information from text

## Task 2: Consolidate count tokens

Return a **single answer** per text, given subject-predicate pair

1. Sum up compositional cues
2. Select prediction per type
3. Rank mention types

# Count information from text

## Task 2: Consolidate count tokens

Return a single answer per text, given subject-predicate pair

### 1. Sum up compositional cues

6

Jolie brought her **six** children: **twins** , **one** daughter **and** **three** adopted children to the gala.

---

**Subject:** Angelina Jolie

**Predicate:** number\_of\_children

# Count information from text

## Task 2: Consolidate count tokens

Return a single answer per text, given subject-predicate pair

1. Sum up compositional cues
2. **Select prediction per type**

**6 (cardinal)**

**6 (cardinal)**

Jolie brought her **six** children: **twins** , **one** daughter **and three** adopted children to the gala.

---

**Subject:** Angelina Jolie

**Predicate:** number\_of\_children

**6 (cardinal)**

# Count information from text

## Task 2: Consolidate count tokens

Return a single answer per text, given subject-predicate pair

1. Sum up compositional cues
2. Select prediction per type
3. **Rank mention types**

cardinal	>>	number-related terms	>>	ordinals	>>	indefinite article
two children	>>	twins	>>	second child	>>	a child



# Count information from text

## Task 2: Consolidate count tokens

Return a single answer per text, given subject-predicate pair

1. Sum up compositional cues
2. Select prediction per type
3. **Rank mention types**

cardinal	>>	number-related terms	>>	ordinals	>>	indefinite article
two children	>>	twins	>>	second child	>>	a child

Jolie brought her **six** children: **twins** , **one** daughter **and** **three** adopted children to the gala.

**Subject:** Angelina Jolie

**Predicate:** number\_of\_children

**6 (cardinal)**

# Count information from text

**Training data generation:** Incompleteness-aware distant supervision

**Input:** KB, count predicate P

**Output:**

- all subjects S and the count
- all sentences about S containing cardinal mentions similar to the KB count

# Count information from text

**Training data generation:** Incompleteness-aware distant supervision

**Input:** KB, count predicate P

**Output:**

- all subjects S and the count
- all sentences about S containing cardinal mentions similar to the KB count

all counts tokens



# Count information from text

**Training data generation:** Incompleteness-aware distant supervision

**Input:** KB, count predicate P

**Output:**

- all subjects S and the count
- all sentences about S containing cardinal mentions similar to the KB count

all counts tokens

The diagram consists of a rectangular box on the left containing the text 'all counts tokens'. Two arrows originate from the right side of this box. The first arrow points diagonally upwards and to the right, ending at the word 'cardinal' in the second bullet point of the 'Output:' section. The second arrow points diagonally upwards and to the right, ending at the word 'similar' in the same bullet point. The words 'cardinal mentions' and 'similar' are circled in red in the original image.

**+ve:** equal to or representative of KB count  
**-ve :** otherwise and all non-numerals  
Ignore: candidate counts > KB counts

# Count information from text

## Ground Truth

Use KB information as Ground Truth


## Challenges

KB incompleteness negatively impacts training quality

## Solution

Consider only popular KB entities

Set upper bound for predicate count value = 99<sup>th</sup> percentile of KB predicate value distribution



# Count information from text

## Challenge

Counting cardinality when it is Zero

# Count information from text

## Challenge

Counting cardinality when it is Zero

## Solution

Focus on

- i) Negation determiners: 'no' and 'any'
- ii) Non-existence-proving adverbs: 'without' and 'never'

# Count information from text

## Challenge

Counting cardinality when it is Zero

## Solution

Focus on

- i) Negation determiners: 'no' and 'any'
- ii) Non-existence-proving adverbs: 'without' and 'never'

No training - Labelling only when applying models

1. Text preprocessing 

<i>They <u>didn't</u> have <u>any</u> children</i>	→ <i>They have <u>no</u> children</i>
<i>He has <u>never</u> been married</i>	→ <i>He has been married <u>0</u> times</i>
<i>The marriage was <u>without</u> children</i>	→ <i>The marriage was with <u>no</u> children.</i>
2. Textual occurrences of 'no' and '0' → CARDINAL (0)



# Count information from text

Relation	Baseline [22]			CINEX-CRF			CINEX-CRF (per type)					
	P	Cov	MAE	P	Cov	MAE	Cardinals		Numt.+Art.		Ordinals	
							P	Contr	P	Contr	P	Contr
containsWork	42.0	<b>29.0</b>	3.7	<b>49.2</b>	<b>29.0</b>	<b>2.6</b>	55.0	33.9	62.5	40.7	20.0	25.4
hasMember	11.8	6.0	3.8	<b>64.3</b>	<b>18.0</b>	<b>1.2</b>	62.5	28.6	65.0	71.4	0	0
containsAdmin	51.8	14.5	7.3	<b>78.6</b>	<b>22.0</b>	<b>1.7</b>	85.7	87.5	33.3	10.7	0	1.8
hasChild	37.0	<b>22.0</b>	<b>2.2</b>	<b>50.0</b>	19.5	2.3	67.3	70.5	6.3	20.5	14.3	9.0
hasSpouse	26.8	11.0	1.3	<b>58.1</b>	<b>12.5</b>	<b>0.5</b>	75.0	18.6	43.8	37.2	63.2	44.2
hasZeroChild				92.3	18.8	-						
hasZeroSpouse				71.9	13.7	-						

Performance of CINEX in consolidation of counting quantifier mentions on Wikidata.

Paramita Mirza, Simon Razniewski, Fariz Darari, Gerhard Weikum

[Enriching Knowledge Bases with Quantifiers](#)

International Semantic Web Conference (ISWC) 2018.

1. Count information for recall assessment
2. How can we extract count information from text?
3. Variants of count information in KB
4. How much count information is accounted for?
5. Counts for KB curation

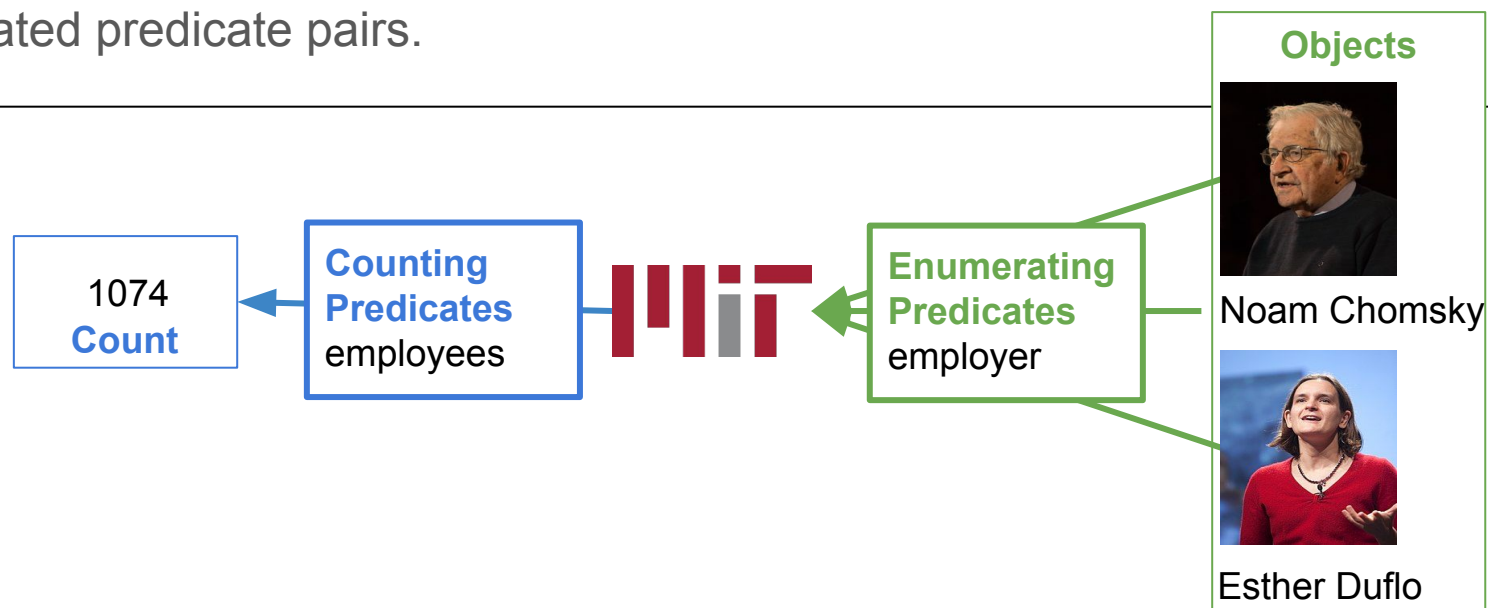
# Count information in KB

**Problem:** Identification of semantically related count predicates

**Input:**

- a set of KB triples  $(s,p,o)$
- and its inverse predicate triples  $(s,p^{-1},o)$

**Task:** Determine counting and enumerating predicates and semantically related predicate pairs.



# Count information in KB

**Task 1:** Identification of the count predicates - **counting** and **enumerating**

# Count information in KB

**Task 1:** Identification of the count predicates - **counting** and **enumerating**

List of frequent KB predicates

academic\_staff, staff,  
faculty

number\_of\_children

wins, doubles\_titles,  
singles\_titles

...

work\_institution<sup>-1</sup>, workplace<sup>-1</sup>,  
work\_institutions<sup>-1</sup>

child

gold<sup>-1</sup>

# Count information in KB

Task 1: Identification of the two variants of count predicates

## Counting Predicates

academic\_staff, staff,  
faculty

number\_of\_children

...

wins, doubles\_titles,  
singles\_titles

## Enumerating Predicates

work\_institution<sup>-1</sup>, workplace<sup>-1</sup>,  
work\_institutions<sup>-1</sup>

child

...

gold<sup>-1</sup>

# Count information in KB

**Task 1:** Identification of the two variants of count predicates

## Counting Predicates

academic\_staff, staff, faculty      number\_of\_children      ...      wins, doubles\_titles, singles\_titles

## Enumerating Predicates

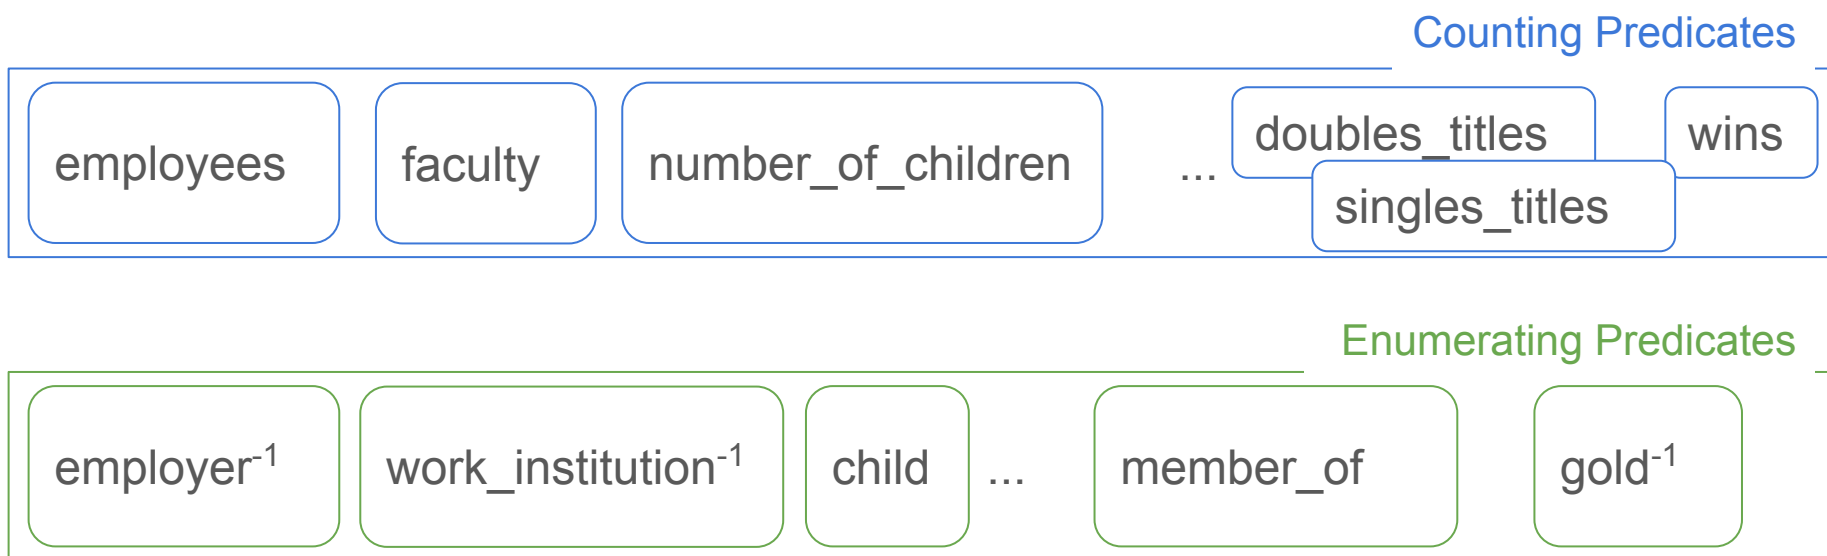
work\_institution<sup>-1</sup>, workplace<sup>-1</sup>, work\_institutions<sup>-1</sup>      child      ...      gold<sup>-1</sup>

Supervised Classification using:

- **Textual Features** - count predicates are more often used in singular form
- **Type Information** - classes of subject and objects
- **KB statistics** - #objects per subject, datatype distribution of the objects

# Count information in KB

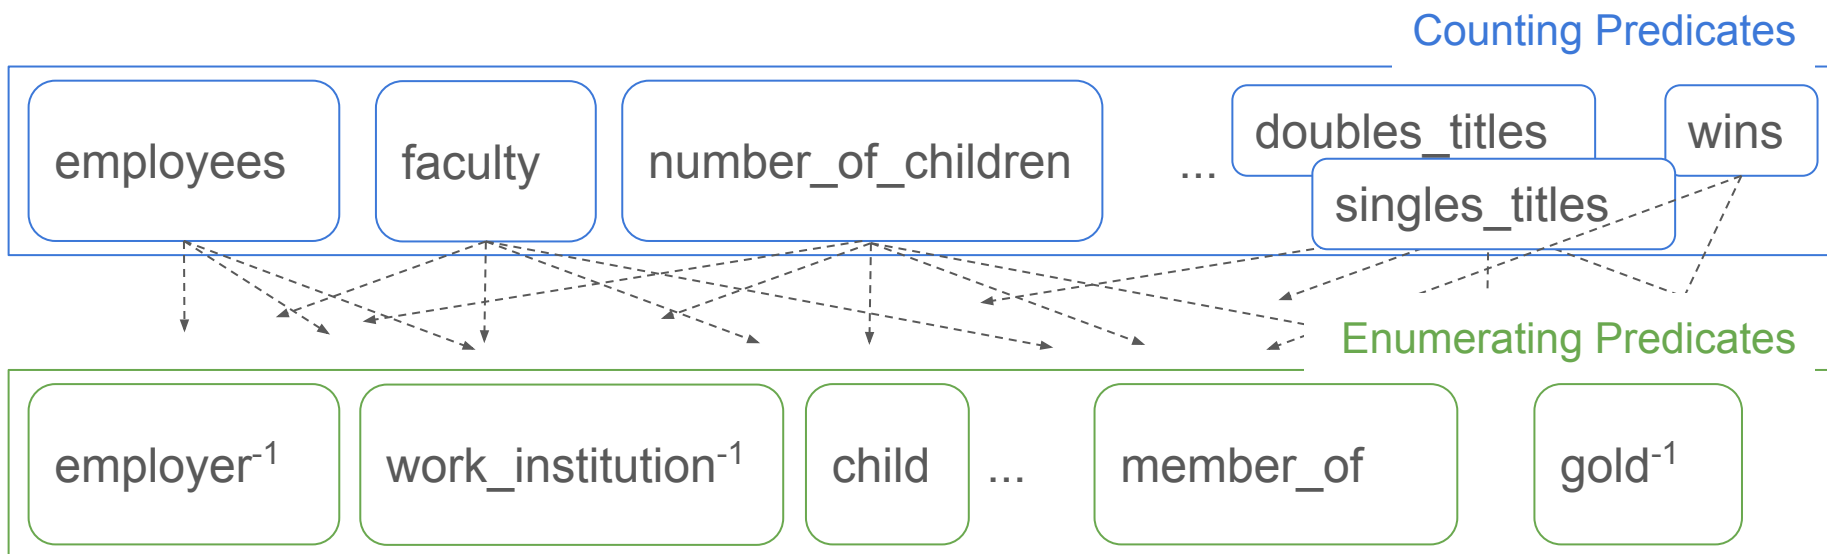
**Task 2:** Align pairs of counting and enumerating predicates






# Count information in KB

**Task 2:** Align pairs of counting and enumerating predicates




# Count information in KB



Heuristics used for the predicate pair  $(e, c)$ , where  $e$  stores entities and  $c$  counts.

1. Predicate pair co-occurrences - #subjects  $e$  and  $c$  co-occur


# Count information in KB



Heuristics used for the predicate pair  $(e, c)$ , where  $e$  stores entities and  $c$  counts.

1. Predicate pair co-occurrences - #subjects  $e$  and  $c$  co-occur
2. Value distribution - number of objects of  $e$  compared to count in  $c$


# Count information in KB



Heuristics used for the predicate pair  $(e, c)$ , where  $e$  stores entities and  $c$  counts.

1. Predicate pair co-occurrences - #subjects  $e$  and  $c$  co-occur
2. Value distribution - number of objects of  $e$  compared to count in  $c$ 
  - a. is it equal for all subjects?
  - b. is there any correlation?

# Count information in KB



Heuristics used for the predicate pair  $(e, c)$ , where  $e$  stores entities and  $c$  counts.

1. Predicate pair co-occurrences - #subjects  $e$  and  $c$  co-occur
2. Value distribution - number of objects of  $e$  compared to count in  $c$ 
  - a. is it equal for all subjects?
  - b. is there any correlation?
2. Linguistic similarity - do  $e$  and  $c$  talk share topical similarity?

# Count information in KB

**Training data generation:** Crowd-sourced annotation of randomly selected predicate subsets

**Challenges:** KB predicates rarely have clean values

# Count information in KB

**Training data generation:** Crowd-sourced annotation of randomly selected predicate subsets

**Challenges:** KB predicates rarely have clean values

- Cannot rely only on #triples per subject for enumerating predicates
- Integer value for a predicate  $\rightarrow$  Counting predicate (seat number, codes, IDs)
- Need for human in the loop

# Count information in KB

**Training data generation:** Crowd-sourced annotation of randomly selected predicate subsets

**Challenges:** KB predicates rarely have clean values

- Cannot rely only on #triples per subject for enumerating predicates
- Integer value for a predicate  $\rightarrow$  Counting predicate (seat number, codes, IDs)
- Need for human in the loop

**Input:** Predicate P, 5 KB triples per predicate

**Output:**

- Graded relevance score for each P
- **+ve:** Average score from 3 users is between [0.6, 1.0]
- **-ve:** Average score between [0, 0.4]



# Count information in KB

**Training data generation:** Crowd-sourced annotation of randomly selected predicate subsets

**Challenges:** KB predicates rarely have clean values

- Cannot rely only on #triples per subject for enumerating predicates
- Integer value for a predicate  $\rightarrow$  Counting predicate (seat number, codes, IDs)
- Need for human in the loop

**Input:** Predicate P, 5 KB triples per predicate

**Output:**

- Graded relevance score for each P
- +ve: Average score from 3 users is between [0.6, 1.0]
- -ve: Average score between [0, 0.4]

to weed out predicates without clear polarity

# Count information in KB

**Ground truth data:** Crowd-sourced annotation of the top enumerating (counting) predicates aligned to randomly selected counting (enumerating) predicates

**Input:** Counting predicate C and top aligned predicates of the other set ( $E_1, E_2, \dots$ ) returned by all heuristics.

faculty\_size

work\_institution<sup>-1</sup>

works\_at<sup>-1</sup>

employer<sup>-1</sup>

# Count information in KB

**Ground truth data:** Crowd-sourced annotation of the top enumerating (counting) predicates aligned to randomly selected counting (enumerating) predicates

**Input:** Counting predicate  $C$  and top aligned predicates of the other set  $(E_1, E_2, \dots)$  returned by all heuristics.

faculty\_size

work\_institution<sup>-1</sup>

works\_at<sup>-1</sup>

**Output:**

- Graded relevance score for each pair  $(C, E_1), (C, E_2), \dots$
- Determine top-3 aligned predicates for  $C$

employer<sup>-1</sup>

Repeat this for enumerating predicates to get their top-3 aligned counting predicates.

# Count information in KB

Model	Recall	Precision	F1
Random	40.6	40.6	40.6
Logistic	<b>55.6</b>	51.7	53.5
Prior	<b>55.6</b>	51.0	53.5
Lasso	51.1	<b>59.6</b>	<b>55.0</b>
Neural	53.0	49.6	51.2

Model	Recall	Precision	F1
Random	12.8	12.8	12.8
Logistic	51.2	19.0	27.7
Prior	48.7	20.2	28.5
Lasso	<b>71.7</b>	<b>23.3</b>	<b>35.1</b>
Neural	35.8	20.8	26.3

Scores for predicting i) Enumerating  
ii) Counting predicates

Metric	Counting		Enumerating	
	@1	@3	@1	@3
<i>Absolute</i>	0.71	0.56	0.62	0.63
<i>Jaccard</i>	0.76	0.61	0.69	0.67
<i>Conditional<sub>C</sub></i>	0.71	0.56	0.68	0.67
<i>Conditional<sub>E</sub></i>	0.76	0.68	0.62	0.63
<i>P'wiseMI</i>	0.73	0.58	0.71	0.70
<i>P'fectMR</i>	0.70	0.57	0.73	0.72
<i>Correlation</i>	0.77	0.69	0.62	0.61
<i>P'tileVM</i>	0.72	0.57	0.65	0.65
<i>CosineSim</i>	0.79	0.61	0.74	0.73
<b><i>Combined</i></b>	<b>0.84</b>	<b>0.67</b>	<b>0.75</b>	<b>0.75</b>

NDCG scores for predicate alignment

Shrestha Ghosh, Simon Razniewski, Gerhard Weikum

[Uncovering Hidden Semantics of Set Information in Knowledge Bases](#)

Journal of Web Semantics (JWS) 2020.

1. Count information for recall assessment
2. How can we extract count information from text?
3. Variants of count information in KB
4. How much count information is accounted for?
5. Counts for KB curation

# How much count information is accounted for?

## Counts from text

173k new count facts increasing KB knowledge by **77%**

from just 4 Wikidata properties across 10 classes

2,205 negative assertions

2.5M new count facts increasing KB knowledge by **28.3%**

# How much count information is accounted for?

## Counts from text

173k new count facts increasing KB knowledge by **77%**

from just 4 Wikidata properties across 10 classes

2,205 negative assertions

2.5M new count facts increasing KB knowledge by **28.3%**

for the predicates: *hasSpouse* and *hasChild*

# How much count information is accounted for?

## Counts from text

173k new count facts increasing KB knowledge by **77%**

from just 4 Wikidata properties across 10 classes

2,205 negative assertions

2.5M new count facts increasing KB knowledge by **28.3%**

for the predicates: *hasSpouse* and *hasChild*

for 110 Wikidata properties-class pairs

Paramita Mirza, Simon Razniewski, Fariz Darari, Gerhard Weikum

[Enriching Knowledge Bases with Quantifiers](#)

International Semantic Web Conference (ISWC) 2018.



# How much count information is accounted for?

<b>KB</b>	<b>Enumerating</b>
DBpedia-raw	4,090
DBpedia mapped	308
Wikidata-truthy	203
Freebase	7,614
<b>Total</b>	<b>12,215</b>

Number of predicted **enumerating** KB predicates

# How much count information is accounted for?

KB	Enumerating	Counting
DBpedia-raw	4,090	5,853
DBpedia mapped	308	898
Wikidata-truthy	203	1,067
Freebase	7,614	1,687
Total	12,215	9,505

Number of predicted [counting](#) KB predicates

Shrestha Ghosh, Simon Razniewski, Gerhard Weikum  
[Uncovering Hidden Semantics of Set Information in Knowledge Bases](#)  
Journal of Web Semantics (JWS) 2020.

# How much count information is accounted for?

Number of predicted count predicates and KB alignments

KB	Enumerating	Counting	Alignments
DBpedia-raw	4,090	5,853	3,703
DBpedia mapped	308	898	270
Wikidata-truthy	203	1,067	31
Freebase	7,614	1,687	274
Total	12,215	9,505	4,278

Quite a low number of alignments: indicative of KB sparsity

Shrestha Ghosh, Simon Razniewski, Gerhard Weikum

[Uncovering Hidden Semantics of Set Information in Knowledge Bases](#)

Journal of Web Semantics (JWS) 2020.

# How much count information is accounted for?

## Open questions and challenges

- #alignments  $\ll$  #counting and #enumerating predicates
  - unaligned count predicate  $\rightarrow$  scope for new predicates
- Clustering similar predicates (faculty  $\leftrightarrow$  staff size)
  - staff size exists for an entity instead of faculty, then use it
- Cardinality extractors from text individually trained for each predicate
- Enumeration for static (children, spouses) vs dynamic classes (population, books)

1. Count information for recall assessment
2. How can we extract count information from text?
3. Variants of count information in KB
4. How much count information is accounted for?
5. Counts for KB curation

# Counts for KB curation

Entity Noam Chomsky: Amer Set Predicate P40: child (1)

!! Hope the results satisfy your curiosity!

Noam Chomsky	P40: child	Aviva Chomsky	
--------------	------------	---------------	--

**Related Counting Predicates**

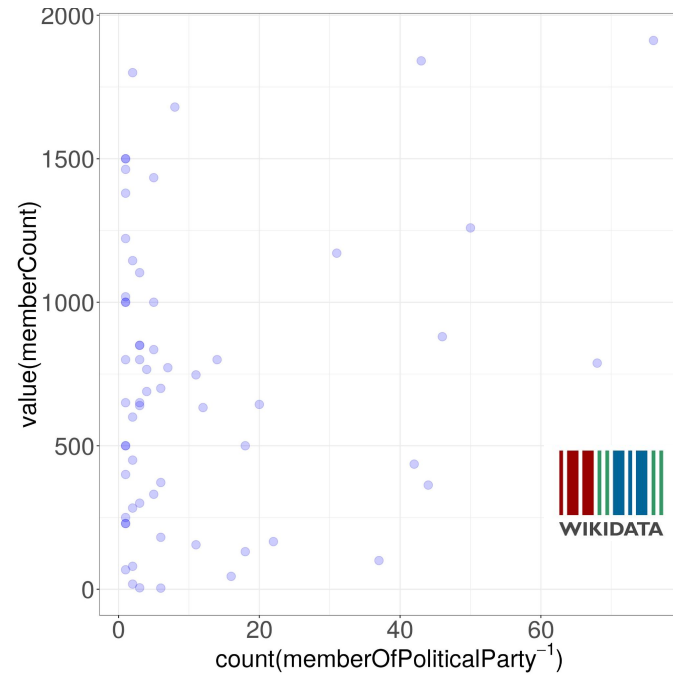
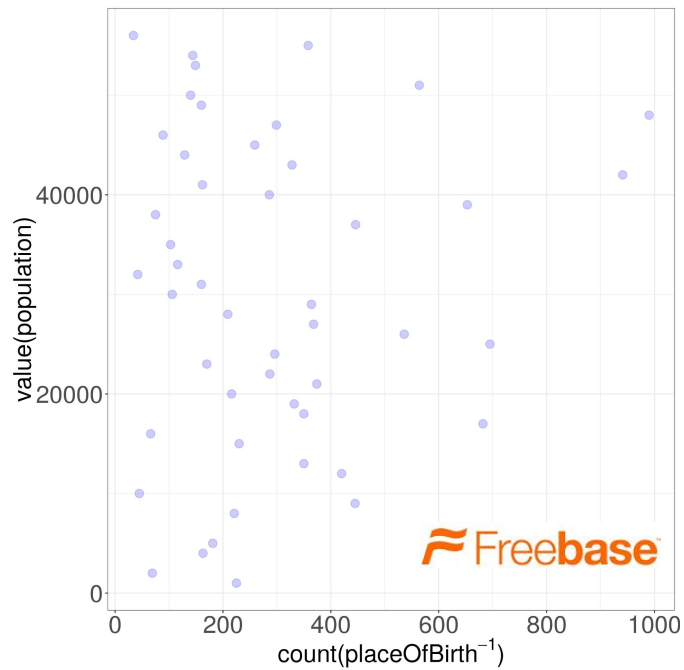
Noam Chomsky	P1971: number of children	(no instantiations)	
--------------	---------------------------	---------------------	--

KB inconsistencies are highlighted

<https://counqer.mpi-inf.mpg.de/spo>

# Counts for KB curation

Value distribution of aligned predicates show incompleteness



# Counts for KB curation

Enhanced KB question answering

**CounQER**

Entity: Roger Federer    Set Predicate: dbp: gold

!! There exist no instantiated facts on the queried set predicate.

**Related Counting Predicates**

Roger Federer	dbp: wins	(no instantiations)	<a href="#">↗</a>
Roger Federer	dbp: doublestitles	8	<a href="#">↗</a>
Roger Federer	dbp: singlestitles	88	<a href="#">↗</a>
Roger Federer	dbp: totalpodiums	(no instantiations)	<a href="#">↗</a>
Roger Federer	dbp: doublesrecord	(no instantiations)	<a href="#">↗</a>

No answer to the original query on enumerating predicate

Related count answers obtained from aligned count predicates

Count predicates which could potentially give more information



# Takeaway: Counts from text and KB

1. Count information for recall assessment
  - Counts and entities benefit from each other
  - KB mixes counts with standard facts
  - Counts can improve KB recall
2. Count information in text
  - is linguistically diverse (cardinals, ordinals, ..)
  - used to get the #objects for a given subject and predicate
3. Count information in KBs
  - can be identified by supervised classification
  - occurs as semantically related counting and enumerating predicates
4. KB curation using counts
  - highlights inconsistencies
  - gives value distribution of aligned predicates
  - can enhance KB question answering