Representing a concept by the distribution of names of its instances

Matthijs Westera, Gemma Boleda and Sebastian Padó

Universitat Pompeu Fabra Barcelona

Computational Linguistics and Linguistic Theory

Universität Stuttgart
Institut für Maschinelle Sprachverarbeitung

Representing a concept by the distribution of names of its instances

Abhijeet Gupta &
Matthijs Westera, Gemma Boleda and Sebastian Padó

Universitat Pompeu Fabra Barcelona
Computational Linguistics and Linguistic Theory
Universität Stuttgart
Institut für Maschinelle Sprachverarbeitung

Interest in Distributional Semantics (etc.)

Interest in Distributional Semantics (etc.)

- Relation to formal semantics;

Interest in Distributional Semantics (etc.)

- Relation to formal semantics;
- Relevance to experimental linguistics;

Interest in Distributional Semantics (etc.)

- Relation to formal semantics;
- Relevance to experimental linguistics;
- Relation between language and the world.

Interest in Distributional Semantics (etc.)

- Relation to formal semantics;
- Relevance to experimental linguistics;
- Relation between language and the world.

Language and the world

Language and the world

… that dog ate my shoe …

Language and the world

… that dog ate my shoe …

… a young dog is called a puppy …

Language and the world

… that dog ate my shoe …

… a young dog is called a puppy …

… every cat ate too much …

Language and the world

… that dog ate my shoe …

… a young dog is called a puppy …

… every cat ate too much …

… when my cat was young she …

Language and the world

… that dog ate my shoe …

… a young dog is called a puppy …

… every cat ate too much …

… when my cat was young she …

Language and the world

… that dog ate my shoe …

… a young dog is called a puppy …

… every cat ate too much …

… when my cat was young she …

Language and the world

… that dog ate my shoe …

… a young dog is called a puppy …

… every cat ate too much …

… when my cat was young she …

Language and the world

… that dog ate my shoe …

… a young dog is called a puppy …

… every cat ate too much …

… when my cat was young she …

… that dog ate my shoe …

… a young dog is called a puppy …

… every cat ate too much …

… when my cat was young she …

Distributional Semantics (DS)

… that dog ate my shoe …

… a young dog is called a puppy …

… every cat ate too much …

… when my cat was young she …

Distributional Semantics (DS)

… that dog ate my shoe …

… a young dog is called a puppy …

… every cat ate too much …

… when my cat was young she …

Distributional Semantics (DS)

… that dog ate my shoe …

… a young dog is called a puppy …

… every cat ate too much …

… when my cat was young she …

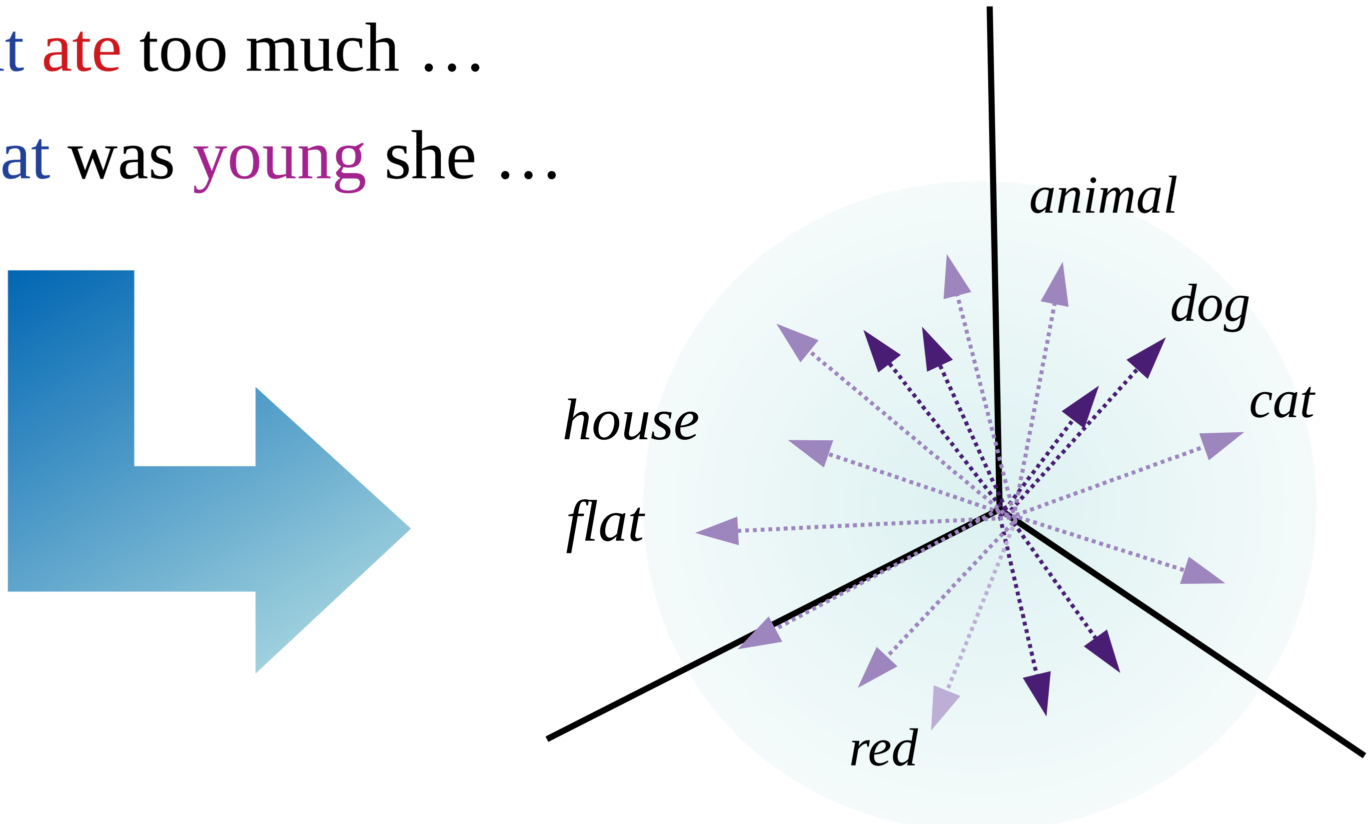

Westera & Boleda (2019, IWCS):

Distributional Semantics as a model of concepts?

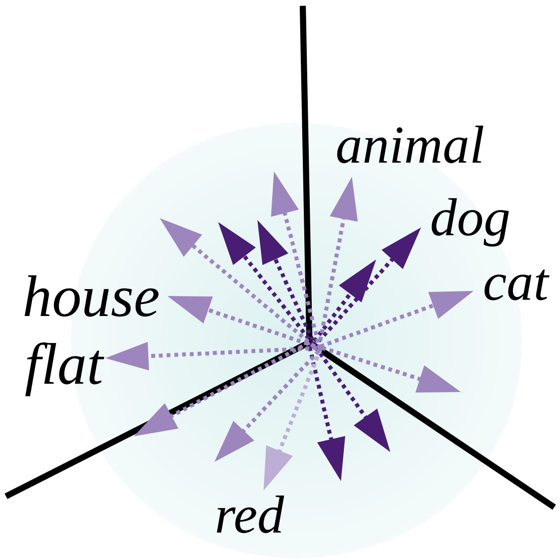

Westera & Boleda (2019, IWCS):

Distributional Semantics as a model of concepts?

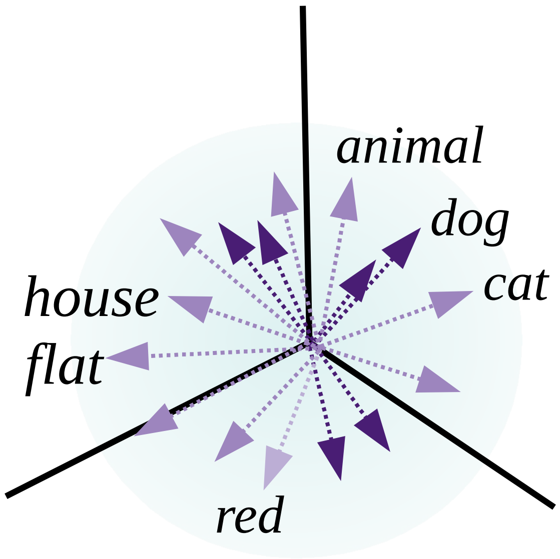

- The vectors of DS are *abstractions over occurrences*.

Westera & Boleda (2019, IWCS):
Distributional Semantics as a model of concepts?

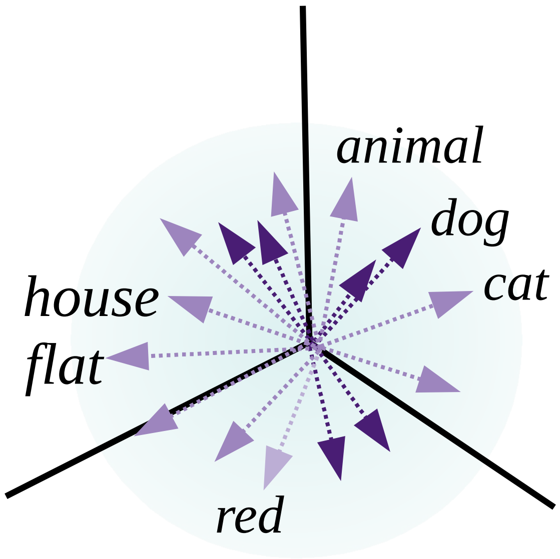

- The vectors of DS are *abstractions over occurrences*.
- And so are *concepts* (e.g., Piaget).

Westera & Boleda (2019, IWCS):

Distributional Semantics as a model of concepts?

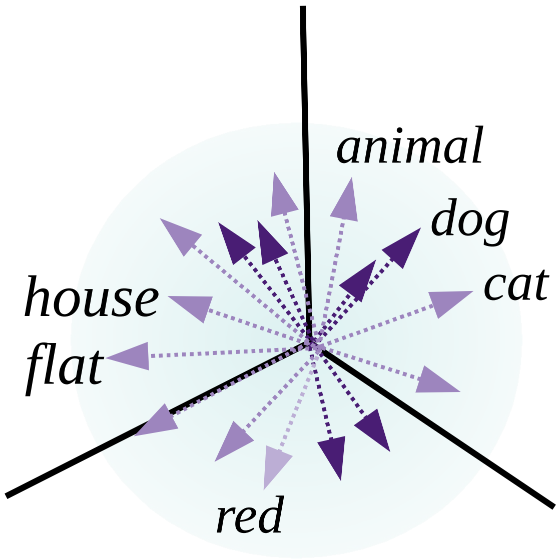

- The vectors of DS are *abstractions over occurrences*.
- And so are *concepts* (e.g., Piaget).

But what sort of concepts does DS model?

Westera & Boleda (2019, IWCS):

Distributional Semantics as a model of concepts?

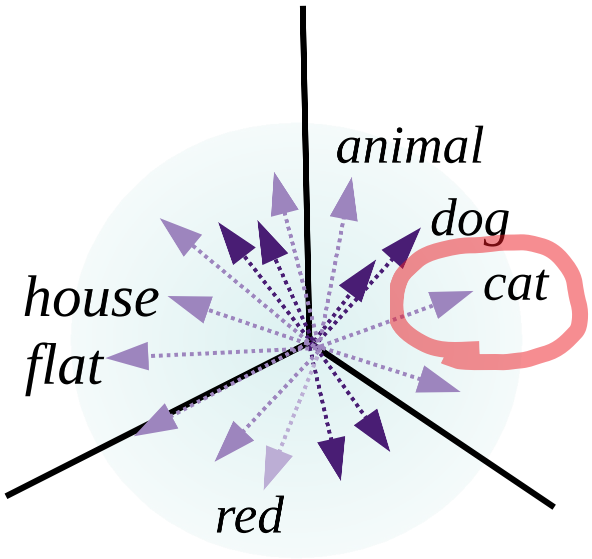

- The vectors of DS are *abstractions over occurrences*.
- And so are *concepts* (e.g., Piaget).

But what sort of concepts does DS model?

Westera & Boleda (2019, IWCS):

Distributional Semantics as a model of concepts?

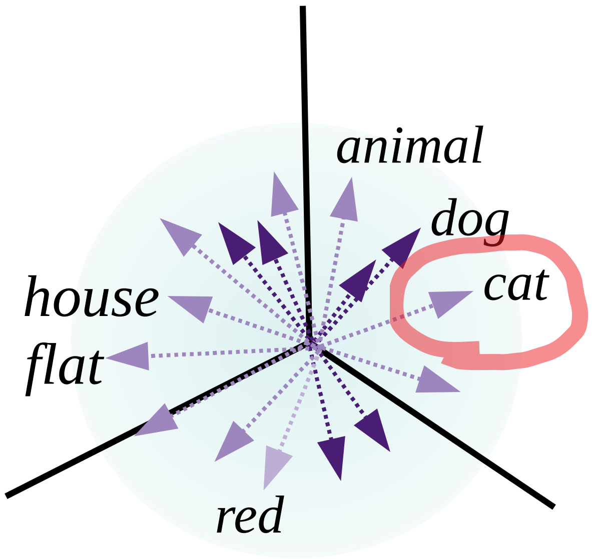

- The vectors of DS are *abstractions over occurrences*.
- And so are *concepts* (e.g., Piaget).

But what sort of concepts does DS model?

Westera & Boleda (2019, IWCS):

Distributional Semantics as a model of concepts?

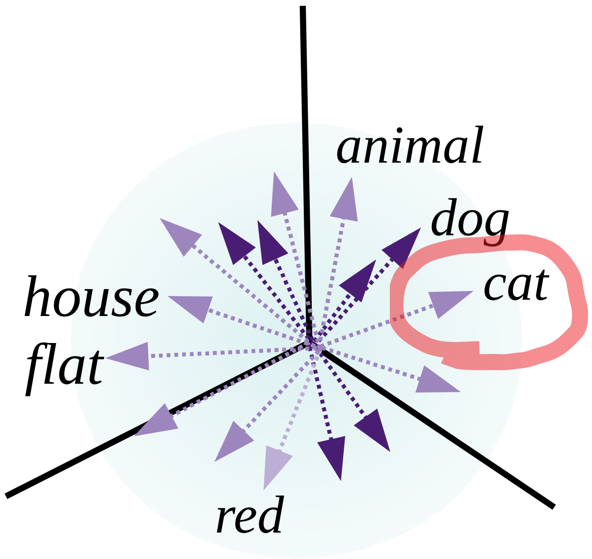

- The vectors of DS are *abstractions over occurrences*.
- And so are *concepts* (e.g., Piaget).

But what sort of concepts does DS model?

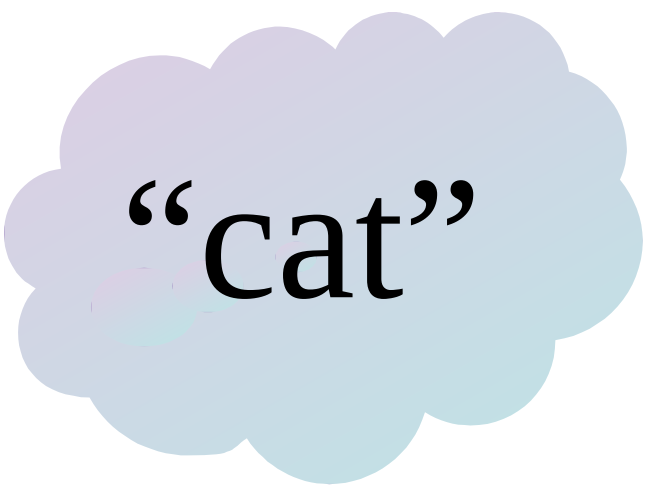

Westera & Boleda (2019, IWCS):

Distributional Semantics as a model of concepts?

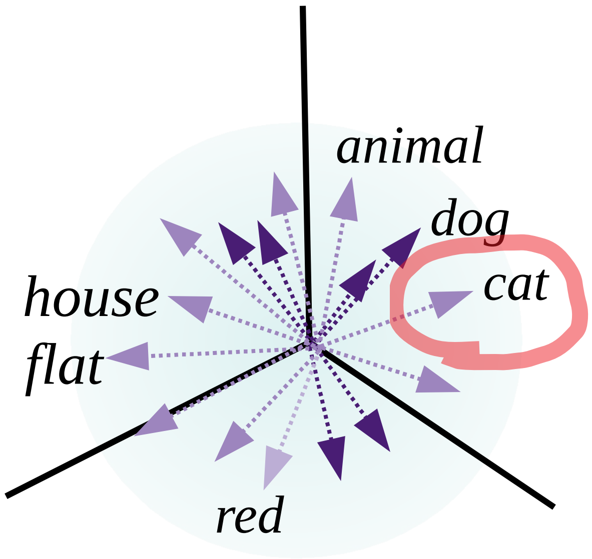

- The vectors of DS are *abstractions over occurrences*.
- And so are *concepts* (e.g., Piaget).

But what sort of concepts does DS model?

Westera & Boleda (2019, IWCS):

Should Distributional Semantics account for *entailment*?

Westera & Boleda (2019, IWCS):

Should Distributional Semantics account for *entailment*?

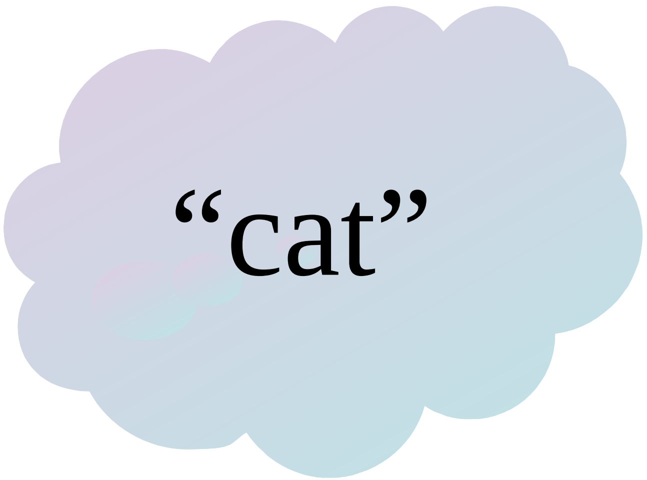

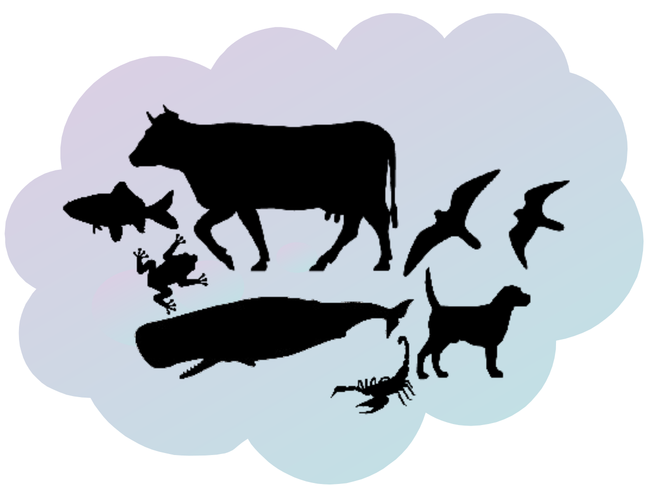

Westera & Boleda (2019, IWCS):

Should Distributional Semantics account for *entailment*?

“cat”

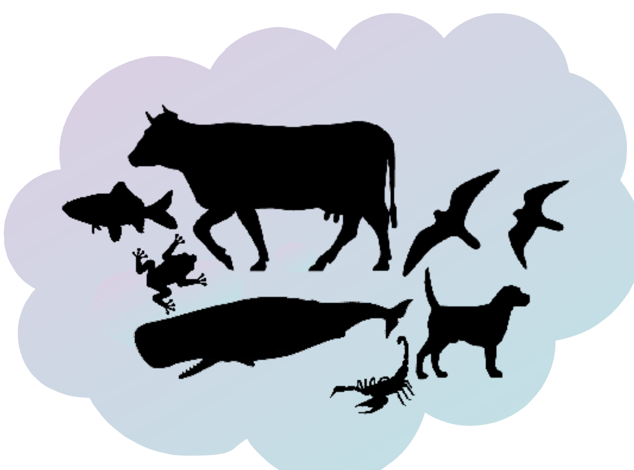

“animal”

Westera & Boleda (2019, IWCS):

Should Distributional Semantics account for *entailment*?

“cat”

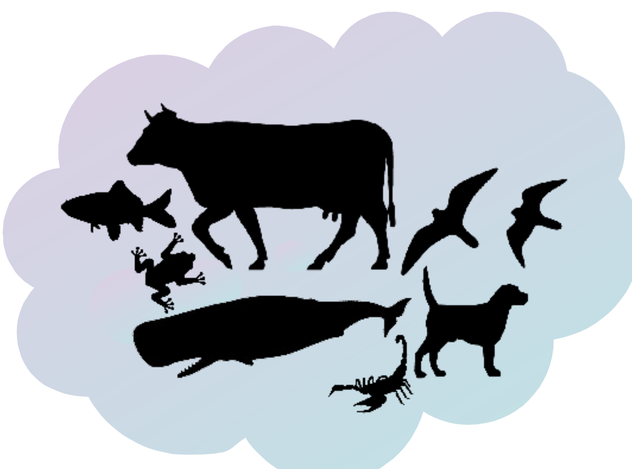

“animal”

Westera & Boleda (2019, IWCS):

Should Distributional Semantics account for *entailment*?

No.

“cat”

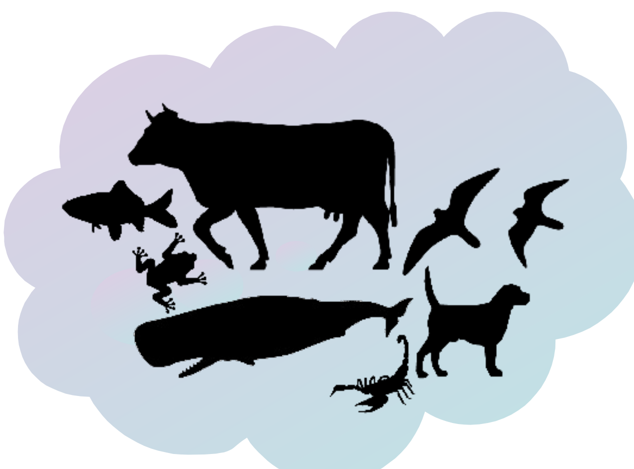

“animal”

“cat”

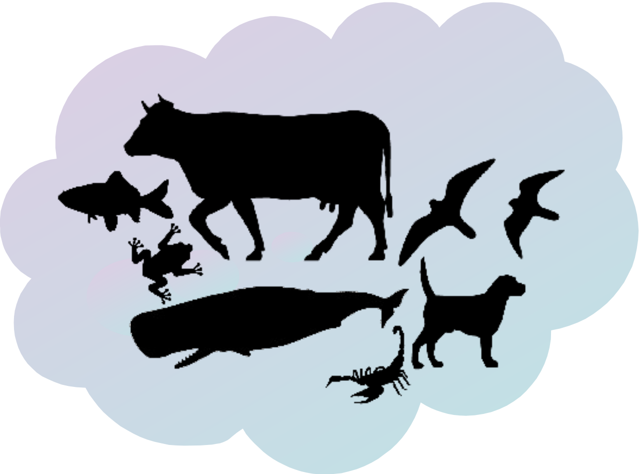

“animal”

Language and the world are not perfectly aligned

“cat”

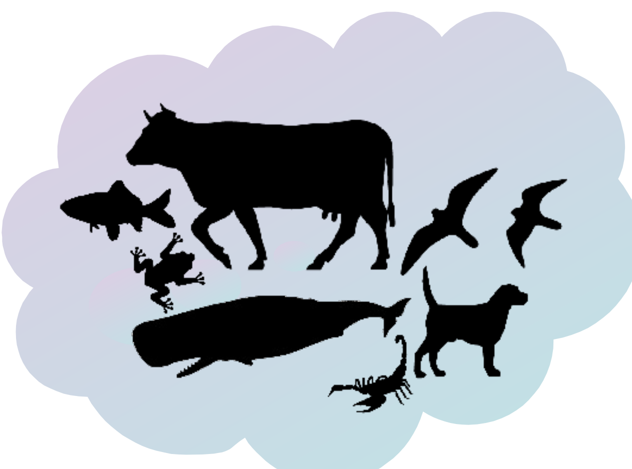

“animal”

Language and the world are not perfectly aligned

~ "cat"

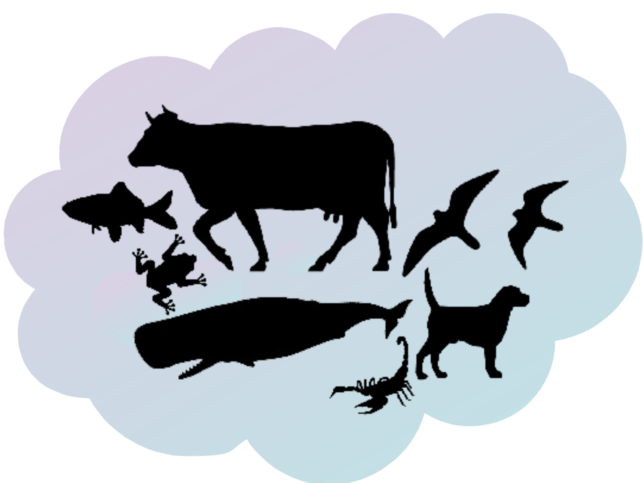

~ "animal"

Language and the world are not perfectly aligned

Language and the world are not perfectly aligned

- This is not (just) a technical challenge, but *interesting*.

Language and the world are not perfectly aligned

- This is not (just) a technical challenge, but *interesting*.
- Are some parts of language closer to the world than other parts? Does this show in DS? Can we exploit this?

Language and the world are not perfectly aligned

- This is not (just) a technical challenge, but *interesting*.
- Are some parts of language closer to the world than other parts? Does this show in DS? Can we exploit this?

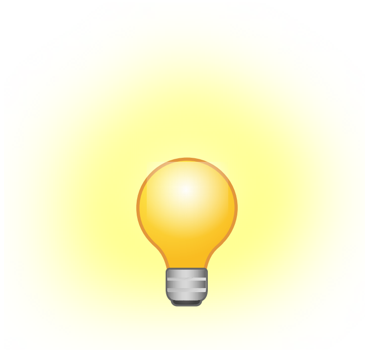

Some expressions are used more rigidly than others...

Language and the world are not perfectly aligned

- This is not (just) a technical challenge, but *interesting*.
- Are some parts of language closer to the world than other parts? Does this show in DS? Can we exploit this?

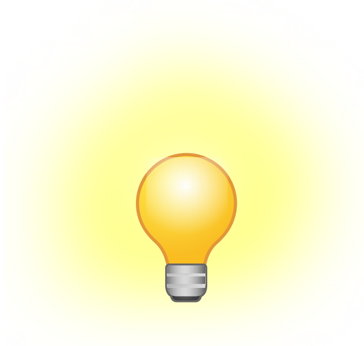

Some expressions are used more rigidly than others... (Kripke, '80)

Approach

Approach

- Let’s compare two kinds of representations of category concepts:

Approach

- Let’s compare two kinds of representations of category concepts:
  - **Predicate-based:**
    Word vector of a predicate that is used to denote the category.

Approach

- Let’s compare two kinds of representations of category concepts:
  - **Predicate-based:**
    Word vector of a predicate that is used to denote the category.
  - **Name-based:**
    Centroid of the word vectors of names of instances of the category.

Approach

- Let’s compare two kinds of representations of category concepts:
  - **Predicate-based:**
    Word vector of a predicate that is used to denote the category.

    E.g., for Scientist, the word vector of "scientist"

  - **Name-based:**
    Centroid of the word vectors of names of instances of the category.

Approach

- Let's compare two kinds of representations of category concepts:
  - **Predicate-based:**
    Word vector of a predicate that is used to denote the category.

    E.g., for Scientist, the word vector of "scientist"
  - **Name-based:**
    Centroid of the word vectors of names of instances of the category.

E.g., the mean of vectors for "Albert Einstein", "Emmy Noether", ...

Approach

- Let’s compare two kinds of representations of category concepts:
  - **Predicate-based:**
    Word vector of a predicate that is used to denote the category.

    E.g., for Scientist, the word vector of "scientist"
  - **Name-based:**
    Centroid of the word vectors of names of instances of the category.

  E.g., the mean of vectors for "Albert Einstein", "Emmy Noether", …
- Evaluation against human judgments of category relatedness.

Representing a concept by the distribution of names of its instances

Abhijeet Gupta &
Matthijs Westera, Gemma Boleda and Sebastian Padó

Universitat Pompeu Fabra Barcelona
Computational Linguistics and Linguistic Theory
Universität Stuttgart
Institut für Maschinelle Sprachverarbeitung

Existing data/model we use

Existing data/model we use

- The **Instantiation** dataset (Boleda, Gupta, and Padó, 2017, EACL):
  - e.g., *<Emmy Noether, scientist>*, *<Edinburgh, capital>*

Existing data/model we use

- The **Instantiation** dataset (Boleda, Gupta, and Padó, 2017, EACL):
  - e.g., *<Emmy Noether, scientist>*, *<Edinburgh, capital>*
  - derived from WordNet's 'instance hyponym' relation.

Existing data/model we use

- The **Instantiation** dataset (Boleda, Gupta, and Padó, 2017, EACL):
  - e.g., <*Emmy Noether*, *scientist*>, <*Edinburgh*, *capital*>
  - derived from WordNet's 'instance hyponym' relation.
- We focus on the 159 categories that have at least 5 entities.

Existing data/model we use

- The **Instantiation** dataset (Boleda, Gupta, and Padó, 2017, EACL):
  - e.g., <*Emmy Noether*, *scientist*>, <*Edinburgh*, *capital*>
  - derived from WordNet's 'instance hyponym' relation.
- We focus on the 159 categories that have at least 5 entities.

- As DS representations of the entities' names and categories' predicates we use the *Google News* embeddings (Mikolov, Sutskever, et al., 2013, ANIPS).

Evaluation: gathering human judgments

Evaluation: gathering human judgments

Following Bruni, Tran and Baroni's MEN benchmark (2012, JAIR):

Evaluation: gathering human judgments

Following Bruni, Tran and Baroni's MEN benchmark (2012, JAIR):

- We semi-randomly sampled 1000 category pairs (out of 12.5K).

Evaluation: gathering human judgments

Following Bruni, Tran and Baroni's MEN benchmark (2012, JAIR):

- We semi-randomly sampled 1000 category pairs (out of 12.5K).
- **'Comparative' task**: *which pair of categories are more related to each other?*

Evaluation: gathering human judgments

Following Bruni, Tran and Baroni’s MEN benchmark (2012, JAIR):

- We semi-randomly sampled 1000 category pairs (out of 12.5K).
- **‘Comparative’ task**: *which pair of categories are more related to each other?*
- Also same way of computing aggregated ‘relatedness’ scores.

Crowdsource task

In this HIT you will see 70 items like the following, each presenting two pairs of categories:

Which pair of categories are more related to each other?

1. wheel ↔ car

2. building ↔ crane *(type of bird)*

Main result

Main result

- Spearman (ranking) correlations between:

Main result

- Spearman (ranking) correlations between:
  - cosine similarities from Name-based / Predicate-based

and

  - aggregate scores from our human judgments

Main result

- Spearman (ranking) correlations between:
  - cosine similarities from Name-based / Predicate-based

    and

  - aggregate scores from our human judgments

- Result:
  - Predicate-based: 0.56

Main result

- Spearman (ranking) correlations between:
  - cosine similarities from Name-based / Predicate-based

and

  - aggregate scores from our human judgments

- Result:
  - Predicate-based: 0.56
  - Name-based: **0.74**

Artist’s impression

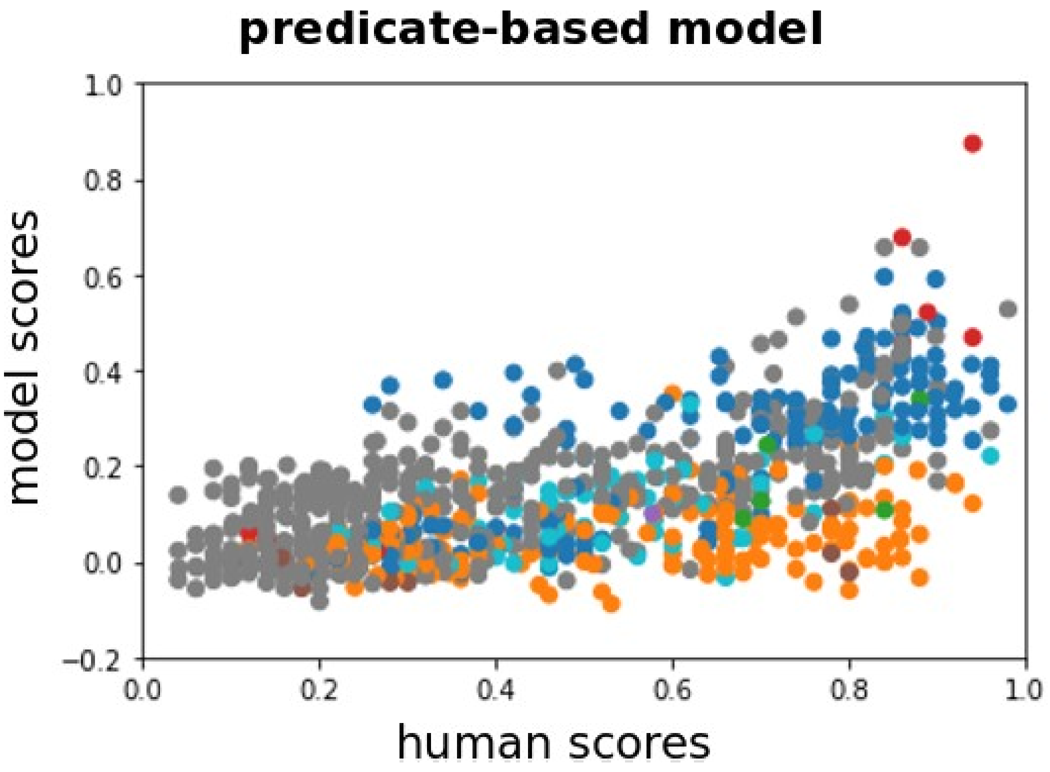

Artist’s impression

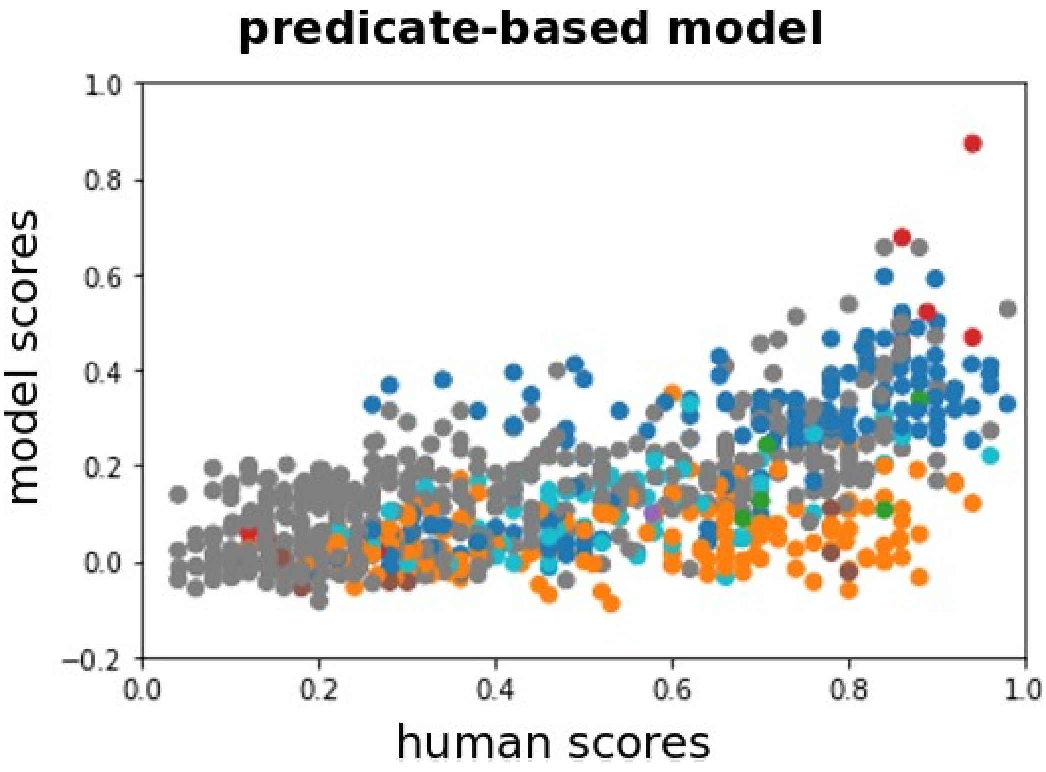

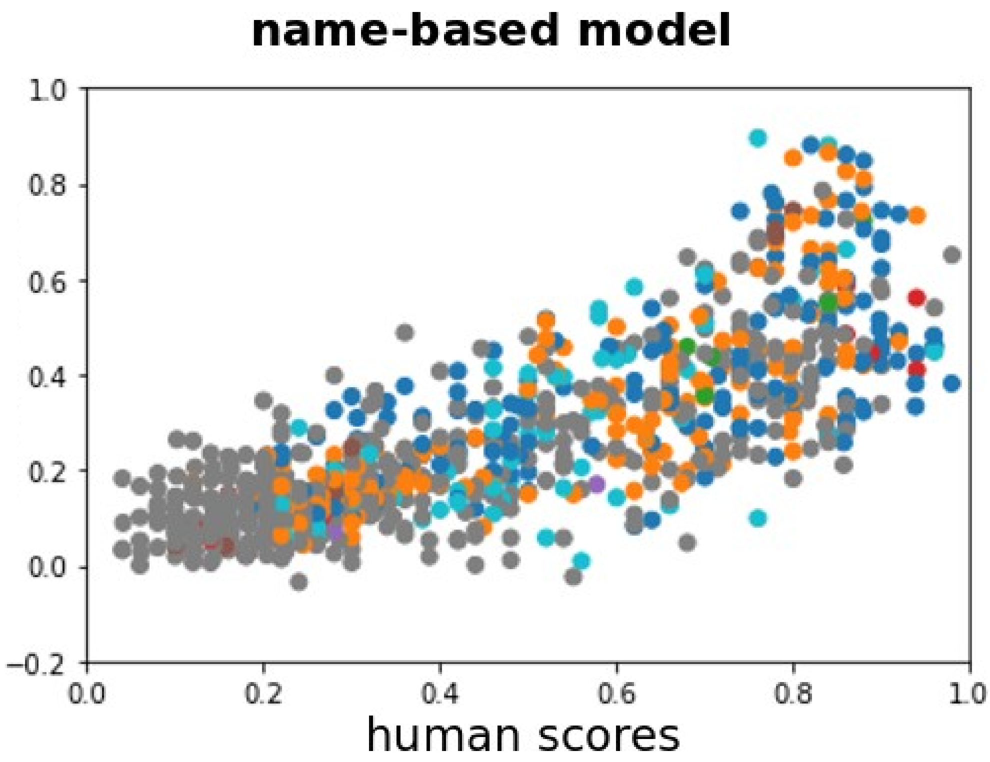

How many names do we need?

How many names do we need?

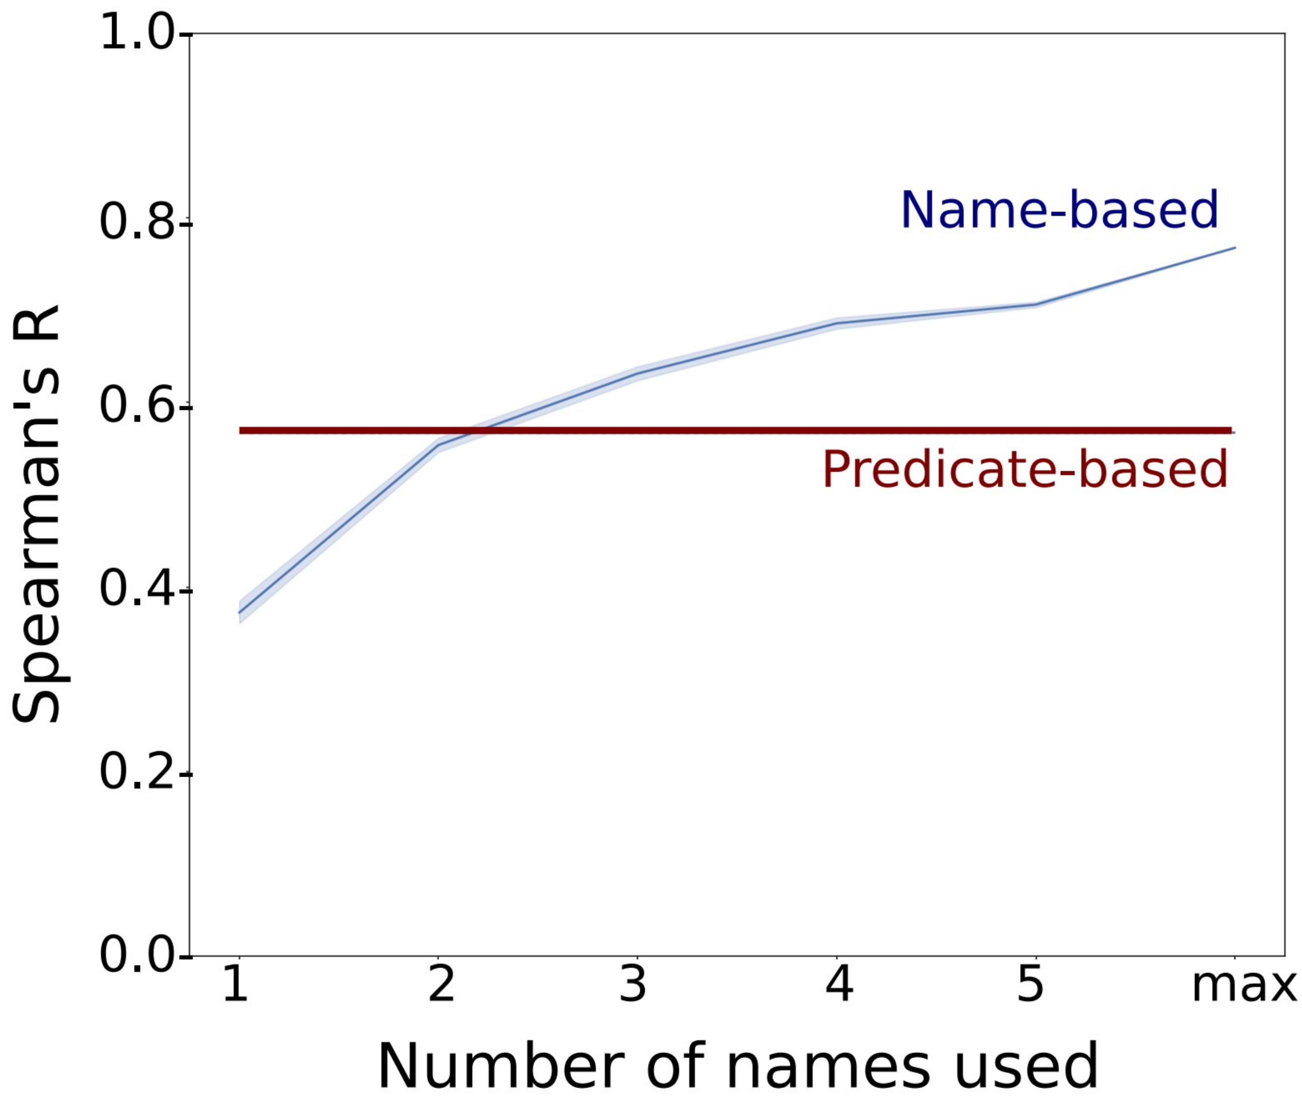

How many names do we need?

Surprisingly few!

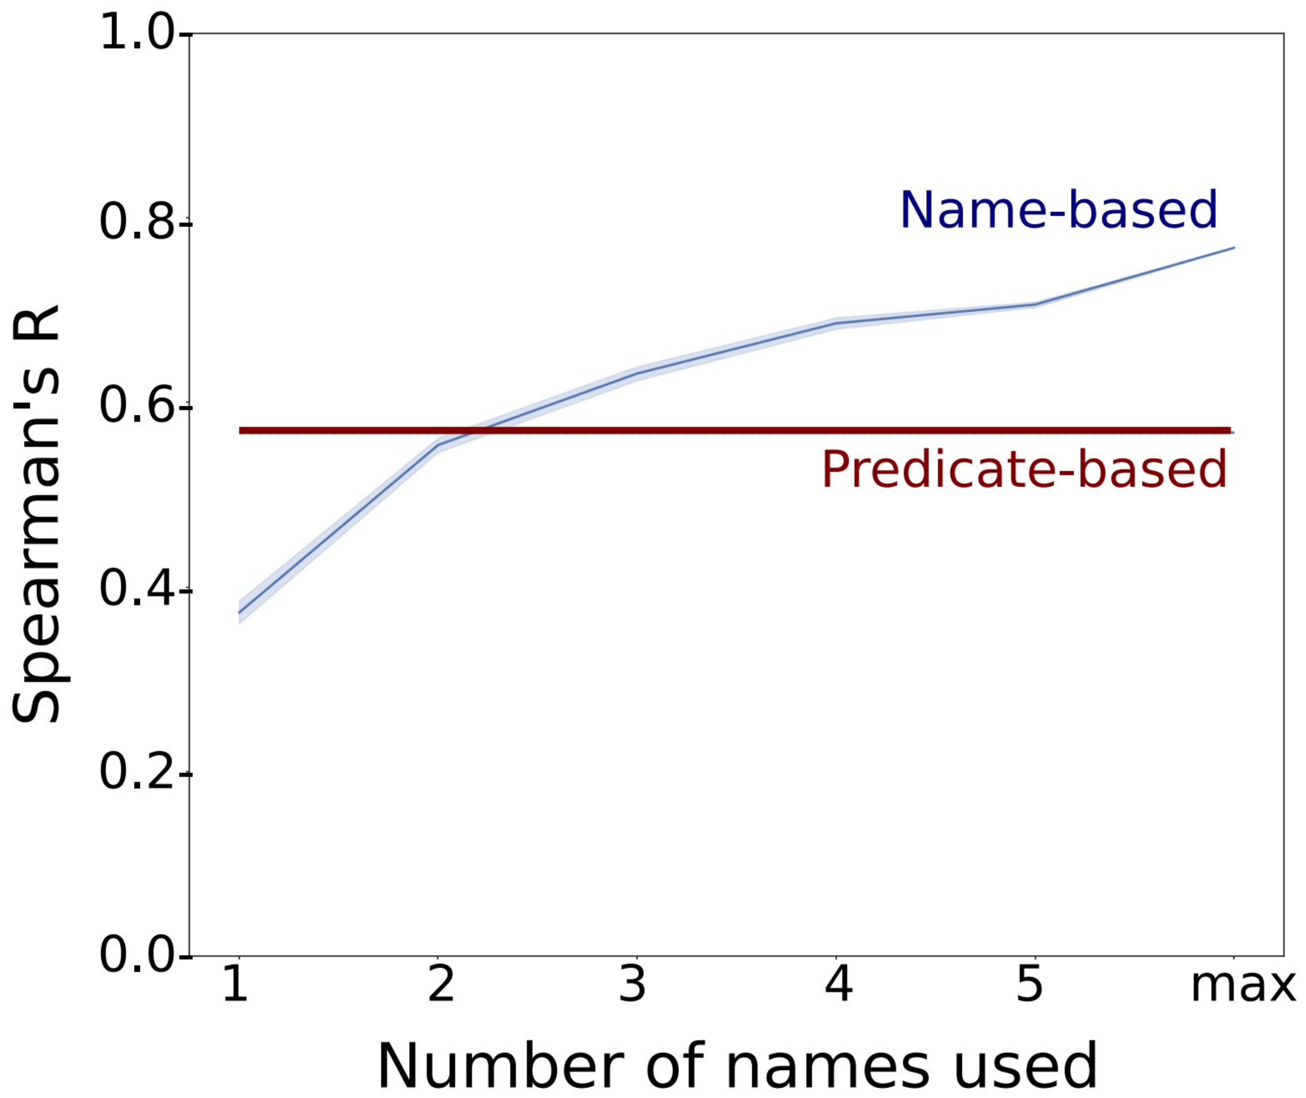

Entities need to be *representative*

Entities need to be *representative*

- E.g., the Name-based model overestimates *surgeon* ~ *siege*...

Entities need to be *representative*

- E.g., the Name-based model overestimates *surgeon* ~ *siege*...
- Instances of surgeon in the Instantiation dataset:
  - William Cowper
  - James Parkinson
  - Alexis Carrel
  - Walter Reed
  - William Beaumont
  - Joseph Lister

Entities need to be *representative*

- E.g., the Name-based model overestimates *surgeon* ~ *siege*...
- Instances of surgeon in the Instantiation dataset:
  - William Cowper
  - James Parkinson
  - Alexis Carrel
  - Walter Reed ← Involved in WW1
  - William Beaumont
  - Joseph Lister

Entities need to be *representative*

- E.g., the Name-based model overestimates *surgeon* ~ *siege*...
- Instances of surgeon in the Instantiation dataset:
  - William Cowper
  - James Parkinson
  - Alexis Carrel
  - Walter Reed ← Involved in WW1
  - William Beaumont
  - Joseph Lister

Members of US military corps

Entities need to be *representative*

- E.g., the Name-based model overestimates *surgeon* ~ *siege*...
- Instances of surgeon in the Instantiation dataset:
  - William Cowper
  - James Parkinson ← Wrote "the siege of chester" (?)
  - Alexis Carrel
  - Walter Reed ← Involved in WW1
  - William Beaumont
  - Joseph Lister

Members of US military corps

Discussion

Discussion

Discussion

- **Main finding:**

Discussion

- **Main finding:**
  - Name-based representations of category concepts align better with ‘the world’ than Predicate-based representations.

Discussion

- **Main finding:**
  - Name-based representations of category concepts align better with ‘the world’ than Predicate-based representations.
  - Even a small number of (representative) names can be enough.

Discussion

- **Main finding:**
  - Name-based representations of category concepts align better with ‘the world’ than Predicate-based representations.
  - Even a small number of (representative) names can be enough.
- **Outlook:**
  - Not every category has *named* instances...

Discussion

- **Main finding:**
  - Name-based representations of category concepts align better with ‘the world’ than Predicate-based representations.
  - Even a small number of (representative) names can be enough.

- **Outlook:**
  - Not every category has *named* instances...
  - NLP relevance? Vs. sense disambiguation? Contextualized word embeddings (ELMo, BERT, …)?

Discussion

- **Main finding:**
  - Name-based representations of category concepts align better with ‘the world’ than Predicate-based representations.
  - Even a small number of (representative) names can be enough.

- **Outlook:**
  - Not every category has *named* instances...
  - NLP relevance? Vs. sense disambiguation? Contextualized word embeddings (ELMo, BERT, …)?
  - Cognitive relevance? E.g., prototype theory?

Acknowledgments

This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 715154). This paper reflects the authors' view only, and the EU is not responsible for any use that may be made of the information it contains.

European Research Council
Established by the European Commission

Universitat Pompeu Fabra Barcelona

Computational Linguistics and Linguistic Theory

Image sources

https://ui-ex.com/explore/whale-transparent-dark/

https://commons.wikimedia.org/wiki/File:Cowicon.svg

https://commons.wikimedia.org/wiki/File:Bird_1010720_drawing.svg

https://commons.wikimedia.org/wiki/File:Dog_silhouette.svg

https://commons.wikimedia.org/wiki/File:Cat_silhouette_darkgray.svg

https://commons.wikimedia.org/wiki/File:Frog_(example).svg

https://commons.wikimedia.org/wiki/File:PeregrineFalconSilhouettes.svg

https://commons.wikimedia.org/wiki/File:Common_goldfish_silhouette.svg

https://commons.wikimedia.org/wiki/File:Six_weeks_old_cat_(aka).jpg

https://nl.m.wikipedia.org/wiki/Bestand:Kooikerhondje_puppy.jpg

https://nl.m.wikipedia.org/wiki/Bestand:Golden_Retriever_eating_crust_of_pizza.jpg

https://commons.wikimedia.org/wiki/File:Cat-eating-prey.jpg

Where are predicates and names, anyway?

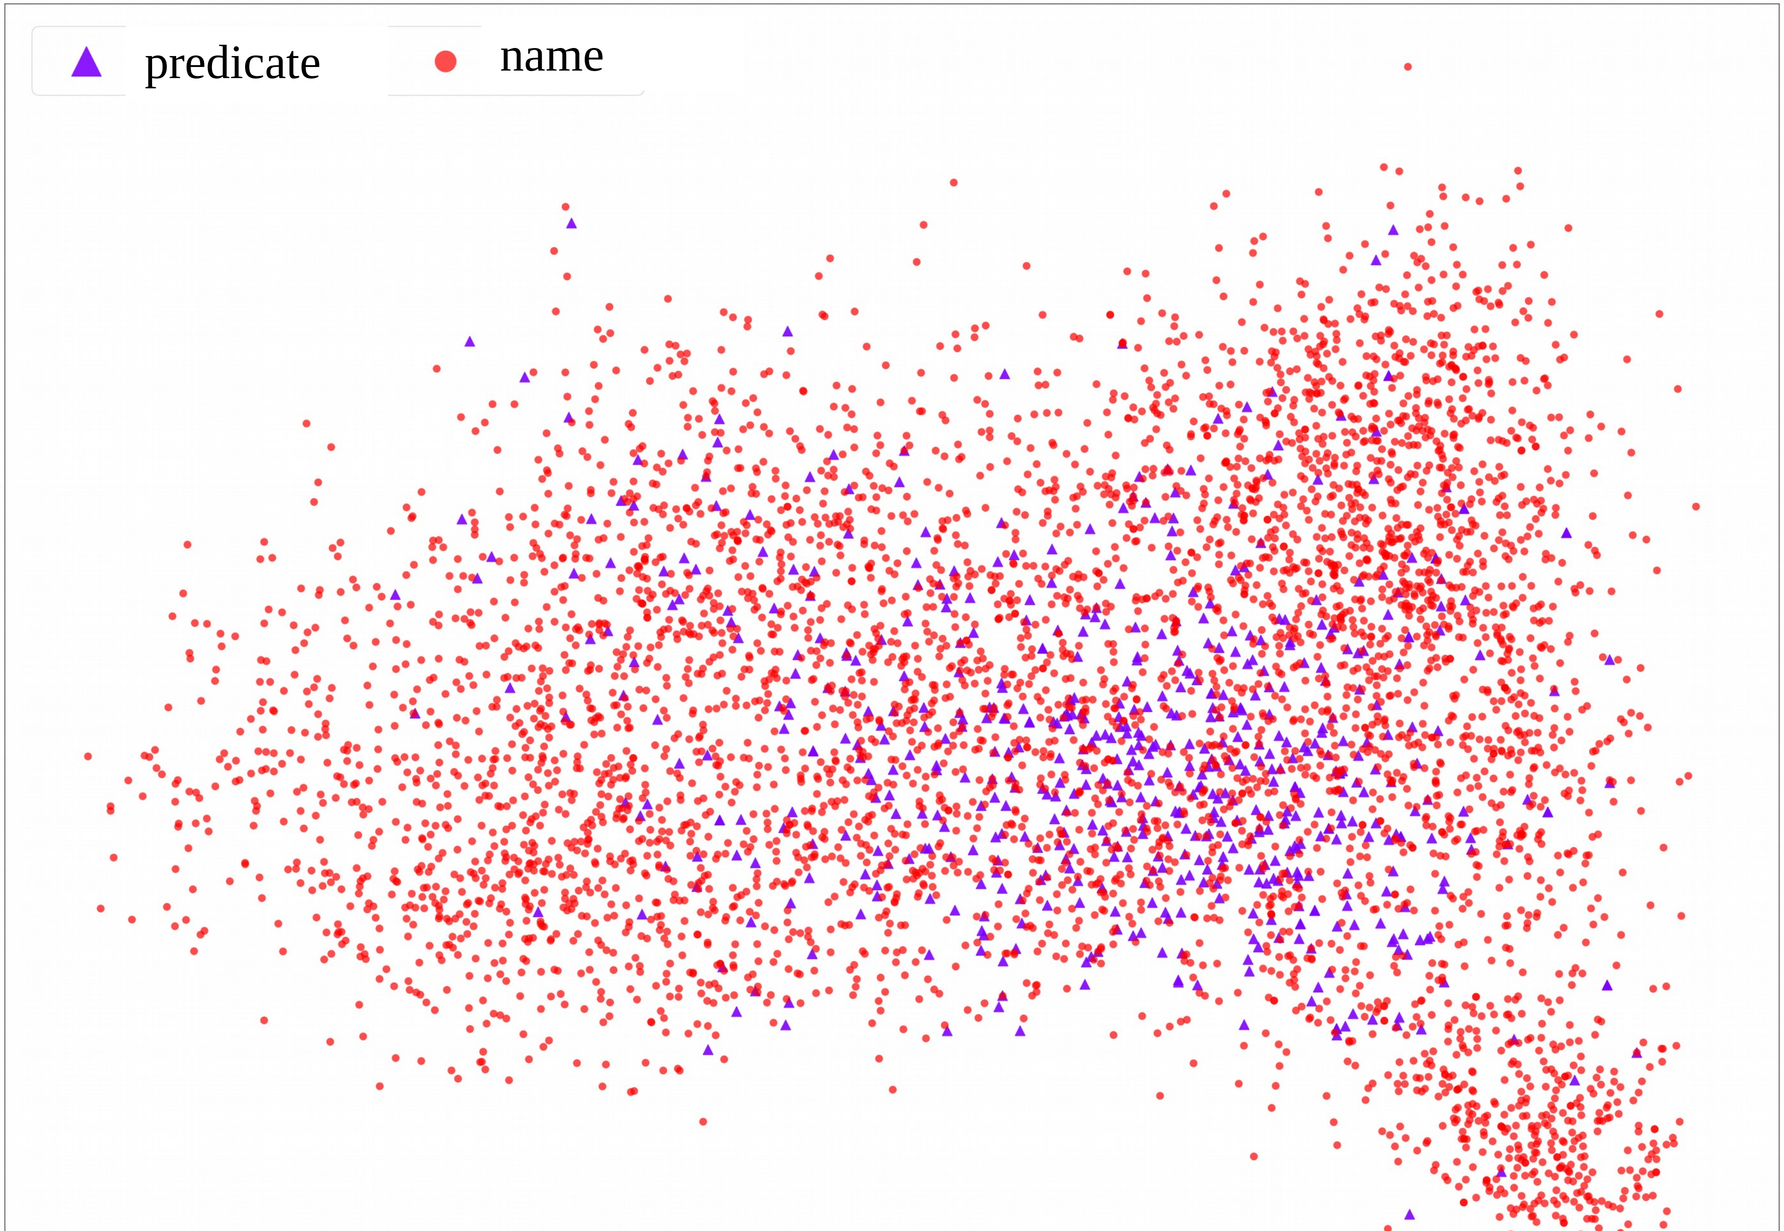

Where are predicates and names, anyway?

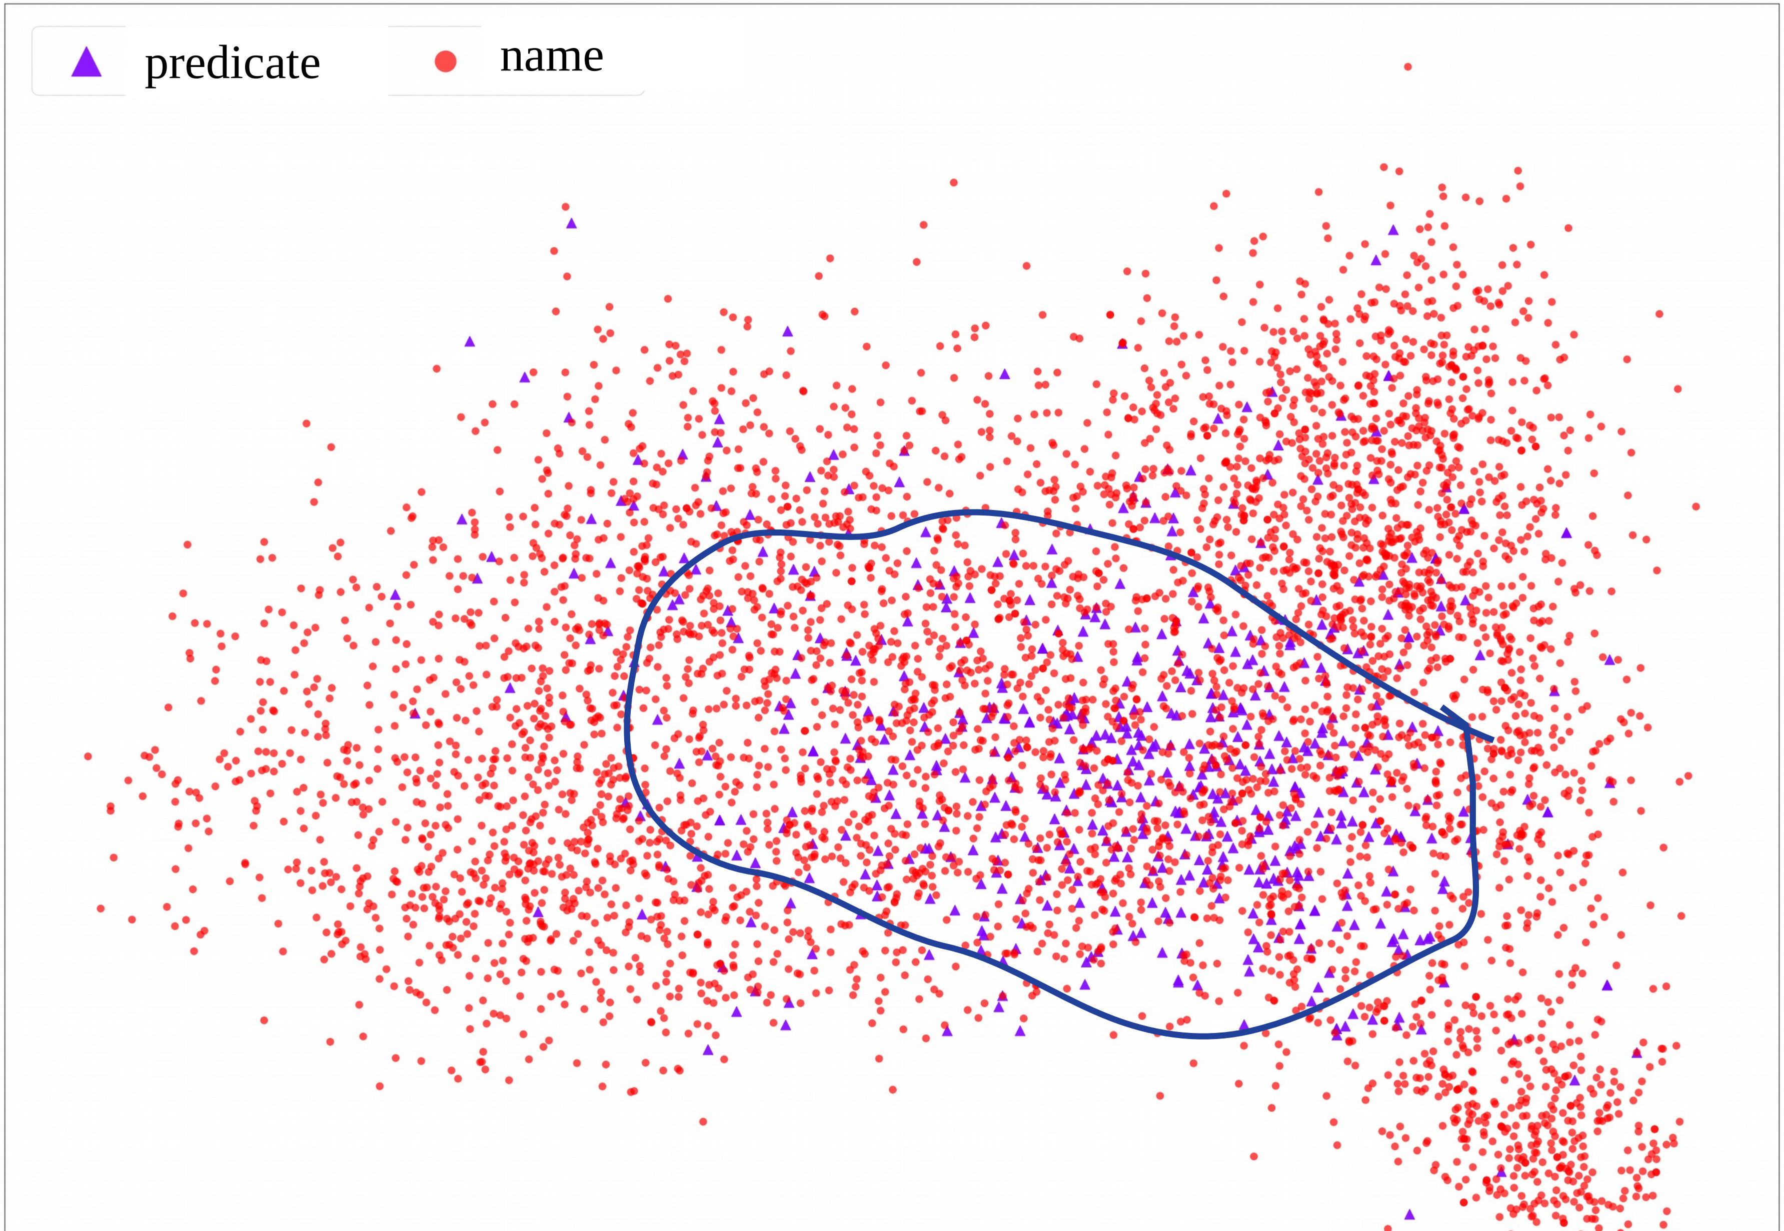

Crowdsource task

In this HIT you will see 70 items like the following, each presenting two pairs of categories:

Which pair of categories are more related to each other?

1. wheel ↔ car

2. building ↔ crane *(type of bird)*

Crowdsource task

In this HIT you will see 70 items like the following, each presenting two pairs of categories:

Which pair of categories are more related to each other?

1. wheel ↔ car

2. building ↔ crane *(type of bird)*

Crowdsource task instructions

Be careful: words in this HIT can sometimes refer to multiple categories! For instance, "crane" could mean a lifting machine or a type of bird. In this case, we mean the type of bird, and you should answer accordingly.

In this example you would probably choose pair 1, because the categories *wheel* and *car* seem more closely related than the categories *building* and (the type of bird!) *crane*.

Don't know the meaning of a word? Use your mouse to hover over a word to see its definitions.

Crowdsource task instructions

Be careful: words in this HIT can sometimes refer to multiple categories! For instance, "crane" could mean a lifting machine or a type of bird. In this case, we mean the type of bird, and you should answer accordingly.

In this example you would probably choose pair 1, because the categories *wheel* and *car* seem more closely related than the categories *building* and (the type of bird!) *crane*.

Don't know the meaning of a word? Use your mouse to hover over a word to see its definitions.

Why definitions?

Why definitions?

- The same words can often be used to denote various categories.

Why definitions?

- The same words can often be used to denote various categories.
- To properly evaluate the Name-based approach, the human judgments should be about the categories as intended by the Instantiation dataset we use.

Why definitions?

- The same words can often be used to denote various categories.
- To properly evaluate the Name-based approach, the human judgments should be about the categories as intended by the Instantiation dataset we use.
- (Would be good practice more generally – e.g., vs. the *good subject effect.*)

Why definitions?

- The same words can often be used to denote various categories.
- To properly evaluate the Name-based approach, the human judgments should be about the categories as intended by the Instantiation dataset we use.
- (Would be good practice more generally – e.g., vs. the *good subject effect.*)

- This may give the Predicate-based approach a disadvantage…

Why definitions?

- The same words can often be used to denote various categories.
- To properly evaluate the Name-based approach, the human judgments should be about the categories as intended by the Instantiation dataset we use.
- (Would be good practice more generally – e.g., vs. the *good subject effect.*)

- This may give the Predicate-based approach a disadvantage…
  - but this disadvantage is not an *unfair* one.

A closer look per ontological domain

A closer look per ontological domain

Predicate-based:

	location	person	object	OTHER
location	**.49**	.46	.44	-
person	-	.68	-	-
object	-	**.49**	.25	-
OTHER	**.52**	.78	.24	.64

A closer look per ontological domain

Predicate-based:

	location	person	object	OTHER
location	**.49**	.46	.44	-
person	-	.68	-	-
object	-	**.49**	.25	-
OTHER	**.52**	.78	.24	.64

A closer look per ontological domain

Predicate-based:

	location	person	object	OTHER
location	**.49**	.46	.44	-
person	-	.68	-	-
object	-	**.49**	.25	-
OTHER	**.52**	.78	.24	.64

Name-based:

	location	person	object	OTHER
location	.35	.44	**.79**	-
person	-	.67	-	-
object	-	.30	**.83**	-
OTHER	.37	.75	**.78**	.64

A closer look per ontological domain

Predicate-based:

	location	person	object	OTHER
location	**.49**	.46	.44	-
person	-	.68	-	-
object	-	**.49**	.25	-
OTHER	**.52**	.78	.24	.64

Name-based:

	location	person	object	OTHER
location	.35	.44	**.79**	-
person	-	.67	-	-
object	-	.30	**.83**	-
OTHER	.37	.75	**.78**	.64

Non-representative instances of ‘object’ categories

- capital: belfast, bridgetown, camelot, cardiff, edinburgh, george_town
- colony: cayman_islands, connecticut, delaware, demerara, georgia, gibraltar, maryland, massachusetts_bay_colony, new_amsterdam, new_hampshire, new_jersey, new_netherland, new_york, north_carolina, pennsylvania, plymouth_colony, rhode_island, rock_of_gibraltar, south_carolina, virginia (most entities used to be colonies, but no longer are.)
- region: achaea, far_east, french_west_indies, kennelly-heaviside_layer, occident, old_world, rand, transylvania, west, witwatersrand
- district: acadia, acre, american_samoa, aragon, attica, boeotia, castilla, catalonia, darfur, east_malaysia, galloway, kwazulu-natal, lake_district, louisiana_purchase, mount_athos, north_borneo, northern_mariana_islands, northern_territory, northwest_territories, nunavut, palatinate, papal_states, sarawak, yukon (I suspect US people will interpret ’district’ as a part of a city, rather than a part of a country?)

A closer look per ontological domain

Predicate-based:

	location	person	object	OTHER
location	**.49**	.46	.44	-
person	-	.68	-	-
object	-	**.49**	.25	-
OTHER	**.52**	.78	.24	.64

Name-based:

	location	person	object	OTHER
location	.35	.44	**.79**	-
person	-	.67	-	-
object	-	.30	**.83**	-
OTHER	.37	.75	**.78**	.64

A closer look per ontological domain

Predicate-based:

	location	person	object	OTHER
location	.54	.46	.44	-
person	-	.68	-	-
object	-	**.49**	.25	-
OTHER	**.52**	.78	.24	.64

Name-based:

	location	person	object	OTHER
location	**.64**	.44	**.79**	-
person	-	.67	-	-
object	-	.30	**.83**	-
OTHER	.37	.75	**.78**	.64