What's in a name? or, what recent results in computational semantics tell us about ambiguity

Gemma Boleda
Universitat Pompeu Fabra
Joint work in progress with Sebastian Padó and Matthijs Westera

Ambiguity: Perspectives on Representation and Resolution Workshop
ESSLLI 2018, 610 August 2018, Sofia, Bulgaria

European Research Council
Established by the European Commission
Supporting top researchers from anywhere in the world

European Commission
Horizon 2020
European Union funding for Research & Innovation

Universitat Pompeu Fabra Barcelona

Acknowledgements

- Co-authors in cited papers, esp. Abhijeet Gupta
- AMORE team: Laura Aina, Kristina Gulordava, Carina Silberer, Ionut-Teodor Sorodoc

- This project has received funding from the European Research Council (ERC) under the European Unions Horizon 2020 research and innovation programme (grant agreement No 715154), AMORE: A distributional Model Of Reference to Entities).

The issue

- words vs. ontological categories: *dog* vs. DOG

The issue

- words vs. ontological categories: *dog* vs. DOG

The (naive) referentialist view

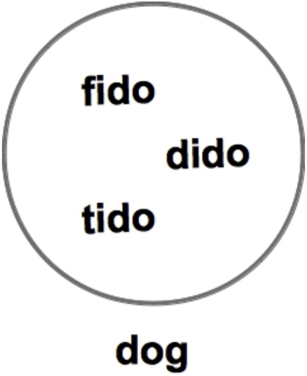

$$\lambda x.dog(x)$$

The (naive) referentialist view

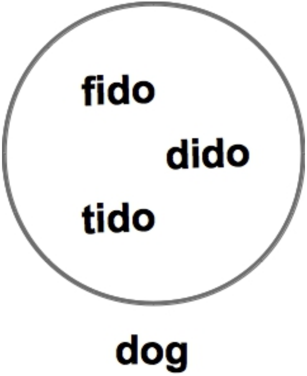

$$\lambda x.p786(x)$$

how do we define the set 'dog'?

- DOG?
- whatever humans call *dog*?

Empirical issues

- also: referentialist very quickly runs into empirical problems
- e.g. composition: *red car / face / wine* (Kennedy & McNally 2010, Boleda 2007, Bruni et al. 2012)
- $\lambda x.red(x)$?
- red_1, red_2, red_3?

Empirical issues

- also: referentialist very quickly runs into empirical problems
- e.g. composition: *red car / face / wine* (Kennedy & McNally 2010, Boleda 2007, Bruni et al. 2012)
- $\lambda x.red(x)$?
- red_1, red_2, red_3?
- NO! (Pustejovsky 1995, Kilgarriff 1995, a.o.)

But not everything goes

> 'There's glory for you!' 'I don't know what you mean by "glory,"' Alice said. Humpty Dumpty smiled contemptuously. 'Of course you don't — till I tell you. I meant "there's a nice knock-down argument for you!"' 'But "glory" doesn't mean "a nice knock-down argument,"' Alice objected. 'When *I* use a word,' Humpty Dumpty said in rather a scornful tone, 'it means just what I choose it to mean — neither more nor less.' 'The question is,' said Alice, 'whether you CAN make words mean so many different things.' 'The question is,' said Humpty Dumpty, 'which is to be master — that's all.'

Lewis Carroll, *Through the Looking Glass*.

Today

Empirical take on the relationship between words and ontological categories; evidence from computational semantics

What does this have to do with ambiguity?

- many ambiguities arise from the (often implicit)
  - *(controversial?)* equation between words' meanings and ontological categories
  - *(softer)* privileged role of ontological information in delimiting meanings

What does this have to do with ambiguity?

- many ambiguities arise from the (often implicit)
  - *(controversial?)* equation between words' meanings and ontological categories
  - *(softer)* privileged role of ontological information in delimiting meanings
- *Frodo*: fictional character, flesh and blood hobbit? (Semeijn, BRIDGE workshop, ESSLLI 2018)

What does this have to do with ambiguity?

- many ambiguities arise from the (often implicit)
  - *(controversial?)* equation between words' meanings and ontological categories
  - *(softer)* privileged role of ontological information in delimiting meanings
- *Frodo*: fictional character, flesh and blood hobbit? (Semeijn, BRIDGE workshop, ESSLLI 2018)
- *fall*: *The glass fell / Christmas day falls on a Sunday*; "64 attested truth conditionally non equivalent readings of the verb fall" (Zeevat, Ambiguity workshop, ESSLLI 2018)

What does this have to do with meaning?

- semantic change (Sagi et al. GEMS 2009:106):
  - broadening: "Late Old English *docga* 'a (specific) powerful breed of dog' $\rightarrow$ *dog* 'any member of the species *Canis familiaris*"
  - narrowing: "Old English *deor* 'animal' $\rightarrow$ *deer*"

  $\rightarrow$ subclass - superclass changes in an ontological hierarchy
- question of meaning linked to question of how humans categorize the world

Exploring categories with distributional semantics

Gupta, Boleda, Padó submitted; https://arxiv.org/abs/1808.01662

- entities, represented from names (4,750)
- predicates expressing ontological categories (common nouns or phrases; 577)
- instantiation relations between entities and categories extracted from knowledge base:

  Nile - RIVER
  George Washington - PRESIDENT OF THE U.S.
  Boccaccio - POET
  Iliad - EPOS
  Alamo - SIEGE

Exploring categories with distributional semantics

Gupta, Boleda, Padó submitted; https://arxiv.org/abs/1808.01662

Plan

- (I) brief general exploration
- (II) experiment: find out whether an entity instantiates a category
  Nile - RIVER?
- based on two definitions of "category":
  - predicate-based:
    RIVER = *river*
  - instance-based:
    RIVER = average of *Nile*, *Danube*, *Mekong*, *Orinoco*, ...
- method: Machine Learning
- representation: distributional semantics

Brief intro to Distributional semantics

Aka vector-space semantics, related to Neural Networks / deep learning

(See Stefan Evert's course this week at ESSLLI for more!)

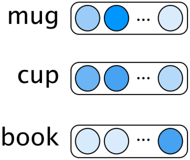

Brief intro to Distributional semantics

Aka vector-space semantics, related to Neural Networks / deep learning

(See Stefan Evert's course this week at ESSLLI for more!)

```
likely) mug of bourbon in hand. Some
stewed milk into a heavy mug, granules of
holding his coffee mug cupped in his hands.
drained his mug, dropping it over his
tablespoons of coffee and a single mug of
milk into the mug plus four spoons of sugar
placing the empty mug on the floor
picking up my mug with one hand and
followed by a very hot mug of tea into which
from time to time to drink a mug of tea. The
briefed, relax over a mug of tea and a
cake and cheese and a mug of strong, black
then we had a mug of cocoa and a gingerbread
and a white mug with a blurred inscription.
was carrying a mug of tea and
```

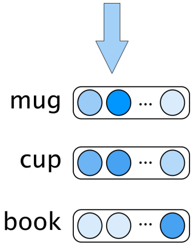

Brief intro to Distributional semantics

Aka vector-space semantics, related to Neural Networks / deep learning

(See Stefan Evert's course this week at ESSLLI for more!)

```
likely) mug of bourbon in hand. Some
stewed milk into a heavy mug, granules of
holding his coffee mug cupped in his hands.
drained his mug, dropping it over his
tablespoons of coffee and a single mug of
milk into the mug plus four spoons of sugar
placing the empty mug on the floor
picking up my mug with one hand and
followed by a very hot mug of tea into which
from time to time to drink a mug of tea. The
briefed, relax over a mug of tea and a
cake and cheese and a mug of strong, black
then we had a mug of cocoa and a gingerbread
and a white mug with a blurred inscription.
was carrying a mug of tea and
```

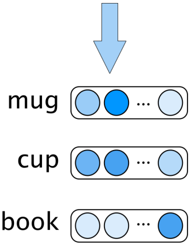

reasonable proxy for lexical semantics (shown in a lot of work in Cognitive Science, Computational Linguistics)

Words as vectors

Landauer and Dumais 1997, Baroni and Lenci 2010, Erk 2012

	runs	*sleeps*
dog	1	4
cat	1	5
car	4	0

Words as vectors

Landauer and Dumais 1997, Baroni and Lenci 2010, Erk 2012

	runs	*sleeps*
dog	1	4
cat	1	5
car	4	0

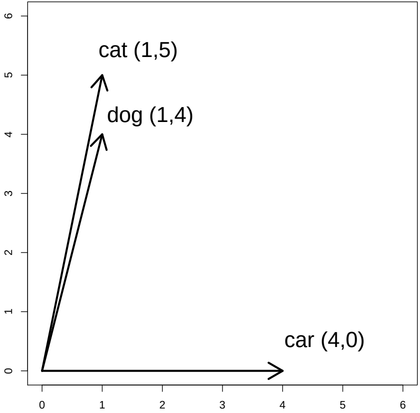

Based on material by Marco Baroni

The ugly reality

dog 0.0067840000000000001 -0.083669999999999994 -0.0276
0.15977 -0.051539000000000001 0.25880999999999998 0.128604
0.043097000000000003 0.022886 0.16512099999999999
-0.19905800000000001 -0.11175599999999999 0.011864
-0.20073099999999999 0.168099 -0.146171 0.244815
-0.31515599999999999 0.012591 -0.099188999999999999
0.011284000000000001 0.15019299999999999
0.075329999999999994 -0.23896400000000001
0.032051999999999997 0.24129900000000001 0.058816
-0.38864799999999999 0.099677000000000002 0.183504
-0.018511 0.123728 0.19941200000000001 -0.191748
-0.019918000000000002 -0.101323 -0.029946
-0.0053169999999999997 -0.007123 0.082957000000000003
-0.087373000000000006 0.272984 0.026393 0.124167 0.231517
-0.242756 -0.173259 -0.089765999999999999 0.204042 -0.017602
...

Representation of *dog* in the space of Baroni et al. (2014).

Words as vectors

Landauer and Dumais 1997, Baroni and Lenci 2010, Erk 2012

	runs	*legs*
dog	1	4
cat	1	5
car	4	0

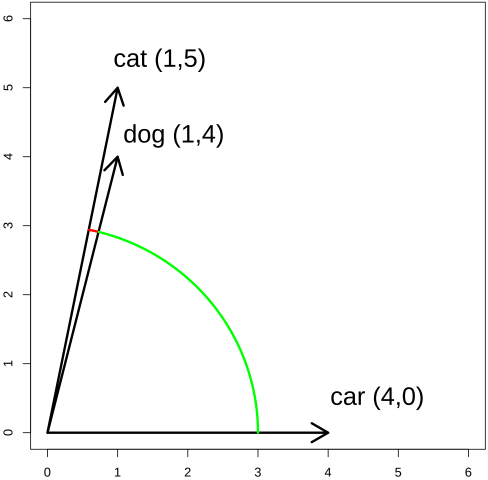

Words as vectors

Landauer and Dumais 1997, Baroni and Lenci 2010, Erk 2012

	runs	*legs*
dog	1	4
cat	1	5
car	4	0

similarity:

- *dog* - *cat*: 0.99
- *dog* - *car*: 0.20

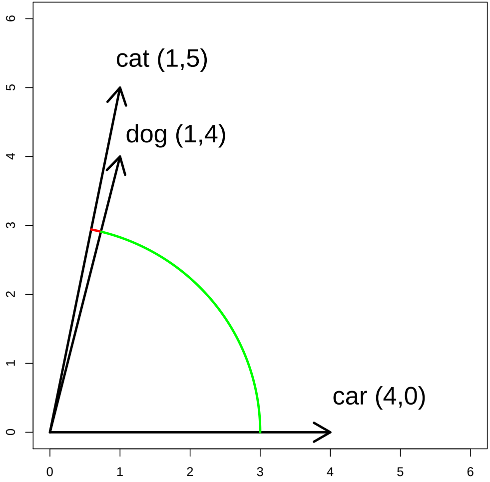

Words as vectors

Landauer and Dumais 1997, Baroni and Lenci 2010, Erk 2012

	runs	*legs*
dog	1	4
cat	1	5
car	4	0

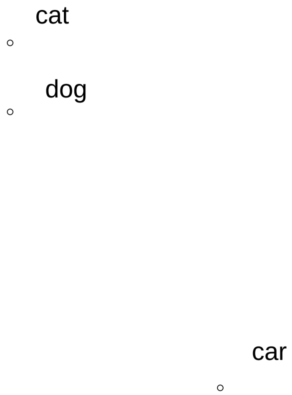

similarity:

- *dog* - *cat*: 0.99
- *dog* - *car*: 0.20

Words as vectors

Landauer and Dumais 1997, Baroni and Lenci 2010, Erk 2012

	runs	*legs*
dog	1	4
cat	1	5
car	4	0

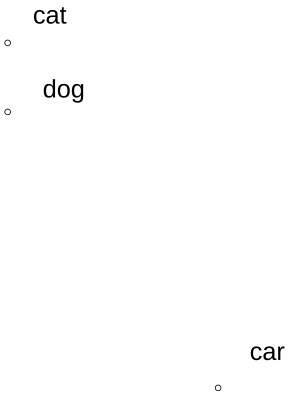

similarity:

- *dog* - *cat*: 0.99
- *dog* - *car*: 0.20

nearest neighbor

Words as vectors

Multidimensional; continuous; distributed; not (only) truth-conditional

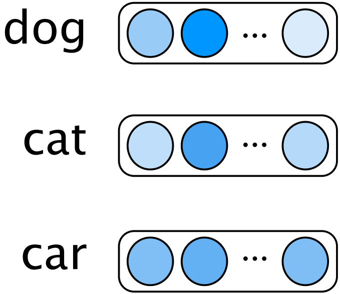

Meaning in distributional semantics
Boleda and Erk 2015

man
woman
gentleman
gray-haired
boy
person
lad
men
girl

Words most similar to *man* in Baroni et al. (2014)

Meaning in distributional semantics

Boleda and Erk 2015

man
woman
gentleman
gray-haired
boy
person
lad
men
girl

+HUMAN

Words most similar to *man* in Baroni et al. (2014)

Meaning in distributional semantics

Boleda and Erk 2015

man
woman
gentleman
gray-haired
boy
person
lad
men
girl

+HUMAN +MALE

Words most similar to *man* in Baroni et al. (2014)

Meaning in distributional semantics

Boleda and Erk 2015

man
woman
gentleman
gray-haired
boy
person
lad
men
girl

+HUMAN +MALE +ADULT

Words most similar to *man* in Baroni et al. (2014)

Meaning in distributional semantics

Boleda and Herbelot 2016

man	chap	lad	dude	guy
woman	bloke	boy	freakin'	bloke
gentleman	guy	bloke	woah	chap
gray-haired	lad	scouser	dorky	doofus
boy	fella	lass	dumbass	dude
person	man	youngster	stoopid	fella

Words most similar to *man*, *chap*, *lad*, *dude*, *guy* in Baroni et al. (2014).

Meaning in distributional semantics

Boleda and Herbelot 2016

man	chap	lad	dude	guy
woman	bloke	boy	freakin'	bloke
gentleman	guy	bloke	woah	chap
gray-haired	lad	scouser	dorky	doofus
boy	fella	lass	dumbass	dude
person	man	youngster	stoopid	fella

Words most similar to *man*, *chap*, *lad*, *dude*, *guy* in Baroni et al. (2014).

Exploring categories with distributional semantics

- distributional semantics has been chiefly used to explore predicates
- we use it also for entities

- some evidence that they are also reasonable
  - Mikolov et al. 2013, Gupta et al. EMNLP 2015, Herbelot and Vecchi EMNLP 2015, Boleda et al. EACL 2017

Exploring categories with distributional semantics

Gupta, Boleda, Padó submitted; https://arxiv.org/abs/1808.01662

- entities
- predicates
- instantiation relations between entities and categories

  Nile - RIVER
  George Washington - PRESIDENT OF THE U.S.
  Boccaccio - POET
  Iliad - EPOS
  Alamo - SIEGE

Names and predicates in distributional space

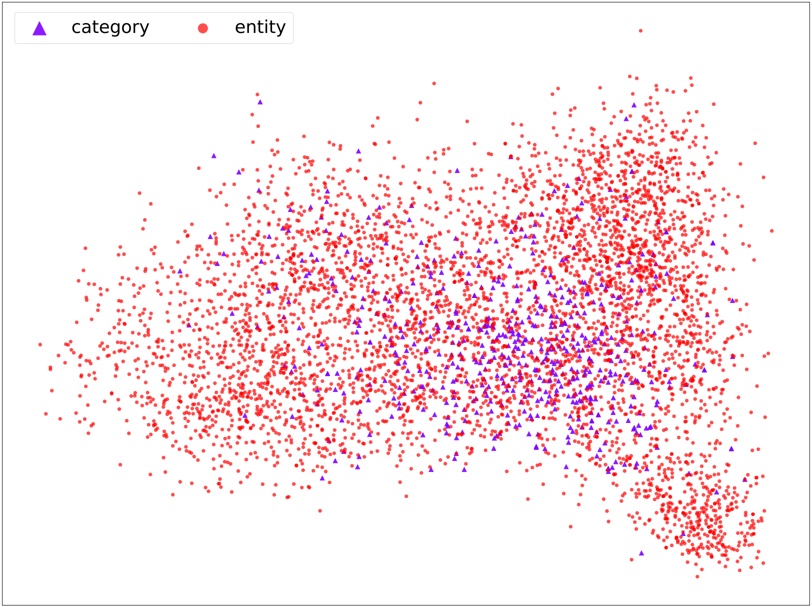

Names and predicates in distributional space

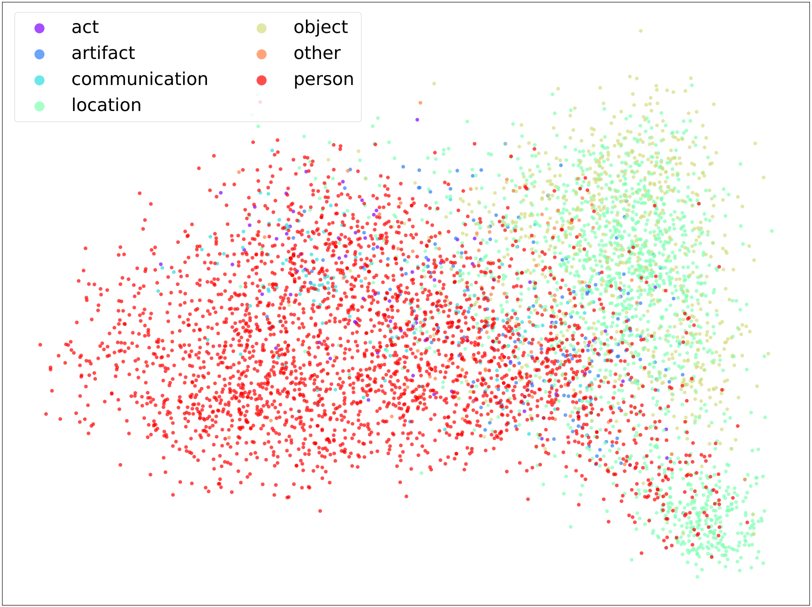

Names and predicates in distributional space

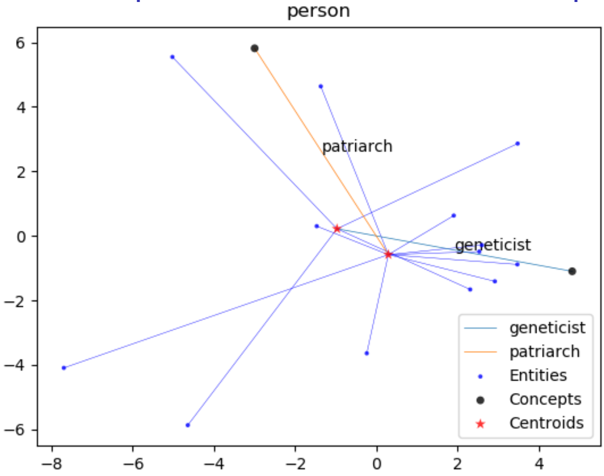

- entity-predicate (*Nile* - *river*): 0.16
- entity-entity, within a category (*Nile* - *Po*): **0.22**

Experiments

Task: decide if an entity instantiates a category:

Nile - RIVER

Method: Machine Learning

3-slide intro to Machine Learning (I)

Task: decide if a student passes the course

- Training phase: expose computer to examples
  - representation: *grade, pass?*

    9.1, yes

    2.2, no

    4.5, no

    8.3, yes

    3.4, no

    ...
  - computer learns to associate the representations of an entity with whether it passed

→ Output: computational model:

default: no; if grade $>$ *4.5: yes*

3-slide intro to Machine Learning (II)

Task: decide if a student passes the course

- Training phase: expose computer to examples, build computational model

  default: no; if grade $>$ 4.5: yes

- Testing phase: ask computer to make decisions for new examples
  1) applies computational model

  7.1?

3-slide intro to Machine Learning (II)

Task: decide if a student passes the course

- Training phase: expose computer to examples, build computational model

  default: no; if grade $>$ 4.5: yes

- Testing phase: ask computer to make decisions for new examples
  1) applies computational model

  7.1? yes
  2.2?

3-slide intro to Machine Learning (II)

Task: decide if a student passes the course

- Training phase: expose computer to examples, build computational model

  default: no; if grade $>$ 4.5: yes

- Testing phase: ask computer to make decisions for new examples

  1) applies computational model

     7.1? yes
     2.2? no
     4.6?

3-slide intro to Machine Learning (II)

Task: decide if a student passes the course

- Training phase: expose computer to examples, build computational model

  default: no; if grade $>$ 4.5: yes

- Testing phase: ask computer to make decisions for new examples
  1) applies computational model

     7.1? yes
     2.2? no
     4.6? yes

3-slide intro to Machine Learning (II)

Task: decide if a student passes the course

- Training phase: expose computer to examples, build computational model

  default: no; if grade $>$ 4.5: yes

- Testing phase: ask computer to make decisions for new examples
  1) applies computational model
  2) (for research/development) compares to ground truth

     7.1? yes

     2.2? no

     4.6? yes

3-slide intro to Machine Learning (II)

Task: decide if a student passes the course

- Training phase: expose computer to examples, build computational model

  default: no; if grade $>$ 4.5: yes

- Testing phase: ask computer to make decisions for new examples
  1) applies computational model
  2) (for research/development) compares to ground truth

     7.1? yes yes OK!

     2.2? no no OK!

     4.6? yes no **OOPS**

3-slide intro to Machine Learning (II)

Task: decide if a student passes the course

- Training phase: expose computer to examples, build computational model

  default: no; if grade $>$ 4.5: yes

- Testing phase: ask computer to make decisions for new examples
  1) applies computational model
  2) (for research/development) compares to ground truth

     7.1? yes yes OK!

     2.2? no no OK!

     4.6? yes no **OOPS**

  $\rightarrow$ see performance, analyze mistakes, ...
  - 2 out of 3 testing examples are right; accuracy = 66%

3-slide intro to Machine Learning (III)

- our example: 1 feature (*grade*)
- real-life Machine Learning: dozens to hundreds of thousands of features
  - here: distributional semantic representations
- much more complex models
  - logistic regression
  - deep learning models
  - ...
- neither representations nor computational models directly inspectable :/
- researchers need to do work to make both interpretable

Experiments

Task: decide if an entity instantiates a category:

Nile - RIVER

Method: Machine Learning

Experiments

Task: decide if an entity instantiates a category:

Nile - RIVER

Method: Machine Learning

- Training phase: expose computer to examples

  George Washington - PRESIDENT OF THE U.S., yes

  Nile - SEA, no
- Representation: Distributional semantics
- Testing phase: ask computer to decide whether new examples are in an instantiation relationship

  Boccaccio - POET?
- we reuse neither entities nor categories

Experiments: Two definitions of "category"

- based on predicates:
  $\textsc{River} = \textit{river}$
- based on instances:
  $\textsc{River}$ = average of *Nile*, *Danube*, *Mekong*, *Orinoco*, ...

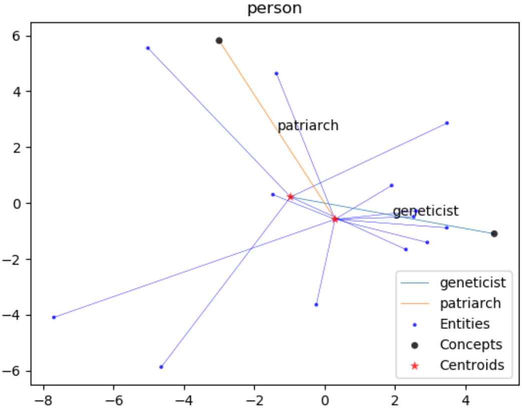

Results

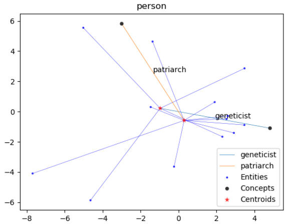

	predicates	instances
Nile - LAWYER (easy)	0.85	**0.90**
Nile - SEA (hard)	0.67	**0.79**

Results

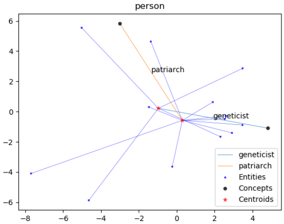

	predicates	instances
Nile - LAWYER (easy)	0.85	**0.90**
Nile - SEA (hard)	0.67	**0.79**

→ Instantiation is easier to model if we define categories in terms of their **instances** than in terms of **predicates** lexicalizing those categories.

Discussion

- unlink words from ontological categories?

$\rightarrow$ against referentialist view

Discussion

- however, confounder: entity representations from names only: *George Washington*, *Nile*
- is the observed difference due to the proper noun / common noun distinction instead?

→ test 1: build representations based on all kinds of referring expressions

→ test 2: repeat experiment with hypernymy (work in progress)

Converging evidence

Failure of the Distributional Inclusion Hypothesis

Lexical entailment

- e.g. $\forall x, dog_1(x) \models animal_2(x)$

Geffet and Dagan (2005) translate into Distributional Inclusion Hypothesis:

- *animal* can occur in all the contexts in which *dog* can occur, plus some contexts in which *dog* cannot:
  - *The animal/dog slept*
  - *animal/#dog rights*

$\rightarrow$ If so, we can use (distributional) feature inclusion as a clue to lexical entailment.

Failure of the Distributional Inclusion Hypothesis

Roller, Erk, Boleda COLING 2014

Roller et al. (2014): hypernymy detection

- method: Machine Learning
  - Training phase: *dog - animal*, *dog - plant*
  - Testing phase: *cat - mammal*?
- unsupervised method: use the Distributional Inclusion Hypothesis directly
  - say "yes" if *animal* features are included in *dog* features
    - *The animal/dog slept*
    - *animal/#dog rights*

$\rightarrow$ doesn't work

Failure of the Distributional Inclusion Hypothesis

Roller, Erk, Boleda COLING 2014

- supervised method: we let the machine figure out how to use the representations
  - what is the relationship between the representations of *animal* and *dog*?
  - → works, but doesn't use feature inclusion
  - domain-dependent features: occurrence with *extinct, fauna, species*, ...
  - computer learns that certain features characterize hypernyms in a domain, not the hypernymy relation (Levy et al. 2015)

→ The contexts of use of a hypernym are not a superset of the contexts of use of a hypernym.

Discussion

- hyper/hyponyms are not interchangeable
- we could truthfully say *the animal sleeps* whenever we say *the dog sleeps*, but we just don't
- (hypothesis) we use *dog* and *animal* for different facets of entities

Discussion

- hyper/hyponyms are not interchangeable
- we could truthfully say *the animal sleeps* whenever we say *the dog sleeps*, but we just don't
- (hypothesis) we use *dog* and *animal* for different facets of entities
- unlink words from ontological categories?
- words have a life of their own

→ against referentialist view

- but then, how to escape all the philosophical problems that led to referentialism in the first place?

Discussion

- unlink words from ontological categories...

Discussion

- unlink words from ontological categories. . .
- . . . but not completely
- replace "meaning" by appplicability conditions of a word for a referent
- words don't mean, people mean (Grice and many others)

Discussion

- unlink words from ontological categories...
- ...but not completely
- replace "meaning" by appplicability conditions
  of a word for a referent
- words don't mean, people mean
  (Grice and many others)
- appplicability conditions driven by
  - past uses of a word
    including reference acts and context
  - analogical processes
    (Hofstaedter, "Analogy as the core of cognition")

$\rightarrow$ "is this an instance of the same type?"
"do applicability conditions apply?"

Discussion

- distributional semantics provide operational representation of past uses of a word (Westera and Boleda, in prep.) and a (crude) means to model analogy
- recent deep learning extensions also allow us to model the utterance context, too (see e.g. work on caption generation)

→ some of the pitfalls that motivated the referentialist move have been removed, such that it is now possible to examine applicability conditions empirically

Discussion

- Aina (2018)
- production view: given instance (<object, context>*as we perceive it*), determine which word to use

- predicts: the better the match between instance and linguistic expression, the more successful the linguistic expression (easier to remember, more likely to be used again for the same referent, . . .)
- example: the *fridge room*

Discussion

- interpretation view: given instance ($<$word, context $>$), determine which referent is intended
- predicts: the better the match between instance and past uses, the more successful the reference act
- "Meanings, if anything, are types of interpretation of occurrences of types of expressions." (Paul Dekker, Ambiguity workshop, ESSLLI, Sofia)

What's in a name? or, what recent results in computational semantics tell us about ambiguity

Gemma Boleda
Universitat Pompeu Fabra
Joint work in progress with Sebastian Padó and Matthijs Westera

Ambiguity: Perspectives on Representation and Resolution Workshop
ESSLLI 2018, 610 August 2018, Sofia, Bulgaria