Lexicon, Syntax, Semantics IIb: Modeling Meaning

Distributional Semantics II

Eva Maria Vecchi

Center for Information and Language Processing
LMU Munich

June 4, 2020

Final Exams

- Final exam schedule will be posted on course website by next week
- **Reminder**: Literature and research materials need to be selected and approved by tomorrow, **June 5th**
- **Update:** I ask that you submit the slides the Monday before your presentation!

Outline

Distributional similarity as semantic similarity

- DSMs interpret semantic similarity as a **quantitative notion**
  - if **a** is closer to **b** than to **c** in the distributional vector space, then a is more semantically similar to b than to c
- Different from **categorical** nature of most theoretical accounts
  - often expressed in terms of semantic classes and relations
- But it is not clear a priori what exactly makes two words or concepts "semantically similar" according to a DSM
  - may also depend on parameter settings

Semantic similarity and relatedness

1. **Attributional similarity** – two words sharing a large number of salient features (attributes)
   - synonymy (*car*/*automobile*)
   - hyperonymy (*car*/*vehicle*)
   - co-hyponomy (*car*/*van*/*truck*)

Semantic similarity and relatedness

1. **Attributional similarity** – two words sharing a large number of salient features (attributes)
   - synonymy (*car/automobile*)
   - hyperonymy (*car/vehicle*)
   - co-hyponomy (*car/van/truck*)

2. **Semantic relatedness** (Budanitsky & Hirst, 2006) – two words are semantically associated without necessarily being similar
   - function (*car/drive*)
   - meronymy (*car/tyre*)
   - location (*car/road*)
   - attribute (*car/fast*)

Semantic similarity and relatedness

1. **Attributional similarity** – two words sharing a large number of salient features (attributes)
   - synonymy (*car*/*automobile*)
   - hyperonymy (*car*/*vehicle*)
   - co-hyponomy (*car*/*van*/*truck*)
2. **Semantic relatedness** (Budanitsky & Hirst, 2006) – two words are semantically associated without necessarily being similar
   - function (*car*/*drive*)
   - meronymy (*car*/*tyre*)
   - location (*car*/*road*)
   - attribute (*car*/*fast*)
3. **Relational similarity** (Turney, 2006) – similar relation between pairs of words (analogy)
   - *policeman*:*gun* :: *teacher*:*book*
   - *mason*:*stone* :: *carpenter*:*wood*
   - *traffic*:*street* :: *water*:*riverbed*

DSMs and semantic similarity

- DSMs are thought to represent **paradigmatic** similarity
  - words that tend to occur in the same contexts
- Words that share many contexts will correspond to concepts that share many attributes (**attributional similarity**), i.e. concepts that are **taxonomically/ontologically similar**
  - synonyms (*rhino*/*rhinoceros*)
  - antonyms and values on a scale (*good*/*bad*)
  - co-hyponyms (*rock*/*jazz*)
  - hyper- and hyponyms (*rock*/*basalt*)
- Taxonomic similarity is seen as the **fundamental semantic relation** organising the vocabulary of a language, allowing categorization, generalization and inheritance

Evaluation of (attributional) similarity

- **Synonym Identification**
  - TOEFL test (Landauer & Dumais, 1997)

Evaluation of (attributional) similarity

- **Synonym Identification**
  - TOEFL test (Landauer & Dumais, 1997)
- **Approximating semantic similarity** judgments
  - RG norms (Rubenstein & Goodenough, 1965)
  - WordSim-353 (Finkelstein et al., 2002)
  - MEN (Bruni et al., 2014), SimLex-999 (Hill et al., 2015)

Evaluation of (attributional) similarity

- **Synonym Identification**
  - TOEFL test (Landauer & Dumais, 1997)
- **Approximating semantic similarity** judgments
  - RG norms (Rubenstein & Goodenough, 1965)
  - WordSim-353 (Finkelstein et al., 2002)
  - MEN (Bruni et al., 2014), SimLex-999 (Hill et al., 2015)
- **Noun categorization**
  - ESSLLI 2008 dataset
  - AP (Almuhareb & Poesio, 2006)

Evaluation of (attributional) similarity

- **Synonym Identification**
  - TOEFL test (Landauer & Dumais, 1997)
- **Approximating semantic similarity** judgments
  - RG norms (Rubenstein & Goodenough, 1965)
  - WordSim-353 (Finkelstein et al., 2002)
  - MEN (Bruni et al., 2014), SimLex-999 (Hill et al., 2015)
- **Noun categorization**
  - ESSLLI 2008 dataset
  - AP (Almuhareb & Poesio, 2006)
- **Semantic Priming**
  - Hodgson dataset (Padó & Lapata, 2007)
  - Semantic Priming Project (Hutchison et al., 2013)

Evaluation of (attributional) similarity

- **Synonym Identification**
  - TOEFL test (Landauer & Dumais, 1997)
- **Approximating semantic similarity** judgments
  - RG norms (Rubenstein & Goodenough, 1965)
  - WordSim-353 (Finkelstein et al., 2002)
  - MEN (Bruni et al., 2014), SimLex-999 (Hill et al., 2015)
- **Noun categorization**
  - ESSLLI 2008 dataset
  - AP (Almuhareb & Poesio, 2006)
- **Semantic Priming**
  - Hodgson dataset (Padó & Lapata, 2007)
  - Semantic Priming Project (Hutchison et al., 2013)
- **Analogies** & **semantic relations** (similarity vs. relatedness)
  - Google (Mikolov et al., 2013b), BATS (Gladkova et al., 2016)
  - BLESS (Baroni & Lenci, 2011), CogALex (Santus et al., 2016)

The TOEFL synonym task

- The TOEFL dataset (80 items)
  - Target: *levied*
    Candidates: *believed*, *correlated*, ***imposed***, *requested*
  - Target: *fashion*
    Candidates: *craze*, *fathom*, ***manner***, *ration*
- DSMs and TOEFL
  1. take vectors of the target ($\mathbf{t}$) and of the candidates ($\mathbf{c_1} \dots \mathbf{c_n}$)
  2. measure the distance between $\mathbf{t}$ and $\mathbf{c_i}$, with $1 \leq i \leq n$
  3. select $\mathbf{c_i}$ with the shortest distance in space from $\mathbf{t}$

Humans vs. machines on the TOEFL task

- Average foreign test taker: 64.5%

Humans vs. machines on the TOEFL task

- Average foreign test taker: 64.5%
- Macquarie University staff (Rapp, 2004):
  - Average of 5 non-natives: 86.75%
  - Average of 5 natives: **97.75**%

Humans vs. machines on the TOEFL task

- Average foreign test taker: 64.5%
- Macquarie University staff (Rapp, 2004):
  - Average of 5 non-natives: 86.75%
  - Average of 5 natives: **97.75**%
- Distributional semantics
  - Classic LSA (Landauer & Dumais, 1997): 64.4%
  - Padó & Lapata's (2007) dependency-based model: 73.0%
  - Distributional memory (Baroni & Lenci, 2010): 76.9%
  - Rapp's (2004) SVD-based model, lemmatized BNC: 92.5%
  - Bullinaria & Levy (2012) carry out aggressive parameter optimization: 100.0%

Semantic similarity judgments

- Rubenstein & Goodenough (1965) collected similarity ratings for 65 noun pairs from 51 subjects on a 0-4 scale

w_1	w_2	avg. rating
car	*automobile*	3.9
food	*fruit*	2.7
cord	*smile*	0.0

- DSMs vs. Rubenstein & Goodenough
  - for each test pair (w_1, w_2), take vectors $\mathbf{w_1}$ and $\mathbf{w_2}$
  - measure the distance (e.g. cosine) between $\mathbf{w_1}$ and $\mathbf{w_2}$
  - measure (Pearson) correlation between vector distances and R&G average judgments (Padó & Lapata, 2007)

Semantic similarity judgments

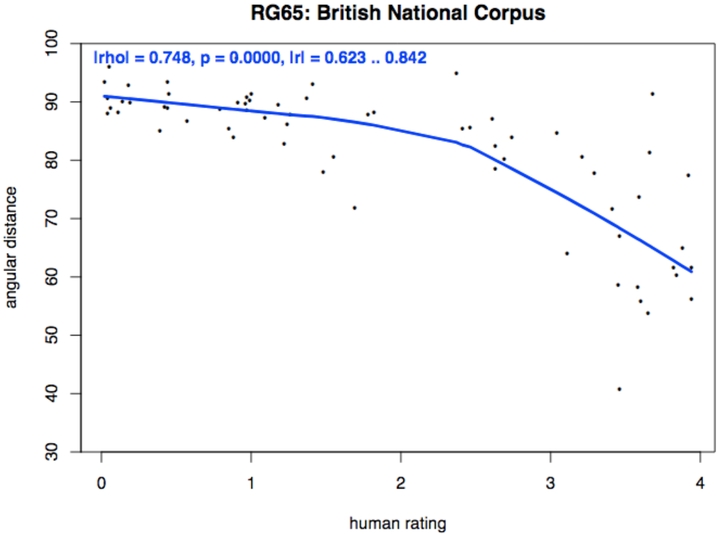

Semantic similarity judgments: results

Results on RG65 task

- Padó & Lapata's (2007) dependency-based model: 0.62
- Dependency-based on Web corpus (Herdağdelen et al., 2009)
  - without SVD reduction: 0.69
  - with SVD reduction: 0.80
- Distributional memory (Baroni & Lenci, 2010): 0.82
- Salient Semantic Analysis (Hassan & Mihalcea, 2011): 0.86

Semantic Priming

- Hearing/reading a "related" prime facilitates access to a target in various psycholinguistic tasks (naming, lexical decision, reading)
  - e.g. the word *pear* is recognized faster if heard/read after *apple*
- Hodgson (1991) single word lexical decision task, 136 prime-target pairs (cf. Padó & Lapata, 2007)
  - similar amounts of priming found for different semantic relations between primes and targets (circa 23 pairs per relation)
    - synonyms (synonym): *to dread/to fear*
    - antonyms (antonym): *short/tall*
    - coordinates (coord): *train/truck*
    - super- and subordinate pairs (supersub): *container/bottle*
    - free association pairs (freeass): *dove/peace*
    - phrasal associates (phrasacc): *vacant/building*

Semantic Priming

- DSMs and semantic priming
  1. for each related prime-target pair, measure cosine-based similarity between items (e.g., *to dread*/*to fear*)
  2. to estimate **unrelated primes**, take average of cosine-based similarity of target with other primes from same semantic relation (e.g., *to value*/*to fear*)
  3. similarity between related items should be significantly higher than average similarity between unrelated items

Semantic Priming

- DSMs and semantic priming
  1. for each related prime-target pair, measure cosine-based similarity between items (e.g., *to dread/to fear*)
  2. to estimate **unrelated primes**, take average of cosine-based similarity of target with other primes from same semantic relation (e.g., *to value/to fear*)
  3. similarity between related items should be significantly higher than average similarity between unrelated items

- Significant effects ($p < .01$) for all semantic relations
  - strongest effects for synonyms, antonyms & coordinates

Semantic Priming

- DSMs and semantic priming
  1. for each related prime-target pair, measure cosine-based similarity between items (e.g., *to dread*/*to fear*)
  2. to estimate **unrelated primes**, take average of cosine-based similarity of target with other primes from same semantic relation (e.g., *to value*/*to fear*)
  3. similarity between related items should be significantly higher than average similarity between unrelated items

- Significant effects ($p < .01$) for all semantic relations
  - strongest effects for synonyms, antonyms & coordinates

- Alternative: **classification** task
  - given target and two primes, identify related prime ($\rightarrow$ multiple choice like TOEFL)

Evaluation Strategies

DSM evaluation in published studies

- **One model, many tasks** (Padó & Lapata 2007; Baroni & Lenci 2010; Pennington et al. 2014)
  - A novel DSM is proposed, with specific features & parameters
  - This DSM is tested on a range of different tasks (e.g. TOEFL, priming, semantic clustering)

Evaluation Strategies

DSM evaluation in published studies

- **One model, many tasks** (Padó & Lapata 2007; Baroni & Lenci 2010; Pennington et al. 2014)
  - A novel DSM is proposed, with specific features & parameters
  - This DSM is tested on a range of different tasks (e.g. TOEFL, priming, semantic clustering)

- **Incremental tuning of parameters** (Bullinaria & Levy 2007, 2012; Kiela & Clark 2014; Polajnar & Clark 2014)
  - Several parameters (e.g., scoring measure, distance metric, dimensionality reduction)
  - Many tasks (e.g. TOEFL, semantic & syntactic clustering)
  - Varying granularity of parameter settings
  - One parameter (sometimes two) varied at a time, with all other parameters set to fixed values or optimized for each setting
  - Optimal parameter values are determined sequentially

(Evert et al, tutorial at ESSLLI 2018)

Recommended Readings

- Bullinaria, John A. and Levy, Joseph P. (2007). Extracting semantic representations from word co-occurrence statistics: A computational study. *Behavior Research Methods*, **39**(3), 510-526.

- Bullinaria, John A. and Levy, Joseph P. (2012). Extracting semantic representations from word co-occurrence statistics: Stop-lists, stemming and SVD. *Behavior Research Methods*, **44**(3), 890-907.

- Lapesa, Gabriella and Evert, Stefan (2014). A large scale evaluation of distributional semantic models: Parameters, interactions and model selection. *Transactions of the Association for Computational Linguistics*, **2**, 531-545.

Outline

The Meaning of *Watermelon*

- The **watermelon** fruit has a smooth exterior rind (usually green with dark green stripes or yellow spots) and a juicy, sweet interior flesh.
- **Watermelon** not only boosts your "health esteem," but it is has excellent levels of vitamins A and C and a good level of vitamin B6.

The Meaning of *Watermelon*

- The **watermelon** fruit has a smooth exterior rind (usually green with dark green stripes or yellow spots) and a juicy, sweet interior flesh.
- **Watermelon** not only boosts your "health esteem," but it is has excellent levels of vitamins A and C and a good level of vitamin B6.

The Meaning of *New York City*

Multi-Modal Semantics: Motivation

- Semantics requires "grounding"
- Interesting applications at the interface of vision and language
- Better semantic representations for NLP
- Suggested Readings:
  - Bruni et al., 2014
  - Lazaridou et al., 2014
  - Silberer & Lapata, 2010
  - Roller & Schulte im Walde, 2013
  - ... *among others*

Multi-Modal Semantics: Motivation

- The relationship between form and meaning

“violin” <==> 

- How far can we get with textual representations alone?

Language and Vision

- Enrichment of pure textual vectors with **complementary information** coming from perceptual visual features.
  - Bruni et al., Multimodal Distributional Semantics. 2014

Language and Vision

- Enrichment of pure textual vectors with **complementary information** coming from perceptual visual features.
  - Bruni et al., Multimodal Distributional Semantics. 2014

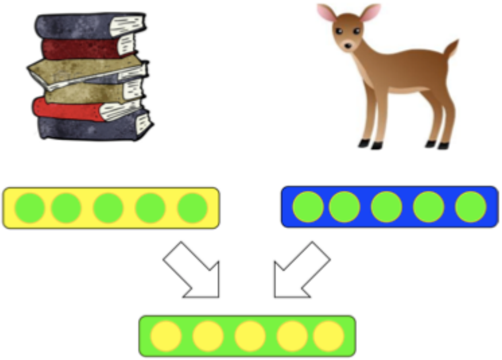

(Elia Bruni, 2014)

Applications

Task 1 Predicting human **semantic relatedness** judgments
$\longrightarrow$ Improved!

Applications

Task 1 Predicting human **semantic relatedness** judgments
$\longrightarrow$ Improved!

Task 2 **Concept categorization**

- i.e. grouping words into classes based on their semantic relatedness
- *car* ISA *vehicle*, *banana* ISA *fruit*

$\longrightarrow$ Improved!

(Elia Bruni, 2014)

Applications

Task 1 Predicting human **semantic relatedness** judgments
$\longrightarrow$ Improved!

Task 2 **Concept categorization**

- i.e. grouping words into classes based on their semantic relatedness
- *car* ISA *vehicle*, *banana* ISA *fruit*

$\longrightarrow$ Improved!

Task 3 Determine the **typical color** of concrete objects

- *cardboard* is brown, *tomato* is red

$\longrightarrow$ Improved!

(Elia Bruni, 2014)

Applications

Task 1 Predicting human **semantic relatedness** judgments

$\longrightarrow$ Improved!

Task 2 **Concept categorization**

- i.e. grouping words into classes based on their semantic relatedness
- *car* ISA *vehicle*, *banana* ISA *fruit*

$\longrightarrow$ Improved!

Task 3 Determine the **typical color** of concrete objects

- *cardboard* is brown, *tomato* is red

$\longrightarrow$ Improved!

Task 4 Distinguish **literal vs. non-literal** usages of color adjectives

- *blue uniform* vs ***blue note***

$\longrightarrow$ Improved!

(Elia Bruni, 2014)

Do pigs fly?

- No, they don't → even though *pig* and *fly* are commonly seen together (idiomatic expression)

Do cats have heads?

A state-of-the-art distributional cat (Baroni et al, 2014)

0.042 seussentennial	0.031 scarer	0.029 ragdoll
0.041 scaredy	0.031 scarer	0.029 purring
0.035 saber-toothed	0.031 repeller	0.029 whiskas
0.034 un-neutered	0.031 miaow	0.029 shorthair
0.034 meow	0.031 sphynx	0.029 scalded
0.034 unneutered	0.031 headbutts	0.029 retranslation
0.033 fanciers	0.031 spay	0.029 feral
0.033 pussy	0.030 fat	0.028 whisker
0.033 pedigreed	0.030 yowling	0.028 silvestris
0.032 sabre-toothed	0.030 flat-headed	0.028 laziest
0.032 tabby	0.030 genzyme	0.028 flap
0.032 civet	0.030 tail-less	0.028 purred
0.032 redtail	0.030 shorthaired	0.028 mummified
0.032 meowing	0.030 longhaired	...
0.032 felis	0.030 short-haired	0.0161 two-headed
0.032 whiskers	0.030 siamese	...
0.032 morphosys	0.030 english/french	0.0092 headless
0.031 meows	0.030 strangling	...
0.031 scratcher	0.030 non-pedigree	0.0021 pilgrim
0.031 black-footed	0.029 sabertooth	0.0021 out
0.031 mouser	0.029 woodpile	0.0021 head
0.031 orinthia	0.029 mewing	...

A state-of-the-art distributional cat (Baroni et al, 2014)

0.042 seussentennial
0.041 scaredy
0.035 saber-toothed
0.034 un-neutered
0.034 meow
0.034 unneutered
0.033 fanciers
0.033 pussy
0.033 pedigreed
0.032 sabre-toothed
0.032 tabby
0.032 civet
0.032 redtail
0.032 meowing
0.032 felis
0.032 whiskers
0.032 morphosys
0.031 meows
0.031 scratcher
0.031 black-footed
0.031 mouser
0.031 orinthia

0.031 scarer
0.031 scarer
0.031 repeller
0.031 miaow
0.031 sphynx
0.031 **head**butts
0.031 spay
0.030 fat
0.030 yowling
0.030 flat-**head**ed
0.030 genzyme
0.030 tail-less
0.030 shorthaired
0.030 longhaired
0.030 short-haired
0.030 siamese
0.030 english/french
0.030 strangling
0.030 non-pedigree
0.029 sabertooth
0.029 woodpile
0.029 mewing

0.029 ragdoll
0.029 purring
0.029 whiskas
0.029 shorthair
0.029 scalded
0.029 retranslation
0.029 feral
0.028 whisker
0.028 silvestris
0.028 laziest
0.028 flap
0.028 purred
0.028 mummified
...
0.0161 two-**head**ed
...
0.0092 **head**less
...
0.0021 pilgrim
0.0021 out
0.0021 **head**
...

World knowledge in language

- Distributional Semantics does not explain how our knowledge of **language** and our knowledge of the **world** interact!
- Model-theoretic semantics?
  - successful at modeling logical phenomena, e.g. quantification
  - set-theoretic interpretation
  - easy to interpret the logical inference of the examples given so far
  - need to integrate model-theoretic semantics, such as quantification

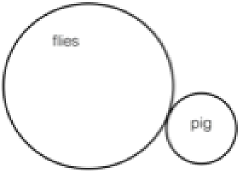

Logical inference:
if Bessy is a pig,
Bessy can't fly

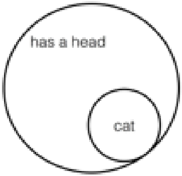

Logical inference:
if Felix is a cat,
Felix has a head

Quantification

"Mice are in the cellar"

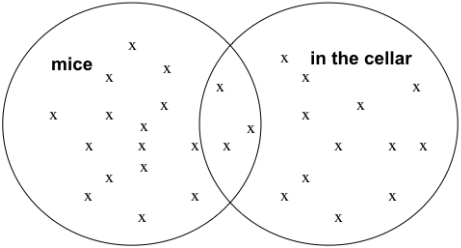

- Quantification intrinsic to most utterances
  - However, rarely explicit in naturally-occurring text
- Reference Act: *some*, *most*, *all* individuals in X do P
- Intuitive process
  - we assume only *some* of all the mice in the world have gathered – despite it not being explicit and despite not having infinite examples of mice in cellars

Modeling quantification

Quantification prerequisite for lexical semantics and inference tasks, e.g.

- hyponomy: *cat* is mammal
  - Without quantification we can do hyponomy, but with it, we can represent the whole scale of set overlap, up to disjointness (Erk, 2014)
- entailment: *most dogs have 4 legs* → *Lassie has 4 legs*
  - quantifier info as, say, features could permit a more direct representation of entailment (Baroni et al, 2012)
- logical inference: *the kouprey is a* MAMMAL
  - speakers have no problem knowing that if x is a *kouprey*, x is a MAMMAL, inference supported by lexical semantics of MAMMAL, which applies the property MAMMAL to all instances of the class

Modeling quantification is not trivial

- uncommon in text (circa 7% of NPs in large corpus)
- account for non-grounded quantification (*all cats are mammals*) and generics (*lions have manes*)
  - even adults make mistakes with generics
- semantics and pragmatics fail to provide an account of models themselves
- quantification highly dependent on speaker's interaction with the world and language
  - lexical semantic vs. world knowledge (e.g. speaker's beliefs about the concepts *bats* and *blind*)
  - pragmatics of quantifier use (e.g. speaker's personal interpretation of quantifiers in context)

From words to worlds

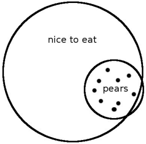

I picked some pears today. They're really nice.

The reporters asked questions at the press conference.

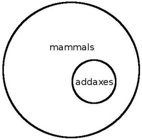

The addax is a mammal.

Distributional and Model-Theoretic Semantics

- Distributional information influences semantic ‘knowledge’
  - e.g. knowing an *alligator* (see Erk, 2015)
  - assume a systematic relation

Distributional and Model-Theoretic Semantics

- Distributional information influences semantic ‘knowledge’
  - e.g. knowing an *alligator* (see Erk, 2015)
  - assume a systematic relation
- Set-theoretic models, like distributions, can be expressed in terms of vectors
  - good approximation of shared intuitions about the world

Distributional and Model-Theoretic Semantics

- Distributional information influences semantic 'knowledge'
  - e.g. knowing an *alligator* (see Erk, 2015)
  - assume a systematic relation
- Set-theoretic models, like distributions, can be expressed in terms of vectors
  - good approximation of shared intuitions about the world
- Distributions can be translated into set-theoretic equivalents
  - assuming supervised learning

Distributional vector space

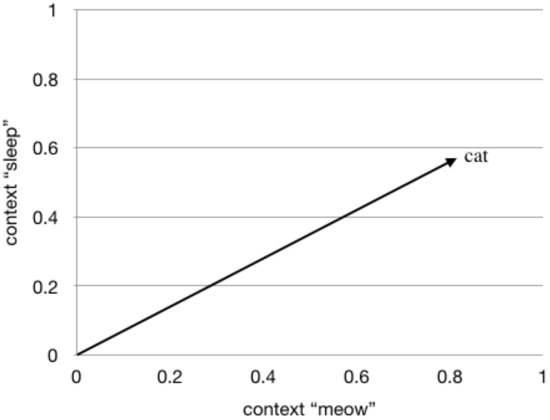

Weight: how lexically characteristic a context is for a target word.

Set-theoretic vector space

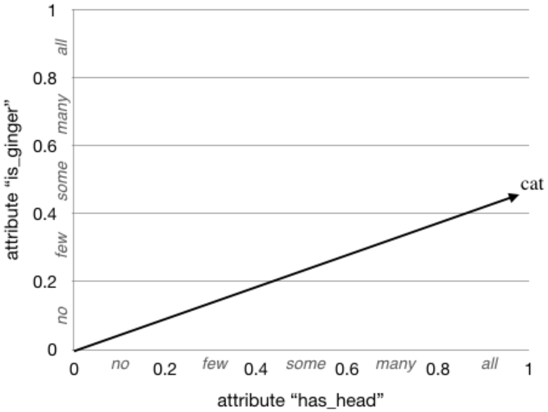

Weight: the set overlap between target and attribute.

Feature Norms

- Human subjects are asked to identify a concept's key attributes

AIRPLANE	SHRIMP	CUCUMBER
flies, 25	is_edible, 19	a_vegetable, 25
has_wings, 20	is_small, 17	eaten_in_salads, 24
used_for_passengers, 15	lives_in_water, 12	is_green, 23
requires_pilots, 11	is_pink, 11	is_long, 15
is_fast, 11	tastes_good, 9	eaten_as_pickles, 12

- McRae Norms (2005)
  - set of feature norms elicited from 725 participants for 541 concepts (7257 concept-feature pairs)

Feature Norms

- Used extensively in psychology but expensive to produce
- Feature norms are more "cognitively sound" than text-based distributional models, and more interpretable (Andrews et al., 2009; Făgărăşan et al., 2015)

	dog	black	book	animal	bread
CAT	4516	3124	1500	2480	1631

	has_fur	has_wheels	an_animal	a_pet	a_weapon
CAT	22	0	21	17	0

From norms to quantified predicates

(Herbelot & Vecchi, 2016)

Concept	*Feature*
ape	is_muscular
	is_wooly
	lives_on_coasts
	is_blind
	flies
tricycle	has_3_wheels
	used_by_children
	is_small
	used_for_transportation
	a_bike

From norms to quantified predicates

(Herbelot & Vecchi, 2016)

Concept	*Feature*	
ape	is_muscular	ALL
	is_wooly	MOST
	lives_on_coasts	SOME
	is_blind	FEW
tricycle	has_3_wheels	ALL
	used_by_children	MOST
	is_small	SOME
	used_for_transportation	FEW

From norms to quantified predicates

(Herbelot & Vecchi, 2016)

Concept	*Feature*		*weight*
ape	is_muscular	ALL	1.0
	is_wooly	MOST	0.95
	lives_on_coasts	SOME	0.35
	is_blind	FEW	0.05
tricycle	has_3_wheels	ALL	1.0
	used_by_children	MOST	0.95
	is_small	SOME	0.35
	used_for_transportation	FEW	0.05

Mapping between spaces

Andrews et al. (2009), Frome et al. (2013), Mikolov et al. (2013), Lazaridou et al. (2014), Făgărăşan et al. (2015), Dinu et al. (2015), etc.

Mapping between spaces

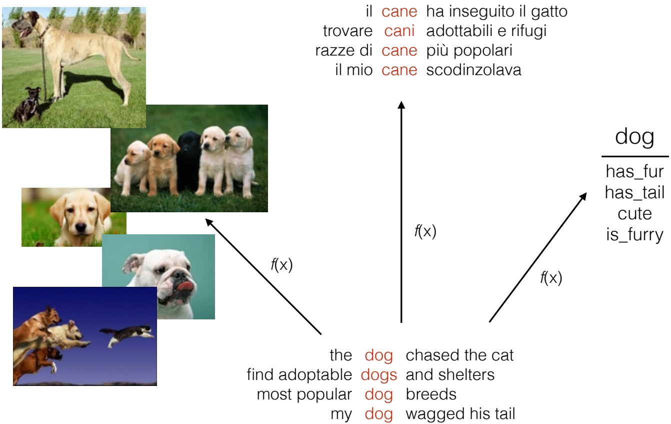

Andrews et al. (2009), Frome et al. (2013), Mikolov et al. (2013), Lazaridou et al. (2014), Făgărăşan et al. (2015), Dinu et al. (2015), etc.

Evaluation

(Herbelot & Vecchi, 2015)

1. Agreement with quantifier annotations
   - correlation between concept values in gold and mapped spaces

Evaluation

(Herbelot & Vecchi, 2015)

1. Agreement with quantifier annotations
   - correlation between concept values in gold and mapped spaces
2. Qualitative vector analysis (error analysis)
   - analysis of highly weighted contexts in mapped model-theoretic space
   - quality of neighborhoods

Evaluation

(Herbelot & Vecchi, 2015)

1. Agreement with quantifier annotations
   - correlation between concept values in gold and mapped spaces
2. Qualitative vector analysis (error analysis)
   - analysis of highly weighted contexts in mapped model-theoretic space
   - quality of neighborhoods
3. Generating quantifiers∗∗
   - map set-theoretic vectors back to natural language quantifiers for subject-predicate pairs

Generating natural language quantifiers

(Herbelot & Vecchi, 2015)

Instance		*Mapped*	*Gold*
raven	a_bird	MOST	ALL
pigeon	has_hair	FEW	NO
elephant	has_eyes	MOST	ALL
crab	is_blind	FEW	FEW
snail	a_predator	NO	NO
octopus	is_stout	NO	FEW
turtle	roosts	NO	FEW
moose	is_yellow	NO	NO
cobra	hunted_by_people	SOME	SOME
snail	forages	FEW	NO
chicken	is_nocturnal	FEW	NO
moose	has_a_heart	MOST	ALL
pigeon	hunted_by_people	NO	FEW
cobra	bites	FEW	MOST

Producing ‘true’ statements with 73% accuracy

Multi-modal semantics: From words to worlds

(Herbelot & Vecchi, 2015)

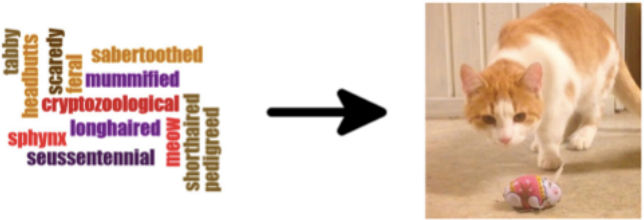

0.042 seussentennial	0.032 tabby	1 walks	1 has-a_tail
0.041 scaredy	0.032 civet	1 purrs	1 has-4_legs
0.035 saber-toothed	0.032 redtail	1 meows	1 an-animal
0.034 un-neutered	0.032 meowing	1 has-eyes	1 a-mammal
0.034 meow	0.032 felis	1 has-a_heart	1 a-feline
0.034 unneutered	0.032 whiskers	1 has-a_head	0.7 is-independent
0.033 fanciers	0.032 morphosys	1 has-whiskers	0.7 eats-mice
0.033 pussy	0.031 meows	1 has-paws	0.7 is-carnivorous
0.033 pedigreed	0.031 scratcher	1 has-fur	0.3 is-domestic
0.032 sabre-toothed	...	1 has-claws	...

Thanks, see you next week!

https://www.vecchi.com/eva/teaching/modelingmeaning.html