
Semantics

Philipp Koehn

9 April 2015

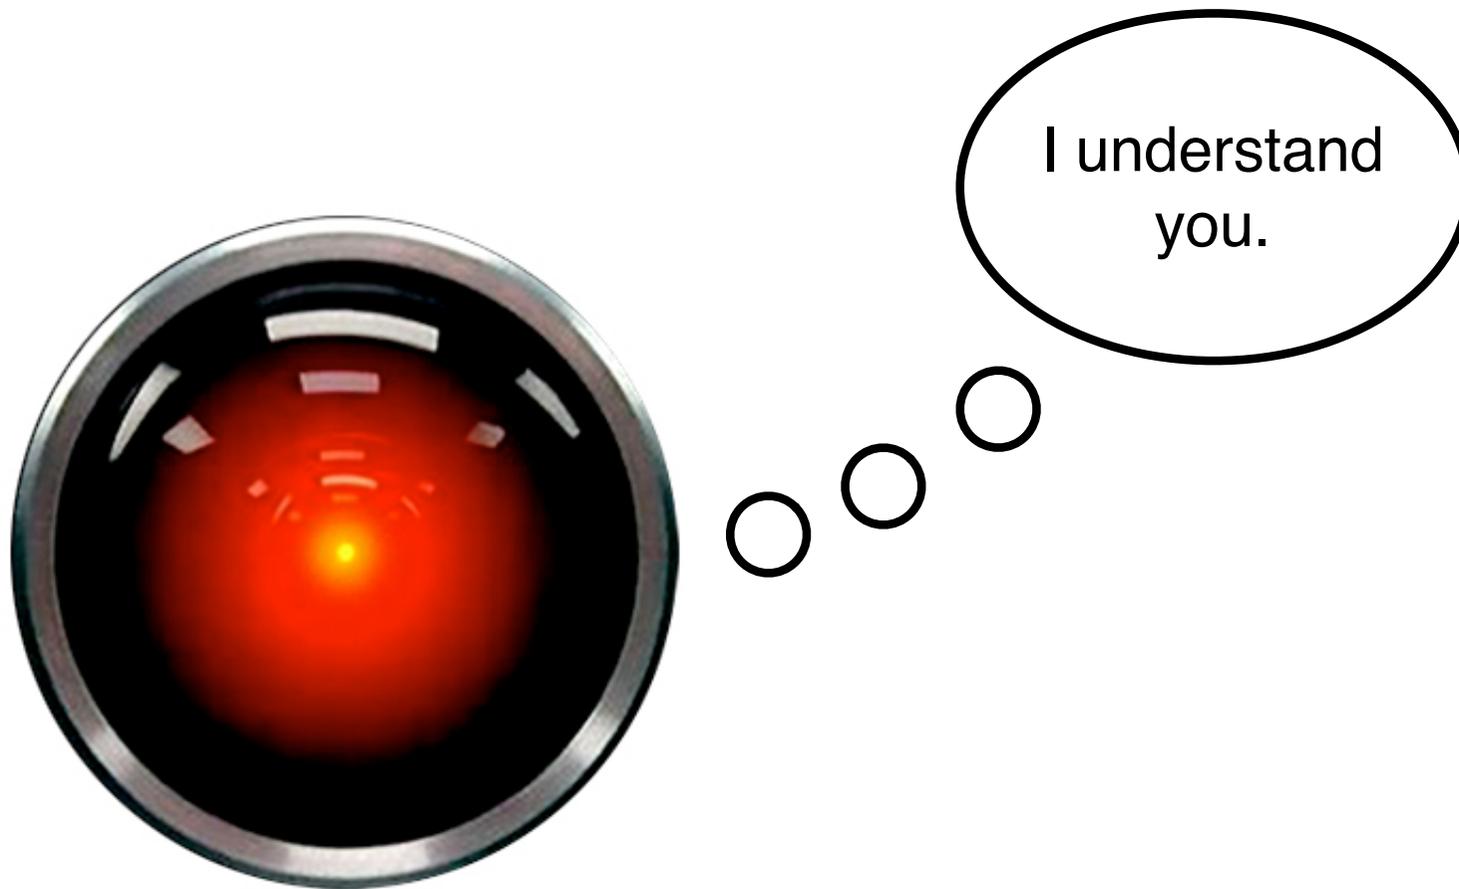


Meaning



- The grand goal of **artificial intelligence**
 - machines that do not mindlessly process data
 - ... but that ultimately **understand** its meaning
- But what is meaning?

Meaningful Machines



A Scale of Understanding?



wisdom



knowledge



data

Eliza



Joseph Weizenbaum wrote the computer program Eliza in 1969 to demonstrate how easily people can be fooled into thinking that machines have some deep understanding.

young woman: Men are all alike.

eliza: In what way?

young woman: They're always bugging us about something specific or other.

eliza: Can you think of a specific example?

young woman: Well, my boyfriend made me come here.

eliza: Your boyfriend made you come here?

young woman: He says I'm depressed much of the time.

eliza: I am sorry to hear you are depressed.

young woman: It's true. I'm unhappy.

eliza: Do you think coming here will help you not to be unhappy?

Online demo: http://www-ai.ijs.si/eliza-cgi-bin/eliza_script

semantic translation problems

Semantic Translation Problems



- Syntactic analysis may be ambiguous

Das Vorhaben verwarf die Kommission .
the plan rejected the commission .

- Both readings (SVO and OSV) are syntactically possible
- But: OSV reading is semantically much more plausible

⇒ Need for semantic model to produce semantically plausible output

Semantic Translation Problems



- Pronominal anaphora

I saw the movie and **it** is good.

- How to translate **it** into German (or French)?
 - **it** refers to **movie**
 - **movie** translates to **Film**
 - **Film** has masculine gender
 - ergo: **it** must be translated into masculine pronoun **er**
- We are not handling this very well [Le Nagard and Koehn, 2010]

Semantic Translation Problems



- Coreference

Whenever I visit my uncle and his daughters,
I can't decide who is my favorite **cousin**.

- How to translate **cousin** into German? Male or female?
- Complex inference required

Semantic Translation Problems



- Discourse

Since you brought it up, I do not agree with you.

Since you brought it up, we have been working on it.

- How to translated *since*? Temporal or conditional?
- Analysis of discourse structure — a hard problem

lexical semantics

Word Senses

- Some words have multiple meanings
- This is called polysemy
- Example: **bank**
 - financial institution: **I put my money in the bank.**
 - river shore: **He rested at the bank of the river.**
- How could a computer tell these senses apart?

Homonym



- Sometimes two completely different words are spelled the same
- This is called a homonym
- Example: **can**
 - modal verb: **You can do it!**
 - container: **She bought a can of soda.**
- Distinction between polysemy and homonymy not always clear

How many senses?

- How many senses does the word **interest** have?
 - She pays 3% **interest** on the loan.
 - He showed a lot of **interest** in the painting.
 - Microsoft purchased a controlling **interest** in Google.
 - It is in the national **interest** to invade the Bahamas.
 - I only have your best **interest** in mind.
 - Playing chess is one of my **interests**.
 - Business **interests** lobbied for the legislation.
- Are these seven different senses? Four? Three?

Wordnet

- Wordnet, a hierarchical database of senses, defines synsets
- According to Wordnet, **interest** is in 7 synsets
 - Sense 1: a sense of concern with and curiosity about someone or something, Synonym: involvement
 - Sense 2: the power of attracting or holding one's interest (because it is unusual or exciting etc.), Synonym: interestingness
 - Sense 3: a reason for wanting something done, Synonym: sake
 - Sense 4: a fixed charge for borrowing money; usually a percentage of the amount borrowed
 - Sense 5: a diversion that occupies one's time and thoughts (usually pleasantly), Synonyms: pastime, pursuit
 - Sense 6: a right or legal share of something; a financial involvement with something, Synonym: stake
 - Sense 7: (usually plural) a social group whose members control some field of activity and who have common aims, Synonym: interest group

Sense and Translation

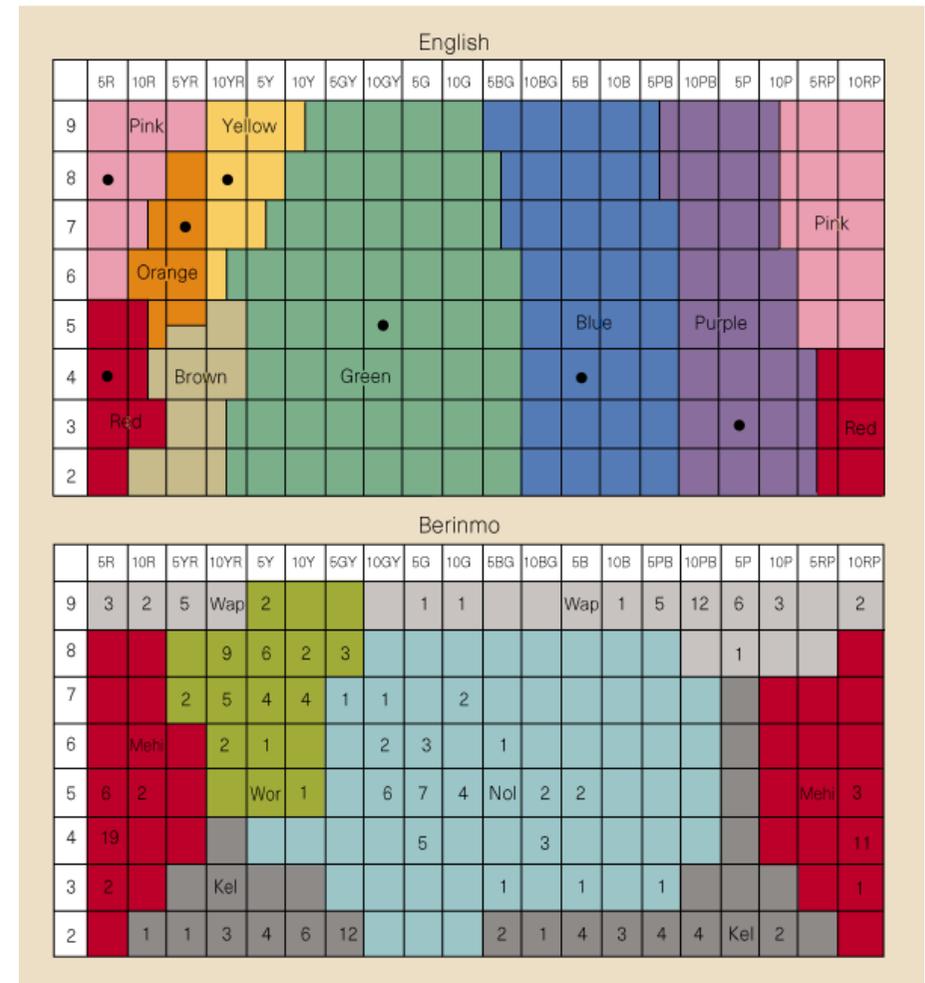
- Most relevant for machine translation:
different translations → different sense
- Example **interest** translated into German
 - **Zins**: financial charge paid for loan (Wordnet sense 4)
 - **Anteil**: stake in a company (Wordnet sense 6)
 - **Interesse**: all other senses

Languages differ

- Foreign language may make finer distinctions
- Translations of **river** into French
 - **fleuve**: river that flows into the sea
 - **rivière**: smaller river
- English may make finer distinctions than a foreign language
- Translations of German **Sicherheit** into English
 - **security**
 - **safety**
 - **confidence**

Overlapping Senses

- Color names may differ between languages
- Many languages have one word for blue and green
- Japanese: **ao**
change early 20th century:
midori (green) and **ao** (blue)
- But still:
 - vegetables are **greens** in English,
ao-mono (blue things) in Japanese
 - “go” traffic light is **ao** (blue)



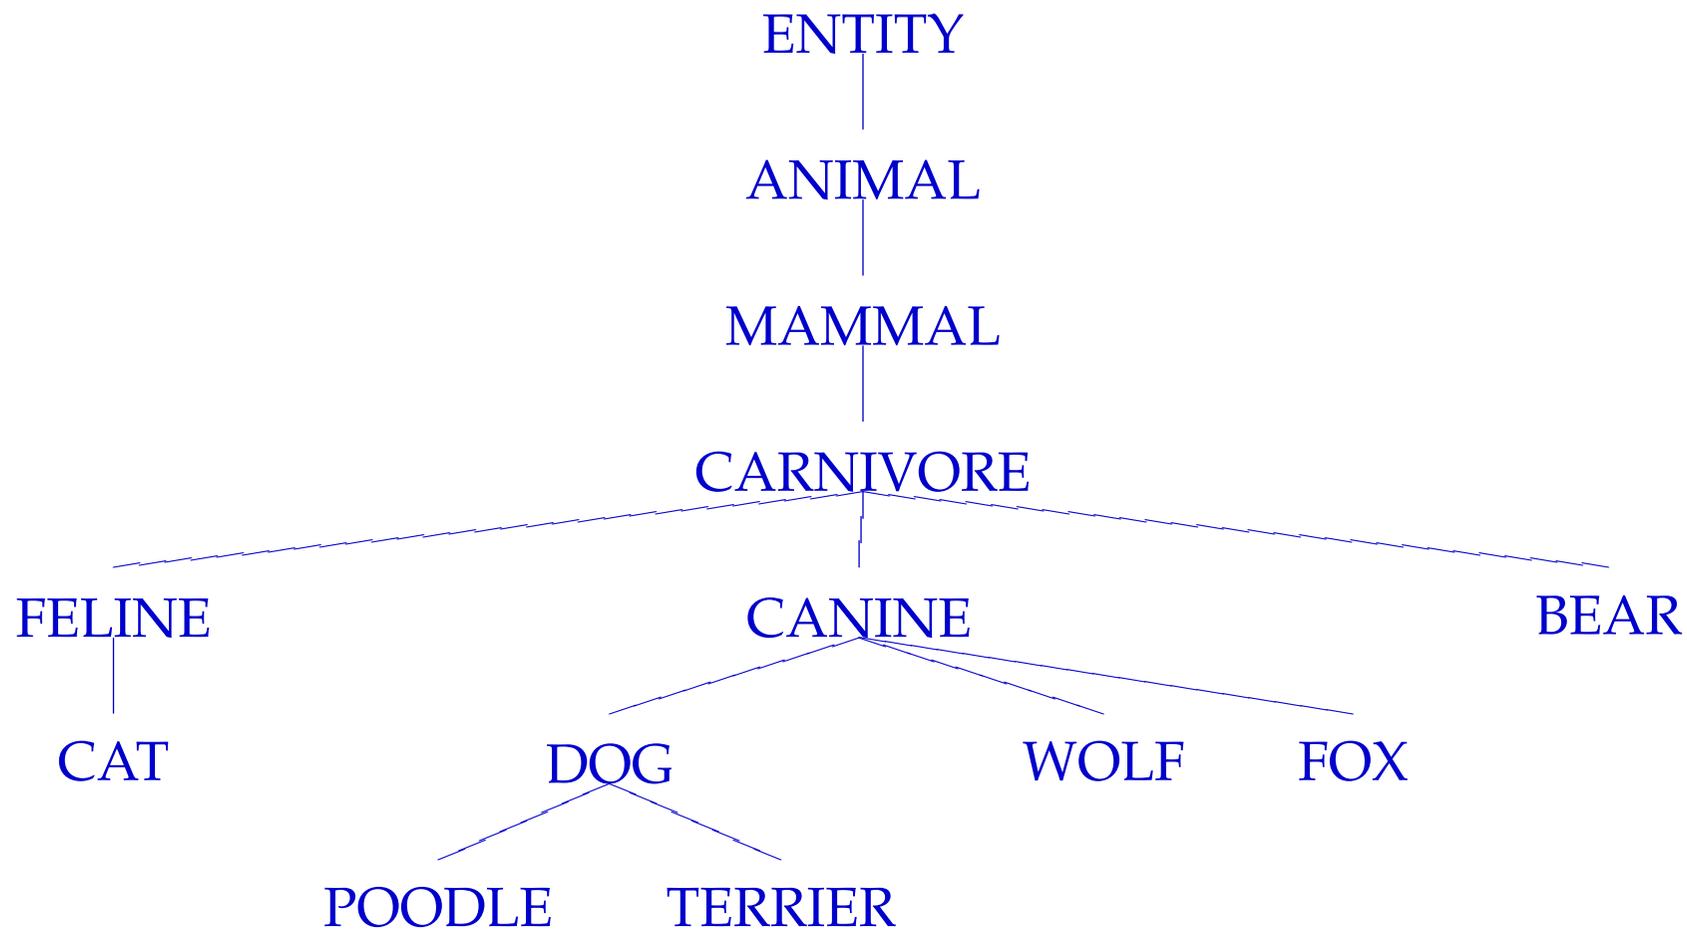
Color names in English and Berinomo (Papua New Guinea)

One last word on senses

- Lot of research in word sense disambiguation is focused on polysemous words with clearly distinct meanings, e.g. **bank**, **plant**, **bat**, ...
- Often meanings are close and hard to tell apart, e.g. **area**, **field**, **domain**, **part**, **member**, ...
 - She is a part of the team.
 - She is a member of the team.
 - The wheel is a part of the car.
 - * The wheel is a member of the car.

Ontology

19



Representing Meaning

- So far: the meaning of **dog** is **DOG** or **dog(x)**
Not much gained here
- Words that have similar meaning should have similar representations
- Composition of meaning

$$\text{meaning}(\text{daughter}) = \text{meaning}(\text{child}) + \text{meaning}(\text{female})$$

- Analogy

$$\text{meaning}(\text{king}) + \text{meaning}(\text{woman}) - \text{meaning}(\text{man}) = \text{meaning}(\text{queen})$$

Distributional Semantics

- Contexts may be represented by a vector of word counts

Example:

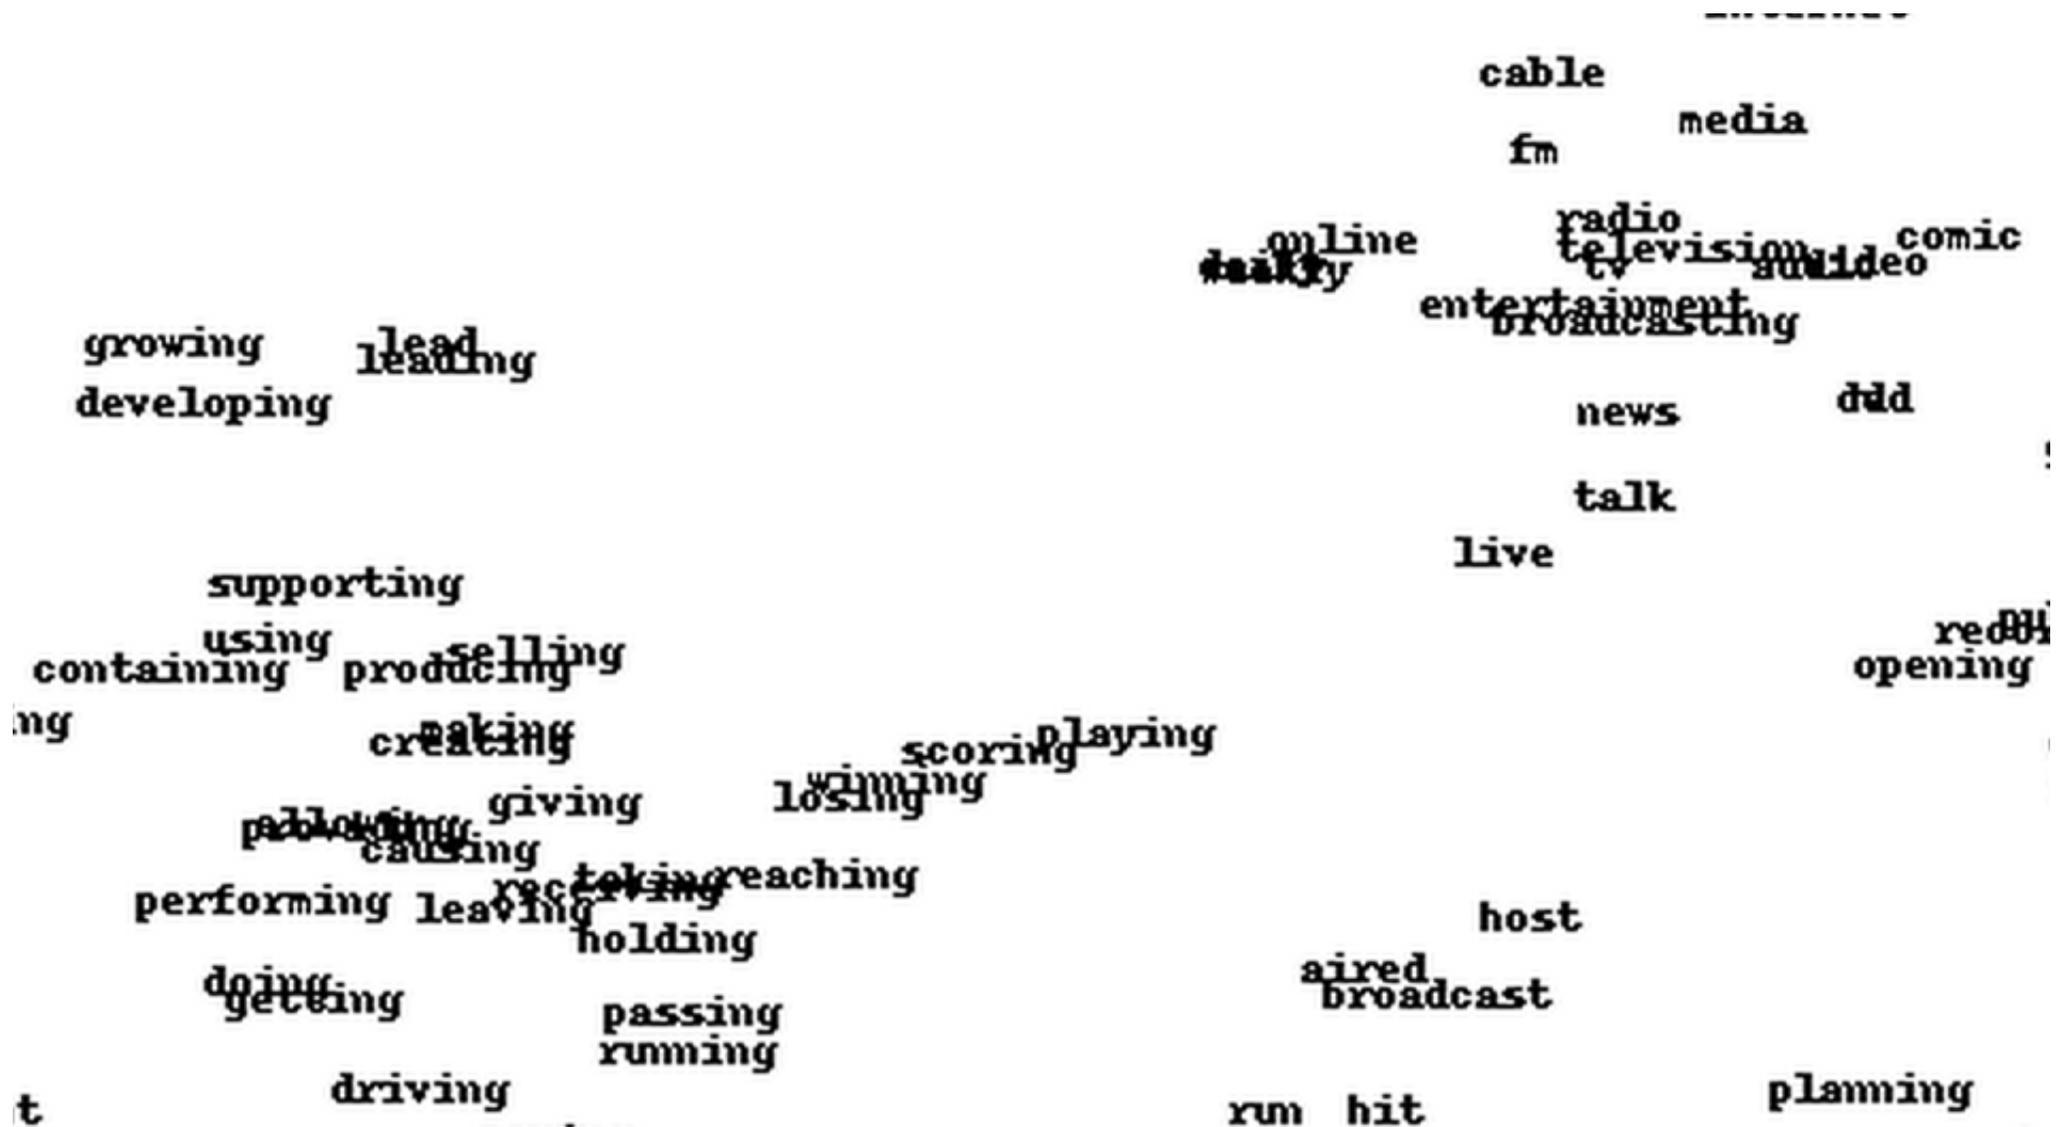
Then he grabbed his new mitt and **bat**, and headed back to the dugout for another turn at bat. Hulet isn't your average baseball player. "It might have been doctoring up a **bat**, grooving a bat with pennies or putting a little pine tar on the baseball. All the players were sitting around the dugout laughing at me."

The word counts normalized, so all the vector components add up to one.

grabbed	(1)	(0.05)
mitt	1	0.05
headed	1	0.05
dugout	2	0.10
turn	1	0.05
average	1	0.05
baseball	2	0.10
player	2	0.10
doctoring	1	0.05
grooving	1	0.05
pennies	1	0.05
pine	1	0.05
tar	1	0.05
sitting	1	0.05
laughing	1	0.05

- Average over all occurrences of word
- Context may also just focus on surrounding words

Word Embeddings



Word Sense Disambiguation

- For many applications, we would like to disambiguate senses
- Supervised learning problem **plant** → **PLANT-FACTORY**
- Features
 - Directly neighboring words
 - * **plant** life
 - * manufacturing **plant**
 - * assembly **plant**
 - * **plant** closure
 - * **plant** species
 - Any content words in a 50 word window
 - Syntactically related words
 - Syntactic role in sense
 - Topic of the text
 - Part-of-speech tag, surrounding part-of-speech tags

- Machine translation models already include the powerful features
 - phrase translation model: condition translation on neighboring words
 - language model: directly neighboring words in target language
- Limited success in adding wider context
 - position-sensitive, syntactic, and local collocational features (Carpuat and Wu, 2007)
 - maximum entropy classifier for surrounding context words (Tamchyna et al., 2014)

subcategorization frames

Verb Subcategorization

- Example

Das Vorhaben verwarf die Kommission .
the plan rejected the commission .

- Propbank

Arg0-PAG: rejecter (vnrole: 77-agent)

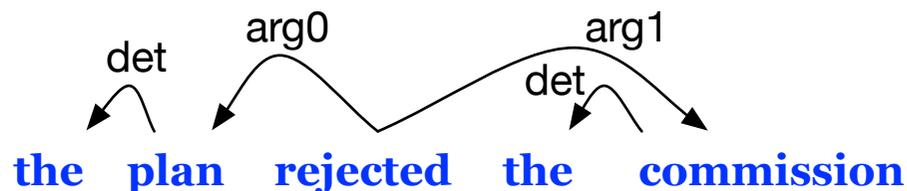
Arg1-PPT: thing rejected (vnrole: 77-theme)

Arg3-PRD: attribute

- Is [plan](#) a typical Arg0 of [reject](#)?

Dependency Parsing

- Dependencies between words



- Can be obtained by
 - dedicated dependency parser
 - CFG grammar with head word rules
- Are dependency relations enough?
 - reject — subj → plan ⇒ bad
 - reject — subj → commission ⇒ good

logical form

- Classical example

Every farmer has a donkey

- Ambiguous, two readings
- Each farmer as its own donkey

$$\forall x: \text{farmer}(x) \exists y: \wedge \text{donkey}(y) \wedge \text{owns}(x,y)$$

- There is only one donkey

$$\exists y: \text{donkey}(y) \wedge \forall x: \text{farmer}(x) \wedge \text{owns}(x,y)$$

- Does this matter for translation? (typically not)

Bigger Issue: Scope

- Example (Knight and Langkilde, 2000)

green eggs and ham

- Only eggs are green

(green eggs) and ham

- Both are green

green (eggs and ham)

- Spanish translations

- Only eggs are green

huevos verdes y jamón

- Also ambiguous

jamón y huevos verdes

- Machine translation should preserve ambiguity

Logical Form and Inference

- Input sentence

Whenever I visit my uncle and his daughters,
I can't decide who is my favorite **cousin**.

- Facts from input sentence

$\exists d: \text{female}(d)$
 $\exists u: \text{father}(d,u)$
 $\exists i: \text{uncle}(u,i)$
 $\exists c: \text{cousin}(i,c)$

- World knowledge

$\forall i,u,c: \text{uncle}(u,i) \wedge \text{father}(u,c) \rightarrow \text{cousin}(i,c)$

- Hypothesis that $c = d$ is consistent with given facts and world knowledge

- Inference

$\text{female}(d) \rightarrow \text{female}(c)$



discourse

Ambiguous Discourse Markers

- Example

Since you brought it up, I do not agree with you.

Since you brought it up, we have been working on it.

- How to translated *since*? Temporal or conditional?

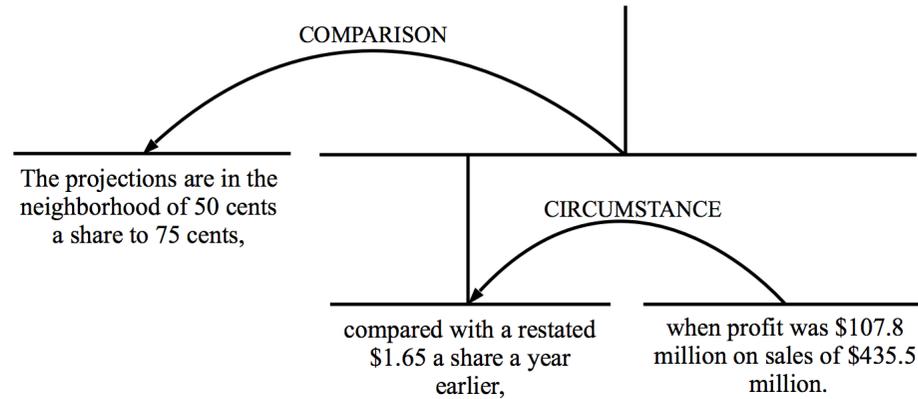
Implicit Discourse Relationships

- English syntactic structure may imply causation

Wanting to go to the other side, the chicken crossed the road.

- This discourse relationship may have to be made explicit in another language

- Discourse relationships,
e.g., Circumstance, Antithesis, Concession, Solutionhood, Elaboration, Background, Enablement, Motivation, Condition, Interpretation, Evaluation, Purpose, Evidence, Cause, Restatement, Summary, ...
- Hierarchical structure



- There is a discourse treebank, but inter-annotator agreement is low

abstract meaning representations

AMR: Towards Interlingua

- Semantic representations of full sentences
- English-oriented
- Builds on Propbank
- Explicit annotation of co-reference
- Some additional semantic relationships (degree, part-of, possessives, etc.)
- Not everything resolved
- Not annotated: tense, plural, passive, focus, and other syntactic properties

Example

He looked at me very gravely , and put his arms around my neck .

(a / and

```
:op1 (l / look-01
      :ARG0 (h / he)
      :ARG1 (i / i)
      :manner (g / grave
               :degree (v / very)))
:op2 (p / put-01
      :ARG0 h
      :ARG1 (a2 / arm
             :part-of h)
      :ARG2 (a3 / around
             :op1 (n / neck
                   :part-of i))))
```

- Abstract meaning representation

```
(l / look-01
  :ARG0 (h / he)
  :ARG1 (i / i)
  :manner (g / grave
           :degree (v / very)))
```

- Possible English sentences

- He looks at me gravely.
- I am looked at by him very gravely.
- He gave me a very grave look.

Directed Acyclic Graphs

- Formally, AMR structures are more complex than trees
- Co-reference \Rightarrow directed acyclic graphs (DAG)
- Processing such DAGs is harder, algorithms are currently developed
- Tasks
 - semantic parsing (English text \rightarrow English AMR)
 - semantic transduction (foreign text \rightarrow English AMR)
 - generation (English AMR \rightarrow English text)
- Active work on algorithms, but no competitive system yet