L98: Introduction to Computational Semantics

Lecture 1: Introduction to Introduction to Computational Semantics

Simone Teufel and Weiwei Sun

Natural Language and Information Processing Research Group
Department of Computer Science and Technology
University of Cambridge

Lent 2021/22

semantics from PIE root *dheie- "to see, look".

meaning from PIE *meino- "opinion, intent", perhaps from root *men- "to think".

PIE = Proto Indo European

Lecture 1: Introduction to Introduction to Computational Semantics

1. Semantics
2. Computational semantics
3. Introduction to computational semantics
4. Introduction to introduction to computational semantics
5. Word senses

Semantics

An example of misleading truth

- **lying**: what is said is false
- **misleading**: what is suggested is false

Immanuel Kant

Computational Semantics

Meta language

a precise representation needs a language.

- natural language, e.g. English
- programming language, e.g. Ruby, Scala
- Math, e.g. matrix
- logic, e.g. λ calculus
- automata, e.g. finite-state machines

Representing word meanings with a natural language

Lexicography, e.g. Cambridge Dictionary (https://dictionary.cambridge.org)

blue

- adjective (COLOUR): of the colour of the sky without clouds on a bright day, or a darker or lighter type of this:
- adjective (SAD): feeling or showing sadness

Representing word *meanings* with vectors (1)

Word embedding, word representations, representation learning, "lexical semantics"

Representing word *meanings* with vectors (1)

Word embedding, word representations, representation learning, "lexical semantics"

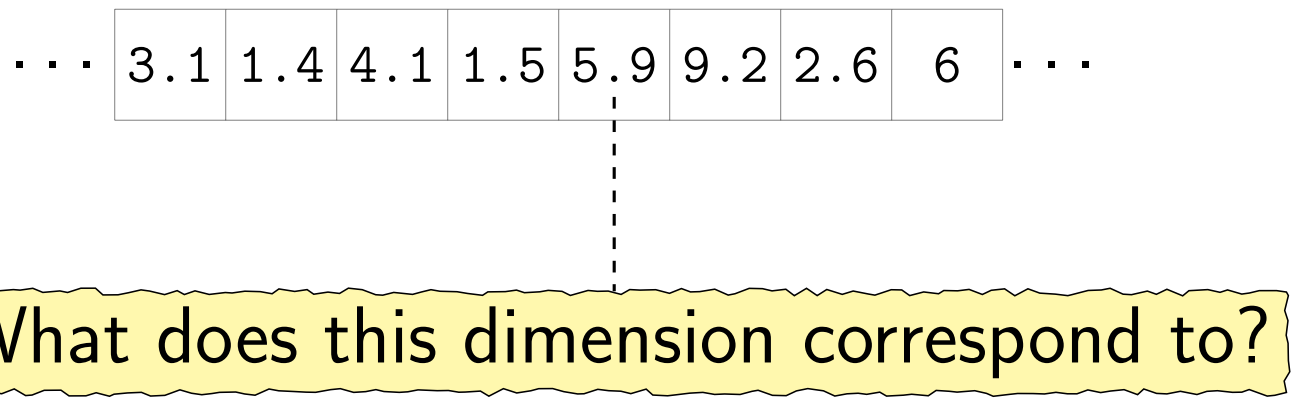

Representing word *meanings* with vectors (2)

Representing word *meanings* with vectors (2)

RGB vector = word meaning??

Cultural effects concerning colour blue

- Russian: subdivision of Western “blue”

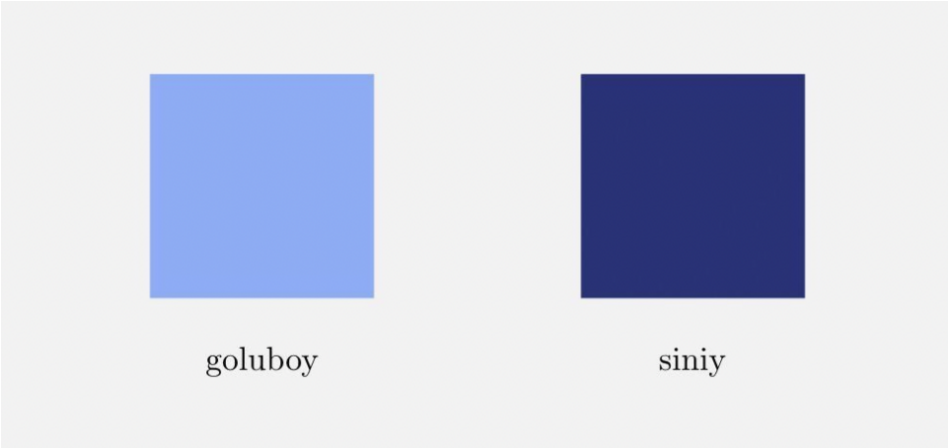

- Japanese: one single word for Western “green” and “blue”: 青

  Similar effect:

#5B8930	萌黄 Moegi "Fresh Onion", listed with yellow
#6B9362	若竹色 Wakatake-iro "Young bamboo color", listed with blue

from `https://en.wikipedia.org/wiki/Blue-green_distinction_in_language`

Introduction to Computational Semantics

Ultimate Goal – understanding a text

- To do so, we need to have knowledge of many things
- Syntax and Semantics – how does the language assemble its meaning-units (locally)?
- Many symbolic NLP courses teach how to assemble meaning from individual words inside a sentence (compositional semantics).
- Individual words' meanings are untreated (left as "atomic").
- Pragmatics – what is left unsaid but can be "calculated" by a human nevertheless? (Not many computational approaches available, but lots of research)

Introduction to Introduction to Computational Semantics

What you can learn here

Answers to questions:

- What are "word meaning", "sentence meaning" and "discourse meaning"?
- Why is there an entire course dedicated to semantics?
- What is the connection to today's practical Natural Language Processing (NLP) tasks?

Course is taught as a mixture of

- Phenomena and Theory
- Automatic methods for recognising/treating the phenomena

Semantic sub-disciplines

- Lexical Semantics (Word senses, Semantic Roles, . . .)

  How can we define and express what individual words mean

- Compositional Semantics (world model, lambda calculus, FOPL, some HOL. . .)

  How basic meaning units are recursively combined

- Pragmatics (one lecture)

- Discourse (one lecture)

Shapeworld: An example

Shapeworld is an environment for testing Visual QA systems created by Alexander Kuhnle (PhD 2020; runner-up to BCS's PhD thesis of the year). It uses a simulated microworld:

A magenta square is to the right of a green shape.

The lowermost green shape is a cross.

A red shape is the same shape as a green shape.

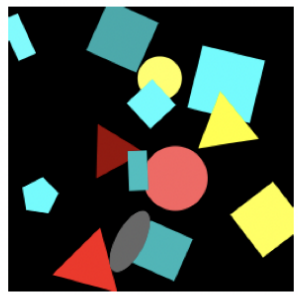

At least half the triangles are red.

More than a third of the shapes are cyan squares.

More than one of the seven cyan shapes is a square.

Green statements are true. Red statements are false.

https://www.cl.cam.ac.uk/techreports/UCAM-CL-TR-942.html

How does Shapeworld know whether something is true?

It has a World model (more about that in Lecture 6):

{ **color**: {**name**: *black*, **shade**: *0.0*}, **noise-stddev**: *0.1*, **size**: *64*, **objects**:
[{ **center**: {**x**: *0.47*, **y**: *0.28*}, **color**: {**name**: *yellow*, **shade**: *-0.24*},
rotation: *0.06*, **shape**: {**name**: *cross*, **extent**: {**x**: *0.10*, **y**: *0.10*}} },
{ **center**: {**x**: *0.49*, **y**: *0.65*}, **color**: {**name**: *red*, **shade**: *0.26*},
rotation: *0.76*, **shape**: {**name**: *cross*, **extent**: {**x**: *0.08*, **y**: *0.08*}} },
{ **center**: {**x**: *0.15*, **y**: *0.91*}, **color**: {**name**: *yellow*, **shade**: *-0.16*},
rotation: *0.27*, **shape**: {**name**: *pentagon*, **extent**: {**x**: *0.09*, **y**: *0.08*}} },
{ **center**: {**x**: *0.80*, **y**: *0.37*}, **color**: {**name**: *red*, **shade**: *-0.12*},
rotation: *0.53*, **shape**: {**name**: *circle*, **extent**: {**x**: *0.12*, **y**: *0.12*}} },
{ **center**: {**x**: *0.92*, **y**: *0.73*}, **color**: {**name**: *yellow*, **shade**: *-0.42*},
rotation: *0.73*, **shape**: {**name**: *cross*, **extent**: {**x**: *0.09*, **y**: *0.09*}} }] }

Images and descriptions are created simulaneously from the same source.

How does Shapeworld know whether something is true?

It has a World model (more about that in Lecture 6):

{ **color**: {**name**: *black*, **shade**: *0.0*}, **noise-stddev**: *0.1*, **size**: *64*, **objects**:
[{ **center**: {**x**: *0.47*, **y**: *0.28*}, **color**: {**name**: *yellow*, **shade**: *-0.24*},
rotation: *0.06*, **shape**: {**name**: *cross*, **extent**: {**x**: *0.10*, **y**: *0.10*}} },
{ **center**: {**x**: *0.49*, **y**: *0.65*}, **color**: {**name**: *red*, **shade**: *0.26*},
rotation: *0.76*, **shape**: {**name**: *cross*, **extent**: {**x**: *0.08*, **y**: *0.08*}} },
{ **center**: {**x**: *0.15*, **y**: *0.91*}, **color**: {**name**: *yellow*, **shade**: *-0.16*},
rotation: *0.27*, **shape**: {**name**: *pentagon*, **extent**: {**x**: *0.09*, **y**: *0.08*}} },
{ **center**: {**x**: *0.80*, **y**: *0.37*}, **color**: {**name**: *red*, **shade**: *-0.12*},
rotation: *0.53*, **shape**: {**name**: *circle*, **extent**: {**x**: *0.12*, **y**: *0.12*}} },
{ **center**: {**x**: *0.92*, **y**: *0.73*}, **color**: {**name**: *yellow*, **shade**: *-0.42*},
rotation: *0.73*, **shape**: {**name**: *cross*, **extent**: {**x**: *0.09*, **y**: *0.09*}} }] }

Images and descriptions are created simulaneously from the same source.

Precise meaning representations supported by methodology

How does Shapeworld generate NL statements?

It translates known relationships, properties and object desriptions into semantic representations (roughly like FOPL):

$$\textit{“to the left of”} \text{ translates to } p = \text{relation}(\text{type} = \text{x-rel}, \text{value} = -1, \text{reference} = r):$$

$$\begin{aligned} p.\text{agree}(e) &:= \exists\, e' \in r.\text{agree}(\cdot) : (e'.\mathbf{x} - e.\mathbf{x}) > \max(\epsilon_{\text{distance}}, |e.\mathbf{y} - e'.\mathbf{y}|) \\ p.\text{disagree}(e) &:= \forall\, e' \in \neg r.\text{disagree}(\cdot) : (e'.\mathbf{x} - e.\mathbf{x}) < -\epsilon_{\text{distance}} \end{aligned}$$

How does Shapeworld generate NL statements?

It then uses an NL generator that can generate language based on the semantic representations:

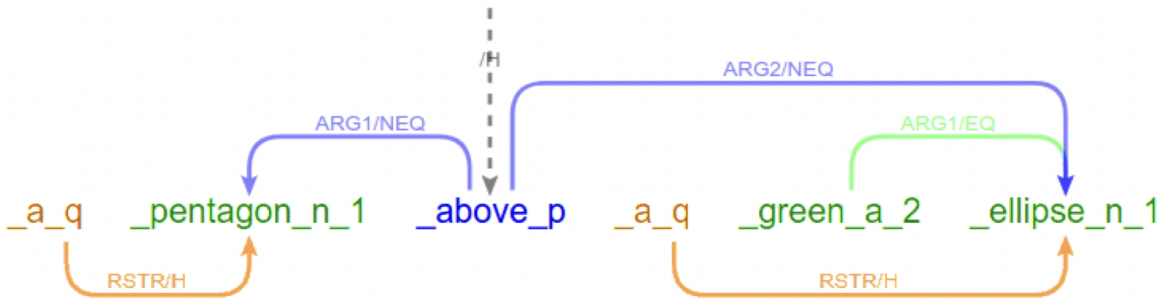


```
[ LTOP: h0 INDEX: e2 RELS: <
    [ _a_q LBL: h4 ARG0:  x3 RSTR: h5 BODY: h6 ]
    [ _pentagon_n_1 LBL: h7 ARG0:  x3 [...]  ]
    [ _above_p LBL: h1 ARG0:  e2 [...]  ARG1:  x3 ARG2:  x8 ]
    [ _a_q LBL: h9 ARG0:  x8 RSTR: h10 BODY: h11 ]
    [ _green_a_2 LBL: h12 ARG0:  e13 [...]  ARG1:  x8 ]
    [ _ellipse_n_1 LBL: h12 ARG0:  x8 [...]  ]
  > HCONS: < h0 ≅ h1, h5 ≅ h7 h10 ≅ h12 > ]
```

"A pentagon is above a green ellipse."

Precise meaning representations supported by methodology

The alternative: Traditional datasets for Visual QA

- What object is shining on the animal?
- What objects is the cat sitting behind?
- How many cats?

- precise?
- world-knowledge free?
- subjective? testable in an experiment?

Word Senses

Lexical semantics: Some topics

- Recognise word senses in text (manually and automatically)
- Describe relations between words (or rather, between word senses)
- Determine how strongly a verb "goes with" its subject or arguments
- Recognise and interpret figurative use of words

Word senses

There are two ways in which a word form can be ambiguous:

- Random historic effects bring two unrelated words together $\rightarrow$ homonymy (same name)
- Senses evolve during language evolution, but there is a connection $\rightarrow$ polysemy (multiple senses)

Polysemy

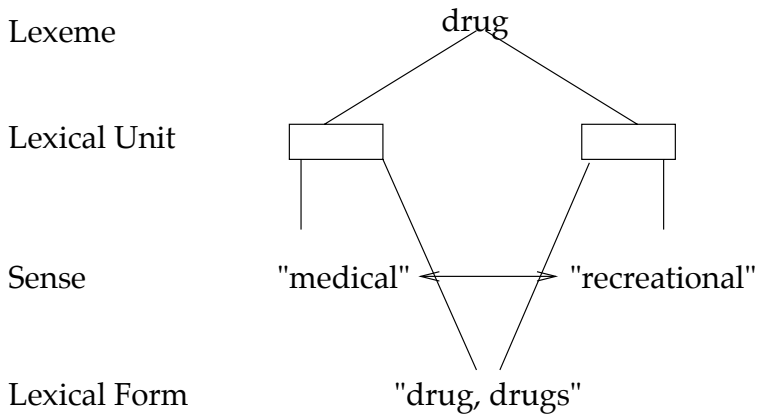

A lexical form corresponding to a single lexeme which has a number of senses is polysemous.

Types of polysemy

Systematic polysemy: two senses are in a systematic semantic relation to each other. This process is productive.

- plant vs food (e.g. wheat)
- content vs physical object (e.g. book)
- rabbit: animal vs meat (e.g. rabbit)
- instrument vs process (e.g. shower)
- unit vs type (e.g. I want that shirt)

Idiosyncratic polysemy

- Has Arthur changed his position? (metaphor)
- The ham sandwich asked for the bill. (situational)

Homonymy

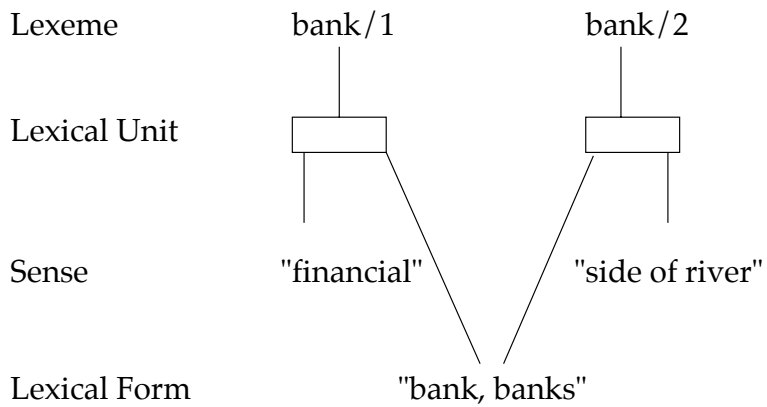

A lexical form corresponding to more than one lexeme, each with their sense(s) is homonymous.

Sense 1 of “bank”

“Arthur reached the bank”

Sense 2 of "bank"

"Arthur reached the bank."

Wordnet as a Meta language

- Wordnet groups word forms into synsets (synonym sets).
- One synset = one sense; this grouping constitutes the senses's definition.
- Homonyms and polysemous word forms are therefore associated with multiple (different) synsets.
- Senses are indicated by slashes and numbers: `interest/1`, `interest/2`...
- Synsets are organized into a hierarchical structure by the use of hyponymy, e.g. *dog* is-a *pet*, *pet* is-a *animal*
- Other relations: meronymy (part-of), paronymy (same stem, morphological variation), antonymy (opposite)

WN example – "interest"

Noun

- S (n) **interest**, involvement (a sense of concern with and curiosity about someone or something) *"an interest in music"*
- S (n) sake, **interest** (a reason for wanting something done) *"for your sake"; "died for the sake of his country"; "in the interest of safety"; "in the common interest"*
- S (n) **interest**, interestingness (the power of attracting or holding one's attention (because it is unusual or exciting etc.)) *"they said nothing of great interest"; "primary colors can add interest to a room"*
- S (n) **interest** (a fixed charge for borrowing money; usually a percentage of the amount borrowed) *"how much interest do you pay on your mortgage?"*
- S (n) **interest**, stake ((law) a right or legal share of something; a financial involvement with something) *"they have interests all over the world"; "a stake in the company's future"*
- S (n) **interest**, interest group (usually plural) a social group whose members control some field of activity and who have common aims) *"the iron interests stepped up production"*
- S (n) pastime, **interest**, pursuit (a diversion that occupies one's time and thoughts (usually pleasantly)) *"sailing is her favorite pastime"; "his main pastime is gambling"; "he counts reading among his interests"; "they criticized the boy for his limited pursuits"*

Verb:

- S (v) **interest** (excite the curiosity of; engage the interest of)
- S (v) concern, **interest**, occupy, worry (be on the mind of) *"I worry about the second Germanic consonant shift"*
- S (v) matter to, **interest** (be of importance or consequence) *"This matters to me!"*

"interest/4" – a closer look

S: (n) **interest** (a fixed charge for borrowing money; usually a percentage of the amount borrowed) "how much interest do you pay on your mortgage?"
direct hyponym / **full hyponym**

- S: (n) compound interest (interest calculated on both the principal and the accrued interest)
- S: (n) simple interest (interest paid on the principal alone)

direct hyponym/ **inherited hypernym** / sister term:

- S: (n) fixed charge, fixed cost, fixed costs (a periodic charge that does not vary with business volume (as insurance or rent or mortgage payments etc.))
  - S: (n) charge (the price charged for some article or service) "the admission charge"
    - S: (n) cost (the total spent for goods or services including money and time and labor)
      - S: (n) outgo, spending, expenditure, outlay (money paid out; an amount spent)
        - S: (n) transferred property, transferred possession (a possession whose ownership changes or lapses)
          - S: (n) possession (anything owned or possessed)
            - S: (n) relation (an abstraction belonging to or characteristic of two entities or parts together)
              - S: (n) abstraction, abstract entity (a general concept formed by extracting common features from specific examples)
                - S: (n) entity (that which is perceived or known or inferred to have its own distinct existence (living or nonliving))

"interest/5" – a closer look

S: (n) **interest**, stake ((law) a right or legal share of something; a financial involvement with something) *"they have interests all over the world"; "a stake in the company's future"*
direct hyponym/ **inherited hypernym** / sister term:

- S: (n) share, portion, part, percentage (assets belonging to or due to or contributed by an individual person or group) "he wanted his share in cash"
  - S: (n) assets (anything of material value or usefulness that is owned by a person or company)
    - S: (n) possession (anything owned or possessed)
      - S: (n) relation (an abstraction belonging to or characteristic of two entities or parts together)
        - S: (n) abstraction, abstract entity (a general concept formed by extracting common features from specific examples)
          - S : (n) entity (that which is perceived or known or inferred to have its own distinct existence (living or nonliving))

interest/4 and interest/5

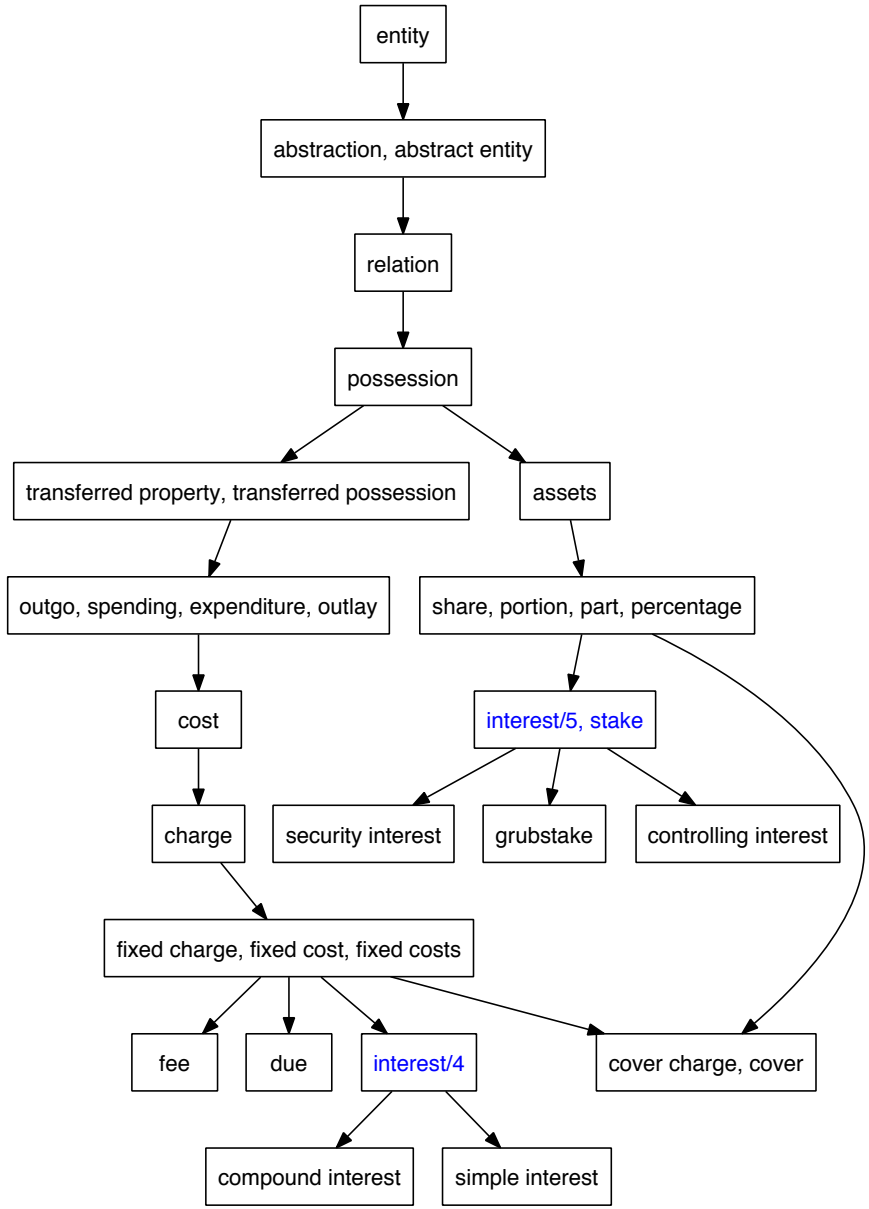

interest/1, interest/2 and interest/3

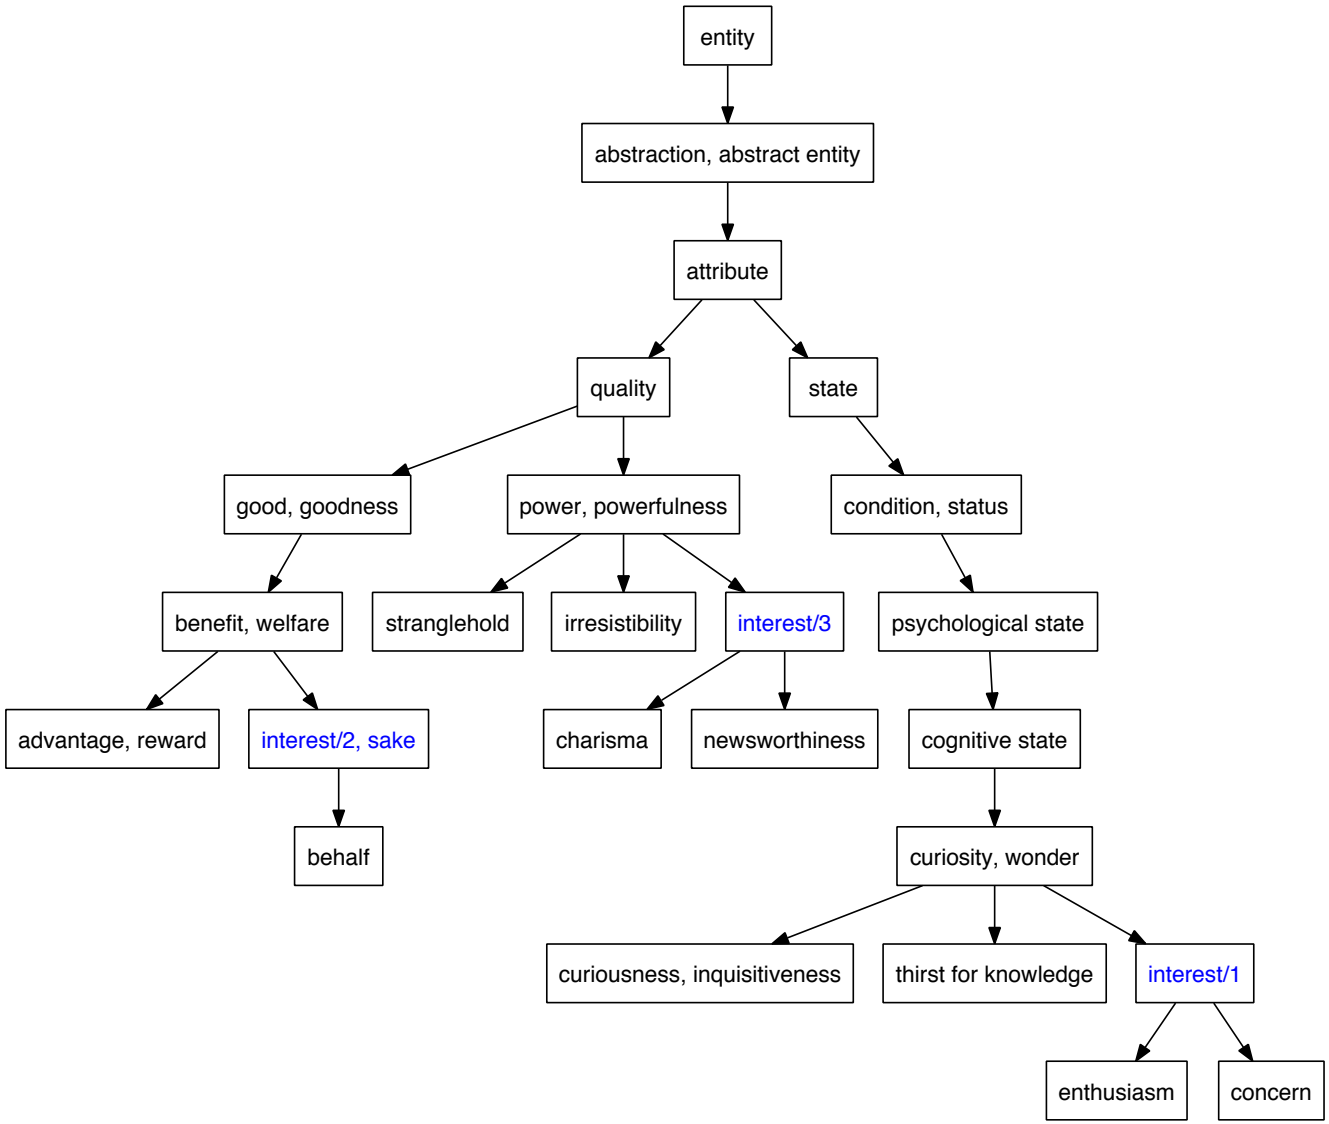

Interest – all senses

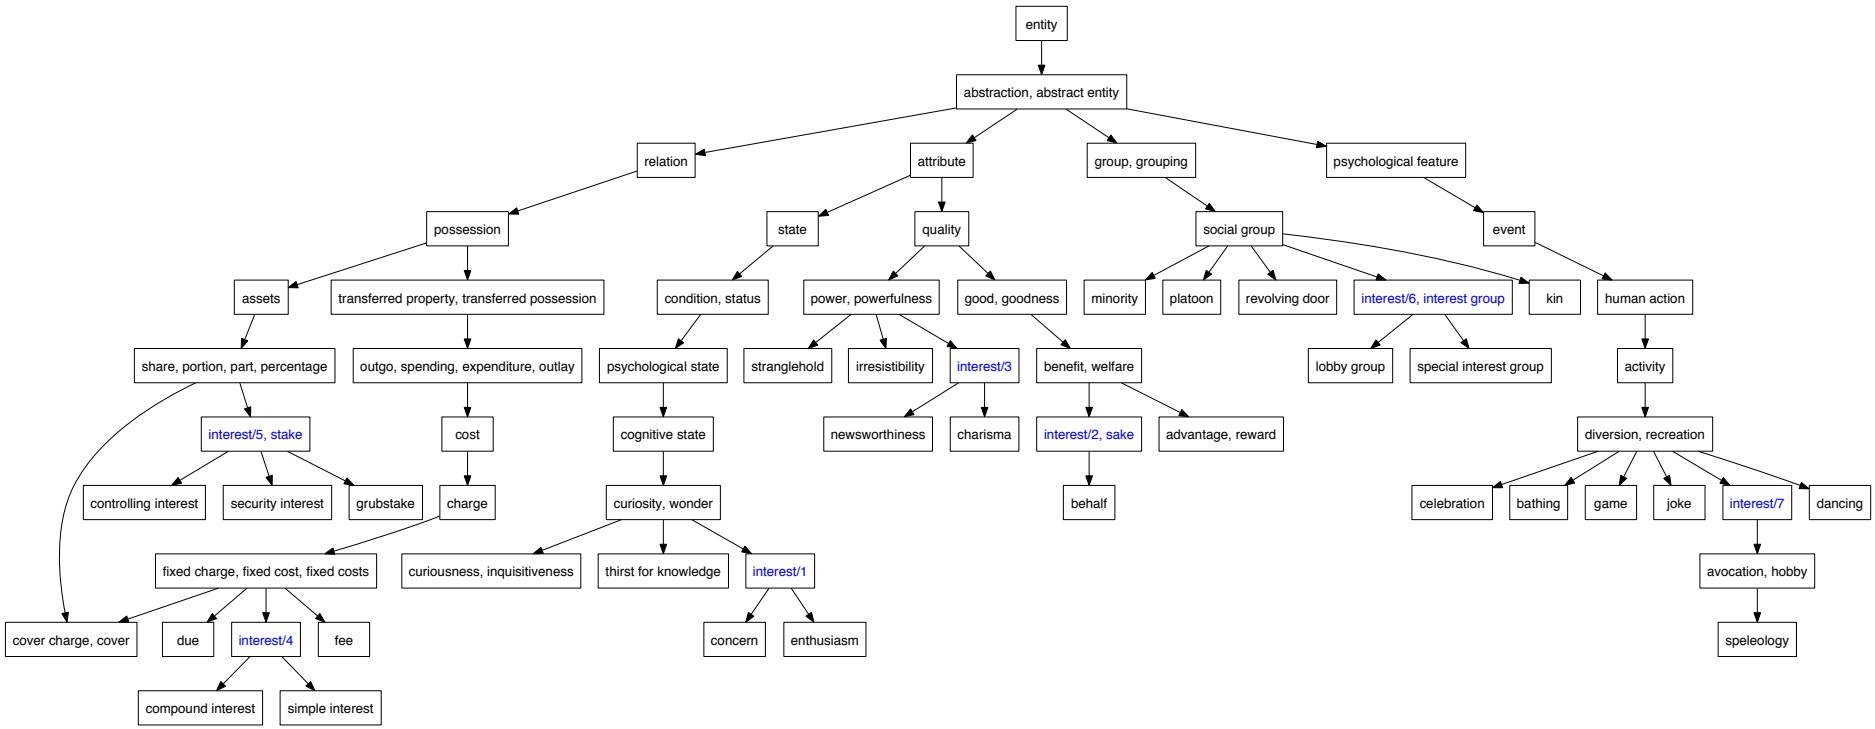

Multilingual aspect of word sense ambiguity

Interest translated into German

- Interesse: curiousness (interest/1)
- Interesse: sake (interest/2)
- Anziehungskraft: attractiveness (interest/3)
- Zins: financial charge (interest/4)
- Anteil: stake in company (interest/5)
- Lobbygruppe: interest group (interest/6)
- Hobby: pastime (interest/7)

Word Senses: Example contexts for *interest*

- *I only have your best* **interest**/_ *in mind.*
- *Primary colours can add* **interest**/_ *to a room.*
- *She pays 3%* **interest**/_ *on the loan.*
- *He showed a lot of* **interest**/_ *in the painting.*
- *Microsoft purchased a controlling* **interest**/_ *in Google.*
- *He said nothing of great* **interest**/_.
- *It is in the national* **interest**/_ *to invade the Bahamas.*
- *Playing chess is one of my* **interests**/_.
- *Business* **interests**/_ *lobbied for the legislation.*

(Invitation to perform Word Sense Disambiguation (WSD) on these examples; simply add WN sense numbers)

Coursework 1

- Perform all word WSD as per instructions on 3 example sentences
- Deadline in one week
- This coursework is ticked