Lexicon, Syntax, Semantics II: Modeling Meaning

Introduction

Eva Maria Vecchi

Center for Information and Language Processing
LMU Munich

April 23, 2020

Scheduling

- Lectures: Thursdays 10:00–12:00, (Zoom, for now)

Scheduling

- Lectures: Thursdays 10:00–12:00, (Zoom, for now)
- Lecture Period: April 23 – July 23, 2020

Scheduling

- Lectures: Thursdays 10:00–12:00, (Zoom, for now)
- Lecture Period: April 23 – July 23, 2020
- Final Exam: July 16/23, 2019

Scheduling

- Lectures: Thursdays 10:00–12:00, (Zoom, for now)
- Lecture Period: April 23 – July 23, 2020
- Final Exam: July 16/23, 2019
- Course Website:
  `www.vecchi.com/eva/teaching/modelingmeaning.html`

Scheduling

- Lectures: Thursdays 10:00–12:00, (Zoom, for now)
- Lecture Period: April 23 – July 23, 2020
- Final Exam: July 16/23, 2019
- Course Website:
  `www.vecchi.com/eva/teaching/modelingmeaning.html`

- My Contact: `evamariavecchi@gmail.com`

What you'll gain from this course

- What is "meaning" in language?
- How do we model it?
  - . . . theoretically and computationally
- Why is it important to understand and study this?
- What obstacles are we facing?
- Can we turn Plato's theories into AI systems? (*eyeroll*)

What you'll gain from this course

- What is "meaning" in language?
- How do we model it?
  - ...theoretically and computationally
- Why is it important to understand and study this?
- What obstacles are we facing?
- Can we turn Plato's theories into AI systems? (*eyeroll*)
- What I expect from you:
  - Read research papers and recent literature
  - Step outside of the box
  - Participate in class exercises and discussion

Final Exam: Research Project

- Research a chosen topic discussed in lecture, submit (and present?) personal project
- Focus on theory and practice
  - **Theory** Scientific literature review
  - **Practice** Implementation of a system, report the design and results
- . . . more details in 3 weeks

Eva Maria Vecchi

- **2005**: B.A. in Linguistics, with minors in Mathematics and Italian from University of Colorado at Boulder
- **2007**: M.Sc. in Computational Linguistics from Georgetown University
- **2005-2007**: Worked at MITRE and The Federation of American Scientists
- **2008-2010**: Research Fellow at Center for National Research at Pisa, Italy
- **2010-2013**: Ph.D. in Cognitive and Brain Sciences at CIMeC, University of Trento, Italy
- **2014-2017**: Post-doctoral researcher at University of Cambridge: Computer Laboratory
- **2017-2018**: Post-doctoral researcher at University of Stuttgart, IMS

Research Interests

- Distributional/Computational Semantics
- Compositionality & Formal Semantics
- Machine Learning in NLP
- Meaning Representation
- Cognitive Computing

Modeling Meaning: ***Why?***

Ultimate Goal: "Understand" language like humans

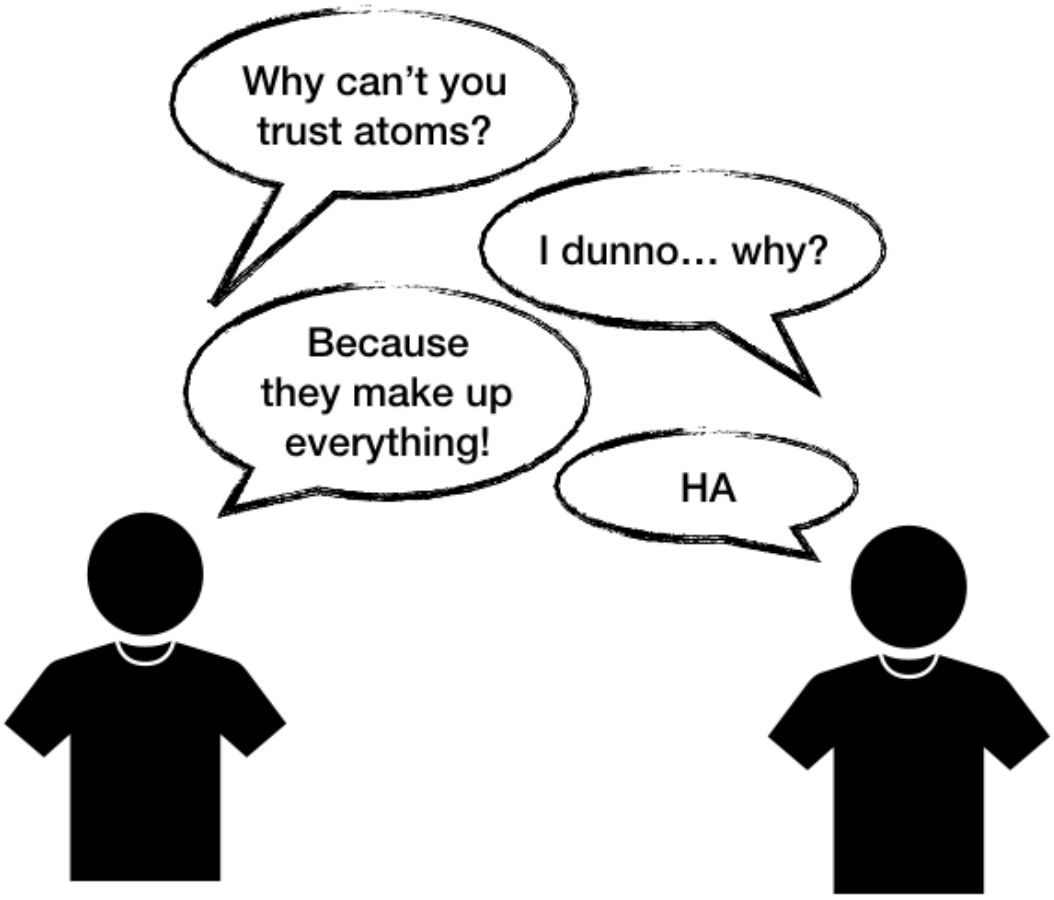

Ultimate Goal: "Understand" language like humans

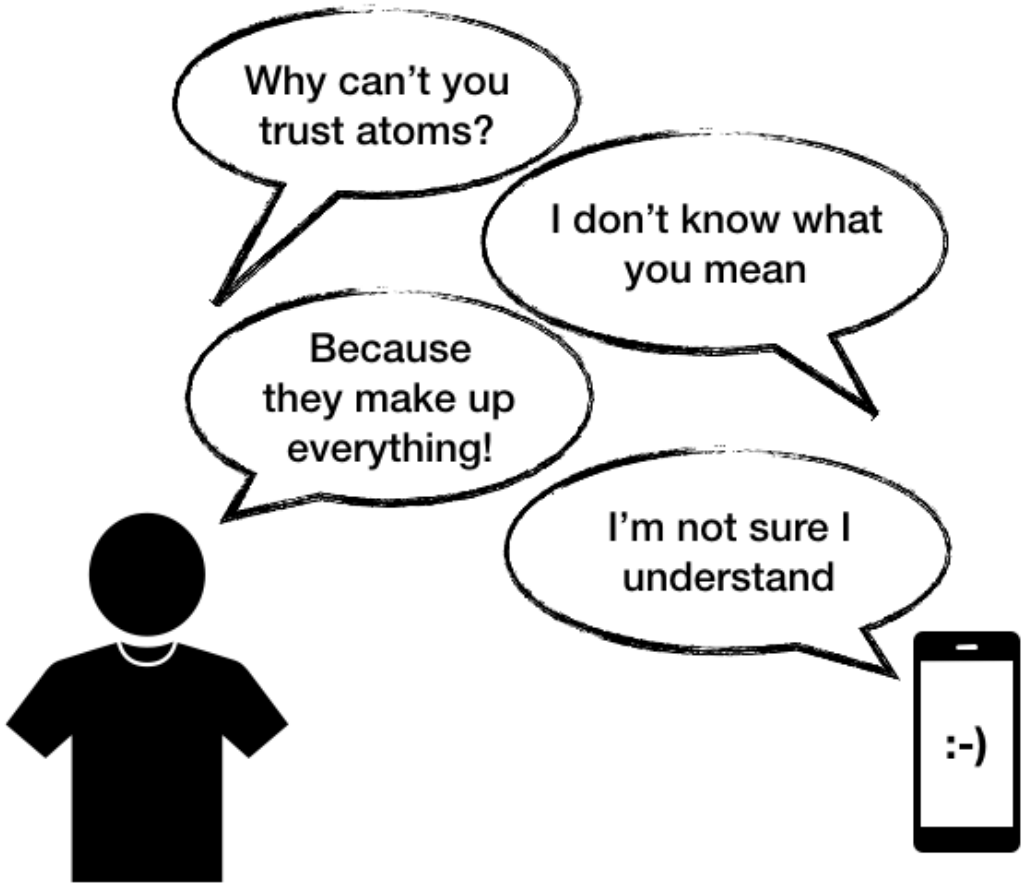

Language: one of the most fundamental human abilities

- Language allows us to speak about complex objects that (we think) **exist in the real world**:
  - *Look at that chair with the velvet back, the one with the flowery English pattern.*
  - *Insulin is a peptide hormone produced by beta cells of the pancreatic islets.*
  - *I'm jealous. It's not that I want that car, but I don't think he should have it either.*
  - *Bring the curd to a boil, let it boil for exactly three minutes whilst gently stirring.*

Spot the phenomena:

Language: one of the most fundamental human abilities

- Language allows us to speak about complex objects that (we think) **exist in the real world**:
  - *Look at that chair with the velvet back, the one with the flowery English pattern.*
  - *Insulin is a peptide hormone produced by beta cells of the pancreatic islets.*
  - *I'm jealous. It's not that I want that car, but I don't think he should have it either.*
  - *Bring the curd to a boil, let it boil for exactly three minutes whilst gently stirring.*

Spot the phenomena: reference to world entity at time t,

Language: one of the most fundamental human abilities

- Language allows us to speak about complex objects that (we think) **exist in the real world**:
  - *Look at that chair with the velvet back, the one with the flowery English pattern.*
  - *Insulin is a peptide hormone produced by beta cells of the pancreatic islets.*
  - *I'm jealous. It's not that I want that car, but I don't think he should have it either.*
  - *Bring the curd to a boil, let it boil for exactly three minutes whilst gently stirring.*

Spot the phenomena: reference to world entity at time t, cultural knowledge,

Language: one of the most fundamental human abilities

- Language allows us to speak about complex objects that (we think) **exist in the real world**:
  - *Look at that chair with the velvet back, the one with the flowery English pattern.*
  - *Insulin is a peptide hormone produced by beta cells of the pancreatic islets.*
  - *I'm jealous. It's not that I want that car, but I don't think he should have it either.*
  - *Bring the curd to a boil, let it boil for exactly three minutes whilst gently stirring.*

Spot the phenomena: reference to world entity at time t, cultural knowledge, vocabulary of speaker,

Language: one of the most fundamental human abilities

- Language allows us to speak about complex objects that (we think) **exist in the real world**:
  - *Look at that chair with the velvet back, the one with the flowery English pattern.*
  - *Insulin is a peptide hormone produced by beta cells of the pancreatic islets.*
  - *I'm jealous. It's not that I want that car, but I don't think he should have it either.*
  - *Bring the curd to a boil, let it boil for exactly three minutes whilst gently stirring.*

Spot the phenomena: reference to world entity at time t, cultural knowledge, vocabulary of speaker, abstract, non-referential concepts,

Language: one of the most fundamental human abilities

- Language allows us to speak about complex objects that (we think) **exist in the real world**:
  - *Look at that chair with the velvet back, the one with the flowery English pattern.*
  - *Insulin is a peptide hormone produced by beta cells of the pancreatic islets.*
  - *I'm jealous. It's not that I want that car, but I don't think he should have it either.*
  - *Bring the curd to a boil, let it boil for exactly three minutes whilst gently stirring.*

Spot the phenomena: reference to world entity at time *t*, cultural knowledge, vocabulary of speaker, abstract, non-referential concepts, reference to events and processes.

Credit: Herbelot 2018

Language: one of the most fundamental human abilities

- Language lets us communicate about things **that don't exist (yet) in the real world**:
  - *Please sit down.* (You haven't yet)

(Credit: Herbelot 2018)

Language: one of the most fundamental human abilities

- Language lets us communicate about things **that don't exist (yet) in the real world**:
  - *Please sit down.* (You haven't yet)
  - *If it snows tomorrow . . .*

(Credit: Herbelot 2018)

Language: one of the most fundamental human abilities

- Language lets us communicate about things **that don't exist (yet) in the real world**:
  - *Please sit down.* (You haven't yet)
  - *If it snows tomorrow . . .*
  - *The unicorn flew over the castle.*

(Credit: Herbelot 2018)

Language: one of the most fundamental human abilities

- Language lets us communicate about things **that don't exist (yet) in the real world**:
  - *Please sit down.* (You haven't yet)
  - *If it snows tomorrow . . .*
  - *The unicorn flew over the castle.*
  - *What if the universe had not 4 but 10 dimensions?*

(Credit: Herbelot 2018)

Language: one of the most fundamental human abilities

- Language lets us communicate about things **that don't exist (yet) in the real world**:
  - *Please sit down.* (You haven't yet)
  - *If it snows tomorrow . . .*
  - *The unicorn flew over the castle.*
  - *What if the universe had not 4 but 10 dimensions?*
  - *Let's have a picnic on Sunday.*

(Credit: Herbelot 2018)

Language: one of the most fundamental human abilities

- Language lets us communicate about things **that don't exist (yet) in the real world**:
  - *Please sit down.* (You haven't yet)
  - *If it snows tomorrow . . .*
  - *The unicorn flew over the castle.*
  - *What if the universe had not 4 but 10 dimensions?*
  - *Let's have a picnic on Sunday.*
- In fact, utterances can be meaningless if *meaning* is taken to be a relation between words and items in the real world.

(Credit: Herbelot 2018)

Language for communication

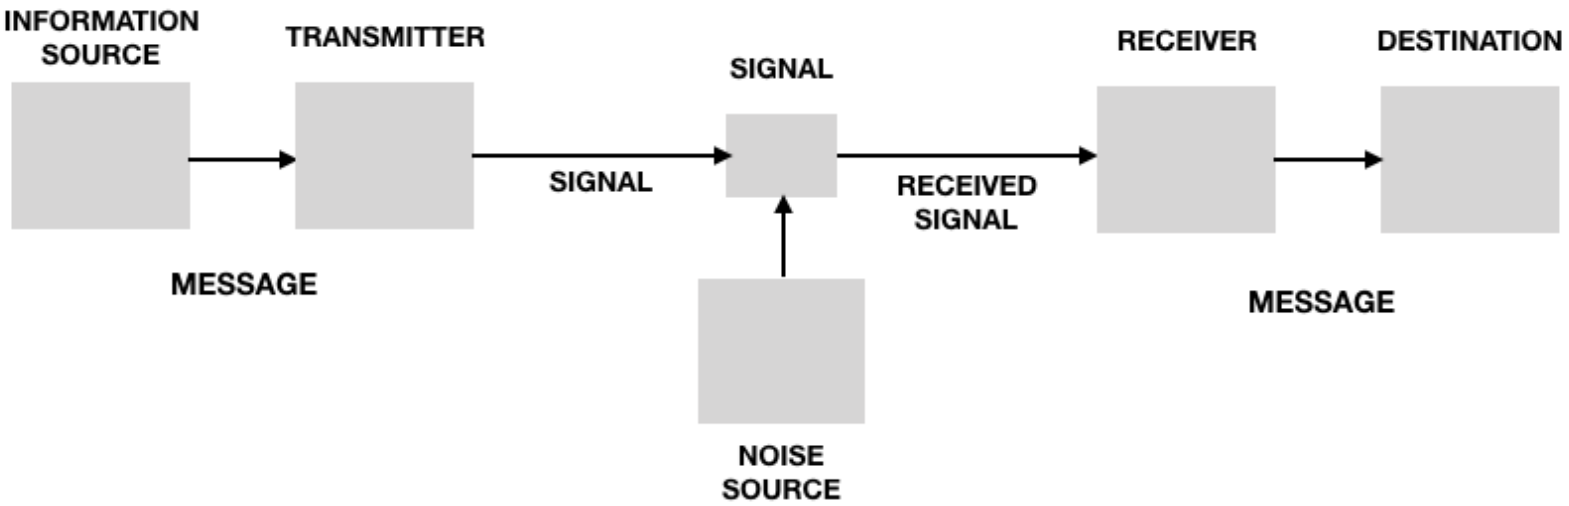

- A message has to be passed in a way that whatever was in the head of the speaker ends up in the head of the hearer.

(Credit: Herbelot 2018)

Language for conceptualization

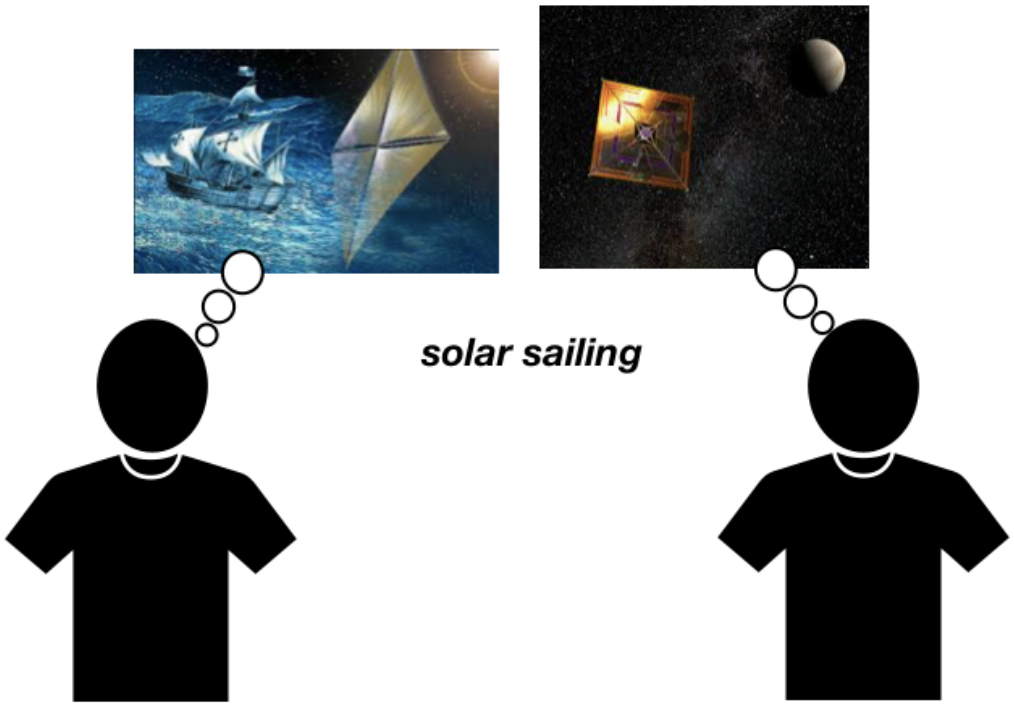

- Composition of existing linguistic constituents lead to new concepts.
- Each speaker has their own conceptual space, made of previous experiences and their own (infinite) creativity.

(Credit: Herbelot 2018)

Ultimate Goal: "Understand" language like humans

But how?

Ultimate Goal: "Understand" language like humans

But how?

1. Define *meaning*: What does it mean to understand language? What are the components of meaning?

Ultimate Goal: "Understand" language like humans

But how?

1. Define *meaning*: What does it mean to understand language? What are the components of meaning?
2. Define and understand phenomena involved in human capacity to process meaning in language

Ultimate Goal: "Understand" language like humans

But how?

1. Define *meaning*: What does it mean to understand language? What are the components of meaning?
2. Define and understand phenomena involved in human capacity to process meaning in language
3. Determine methods that computationally elaborate components of meaning such that they approximate the phenomena of natural language **and** remain computationally-friendly and efficient.

Thought exercise: Defining meaning

1. What is the meaning of the word *ferret* (*das Frettchen*)?

Thought exercise: Defining meaning

1. What is the meaning of the word *ferret* (*das Frettchen*)?
   - "Domesticated albino variety of the polecat [...] bred for hunting rabbits, rats, etc"

Thought exercise: Defining meaning

1. What is the meaning of the word *ferret* (*das Frettchen*)?
   - "Domesticated albino variety of the polecat [...] bred for hunting rabbits, rats, etc"
   - Problem: **Circularity** of defining words with other words

Thought exercise: Defining meaning

1. What is the meaning of the word *ferret* (*das Frettchen*)?
   - "Domesticated albino variety of the polecat [...] bred for hunting rabbits, rats, etc"
   - Problem: **Circularity** of defining words with other words
2. *I dreamt that I was swallowed by a whale.*

Thought exercise: Defining meaning

1. What is the meaning of the word *ferret* (*das Frettchen*)?
   - "Domesticated albino variety of the polecat [...] bred for hunting rabbits, rats, etc"
   - Problem: **Circularity** of defining words with other words
2. *I dreamt that I was swallowed by a whale.*
   - SPEAKER: Believes *whale* is a mammal

Thought exercise: Defining meaning

1. What is the meaning of the word *ferret* (*das Frettchen*)?
   - "Domesticated albino variety of the polecat [...] bred for hunting rabbits, rats, etc"
   - Problem: **Circularity** of defining words with other words
2. *I dreamt that I was swallowed by a whale.*
   - SPEAKER: Believes *whale* is a mammal
   - LISTENER: Believes *whale* is a fish

Thought exercise: Defining meaning

1. What is the meaning of the word *ferret* (*das Frettchen*)?
   - "Domesticated albino variety of the polecat [...] bred for hunting rabbits, rats, etc"
   - Problem: **Circularity** of defining words with other words
2. *I dreamt that I was swallowed by a whale.*
   - SPEAKER: Believes *whale* is a mammal
   - LISTENER: Believes *whale* is a fish
   - Do they have different meanings of *whale*, and do they still understand each other?

Thought exercise: Defining meaning

1. What is the meaning of the word *ferret* (*das Frettchen*)?
   - "Domesticated albino variety of the polecat [...] bred for hunting rabbits, rats, etc"
   - Problem: **Circularity** of defining words with other words
2. *I dreamt that I was swallowed by a whale.*
   - SPEAKER: Believes *whale* is a mammal
   - LISTENER: Believes *whale* is a fish
   - Do they have different meanings of *whale*, and do they still understand each other?
   - Issue of **linguistic** vs. **encyclopedic** knowledge

Thought exercise: Defining meaning

1. What is the meaning of the word *ferret* (*das Frettchen*)?
   - "Domesticated albino variety of the polecat [...] bred for hunting rabbits, rats, etc"
   - Problem: **Circularity** of defining words with other words
2. *I dreamt that I was swallowed by a whale.*
   - SPEAKER: Believes *whale* is a mammal
   - LISTENER: Believes *whale* is a fish
   - Do they have different meanings of *whale*, and do they still understand each other?
   - Issue of **linguistic** vs. **encyclopedic** knowledge
3. *Marvelous weather you have here in Ireland.*

Thought exercise: Defining meaning

1. What is the meaning of the word *ferret* (*das Frettchen*)?
   - "Domesticated albino variety of the polecat [...] bred for hunting rabbits, rats, etc"
   - Problem: **Circularity** of defining words with other words
2. *I dreamt that I was swallowed by a whale.*
   - SPEAKER: Believes *whale* is a mammal
   - LISTENER: Believes *whale* is a fish
   - Do they have different meanings of *whale*, and do they still understand each other?
   - Issue of **linguistic** vs. **encyclopedic** knowledge
3. *Marvelous weather you have here in Ireland.*
   - INSTANCE_1: Cloudless sunny day

Thought exercise: Defining meaning

1. What is the meaning of the word *ferret* (*das Frettchen*)?
   - "Domesticated albino variety of the polecat [...] bred for hunting rabbits, rats, etc"
   - Problem: **Circularity** of defining words with other words
2. *I dreamt that I was swallowed by a whale.*
   - SPEAKER: Believes *whale* is a mammal
   - LISTENER: Believes *whale* is a fish
   - Do they have different meanings of *whale*, and do they still understand each other?
   - Issue of **linguistic** vs. **encyclopedic** knowledge
3. *Marvelous weather you have here in Ireland.*
   - INSTANCE_1: Cloudless sunny day
   - INSTANCE_2: Rain is pouring down outside

Thought exercise: Defining meaning

1. What is the meaning of the word *ferret* (*das Frettchen*)?
   - "Domesticated albino variety of the polecat [...] bred for hunting rabbits, rats, etc"
   - Problem: **Circularity** of defining words with other words
2. *I dreamt that I was swallowed by a whale.*
   - SPEAKER: Believes *whale* is a mammal
   - LISTENER: Believes *whale* is a fish
   - Do they have different meanings of *whale*, and do they still understand each other?
   - Issue of **linguistic** vs. **encyclopedic** knowledge
3. *Marvelous weather you have here in Ireland.*
   - INSTANCE_1: Cloudless sunny day
   - INSTANCE_2: Rain is pouring down outside
   - Contribution of **context** to meaning

Overdoing it...

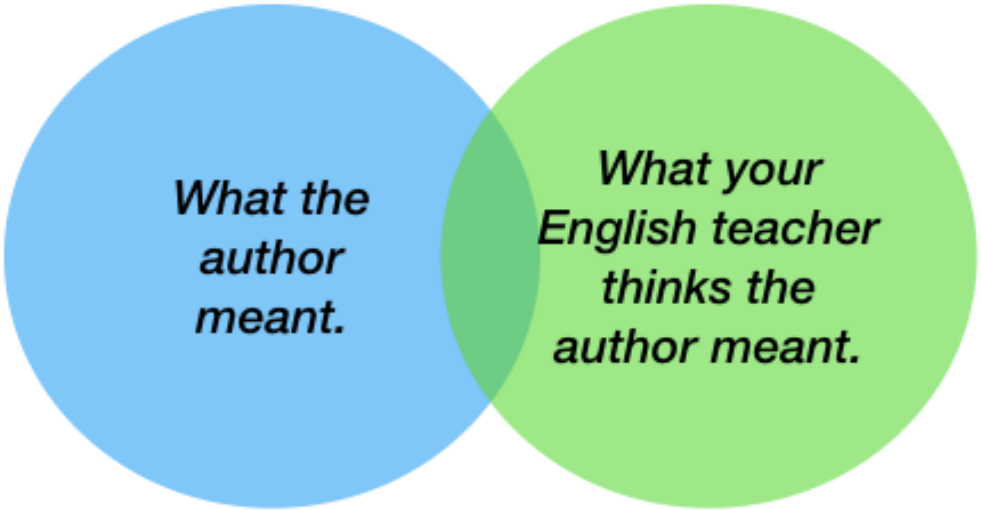

For instance: "The curtains were blue."

What your teacher thinks: "The curtains represent his immense depression and his lack of will to carry on."

*What the author meant: "The curtains were *%$!& blue."*

Courtesy of *The Language Nerds*

Meaning in a nutshell

- Meaning is **extension** / reference: it 'points' at things in the world;

(Credit: Herbelot 2018)

Meaning in a nutshell

- Meaning is **extension** / reference: it 'points' at things in the world;
- Meaning is **intension**: intensions are mappings from possible worlds to extensions;
  - e.g. the properties something has to have to count as a member of a category and be identified by a linguistic expression

(Credit: Herbelot 2018)

Meaning in a nutshell

- Meaning is **extension** / reference: it 'points' at things in the world;
- Meaning is **intension**: intensions are mappings from possible worlds to extensions;
  - e.g. the properties something has to have to count as a member of a category and be identified by a linguistic expression
- Meaning is **conceptual**: linguistic constituents activate cognitive processes involving extra-linguistic features;

(Credit: Herbelot 2018)

Meaning in a nutshell

- Meaning is **extension** / reference: it 'points' at things in the world;
- Meaning is **intension**: intensions are mappings from possible worlds to extensions;
  - e.g. the properties something has to have to count as a member of a category and be identified by a linguistic expression
- Meaning is **conceptual**: linguistic constituents activate cognitive processes involving extra-linguistic features;
- Meaning is **use**: linguistics constituents have certain patterns of use across a community

(Credit: Herbelot 2018)

Modeling Meaning

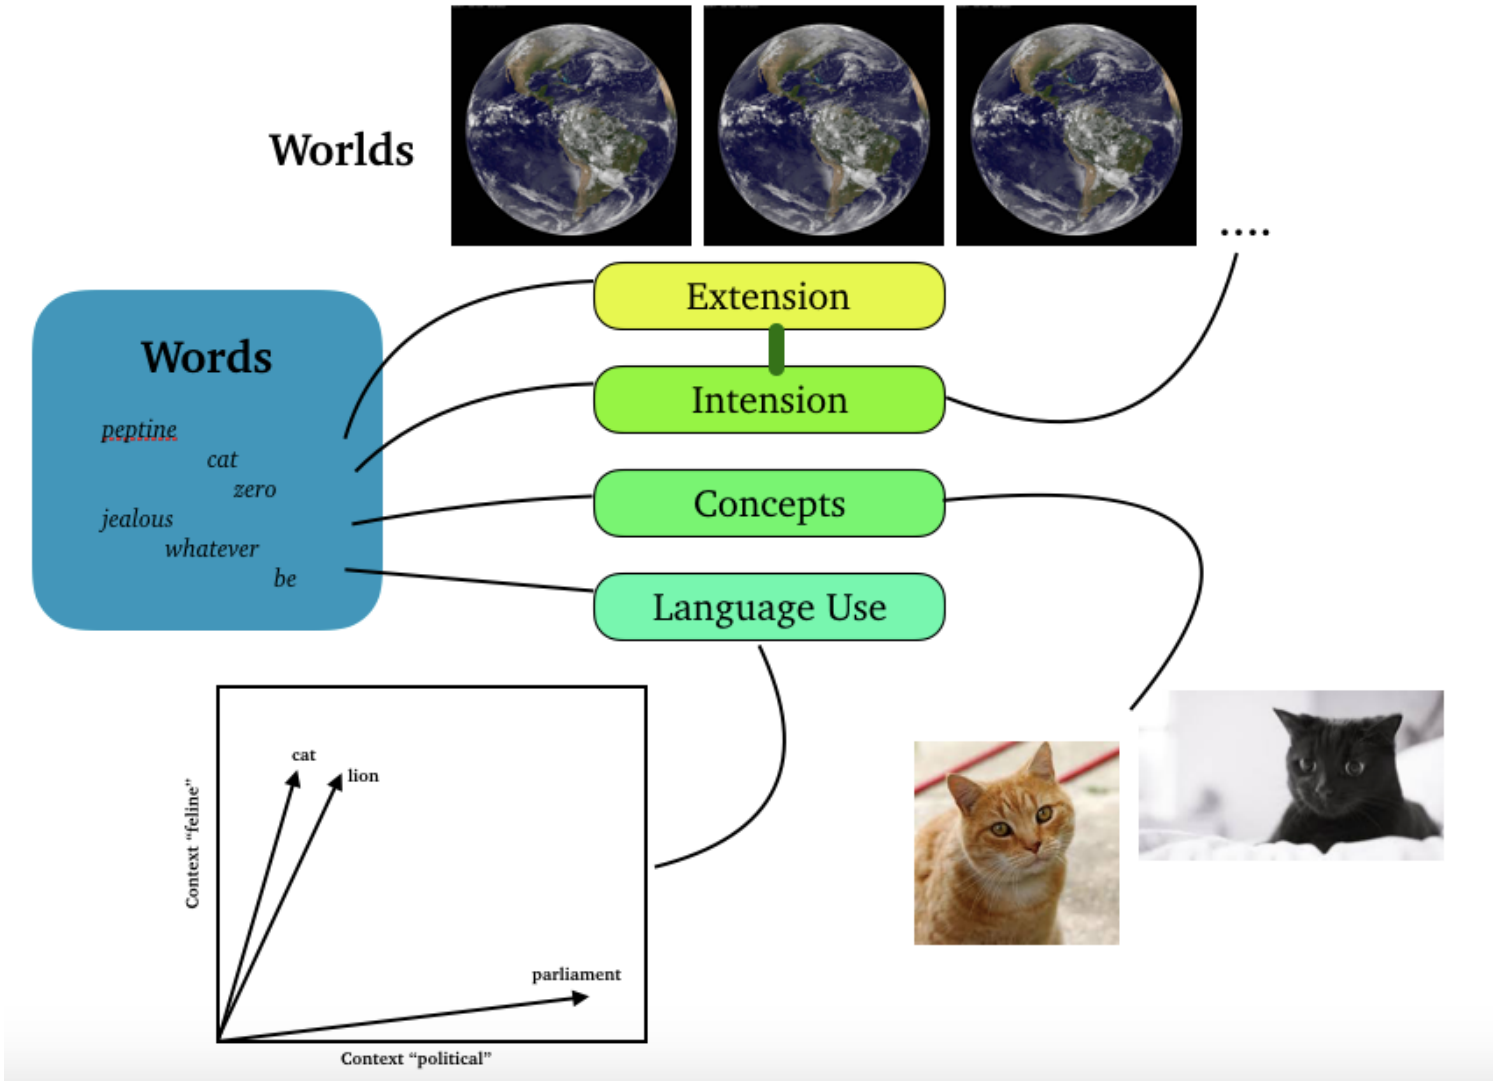

(Credit: Herbelot 2018)

Integration with the real world

- A system that learns extensions must be able to link language to the world.
- What is a bike, a bus? What do *how*, *many* and *how many* mean?
- Formal meaning of bus: *bus'* (the set of buses in a world)
- *bus'* assumes that the speaker holds entities of *bus* in their head (a model), but how do those relate to the actual world? Via perception.

How many bikes are there?	2 2 2	3 4 12
What number is the bus?	48 48 48	4 46 number 6

(Credit: Herbelot 2018)

Constraints: Communication

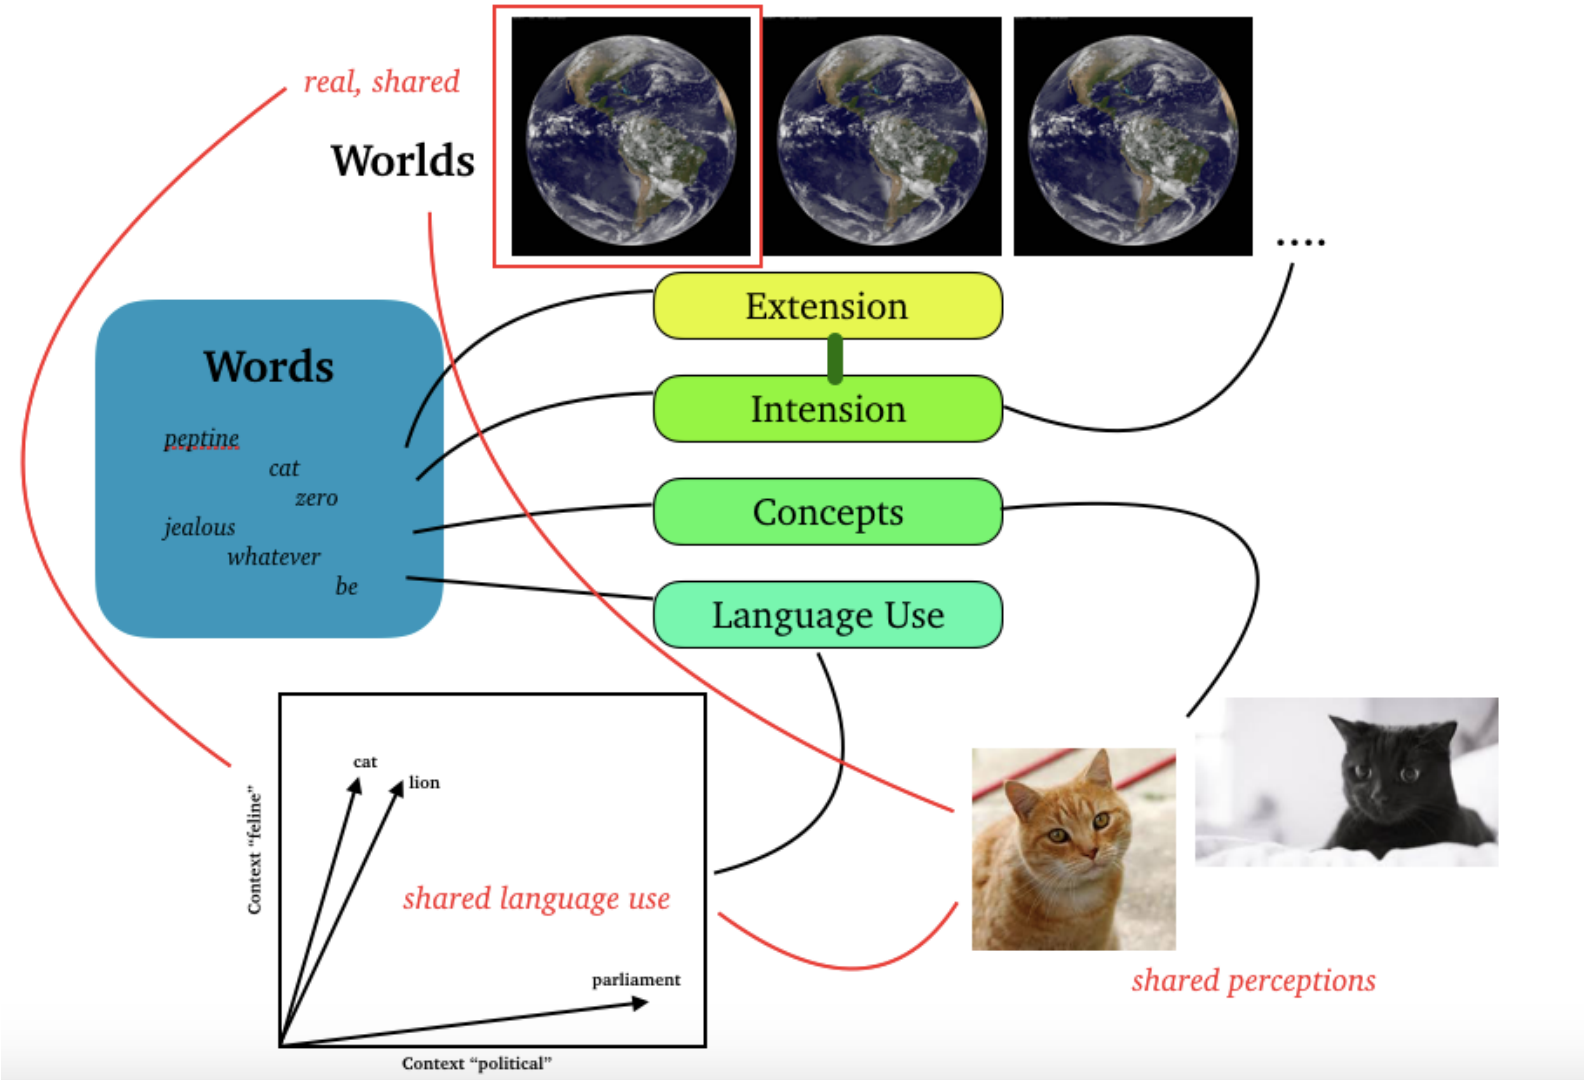

(Credit: Herbelot 2018)

Constraints: Communication

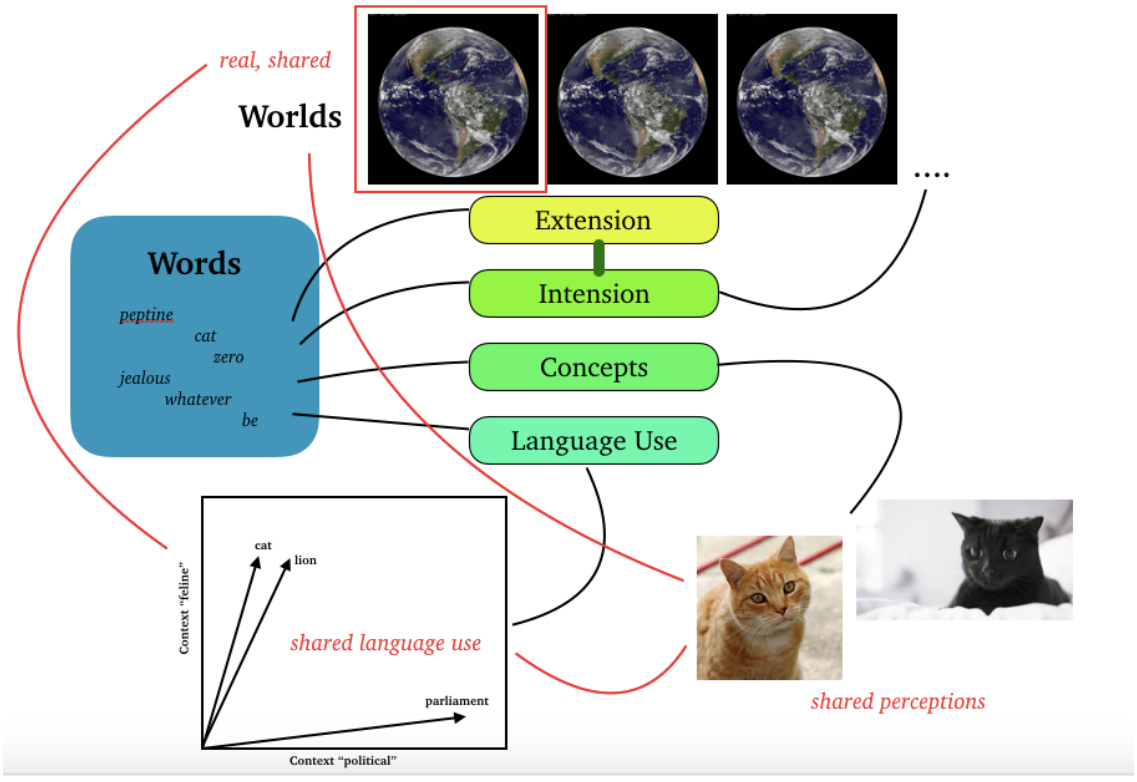

Since we share the real world and (roughly) share our perceptions of it, sharing labels (words) for certain categories allows us to successfully refer. Since we see the same world, we are bound to say the same kind of things about it, so we share language use, which is anchored in perception.

(Credit: Herbelot 2018)

Constraints: Individuality

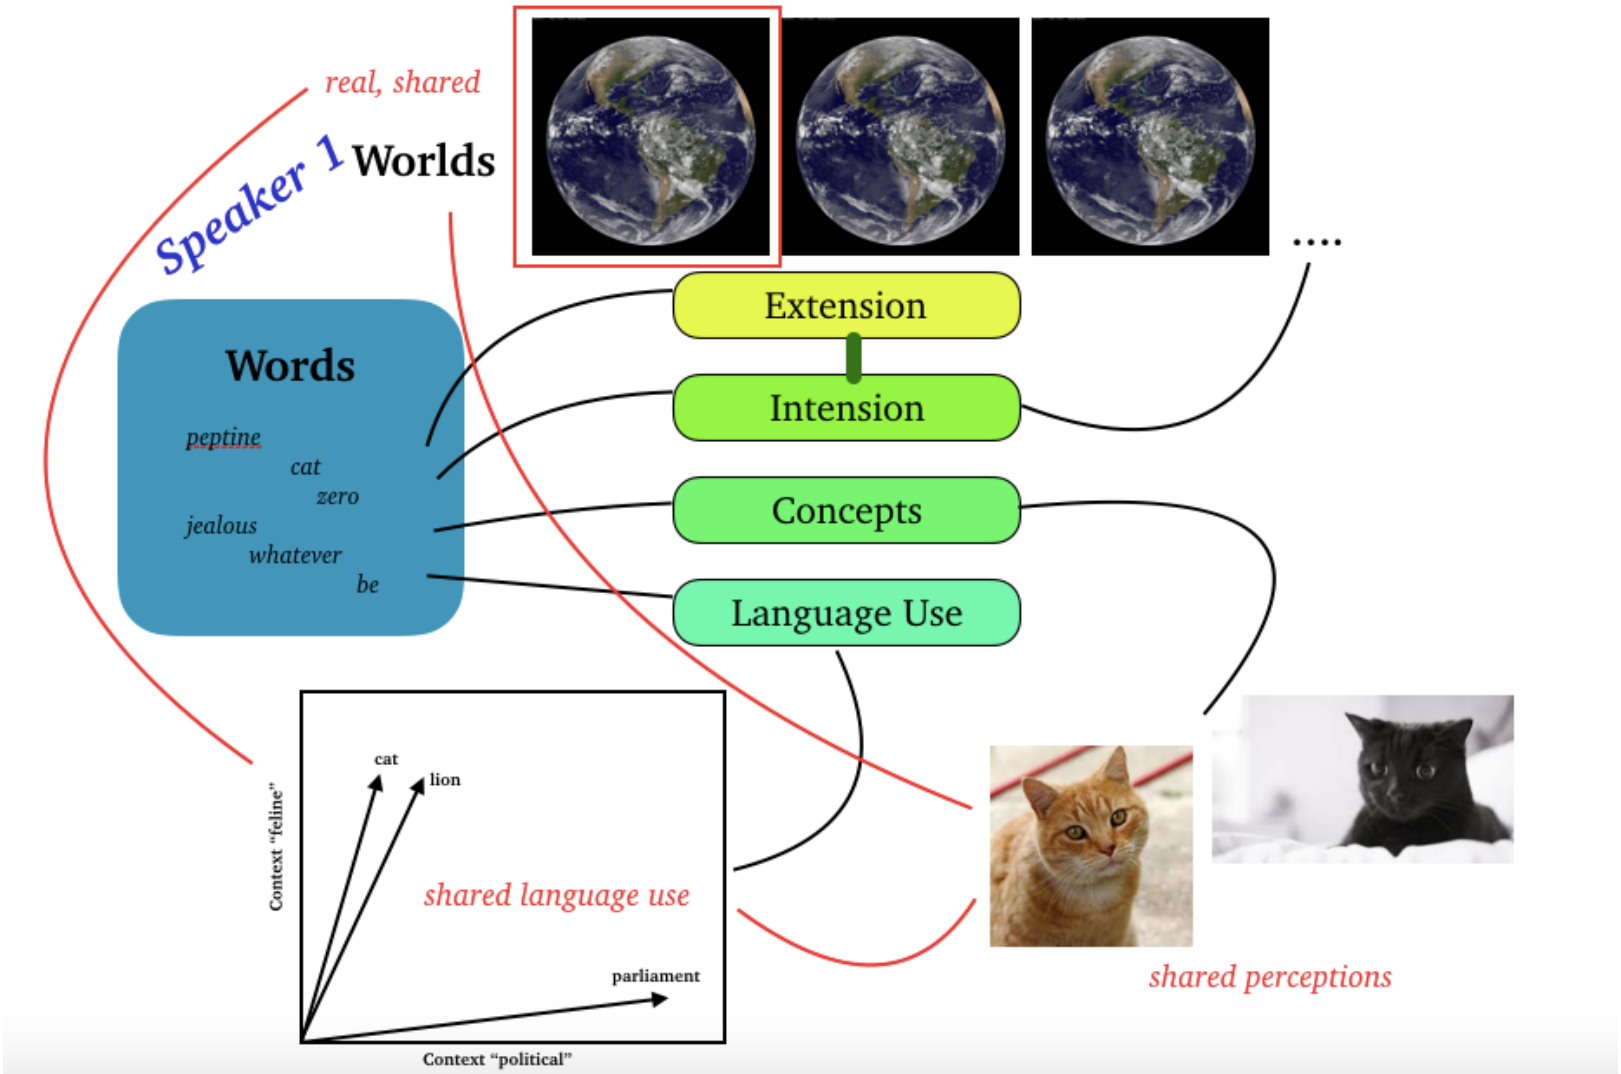

(Credit: Herbelot 2018)

Constraints: Individuality

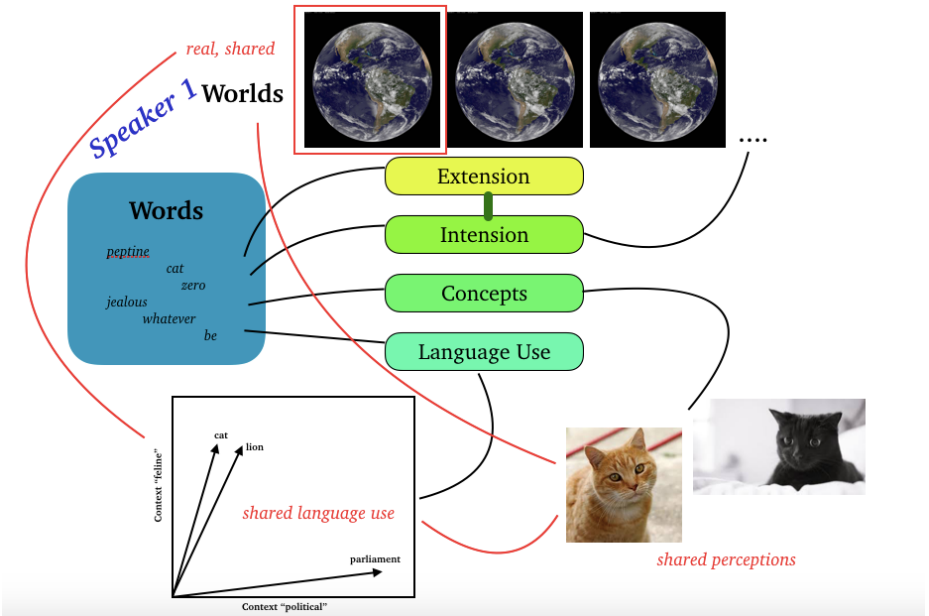

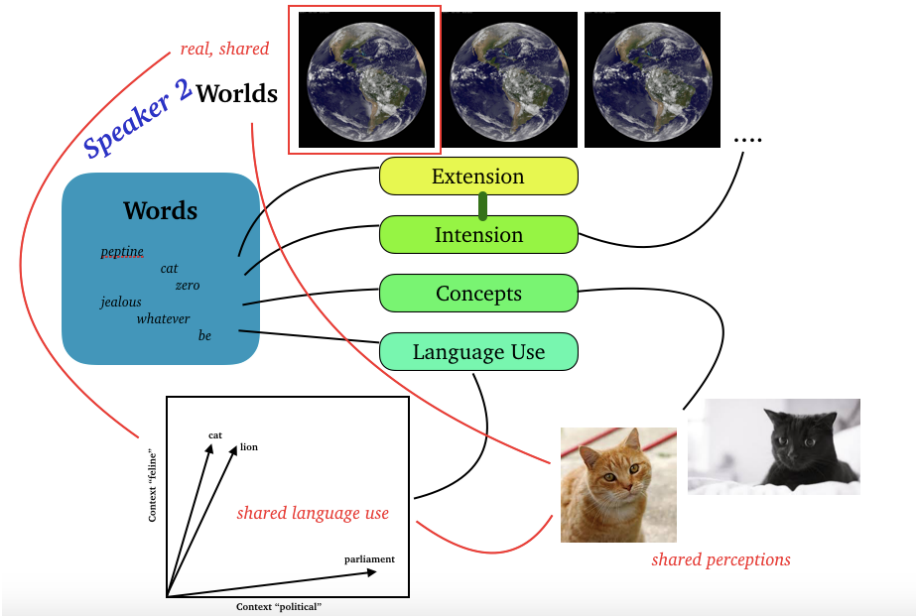

Since we don't share the same experiences of the world, and since concept composition happens in the mind (separate minds), we are bound to have separate conceptual spaces, separate language use, separate extensions, separate possible worlds, separate intensions.

(Credit: Herbelot 2018)

Larger constituents and compositionality

Larger constituents and compositionality

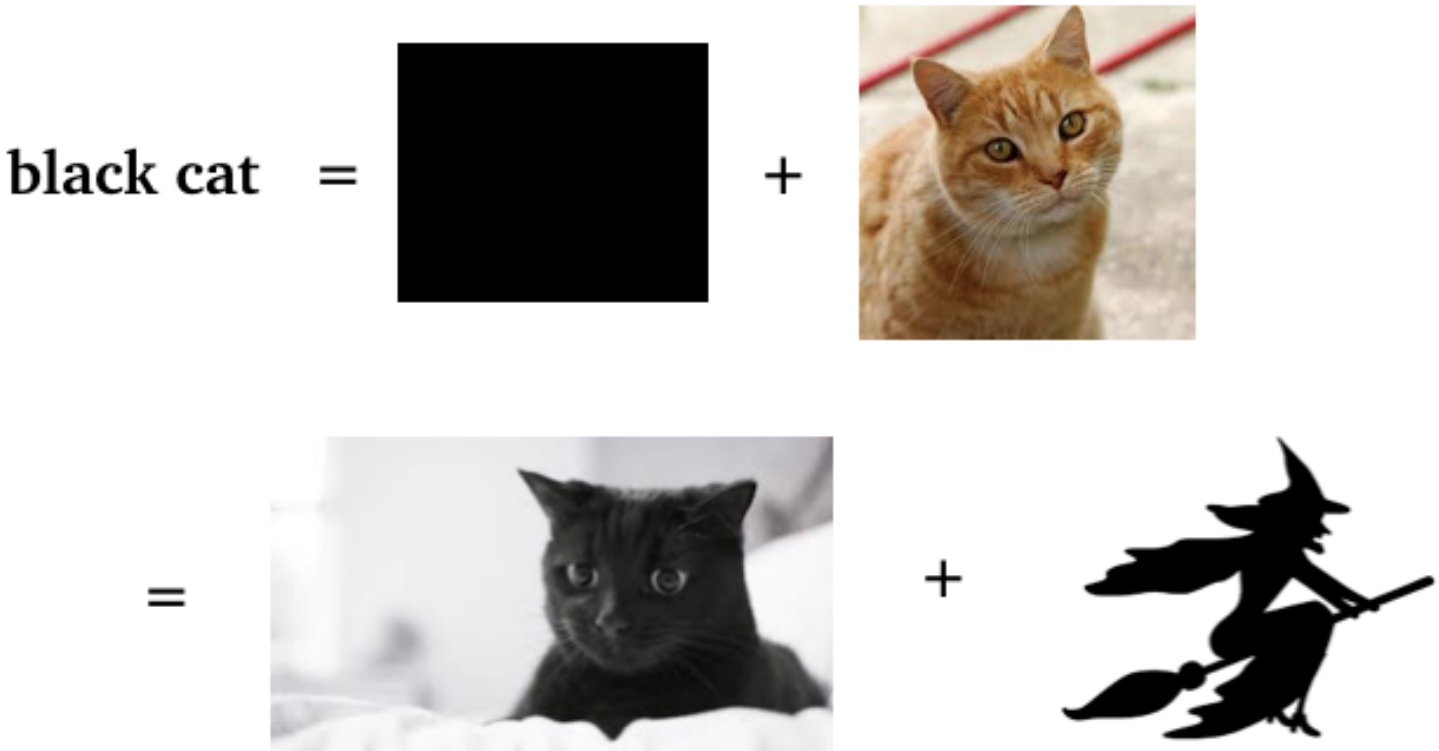

(Credit: Herbelot 2018)

Larger constituents and compositionality

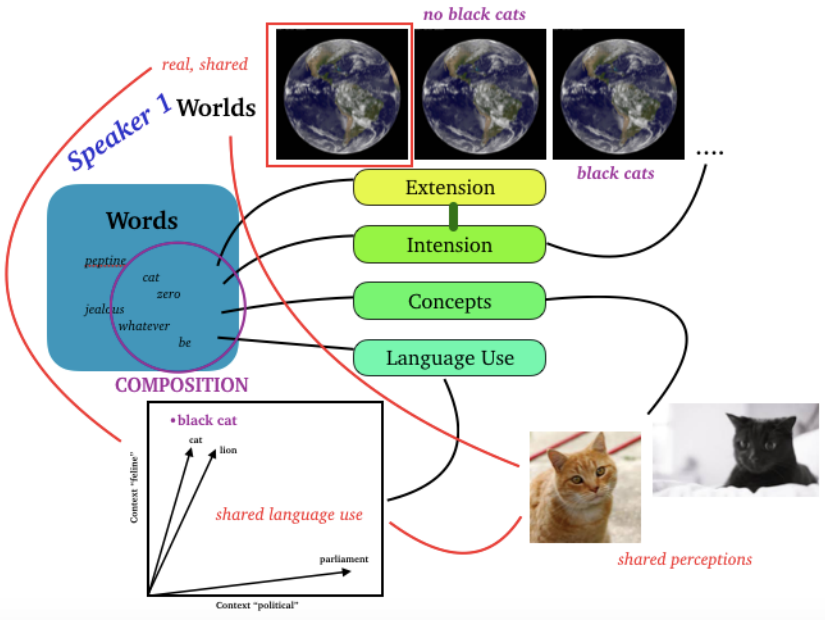

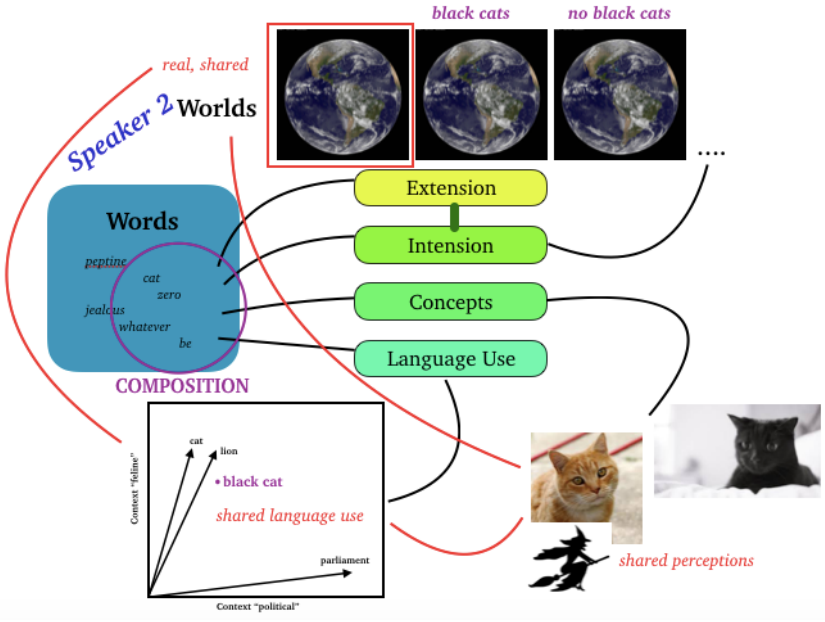

(Credit: Herbelot 2018)

Language in Use

1. Pragmatics
   - How does the *broader* context affect meaning? (situation of utterance, community of speaker, etc)
   - How does a community contribute to the emergence and spread of meanings? (del Tredici & Fernandez, 2017)

Language in Use

Katherine Johnson, *computer* at NASA, 1966

Much less reliable *computer* at NASA, late 1960s

2. Dynamicity of the lexicon
   - The use of *computer* changed meaning over the years.
   - Every utterance changes the meaning to its linguistic constituents.
   - Members of a community are not necessarily fully in sync.

(Credit: Herbelot 2018)

How hard is it to model meaning, then?

(Example from Teufel 2017)

Lecture 1: Introduction

Hurray! You've made it this far :-)

Lecture 2, 3, 4: Semantics and Theory

2. Lexical Semantics
   - How can we define and express what individual words mean?
   - How do they interact with the words "around" them?
   - Problems with pinning down word meaning
   - Implications for NLP tasks

Lecture 2, 3, 4: Semantics and Theory

2. Lexical Semantics
   - How can we define and express what individual words mean?
   - How do they interact with the words "around" them?
   - Problems with pinning down word meaning
   - Implications for NLP tasks

3. Formal/Compositional Semantics
   - Formalizing the meaning of phrases/sentences
   - Entailment
   - Model-theoretic semantics
   - Montague Semantics

Lecture 2, 3, 4: Semantics and Theory

2. Lexical Semantics
   - How can we define and express what individual words mean?
   - How do they interact with the words "around" them?
   - Problems with pinning down word meaning
   - Implications for NLP tasks

3. Formal/Compositional Semantics
   - Formalizing the meaning of phrases/sentences
   - Entailment
   - Model-theoretic semantics
   - Montague Semantics

4. Cognitive Semantics
   - Cognitive Semantics: Introduction to the field and motivation
   - Polysemy, Metaphor and Metonymy
   - Mental Spaces
   - Implications for modeling meaning in NLP: Language, Cognition, and Computation

Lecture 5, 6, 7, 8: Distributional Semantics

5. Distributional Semantics 1
   - Distributional Hypothesis
   - Word spaces
   - Parameterization and evaluation
   - Practical Work: Build your own DSM

Lecture 5, 6, 7, 8: Distributional Semantics

5. Distributional Semantics 1
   - Distributional Hypothesis
   - Word spaces
   - Parameterization and evaluation
   - Practical Work: Build your own DSM

6. Distributional Semantics 2
   - Beyond text: exploring other modalities in semantic representations
   - Introduction to supervised learning (e.g. PLSR)
   - Practical Work: Analysis of multi-modal vectorial representations

Lecture 5, 6, 7, 8: Distributional Semantics

5. Distributional Semantics 1
   - Distributional Hypothesis
   - Word spaces
   - Parameterization and evaluation
   - Practical Work: Build your own DSM

6. Distributional Semantics 2
   - Beyond text: exploring other modalities in semantic representations
   - Introduction to supervised learning (e.g. PLSR)
   - Practical Work: Analysis of multi-modal vectorial representations

7. Compositional Distributional Semantics
   - Representing phrases in semantic spaces
   - Composition functions - dealing with meaning and syntax
   - Evaluation methods
   - Practical Work: Working with phrase vectors

Lecture 5, 6, 7, 8: Distributional Semantics

5. Distributional Semantics 1
   - Distributional Hypothesis
   - Word spaces
   - Parameterization and evaluation
   - Practical Work: Build your own DSM

6. Distributional Semantics 2
   - Beyond text: exploring other modalities in semantic representations
   - Introduction to supervised learning (e.g. PLSR)
   - Practical Work: Analysis of multi-modal vectorial representations

7. Compositional Distributional Semantics
   - Representing phrases in semantic spaces
   - Composition functions - dealing with meaning and syntax
   - Evaluation methods
   - Practical Work: Working with phrase vectors

8. Case Study(ies): Adjectives in CDSMs

Lecture 9, 10: Machine Learning for Meaning Representation

9. ML for Meaning Representation
   - Introduction to Neural Networks
   - Machine Learning for meaning
   - Practical Work: Compare meaning in a "count" vs. "predict" world

Lecture 9, 10: Machine Learning for Meaning Representation

9. ML for Meaning Representation
   - Introduction to Neural Networks
   - Machine Learning for meaning
   - Practical Work: Compare meaning in a "count" vs. "predict" world

10. Ethics in ML for Meaning Representation
    - Ethical issues with ML, bias in models
    - Literature on bias and on de-biasing
    - Practical Work: Visualization of word embeddings for bias detection

Thanks, see you next week!

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