Lecture 2: Overview of the research projects

Katia Shutova

ILLC
University of Amsterdam

3 April 2019

Lecture 2: Overview of the research projects

Topic 1: Modelling meaning variation in context

Topic 2: Discourse information in semantic interpretation

Topic 3: Multitask learning: semantics and NLP applications

Topic 4: Cognitively-driven semantic representations

Modelling meaning variation in context

The children **ran** to the store
If you see this man, **run**!
Service **runs** all the way to Cranbury
She is **running** a relief operation in Sudan
the story or argument **runs** as follows
Does this old car still **run** well?
Interest rates **run** from 5 to 10 percent
Who's **running** for treasurer this year?
They **ran** the tapes over and over again
These dresses **run** small

Polysemy

- homonymy: unrelated word senses. *bank* (raised land) vs *bank* (financial institution)
- *bank* (financial institution) vs *bank* (in a casino): related but distinct senses.
- regular polysemy and sense extension
  - zero-derivation, e.g. *tango* (N) vs *tango* (V), or *rabbit, turkey, halibut* (meat / animal)
  - metaphorical senses, e.g. *swallow* [food], *swallow* [information], *swallow* [anger]
  - metonymy, e.g. he played *Bach*; he drank his *glass*.
- vagueness: *nurse, lecturer, driver*

No clearcut distinctions.

What is metaphor?

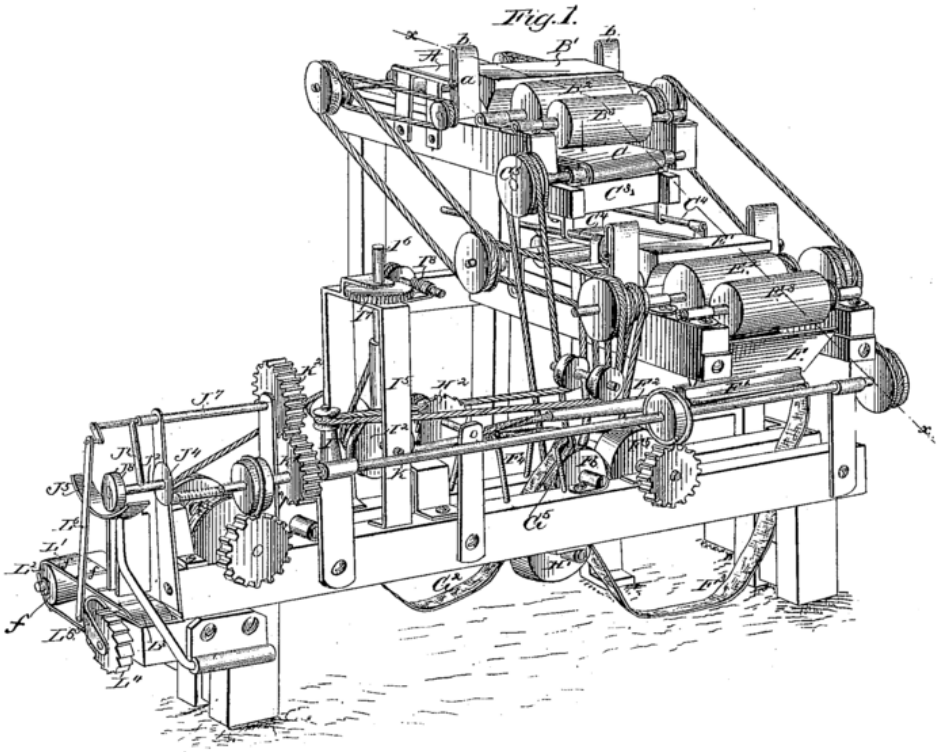

What is metaphor?

"*A political machine*"

"*The wheels of the regime were well oiled and already turning*"

"*Time to mend our foreign policy*"

"*20 Steps towards a Modern, Working Democracy*"

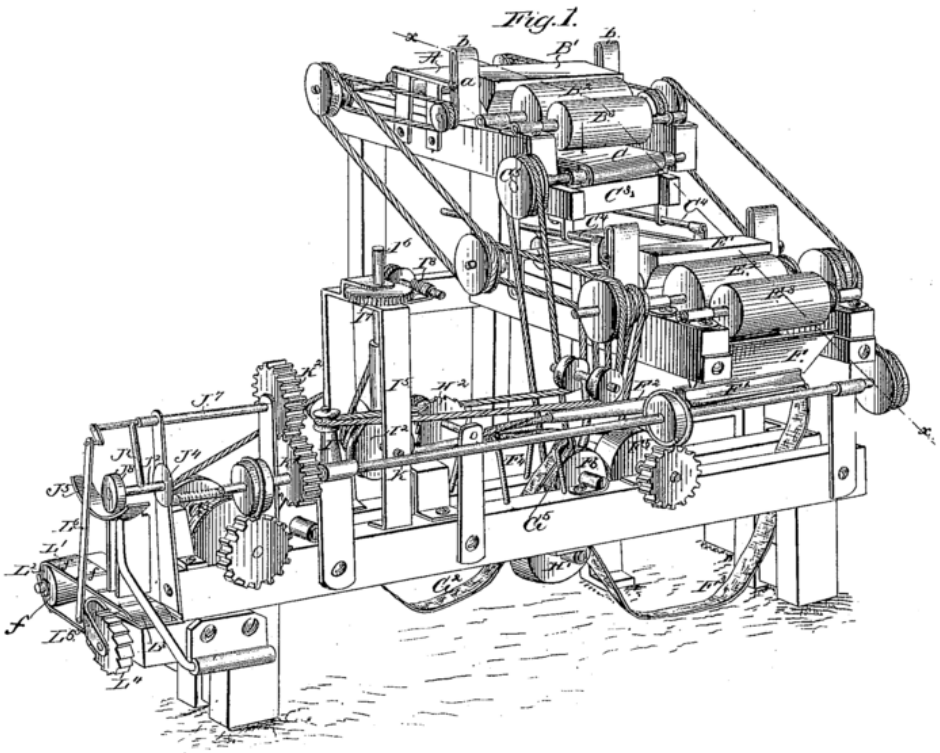

How does it work?

Conceptual Metaphor Theory
(Lakoff and Johnson, 1980)

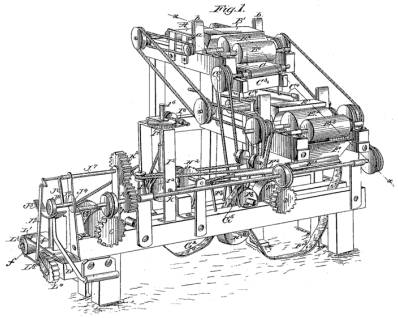

Metaphorical associations between concepts

POLITICALSYSTEM is a *MECHANISM*

target *source*

Cross-domain knowledge projection and inference

Reasoning about the target domain in terms of the properties of the source

Computational metaphor processing tasks

1. Learn metaphorical associations from corpora

   "POLITICAL SYSTEM is a MECHANISM"

2. Identify metaphorical language in text

   "*mend the policy*"

3. Interpret the metaphorical language

   "*mend the policy*" means "improve the policy; address the downsides of the policy"

Research question in this project

- Many models of metaphor and (some of) metonymy exist

But...

- None are integrated with general-purpose representation learning

The goal of this project is to:

Incorporate a model of regular polysemy into representation learning

Methods and experimental setup

Joint learning setup

- Learn sentence representations by training jointly for several tasks:
  - natural language inference
  - metaphor identification
  - metonymy identification
- Experiment with
  - different types of **sentence encoders** and **word representations**
  - different ways to **share parameters** between different tasks

Evaluation

1. Natural language inference task
2. Tasks in the SentEval and GLUE toolkits
3. Word in context (WIC) dataset (10k sentence pairs)

Examples from the WIC dataset:

*There's a lot of trash on the **bed** of the river*
*I keep a glass of water next to my **bed** when I sleep*

FALSE

*We **beat** the competition*
*Agassi **beat** Becker in the tennis championship*

TRUE

Discourse information in semantic interpretation

- Semantic interpretation always takes place in context
- Wider discourse provides an approximation of that context

But...

- There aren't many approaches incorporating discourse information into sentence-level semantic tasks

The goal of this project is to:

Incorporate a model of discourse into semantic tasks

What is discourse?

The (surrounding) text

Canada is warming on average at a rate twice as fast as the rest of the world. The federal government climate report also warns that changes are already evident in many parts of the country and are projected to intensify. Canada's Arctic has seen the deepest impact and will continue to warm at more than double the global rate. The report suggests that many of the effects already seen are probably irreversible. Canada's annual average temperature has increased by an estimated 1.7C (3F) since 1948, when nationwide temperatures were first recorded.

What is discourse?

The (surrounding) text

Canada is warming on average at a rate twice as fast as the rest of the world. The federal government climate report also warns that changes are already evident in many parts of the country and are projected to intensify. Canada's Arctic has seen the deepest impact and will continue to warm at more than double the global rate. The report suggests that many of the effects already seen are probably irreversible. Canada's annual average temperature has increased by an estimated 1.7C (3F) since 1948, when nationwide temperatures were first recorded.

Semantic tasks

Experiment with the following semantic tasks:

- **Metaphor identification**
  - determine if a word is used metaphorically or literally in a given context
- **Word sense disambiguation**
  - determine the sense of a word (e.g. WordNet sense)

  *John **killed** the wolf.*
  *Bill **killed** the project.*
  *Mary **killed** Jane. (at tennis or murdered her?)*

Semantic tasks

- **Semantic role labelling**
  - Semantic roles are types of the semantic relation between a predicate and its arguments.
  - **Task**: automatically label semantic roles

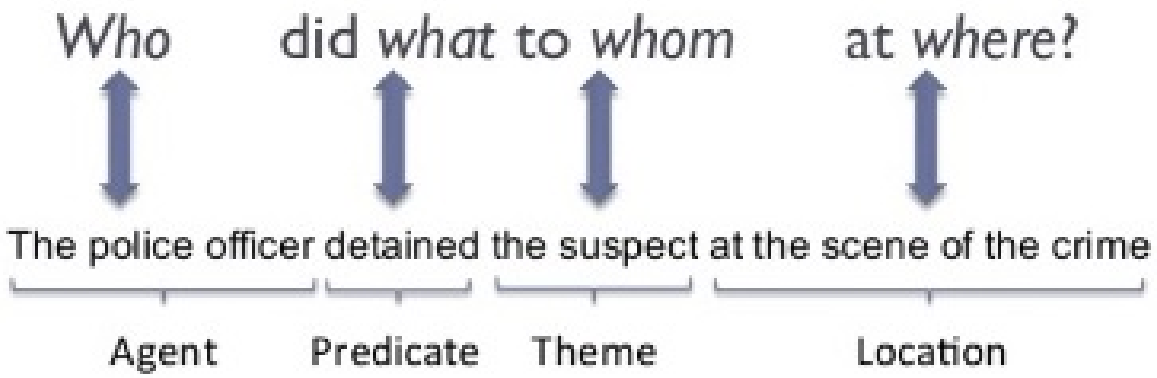

Methods and experiments

- Experiment within a sequence labelling paradigm
- use attention to capture important information in discourse
  - attention at word level
  - attention at sentence level

Why would this work?

- Sentences in discourse are not independent from each other
- Proxy for co-reference (e.g. pronoun) resolution
- Provide a coarse topical structure of text

Multitask learning: semantics and NLP applications

A series of projects focusing on integrating **semantic models** with **NLP applications**.

- Experiment in a multitask learning paradigm
- Train a model to perform two or more tasks at the same time
- which can benefit each other through information sharing

Project 1: Semantic inference and fact checking

Jointly learning natural language inference and fake news detection

The Weekly Meekly

Dinosaurs were lactose-intolerant

By D. Airee Cowse

BRATTLEBORO, VT - According to researchers at the Institute for Milk and Cheese, dinosaurs were unable to digest lactose.

"We think this is a major factor in their extinction, especially for larger dinosaurs like Apatosaurus," said Frank Lee Ronge, who chaired the research project.

Dr. Ronge and his team of scientists examined skeletons of over 200 dinosaurs in collections

scattered across the world, looking for evidence of calcium depletion such as surface pits and malformed bones. "Almost 46% of the samples we examined contained some proof of an inability to find high-quality calcium sources. Even for dinosaurs with lighter masses who relied more on speed than strength, like velociraptors, this inability to process dairy would

- The two tasks have a lot in common
- Both require a wide range of lexical and world knowledge
- and need to determine the veracity of statements

Fact checking: research tasks

Several **experimental paradigms**:

1. Given: Header and body of a news article
   Classes: *agree, disagree, discusses, unrelated*
2. Given: Two news articles
   Classes: *neural, agrees, refutes*
3. Given: One news article
   Classes: *'trusted', 'propaganda', 'hoax', 'satire'*

Example: true or fake?

Ryanair one of Europe's top polluters

Ryanair has become the only airline to be included in a list of Europe's top 10 polluters, according to data from the EU's Transport & Environment group. It is the first time a company that does not run a coal-fired power plant has come near the top of the ranking. Seven plants in Germany and one in both Poland and Bulgaria were on the list. The data said Ryanair's carbon dioxide emissions rose by 6.9% last year, but the firm said it was "Europe's greenest and cleanest airline". In a statement, the carrier added that "passengers travelling on Ryanair have the lowest CO2 emissions per kilometre travelled than any other airline".

Example: true or fake?

Grandmother gives birth to own grandchild

A 61-year-old Nebraskan woman has told of her joy after giving birth to her own grandchild, acting as the surrogate for her son and his husband. Cecile Eledge carried the daughter of her son Matthew Eledge and his husband Elliot Dougherty to term, giving birth to baby Uma Louise last week. Mrs Eledge said she made the offer when her son and Mr Dougherty first said they wanted to start a family.

Methods and experimental setup

Experiment with:

- sentence and document representations
- combining NLI and fake news detection in a multitask learning setup
- different ways to share parameters between the two tasks

Evaluate on the task of **fake news detection**

Project 2: Stance detection and fact checking

Jointly learning stance and fake news detection

The Weekly Meekly

Dinosaurs were lactose-intolerant

By D. Airee Cowse

BRATTLEBORO, VT - According to researchers at the Institute for Milk and Cheese, dinosaurs were unable to digest lactose.

"We think this is a major factor in their extinction, especially for larger dinosaurs like Apatosaurus," said Frank Lee Rouge, who chaired the research project.

Dr. Rouge and his team of scientists examined skeletons of over 200 dinosaurs in collections scattered across the world, looking for evidence of calcium depletion such as surface pits and malformed bones. "Almost 46% of the samples we examined contained some proof of an inability to find high-quality calcium sources. Even for dinosaurs with lighter masses who relied more on speed than strength, like velociraptors, this inability to process dairy would

The task of **stance detection**:

- determine if a text expresses a stance that *agrees* or *disagrees* with a particular statement
- already shown to be informative for fake news detection
- but not in a multitask learning setup

Project 3: Hyperpartisan news detection

Jointly learning to detect metaphor and hyperpartisan news

SHUTDOWN ENDS, SETTING UP CLASH OVER 'DREAMERS'

NEWS ANALYSIS

Hope for Bargain in a Swift Surrender

By CARL HULSE

WASHINGTON — Senate Democrats blinked. But the saving grace for them may be that they did it quickly.

"Should we have ever shut down the government?" asked Senator Joe Manchin III of West Virginia, who played a leading role in finding a way out for Democrats. "Absolutely not."

Over the weekend it became clear that using the shutdown to insist on protections for hundreds of thousands of undocumented immigrants was a serious miscalculation. By abandoning the strategy on its third day, Democrats believe they limited whatever damage there may be

Immigration Issues Unresolved Under Senate Accord

By SHERYL GAY STOLBERG and THOMAS KAPLAN

WASHINGTON — Congress brought an end to a three-day government shutdown on Monday as Senate Democrats buckled under pressure to adopt a short-term spending bill to fund government operations without first addressing the fate of young undocumented immigrants.

The House quickly approved the measure — which will fund the government through Feb. 8 and extend funding for the popular Children's Health Insurance Program for six years — and President Trump signed it on Monday night.

The task of **hyperpartisan news detection**:

- detect highly opinionated news articles
- classify news articles as to where they belong on political spectrum
- classes: *left, left-center, center, right-center, right*

Why is metaphor relevant?

Metaphors we think with: The role of metaphor in reasoning. Thibodeau and Boroditsky (2011)

- investigated how metaphor influences decision-making
- subjects read a text containing metaphors of either
  1. CRIME IS A VIRUS
  2. CRIME IS A BEAST
- then they were asked a set of questions on how to tackle crime in the city
  1. preventive measures
  2. punishment, restraint

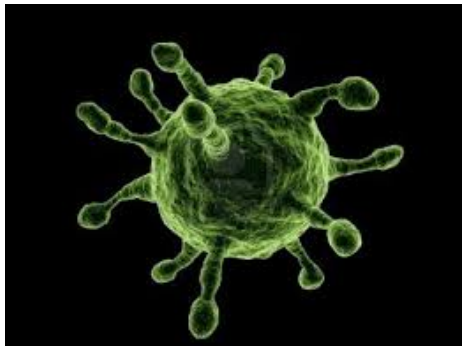

Methods and experimental setup

Hyperpartisan news detection is a document classification task.

- Sentence encoder trained in the metaphor identification task
- (in combination with NLI)
- incorporate this encoder into the document classification architecture
- train jointly.

Evaluate in the hyperpartisan news detection task.

Cognitively-driven semantic representations

Do our semantic models correlate with conceptual representation in the brain?

Can we use cognitive data as a source of bias in training word and sentence representations?

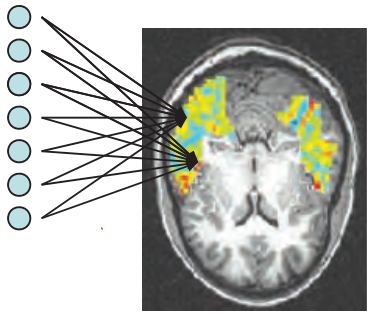

- Previous research has evaluated a range of semantic models in their ability to predict patterns of **human semantic processing**
- found in behavioural and brain imaging data
- with promising results

Datasets

Behavioral data

- free association task
- a dataset for around 5k words (De Deyne et al. 2016)
- elicit associations for a word from humans
  e.g. *summer: hot, sun, ice-cream, holiday* etc.

Brain imaging data

- fMRI neural activation patterns associated with the meaning of words (Mitchell et al, 2008; Pereira et al, 2018)
- fMRI patterns associated with sentence meaning (Pereira et al, 2018)

functional Magnetic Resonance Imaging (fMRI)

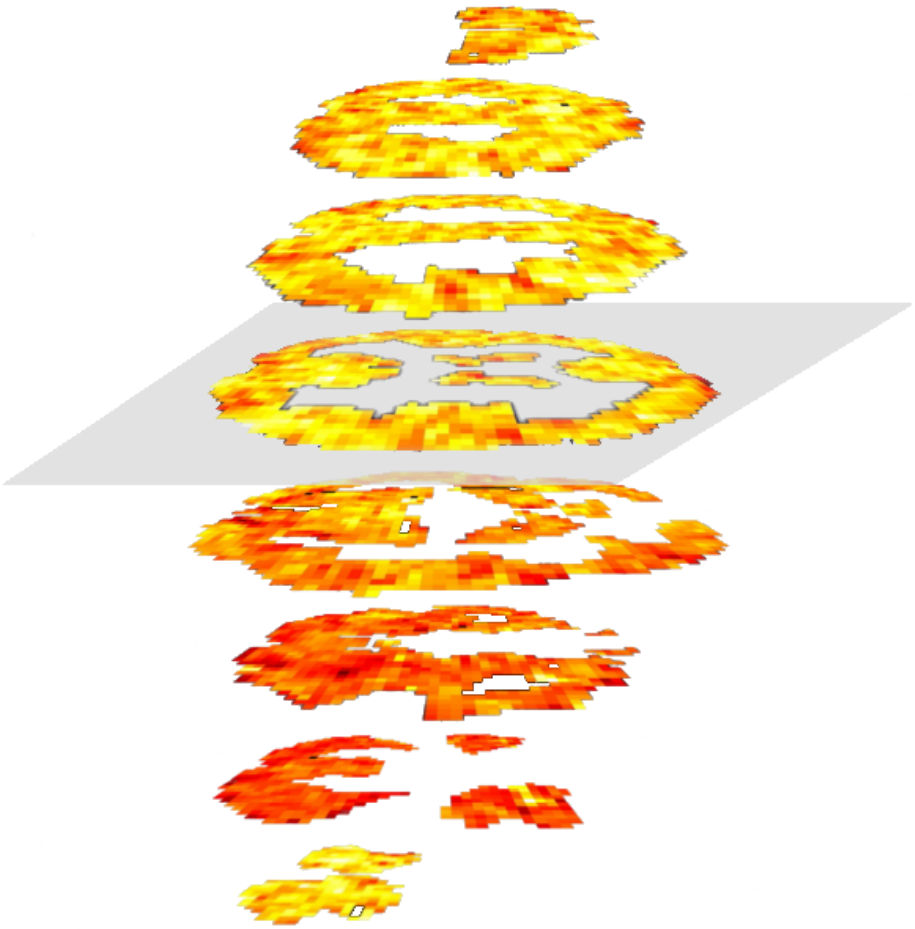

- **Voxel**: a 3x3x6mm^3 cube of brain tissue
- **Voxel value**: intensity of brain activity in that voxel
- **fMRI image**: vector of voxel values (represents brain activation pattern)

Evaluating word representations

- Data: fMRI neural activation patterns associated with the meaning of nouns (Mitchell et al, 2008)
- Task: predict patterns of brain activity
- Semantic models: distributional (window, dependency); skip-gram (window, dependency); visual; multimodal
- Method: linear regression
- Evaluation: leave-two-out cross validation

Promising results, with accuracies above 70%

Research question in this project

Can we use cognitive data as a source of bias in training word and sentence representations?

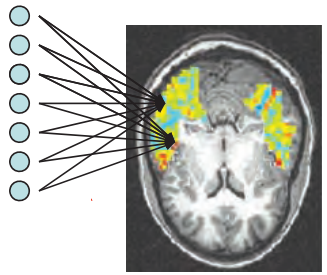

Goals:

- incorporate information about human semantic processing into computational models
- use both behavioral and brain imaging data
- experiment with joint learning
- train a model to learn word / sentence representations,
- and at the same time predict human associations or brain activity

Next steps

1. Sign up to groups on Canvas, by Friday **5 April**
2. Submit your top 2 choices of research topic by **12 April** (1 submission per group)
3. We will assign project topics by Friday, 19 April
4. Ask questions about the projects and get started at the lab session, 23 April.