Advanced Topics in Computational Semantics

Katia Shutova

ILLC
University of Amsterdam

30 March 2020

Taught by...

Katia Shutova
Lecturer
e.shutova@uva.nl

Phillip Lippe
Teaching assistant
phillip.lippe@gmail.com

Verna Dankers
Teaching assistant
vernadankers@gmail.com

Lecture 1: Introduction

Overview of the course

Semantics in wider NLP

Statistical semantics and representation learning

Word representations

Sentence representations

Overview of the course

- Focus on language interpretation and modelling meaning
  - Methods for **learning meaning representations** from linguistic data
  - Analysis of meaning representations learnt
  - Applications
- This is a research seminar
  - Focus on recent progress in the field
  - Lectures
  - You will present and critique research papers
  - and conduct a research project

Overview of the topics

Modelling meaning at different levels

- Word representations
- Compositional semantics and sentence representations
- Modelling meaning variation in context
- Discourse processing, document representations

Overview of the topics

Focus on deep learning and joint learning

- Different neural architectures (e.g. LSTMs, attention, transformers etc.)
- Contextualised representations: ELMo and BERT
- Joint learning at different linguistic levels
- Multitask learning
- Multilingual joint learning
- Few-shot learning and meta-learning

Interdisciplinary topics and applications

- Representation learning and neurocognition of language

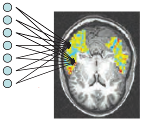

- Modelling social bias and stereotypes in NLP models

- Applications: stance detection and misinformation detection

Assessment

- Presentation and participation (25%)
  - Present 1 paper in class
  - Read and discuss other papers (and submit questions!)
- Practical assignment (25%)
  1. Implement a model of sentence meaning
  2. Evaluate it in a set of NLP tasks
  3. Assessed by presenting results to TAs
  4. **Deadline**: 17 April 2020
- Research project (50%)

No exam!

More information at the first lab session on Tuesday, 31 March.

Research project

- Goal: Investigate a **new research question**
  - Apply the models discussed in the course
  - Perform experiments and analyse results
  - Write a research paper
  - Present the results during the last lecture
- Organisation
  - We will propose projects on several topics – you choose
  - Work in groups of 4
  - **Deadline**: 22 May 2020

It gets even better...

Best Project Presentation Award

Also note:

Course materials and more info:
`https://cl-illc.github.io/semantics-2020`

Piazza for discussions:
`piazza.com/university_of_amsterdam/spring2020/5314atic6y`

Access code: **elmobert**

Contact

- Assignments: Phillip and Verna
- Paper presentations: Katia

Sign up to groups on Canvas by Friday, 3 April.

Video lectures, seminars and labs

Due to COVID-19, the whole course will be taught online

- We will use Zoom for:
- Video lectures (live or recorded)
- Video conference seminars
  - live paper presentations via screen sharing
- Lab sessions with Verna and Phillip via Zoom
- Piazza for questions outside of these sessions

Natural Language Processing

Many popular applications

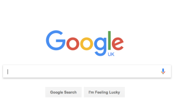


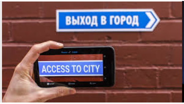

...and the emerging ones

Why is NLP difficult?

Similar strings mean different things, different strings mean the same thing.

- Synonymy: different strings can mean the same thing

  The King's speech *gave the much needed reassurance to his people.*
  His majesty's address *reassured the crowds.*

- Ambiguity: same strings can mean different things

  His majesty's address *reassured the crowds.*
  His majesty's address *is Buckingham Palace, London SW1A 1AA.*

Why is NLP difficult?

Similar strings mean different things, different strings mean the same thing.

- Synonymy: different strings can mean the same thing

  The King's speech *gave the much needed reassurance to his people.*
  His majesty's address *reassured the crowds.*

- Ambiguity: same strings can mean different things

  His majesty's address *reassured the crowds.*
  His majesty's address *is Buckingham Palace, London SW1A 1AA.*

Why is NLP difficult?

Similar strings mean different things, different strings mean the same thing.

- Synonymy: different strings can mean the same thing

  The King's speech *gave the much needed reassurance to his people.*
  His majesty's address *reassured the crowds.*

- Ambiguity: same strings can mean different things

  His majesty's address *reassured the crowds.*
  His majesty's address *is Buckingham Palace, London SW1A 1AA.*

Why is NLP difficult?

Similar strings mean different things, different strings mean the same thing.

- Synonymy: different strings can mean the same thing

  The King's speech *gave the much needed reassurance to his people.*
  His majesty's address *reassured the crowds.*

- Ambiguity: same strings can mean different things

  His majesty's address *reassured the crowds.*
  His majesty's address *is Buckingham Palace, London SW1A 1AA.*

Why is NLP difficult?

Similar strings mean different things, different strings mean the same thing.

- Synonymy: different strings can mean the same thing

  The King's speech *gave the much needed reassurance to his people.*
  His majesty's address *reassured the crowds.*

- Ambiguity: same strings can mean different things

  His majesty's address *reassured the crowds.*
  His majesty's address *is Buckingham Palace, London SW1A 1AA.*

Computational semantics

Computational semantics = Natural language understanding (NLU)

an area of NLP concerned with language interpretation and modelling meaning

1. Lexical semantics: modelling the meaning of words
2. Compositional semantics: modelling the meaning of sentences
3. Discourse processing: modelling larger text passages
4. Pragmatics: modelling meaning in wider situational context (e.g. social meaning)

Statistical semantics

Distributional semantics

- The meaning of a word can be defined by its use
- as a distribution of contexts
- extracted from a text corpus

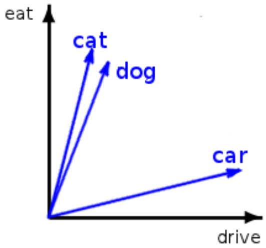

N: dog	N: car
248 bark	493 drive
197 eat	428 park
193 take	317 steal
110 walk	248 stop
101 run	102 break
...	...

Statistical semantics in pre-deep learning era

- Vector space models (dimensionality reduction, SVD etc.)
- Information theoretic approaches
- Supervised learning with hand-engineered features
  - a range of classifiers (SVM, decision trees etc.)
  - features based on lexico-syntactic patterns
  - or lexical resources (such as WordNet)
- Unsupervised learning
  - Clustering

Paradigm shift: representation learning

Deep learning

- dominates the field since ≈2014
- led to performance improvements in many tasks

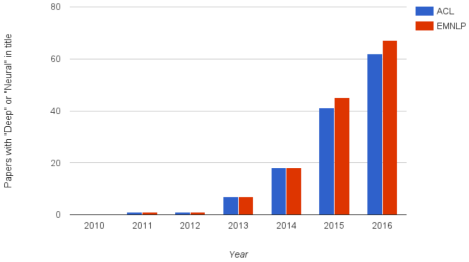

Paradigm shift: representation learning

But why?

- Neural networks have been around for decades.
- What has changed in the way they are applied in NLP?
- Key conceptual innovation:

  *learning **intermediate meaning representations** in the process of end-to-end training for a particular task.*

Paradigm shift: representation learning

But why?

- Neural networks have been around for decades.
- What has changed in the way they are applied in NLP?
- Key conceptual innovation:

> *learning **intermediate meaning representations** in the process of end-to-end training for a particular task.*

Example: sentiment analysis

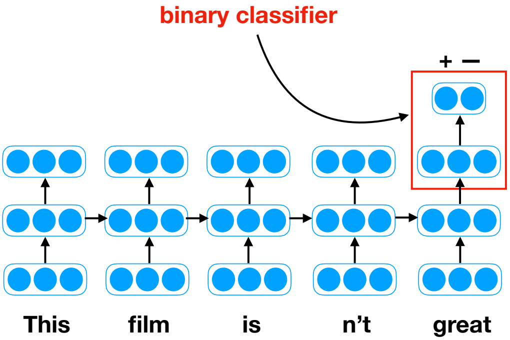

Example: sentiment analysis

Word representations

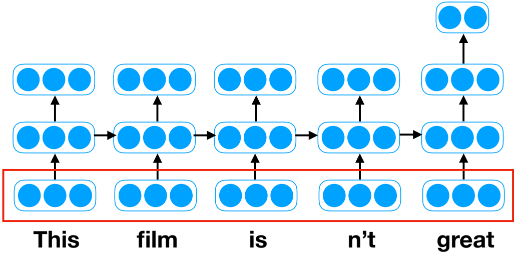

Example: sentiment analysis

Sentence representations

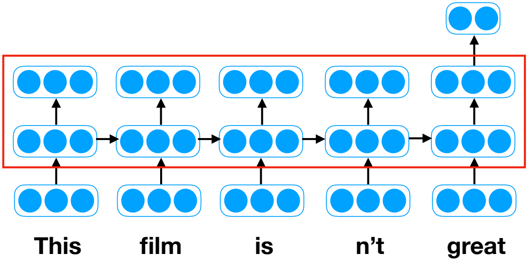

General-purpose word representations

Mikolov et. al. 2013. *Efficient Estimation of Word Representations in Vector Space*.

Skip-gram model:

- Given a word
- predict its neighboring words
- learn word representations in the process

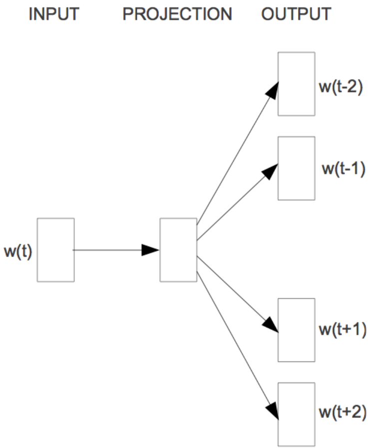

Word embeddings in NLP tasks

- Random initialization, learn as part of task objective
- External initialization (e.g. skip-gram), update as part of task objective
- External initialization, keep fixed

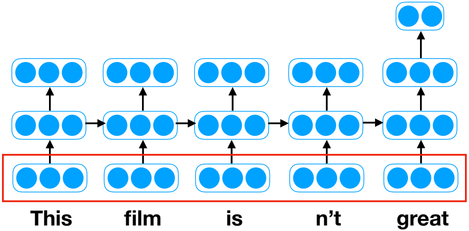

Learning sentence representations

(Long-term) goal:

- a **general-purpose** neural network **sentence encoder**
- which can be applied across diverse NLP tasks.

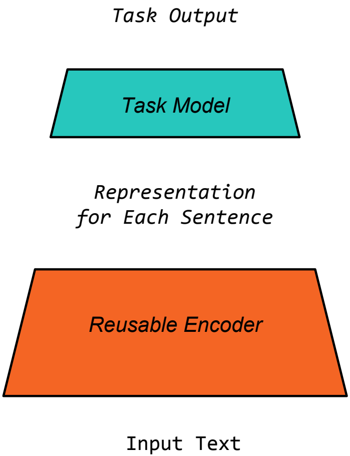

Why is this useful?

1. Improve performance
   - produce **rich semantic representations** for downstream NLP tasks
2. Improve data efficiency
   - provide a model of sentence representation for language understanding tasks which **lack training data**

What can we expect this model to capture?

- Lexical semantics and meaning disambiguation in context
- Word order
- Some syntactic structure
- Semantic composition
- Idiomatic/non-compositional phrase meanings
- Connotation and social meaning.

Sentence representation models

Unsupervised training on **single sentences**:

- Sequence autoencoders (Dai and Le, 2015)
- Paragraph vector (Le and Mikolov, 2015)

Unsupervised training on **running text**:

- SkipThought (Kiros et al., 2015)
- FastSent (Hill et al. 2016)
- Quick Thoughts (Logeswaran and Lee, 2018)

We will look at these models later in the course.

Sentence representation models

Supervised training on **large corpora**:

- Dictionaries (Hill et al. 2015)
- Image captions (Hill et al. 2016)
- Natural language inference data (Conneau et al. 2017)

Learning from dictionary definitions

Hill et al., 2016. *Learning to Understand Phrases by Embedding the Dictionary*

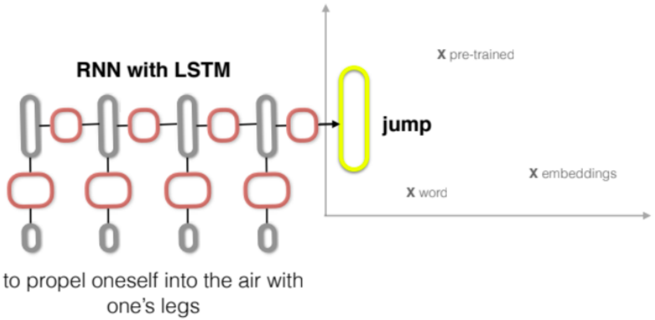

Natural language inference task

Bowman et al, 2015. *A large annotated corpus for learning natural language inference*

- Stanford Natural Language Inference (SNLI) corpus
- 570k sentence pairs
- labeled for entailment, contradiction, and semantic independence

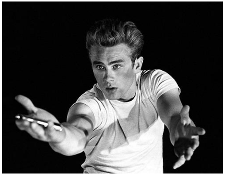

James Byron Dean refused to move without blue jeans

{**entails**, contradicts, neither}

James Dean didn't dance without pants

More NLI examples

A black race car starts up in front of a crowd of people.

A man is driving down a lonely road.

More NLI examples

A black race car starts up in front of a crowd of people.

A man is driving down a lonely road.

CONTRADICTION

More NLI examples

A black race car starts up in front of a crowd of people.

A man is driving down a lonely road.

CONTRADICTION

A soccer game with multiple males playing.

Some men are playing a sport.

More NLI examples

A black race car starts up in front of a crowd of people.

A man is driving down a lonely road.

CONTRADICTION

A soccer game with multiple males playing.

Some men are playing a sport.

ENTAILMENT

General architecture for NLI

Conneau et al, 2017. *Supervised Learning of Universal Sentence Representations from Natural Language Inference Data*

InferSent model

- Siamese architecture (same encoder to represent premise and hypothesis)
- 3-way classification (*entails, contradicts, neither*)

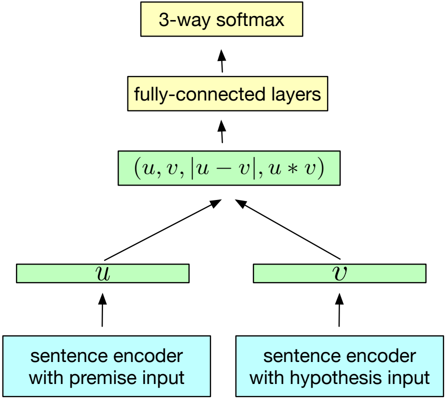

InferSent encoder: BiLSTM with max pooling

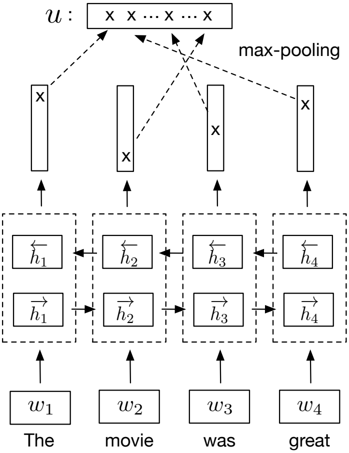

NLI and language understanding

To perform well at NLI, your representations of meaning must handle with the full complexity of compositional semantics...

- Lexical entailment (*cat* vs. *animal*, *cat* vs. *dog*)
- Lexical ambiguity (e.g. *bank, run*)
- Quantification (*all, most, fewer than eight* etc.)
- Modality (*might, should*, etc.)
- Common sense background knowledge

Evaluation framework: SentEval

Conneau and Kiela, 2018. *SentEval: An Evaluation Toolkit for Universal Sentence Representations*

- Formalised an evaluation standard for sentence representations
- Suite of ten tasks
- Software package automatically trains and evaluates per-task classifiers using supplied representations.

SentEval tasks

- Classification tasks:
  - sentiment analysis / opinion polarity
  - subjectivity vs. objectivity
  - question type (e.g. for question answering)
- Natural language inference:
  - several datasets
- Semantic similarity tasks:
  - sentence similarity
  - paraphrasing
  - image caption retrieval

Practical 1

Learning general-purpose sentence representations

- supervised training
- SNLI task
- Implement three variants of the **InferSent** model:
  1. Unidirectional LSTM encoder
  2. Bidirectional (Bi-) LSTM encoder
  3. BiLSTM encoder with max pooling
- Compare to a **baseline** averaging word embeddings
- Evaluate using **SentEval**

Submit a mini-report containing your results and your code
Deadline: 17 April

Research project topics

- Multitask learning
- Few-shot learning and meta-learning
- Multilingual representation learning
- Extending models with human language processing objectives
- Identifying and mitigating gender and racial bias in NLP models

Overview of research projects **next Monday** (get excited!)

Coming next...

On Thursday:

- Seminar: encoding different kinds of semantic information in word embeddings

Next Monday:

- Seminar: learning general-purpose sentence representations

Acknowledgement

Some images were adapted from Sam Bowman and Steve Clark