Natural Language Processing 1 Live Q & A: Discourse and summarisation

Katia Shutova

ILLC University of Amsterdam

2 December 2020

Question 1: Modelling meaning change

Imagine that you are given a dataset of historical texts (1860-2020). It contains 5 parts each covering a particular time period. This type of data captures how language evolves.

How could you use **distributional semantics** or **word embeddings** to model change in word meaning over time?

Question 2: Errors in neural co-reference resolution

Why do these errors happen? And what do they tell us about the model behaviour?

- 3 (The flight attendants) have until 6:00 today to ratify labor concessions. (The **pilots**') union and ground crew did so yesterday.
- (Prince Charles and his new wife Camilla) have jumped across the pond and are touring the United States making (their) first stop today in New York. It's Charles' first opportunity to showcase his new 4 wife, but few Americans seem to care. Here's Jeanie Mowth. What a difference two decades make. (Charles and Diana) visited a JC Penney's on the prince's last official US tour. Twenty years later here's the prince with his new wife.

Question 3: Summarisation

- 1. How can information about discourse nuclei and satellites be exploited in extractive summarisation?
- 2. What would be the benefits of an **abstractive summarisation** model that can distinguish discourse relations?

Question 4: Discussion

- How are the following tasks / methods affected by the differences in genre? Which of them is more sensitive to this difference and why?
 - n-gram language modelling
 - PoS tagging
 - probabilistic syntactic parsing
 - distributional semantics and word embeddings
 - co-reference resolution
- 2. How can you use the above techniques in summarisation?

Advanced course on semantics

Advanced Topics in Computational Semantics (block 5)

- This course is about learning meaning representations
 - Methods for learning meaning representations
 - focus on deep learning (LSTMs, CNNs, transformers)
 - Interpretation of meaning representations learnt
 - Applications
- This is an advanced research seminar
 - Focus on recent progress in the field
 - Lectures
 - You will present and critique recent research papers
 - and conduct a research project (new research question!)

Overview of the topics

For a detailed overview and list of papers see last year's website: https://cl-illc.github.io/semantics-2020/syllabus.html

Modelling meaning at different levels

- Word representations
- Compositional semantics and sentence representations
- Contextualised representations (ELMo and BERT)
- Discourse processing, document representations

Overview of the topics

Focus on deep learning and joint learning

- Different neural architectures (e.g. LSTMs, attention, transformers etc.)
- Joint learning at different linguistic levels
- Multitask learning
- Multilingual joint learning
- Learning from multiple modalities (language and vision)
- Few-shot learning (i.e. learning from a few examples)

Example research projects (from previous years)

- Learning multilingual contextualised representations
- Joint modelling of semantics and discourse
- Multitask learning: semantics in NLP applications
 - stance detection
 - misinformation detection
- Meta-learning for few-shot model adaptation
 - across languages, domains and tasks
- Cognitive properties of meaning representations
 - evaluating learnt representations against brain imaging data

Many of your TAs took it — ask about their experience!