

NEURAL NETWORKS AND LANGUAGE

Research and Development, 5LN714

Johan Sjons,
Department of Linguistics and Philology,
Uppsala University

August 30, 2022

- 1 DNNs
- 2 Language
- 3 The relation between DNNs and language

DNNs

- Deep Neural Networks (DNNs) are powerful learning models
- By now, they have outperformed most other NLP models (e.g., n -gram models, HMMs, Naive Bayes etc)
- It's easy to be swept away by the magnificence of DNNs!

DNNs

- Deep neural networks have mainly three uses in NLP (Linzen and Baroni, 2021)
- In no particular order:
 - ① **As language models** The input is a sequence of words and the output the most likely continuation (from a probability distribution over upcoming words)
 - ② **As classifiers** The input is a sequence of words and the output is a class that that sequence is a member of
 - ③ **As seq2seqs** The input is a sequence of words and the output is another sequence of words in response to this input

DNNs and cognition?

- But what are DNNs models of?
- For example: When we say that we have a language model, what do we mean?
- “We have a model of X.”
- What is X?

- 1 DNNs
- 2 Language
- 3 The relation between DNNs and language

Small and dramatic changes

- *The chicken was ready to eat* (ambiguous)
- *The chicken was ready to be eaten* (ambiguous)
- *The chicken was hard to eat* (unambiguous)
- **The chicken was hard to be eaten* (lacks meaning)

What is language?

- No one knows what language is, but
 - some combinatorial operation can safely be assumed, which
 - yields hierarchical structures
 - DNNs are concerned with linear order

Hierarchical structure

- Consider the sentences in 1 and 2 (Chomsky, 2013):
 - ① *Eagles that fly instinctively swim* (ambiguous)
 - ② *Instinctively, eagles that fly swim* (resolved)
- Linear order is sacrificed, intuitive interpretation as well

Grammaticality and acceptability

- Grammaticality is about whether a sentence is deemed OK or not by a theoretical grammar.
- Acceptability is about humans' judgements about sentences and string patterns.
 - Acceptability (1): A sequence of words sounds good or bad.
 - Acceptability (2): A sequence of words good or bad under an interpretation (reading).
 - He_i loves $John_j$
 - $*He_j$ loves $John_i$

Grammaticality and acceptability

- A sequence of words can be
 - Grammatical and acceptable
 - *The human is taking a walk*
 - Grammatical and unacceptable
 - *? The woman the boy who I love loves loves the boy*
 - This is an example of self-embedding because the relative clause *who I love* is embedded in the relative clause *the boy loves* (which in turn is embedded in the matrix clause *The woman loves the boy*).
 - Ungrammatical and acceptable
 - *? Last year more people visited Rome than I did.*
 - This is an example of a **linguistic illusion**.
 - Ungrammatical and unacceptable
 - ** A the is walk taking human*
 - Such sequences are sometimes called **string patterns** or **word salad**

- 1 DNNs
- 2 Language
- 3 The relation between DNNs and language**

The “and”

- Are there parts in our language machinery that DNNs are good models of?
- Relatedly, what do we bake into our models?
 - Any learner has biases (Mitchell, 1980)
 - Random seed might make a difference

Parsing

- Garden path sentences (locally ambiguous)
 - *The horse raced past the barn fell*
 - *The complex houses married and single soldiers and their families*
 - *The train left at midnight was empty*

Predicting ourselves

- We do predict our own output (Blacfkmer and Mitton, 1991)

Language acquisition

- Children need much less data than DNNs to acquire language
- Children are unsupervised learners for the most part, whereas DNNs are supervised.
- Where do DNNs come in? Statistical learning? (Continuous vs discrete systems)

Speech

- Speech recognition
- Speech synthesis

Summary

- Most seminar papers are focused on syntax (but the most challenging parts probably concern the structures of the DNNs and the experiments carried out)
- You can find your own flavor

References

- Blacfkmer, E. R., and Mitto,n J. L. (1991). Theories of monitoring and the timing of repairs in spontaneous speech. *Cognition*, 39(3), 173–194.
- Chomsky, N. (2013). Problems of projection. *Lingua*, 130, 33–49.
- Linzen, T. and M. Baroni (2021). Syntactic structure from deep learning. *Annual Review of Linguistics* 7, 195–212.
- Mitchell, T. M. (1980). The need for biases in learning generalizations (pp. 184–191). New Jersey: Department of Computer Science, Laboratory for Computer Science Research, Rutgers University.

- Questions?