Climbing Towards NLU: On Meaning, Form, and Understanding in the Age of Data

~ Bender & Koller, ACL2020 ~

Hermann Kroll

Institut für Informationssysteme

Technische Universität Braunschweig

http://www.ifis.cs.tu-bs.de

ifis
Institut für Informationssysteme
Technische Universität Braunschweig

“a system trained only on form has a priori no way to learn meaning”

(B&K, Abstract)

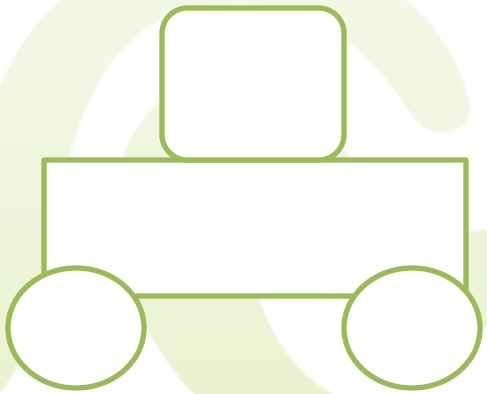

BERTology

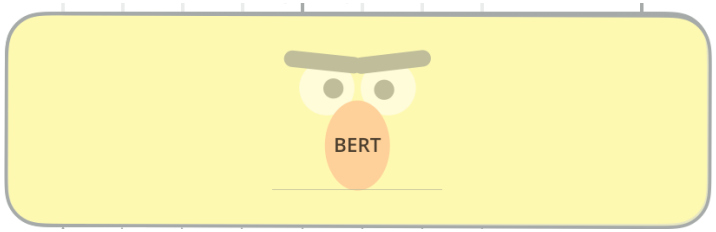

- "In order to train a model that **understands sentence relationships**, we pre-train for a binarized next sentence prediction task." (Devlin et al., 2019)
- "The surprisingly strong ability of these models to **recall factual knowledge** without any fine-tuning demonstrates their potential as unsupervised open-domain QA systems." (Petroni et al., 2019)
- "BERT is a system by which Google's algorithm uses pattern recognition to better **understand how human beings communicate** so that it can return more relevant results for users." - Google

Whereas NLP Experts argue

- "These systems are **still a really long way** from truly understanding running prose."

  (Gary Marcus)

- "Though BERT passed the lab's common-sense test, machines **are still a long way from** an artificial version of a human's common sense."

  (Oren Etzioni)

What is Meaning?

- Meaning is a relation $M \subseteq E \times I$
  - E is the set of all **natural language expressions**
  - I is the set of all **communicative intents**
    - "they may be to convey some information to the other person; or to ask them to do something; or simply to socialize."
    - "the communicative **intent is grounded in the real world** the speaker and listener inhabit together"
- Examples:
  - "The photographer asked me to do a Napoleon for the camera."
  - "Never ask two China trips to the same party."

Let's talk Meaning and Intelligence

- Chinese Room experiment, Searle (1980)
  - "He develops the metaphor of a "system" in which a person **who does not speak** Chinese **answers Chinese questions** by consulting a library of Chinese books according to predefined rules."
  - "From the outside, the system seems like it "understands" Chinese, although in reality **no actual understanding** happens anywhere inside the system."

- "This in turn means that for a human or a machine to learn a language, they must solve what Harnad (1990) calls the **symbol grounding problem**."

Octopus Test

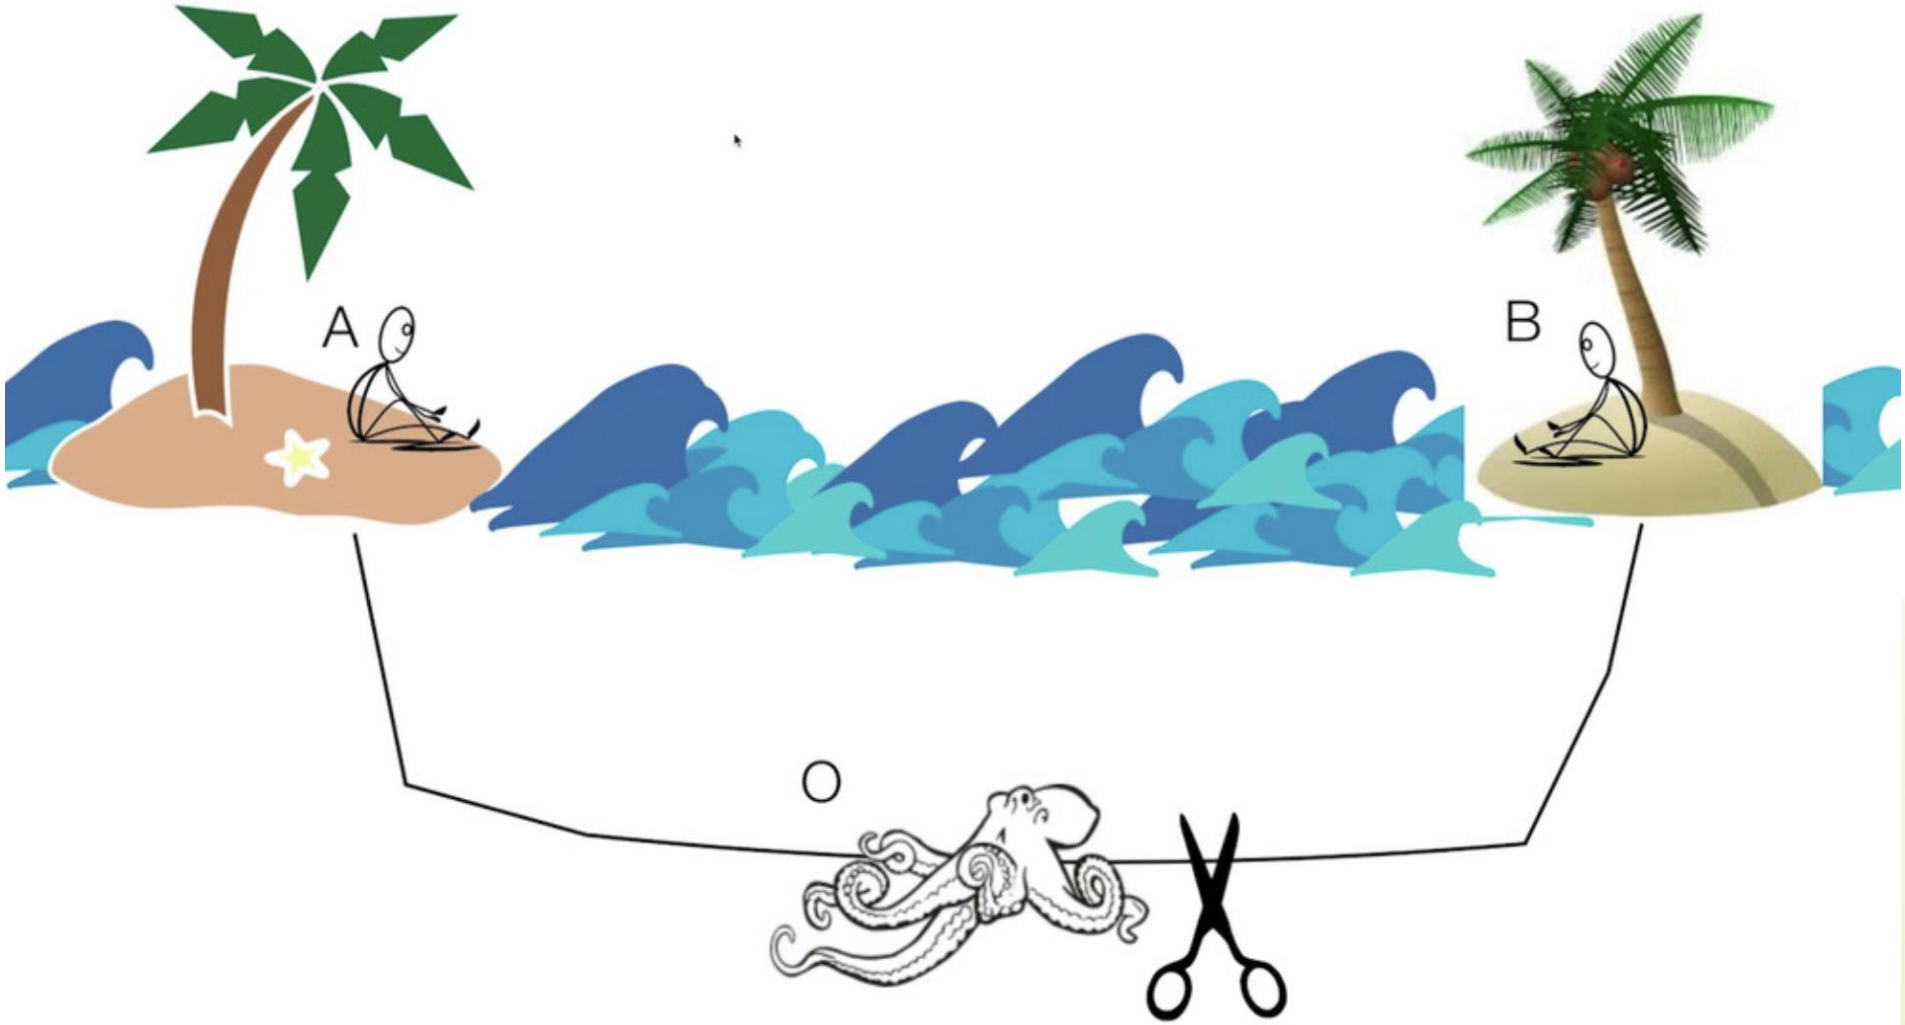

Can the Octopus fool A?

- Suppose a new situation:
  - A constructs a coconut catapult, sends detailed instructions over the wire, and asks for B's thoughts.
  - A is suddenly chased by an angry bear, grabs a couple of sticks to defend herself, and asks B for advice.

- Claim:
  - "a system that is **trained only on form would fail** a sufficiently sensitive test, because **it lacks** the ability to **connect** its **utterances** to the **world**"

Critics

- "The test **is not precise** enough to say which capabilities exactly are lacking" - Blog
- "But couldn't **GPT3** solve all problems?" - Chat

- Answer:
  - "The point isn't really whether O could fool A under what circumstances, but rather to use that thought experiment to **show what is missing in O's** (and thus modern LM's) **input**." - Bender

ifis

Institut für Informationssysteme

Technische Universität Braunschweig

More Thought Experiments

The Java Test

- "Imagine that we were to train an LM on all of the well-formed Java code published on GitHub. The input is only the code. It is not paired with bytecode, nor a compiler, nor sample inputs and outputs for any specific program. We can use any type of LM we like and train it for as long as we like. We then ask the model to execute a sample program, and expect correct program output."

  (Section 5)

Discussion

- Human Language Acquisition
  - Summary of experiments **how human learn to connect language and meaning**
- Distributional Semantics
  - "The lexical similarity relations learned by distributional models trained on text **don't** in themselves **connect any of those words to the world** […]"
  - "From this literature we can see that the slogan "**meaning is use**" (often attributed to Wittgenstein, 1953), refers not to "use" as "distribution in a text corpus" but rather that **language is used in the real world to convey communicative intents to real people**."

Climbing the Hill

- Bottom-Up perspective:
  - Driven by identifying specific research challenges
  - "A scientific result counts as a success if it solves such **a specific challenge**, at least partially."
  - "As long as such successes are frequent and satisfying, there is a general atmosphere of sustained progress."
- Top-Down perspective
  - "[…] the focus is on the remote end goal of offering a **complete, unified theory** for the entire field."

Are we Climbing the Right Hill?

- "Grammar- and knowledge-based methods gave way to statistical methods, and today most research incorporates neural methods. Researchers of each generation felt like they were solving relevant problems and making constant progress, from a bottom-up perspective. However, eventually serious shortcomings of each paradigm emerged, which could not be tackled satisfactorily with the methods of the day, and these methods were seen as obsolete."

What can we do?

- "First, above all, cultivate humility towards language and **ask top-down questions**"
- "Second, be aware of the **limitations of tasks**"
- "Third, value and support the work of carefully creating new tasks"
- "Fourth, evaluate models of meaning **across tasks**."
- "Finally, perform thorough **analysis** of **both errors** and **successes**."

ifis

Institut für Informationssysteme
Technische Universität Braunschweig

Conclusion

Some Guy in the Chat

- "For me the huge contribution here is […] "If someone tells you that their system 'understands' language, put your hand on your wallet and keep it there." Doubly so if they say "like people do" and quadruply so if they make any reference whatsoever to the way children learn language"

~ Someone in the ACL Chat ~

My Conclusion

- Be careful of what **you conclude**:
  - A good precision in task X does not imply that your model understands or has learned anything

- Asking **top-down questions**:
  - They might be frustrating, however, a broader view on your task or your research **helps you** to understand the current limitations
  - It's fair to talk about these limitations, they should not be hidden. It's better to **raise these questions before** someone else will do it

Sources

- Bender, Emily M., and Alexander Koller. "Climbing towards NLU: On meaning, form, and understanding in the age of data." Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics. 2020.

- Blog: https://blog.julianmichael.org/2020/07/23/to-dissect-an-octopus.html

Thank You!

Technische Universität Braunschweig

kroll@ifis.cs.tu-bs.de