

What can we see in deep learning black boxes?

Andrey Kutuzov
University of Oslo
Language Technology Group

May 31, 2018





Yoav Goldberg [Follow](#)

Senior Lecturer at Bar Ilan University. Working on NLP. Recently with Neural Nets. Published a book about it. <http://www.cs.biu.ac.il/~yogo/>

Jun 9, 2017 · 14 min read

An Adversarial Review of “Adversarial Generation of Natural Language”

Or, for fucks sake, DL people, leave language alone and stop saying you solve it.



Yoav Goldberg [Follow](#)

Senior Lecturer at Bar Ilan University. Working on NLP. Recently with Neural Nets. Published a book about it. <http://www.cs.biu.ac.il/~yogo/>

Jun 9, 2017 · 14 min read

An Adversarial Review of "Adversarial Generation of Natural Language"

Or, for fucks sake, DL people, leave language alone and stop saying you solve it.

...human language is magnificent, and complex, and challenging. It has tons of nuances, and corners, and oddities, and surprises.



Yoav Goldberg [Follow](#)

Senior Lecturer at Bar Ilan University. Working on NLP. Recently with Neural Nets. Published a book about it. <http://www.cs.biu.ac.il/~yogo/>
Jun 9, 2017 · 14 min read

An Adversarial Review of "Adversarial Generation of Natural Language"

Or, for fucks sake, DL people, leave language alone and stop saying you solve it.

*...human language is magnificent, and complex, and challenging. It has tons of nuances, and corners, and oddities, and surprises. While natural language processing researchers, and natural language generation researchers, and linguists! who do a lot of the heavy lifting, made some impressive advances towards our understanding of language and how to process it **we are still just barely scratching the surface on this.***



NLP has its problems



NLP has its problems

- ▶ machine learning **models amplifying biases** and discrimination in data [Zhao et al., 2017]

NLP has its problems

- ▶ machine learning **models amplifying biases** and discrimination in data [Zhao et al., 2017]
- ▶ sometimes research success depends on computational power:
 - ▶ '*...do we have enough GPUs?*'





- ▶ People wonder:
 - ▶ *'What are your research questions?'*
 - ▶ *'Just lots of numbers with very small differences?'*



- ▶ People wonder:
 - ▶ *'What are your research questions?'*
 - ▶ *'Just lots of numbers with very small differences?'*
- ▶ Is it a **science** or an **engineering discipline**?



- ▶ People wonder:
 - ▶ *'What are your research questions?'*
 - ▶ *'Just lots of numbers with very small differences?'*
- ▶ Is it a **science** or an **engineering discipline**?
- ▶ Or may be **CL is a science** and **NLP is its application towards empirical problems**?



- ▶ People wonder:
 - ▶ *'What are your research questions?'*
 - ▶ *'Just lots of numbers with very small differences?'*
- ▶ Is it a **science** or an **engineering discipline**?
- ▶ Or may be **CL is a science** and **NLP is its application towards empirical problems**?
- ▶ Motivation for research can be different:



- ▶ People wonder:
 - ▶ *'What are your research questions?'*
 - ▶ *'Just lots of numbers with very small differences?'*
- ▶ Is it a **science** or an **engineering discipline**?
- ▶ Or may be **CL is a science** and **NLP is its application towards empirical problems**?
- ▶ Motivation for research can be different:
 1. trying to provide a **computational explanation for linguistic or psycholinguistic phenomenon**;



- ▶ People wonder:
 - ▶ *'What are your research questions?'*
 - ▶ *'Just lots of numbers with very small differences?'*
- ▶ Is it a **science** or an **engineering discipline**?
- ▶ Or may be **CL is a science** and **NLP is its application towards empirical problems**?
- ▶ Motivation for research can be different:
 1. trying to provide a **computational explanation for linguistic or psycholinguistic phenomenon**;
 2. trying to provide a **working component of a speech or natural language system**.



- ▶ People wonder:
 - ▶ *'What are your research questions?'*
 - ▶ *'Just lots of numbers with very small differences?'*
- ▶ Is it a **science** or an **engineering discipline**?
- ▶ Or may be **CL is a science** and **NLP is its application towards empirical problems**?
- ▶ Motivation for research can be different:
 1. trying to provide a **computational explanation for linguistic or psycholinguistic phenomenon**;
 2. trying to provide a **working component of a speech or natural language system**.
- ▶ Do our top-tier conferences belong to CL or to NLP then?



- ▶ People wonder:
 - ▶ *'What are your research questions?'*
 - ▶ *'Just lots of numbers with very small differences?'*
- ▶ Is it a **science** or an **engineering discipline**?
- ▶ Or may be **CL is a science** and **NLP is its application towards empirical problems**?
- ▶ Motivation for research can be different:
 1. trying to provide a **computational explanation for linguistic or psycholinguistic phenomenon**;
 2. trying to provide a **working component of a speech or natural language system**.
- ▶ Do our top-tier conferences belong to CL or to NLP then?
- ▶ The overwhelming majority of papers are empirical today.



- ▶ People wonder:
 - ▶ *'What are your research questions?'*
 - ▶ *'Just lots of numbers with very small differences?'*
- ▶ Is it a **science** or an **engineering discipline**?
- ▶ Or may be **CL is a science** and **NLP is its application towards empirical problems**?
- ▶ Motivation for research can be different:
 1. trying to provide a **computational explanation for linguistic or psycholinguistic phenomenon**;
 2. trying to provide a **working component of a speech or natural language system**.
- ▶ Do our top-tier conferences belong to CL or to NLP then?
- ▶ The overwhelming majority of papers are empirical today.
- ▶ No final answer yet.



Linguistics is back



Linguistics is back

- ▶ NLP is **re-embracing linguistic structure** now;



Linguistics is back

- ▶ NLP is **re-embracing linguistic structure** now;
- ▶ Linguistics can help reduce the search space;



Linguistics is back

- ▶ NLP is **re-embracing linguistic structure** now;
- ▶ Linguistics can help reduce the search space;
- ▶ It can provide inductive biases (linguistic scaffolding).
- ▶ Even the strongest proponents of purely data-driven approaches acknowledge it;



Linguistics is back

- ▶ NLP is **re-embracing linguistic structure** now;
- ▶ Linguistics can help reduce the search space;
- ▶ It can provide inductive biases (linguistic scaffolding).
- ▶ Even the strongest proponents of purely data-driven approaches acknowledge it;
- ▶ Linguistic structures induced into machine learning systems reduce search space, bringing improvements [Dyer, 2017]:
 - ▶ language is inherently hierarchical → syntactic recency is a preferable inductive bias to sequential recency.
- ▶ **Language is not just sequences of words / characters / bytes.**

References I



Dyer, C. (2017).

Should neural network architecture reflect linguistic structure?

In *Proceedings of the 21st Conference on Computational Natural Language Learning (CoNLL 2017)*, page 1. Association for Computational Linguistics.



Zhao, J., Wang, T., Yatskar, M., Ordonez, V., and Chang, K.-W. (2017).

Men also like shopping: Reducing gender bias amplification using corpus-level constraints.

In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2979–2989. Association for Computational Linguistics.