BERT Rediscovers the Classical NLP Pipeline

Ian Tenney[*1], Dipanjan Das[1], and Ellie Pavlick[1,2]*
[1]Google Research, [2]Brown University

Google AI

Paper

Code

Overview

Question: Does BERT [1] learn linguistic abstractions, or is it just really good at summarizing co-occurrence statistics?

- BERT is a deep model. Do the layers make sequential decisions?
- Is linguistic information localized in different layers of the encoder?

Takeaways:

- Linguistic abstractions appear in a consistent order, with POS tagging in lower layers, followed by parsing, NER, semantic roles, then coreference.
- But, individual decisions don't always follow this: low-level decisions can be revised based on high-level information.

BERT by Layer

For each task τ, train probing classifiers $\{P_\tau^{(\ell)}\}$ for ℓ = 0, 1, ..., L

Scalar Mixing Weights:

ELMo-style: let s_τ = softmax(a_τ), and

$$\mathbf{h}_{i,\tau} = \gamma_\tau \sum_{\ell=0}^{L} s_\tau^{(\ell)} \mathbf{h}_i^{(\ell)}$$

Center-of-gravity:

$$\mathbf{E}_s[\ell] = \Sigma_\ell\, \ell\, s_\tau^{(\ell)} \,/\, \Sigma_\ell\, s_\tau^{(\ell)}$$

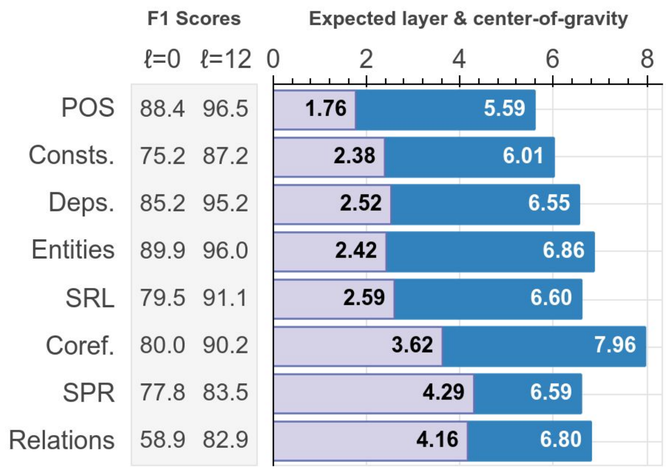

BERT-base (12 layer)

Cumulative Scoring:

$$\Delta_\tau^{(\ell)} = \text{Score}(P_\tau^{(\ell)}) - \text{Score}(P_\tau^{(\ell-1)})$$

Expected layer:

$$\mathbf{E}_\Delta[\ell] = \Sigma_\ell\, \ell\, \Delta_\tau^{(\ell)} \,/\, \Sigma_\ell\, \Delta_\tau^{(\ell)}$$

Tasks appear in a consistent order, reflecting the traditional NLP pipeline!

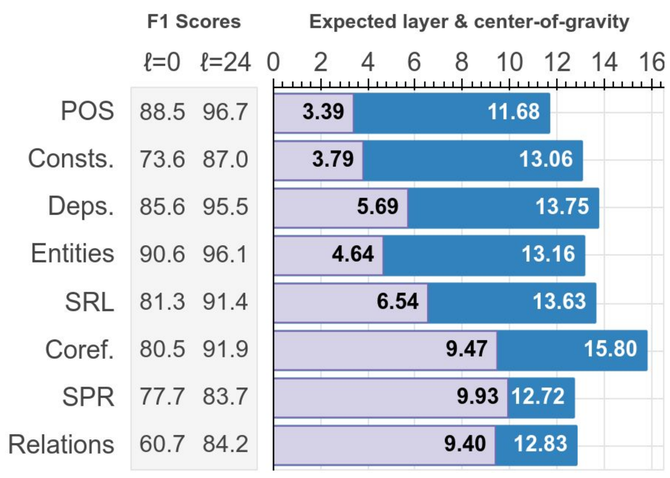

BERT-large (24 layer)

Tracing a Sentence

OntoNotes: τ = {POS, constituents, entities, SRL, coref}

Collect predictions $\{P_\tau^{(\ell)}\}$ for ℓ = 0, 1, ..., L for each task

"he smoked **toronto** in the playoffs with six hits, ... "

Entities: from **GPE → ORG** in layers 3-7

SRL: conclude **ARG1** by layer 2-3

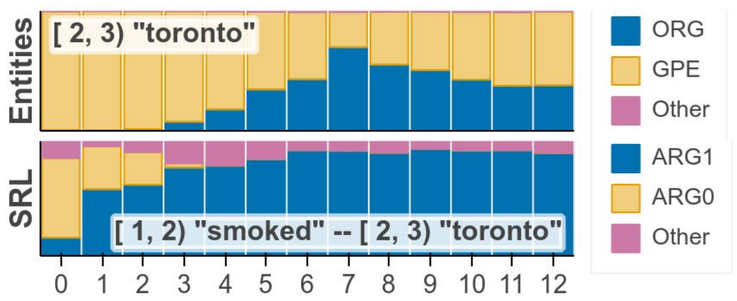

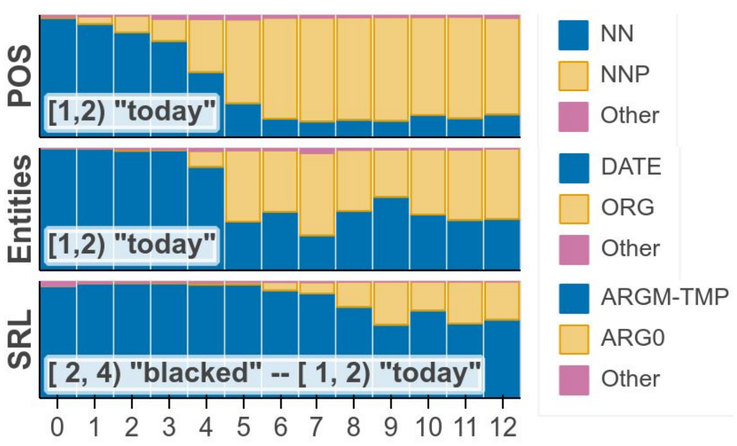

"china **today** blacked out a cnn interview that was critical ... "

POS: from **NN → NNP** in layers 3-5

Entities: from **DATE → ORG** in layer 4-5

SRL: consider **ARG0** from layers 6-9

Edge Probing

Probing suite [2] recasts tasks as edge labeling:

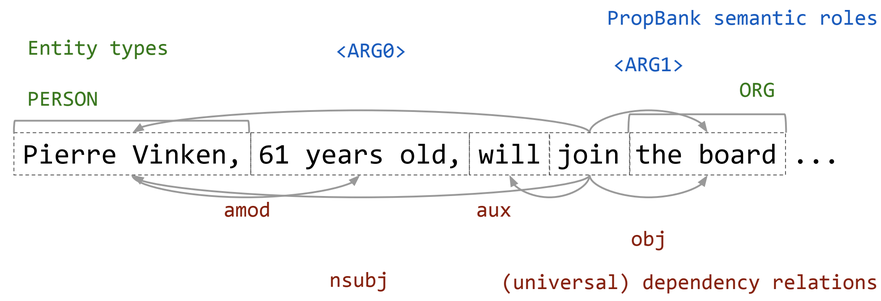

Given contextual vectors E = [e_0, e_1, ..., e_n], predict:

- **Unary:** label(s) for span1 = [i_1, j_1)
- **Binary:** label(s) for (span1 = [i_1, j_1), span2 = [i_2,j_2))

Common classifier model [2] over frozen encoder, with ELMo-style mixing over layers {0,1,...,ℓ}.

Per-layer Contributions

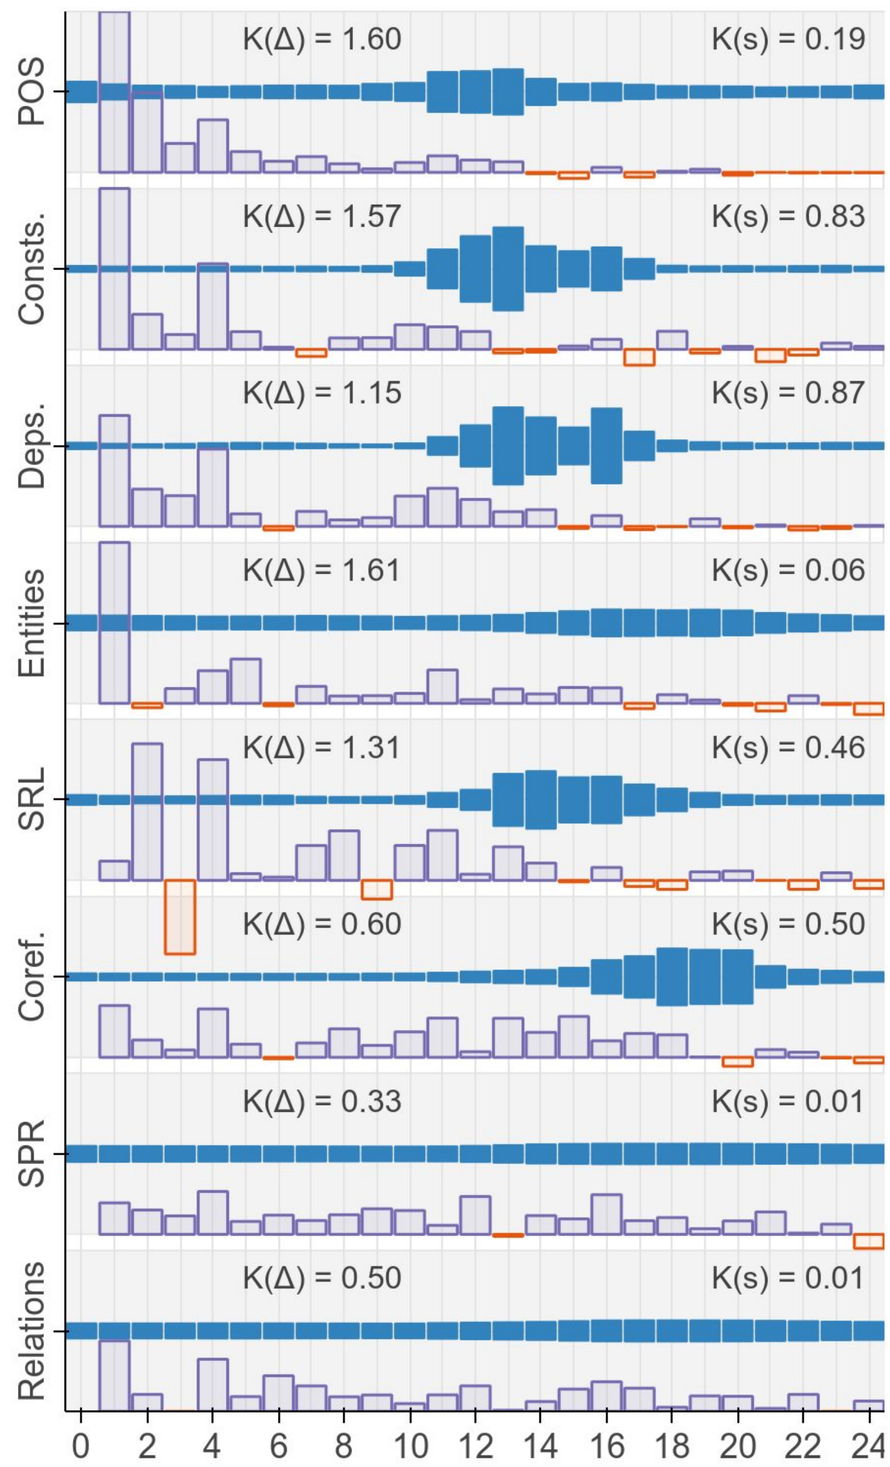

Solid blue: scalar mixing weights (s)

Light purple: relative improvement in F1 score (Δ)

K(*): KL(* | Uniform) over all layers

References

[1] *BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding* (Devlin et al., NAACL 2019)

[2] *What do you learn from context? Probing for sentence structure in contextualized word representations* (Tenney et al., ICLR 2019)