

(D) Pay Attention (1/2) [10 Points]

The meaning of a word depends on its context. For example, in the sentence “The farmer *seeded* the field with corn,” the word *seeded* means “added seeds to.” However, in the sentence “The chef seeded the tomato,” the word *seeded* means “took seeds away from.”

If you were building a model of language, how would you get it to recognize the way that a word’s meaning depends on context? One popular technique for achieving this goal is a mechanism called **attention**. In the way that attention is implemented in current state-of-the-art models of language, the model has a large number of **attention heads**, each of which is denoted by a pair of numbers (for examples, 8-10). When the model processes a sentence, for every pair of words in the sentence, each head calculates the “relatedness” of the two words.

The one wrinkle is that we do not know what exactly “relatedness” should mean, so instead of telling the model how to define “relatedness,” we let the model learn its own definition of relatedness. Recently, computer scientists have started to analyze what these attention heads have learned, and this analysis shows that they often reflect linguistic information! For example, here’s the output of one attention head (head 8-10) when we feed the following sentence into BERT, which is one of the most popular models that uses attention heads:

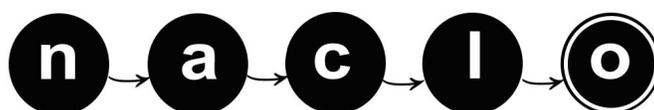
Example: I_1 see $_2$ my $_3$ sister $_4$, but $_5$ she $_6$ can’t $_7$ see $_8$ me $_9$ because $_{10}$ she $_{11}$ is $_{12}$ reading $_{13}$ a $_{14}$ linguistics $_{15}$ book $_{16}$.

Output: $4 \rightarrow 2$, $9 \rightarrow 8$, $16 \rightarrow 13$

This output signifies that head 8-10 connects word #4 (*sister*) to word #2 (*see*), as well as word #9 (*me*) to word #8 (*see*) and word #16 (*book*) to word #13 (*reading*). If you consider what all of those pairs of words have in common, you’ll see that each one is a verb and its direct object: *sister* is the direct object of the first instance of *see*, *me* is the direct object of the second instance of *see*, and *book* is the direct object of *reading*. It appears that head 8-10 has learned to connect verbs to their objects! (Note that these activations are directional; for example, word #2 is not connected to word #4.)

Why is this information useful? If we go back to the example with the verb *seed*, this sort of information can help the model figure out which version of *seed* is being used: If its direct object is something like *field* or *lawn*, then it probably means “to add seeds to;” if its direct object is something like *tomato* or *watermelon*, then it probably means “to take seeds away from.” Of course, one sentence isn’t enough to draw strong conclusions. Instead, computer scientists tend to use a *corpus*, or a database, of example sentences to find patterns in the data. On the next page is a small corpus, the NacloWeb Corpus¹, which has 7 sentences.

¹ Some sentences derived from data in the English Web Treebank.



(D) Pay Attention (2/2)

NacloWeb Corpus

1. My₁ experience₂ with₃ Gelda₄ 's₅ House₆ of₇ Gelbelgarg₈ has₉ been₁₀ extremely₁₁ wonderful₁₂
2. We₁ use₂ Google₃ 's₄ models₅ to₆ delve₇ into₈ the₉ inner₁₀ workings₁₁ of₁₂ language₁₃
3. At₁ this₂ corporation₃ 's₄ meeting₅ people₆ are₇ concerned₈ about₉ the₁₀ company₁₁ 's₁₂ plans₁₃
4. In₁ July₂ we₃ will₄ interview₅ the₆ candidate₇ and₈ review₉ her₁₀ resumé₁₁ again₁₂
5. The₁ platypus₂ is₃ a₄ strange₅ animal₆, with₇ its₈ eggs₉ and₁₀ its₁₁ webbed₁₂ feet₁₃
6. I₁ think₂ that₃ although₄ my₅ NACLO₆ exam₇ was₈ difficult₉ it₁₀ was₁₁ a₁₂ lot₁₃ of₁₄ fun₁₅
7. Linguistics₁ is₂ a₃ beautiful₄ science₅ that₆ provides₇ interdisciplinary₈ insight₉ into₁₀ the₁₁ human₁₂ experience₁₃

Note that the NacloWeb corpus treats the possessive clitic 's as a separate word. (So in Sentence #1, word #5 is 's and word #6 is House.)

In our experiment on the NacloWeb Corpus, we ran each of the corpus' sentences through BERT and recorded the outputs of four attention heads (8-11, 7-6, 9-6, and 5-4). Unfortunately, due to some extremely sloppy experimental procedure, we don't remember which order we ran them through the model; in addition, we forgot to record some data. Your job is to fill in the blanks! Note that some blanks may have more than one connection, and some may have none at all.

Sentence	8-11	7-6	9-6	5-4
Sentence A	12 → 13	5 → 6	14 → 13	10 → 6
Sentence B	(a)	8 → 9, 11 → 12	None	8 → 2, 11 → 2
Sentence C	2 → 3, 10 → 11	4 → 5, 12 → 13	(b)	11 → 3
Sentence D	(c)	4 → 5	12 → 13	None
Sentence E	3 → 4, 11 → 12	(d)	10 → 12	5 → 1
Sentence F	(e)	1 → 2, 5 → 6	7 → 8	None
Sentence G	(f)	(g)	(h)	(i)

D1. Identify sentences A-G. Write one number from 1 to 7 per box.

A		B		C		D	
E		F		G			

D2. Fill in the missing data, labelled (a) to (i), in the table above.

