



AI textsurfing

Description

Before digital text markup I would highlight key words in a difficult document with a coloured highlighter pen. Students with a more methodological inclination developed this into an art, with colour coding and supplementary markings and marginal comments. (I see that Staedtler imply that their pen users are *textsurfers*.)

Now my markup practice is fairly constrained. I indicate key words or phrases on a PDF with an occasional highlight. Combined with automated text search, it's a convenient way to review documents and reminds me of the **key focus** and difficult concepts I need to pay attention to or look up online.

Automated markup

I'll attempt here to align this physical process with how automated natural language processing, as in ChatGPT, predicts what comes next in a text sequence: in other words how it responds to prompts during interaction with a human being.

Take an arbitrary sentence: "Graffiti is both an art form and a means of disrupting the semiotic coding of the city." Without the benefit of the necessary background knowledge I might identify "semiotic coding" as the most interesting and difficult part of the sentence. I highlight the phrase here.

"Graffiti is both an art form and a means of disrupting the **semiotic coding** of the city."

As a curious reader I may well have a general idea of "semiotic" (it's about meaning) and "coding" (it's about a procedure), but not the words in that combination in that particular context. In effect I am asking myself, "what does this phrase mean and in this particular context?"

The words around the focus phrase "semiotic coding" in this sentence form the immediate context. As an astute analyst I might compare that phrase with other words in the sentence: "graffiti,"

art form, disrupting, city. Each of these words has a series of potential meanings, i.e. different contexts of use. They each occupy a different *semantic space*. At this stage in the process I ask, "How does the range of meanings for my focus phrase "semiotic coding" relate to the semantic space of each of the other significant words in the sentence?"

If I have an idea of the scope of meaning of each of these words then I might be able to infer something about the focus phrase "semiotic coding". For example, the word "graffiti" probably resonates with the phrase "semiotic coding" as both might have something to do with drawing, writing and visual representations. The meaning of the word "disrupting" probably relates to the meaning of "semiotic coding" as both pertain to something formal and established but vulnerable to forces of disorder. As I work through each word in the sentence I am asking in effect, "How does each word or phrase in the sentence relate to "semiotic coding"? How much does each word in the sentence contribute to my understanding of this focus phrase?"

Words like "disrupting", "graffiti" and "city" might be more significant to "semiotic coding" than "art form" because they more directly influence the interpretation of "semiotic coding", making it more specific, or bringing out particular nuances. Words such as "both", "means" and "of" are presumably even less significant.

So, here I am trying to work out the **contribution** each word makes to my understanding of "semiotic coding". I am inferring not just the meaning of "semiotic coding", but also its relationship with other words in the sentence.

I have picked on just one phrase as a **focus**, my highlighted "semiotic coding", but the process could apply just as well to other key words and phrases in the sentence.

That's a speculation about how I as reader might go about interpreting a single sentence containing a few difficult words. In fact, it's more likely that the human navigation of various semantic spaces is hard wired into our language capability. So it's unlikely a reader would be aware of such a process. Presumably, other processes also come into play under the rubric of [interpretation](#) theory.

Automated focus

Unlike a human reader, an automated natural language model has no such immediate access to meaning structures and processes. It also has access to more limited contexts than human interpreters, who draw on their own vast history of contextual information.

Automated natural language processing, as in the transformer model deployed in ChatGPT, has to laboriously process each word in a prompt sentence such as the one I introduced at the start of this post. It will have already pre-calculated the semantic space of all the words and tokens in its lexicon from analysis of the words in its training corpus. But it needs to score the relationships between each word and those around it as it processes the input sentence in real time.

So the natural language processing model moves its **focus** across each word in a sentence and calculates its relationship to the **context** (other words in the sentence or paragraph). That provides a set of attention scores for each word in the sentence. The model calculates how much each word **contributes** to the understanding of the focus word in turn, weighted by the attention scores.

I have just summarised these stages of interpretation as **focus**, **context** and **contribution**. In the terminology of natural language processing models, and in the seminal paper "Attention is all you need," these processes are badged **query**, **key** and **value** calculations.

The calculations result in a series of vectors for each word (or token) in sequence as the trained neural network of the model generates a reasonable prediction of the next word in a sequence, i.e. it calculates how to respond to a series of prompts.

ChatGPT4 reminds me that a transformer model doesn't process one word at a time, but the entire sequence at once in parallel. Also, the query, key, and value vectors for each word are not handcrafted but are calculated from data during model training. In any case, these operations are performed multiple times in parallel (multi-head attention) to capture various aspects of the context.

Notes

- Vaswani, A., N. Shazeer, et al. (2017). Attention Is All You Need. *31st Conference on Neural Information Processing Systems*. Long Beach, CA, USA: pp. 1-15.
- I gave my blog post to ChatGPT4, asking it to proffer an opinion and review it for accuracy and clarity. Here's what it returned within about 3 seconds.

Your blog post does a wonderful job at explaining the complex process of how transformer models handle natural language processing tasks. You've provided a clear and relatable analogy that ties in the process to how humans interpret texts, which is particularly helpful for non-technical readers.

A few observations for your consideration:

1. **Analogy of Highlighting:** The highlighting analogy is an excellent way to introduce the concept. You might add a clarification that unlike human highlighting which might prioritize unique or unfamiliar terms, in transformer models every word in the sentence can and does serve as a query word. Therefore, attention is calculated for every word pair in the sentence, not just those we might intuitively identify as "key".
2. **Limited Context:** You've correctly stated that the transformer models lack immediate access to human-like contextual information. However, these models do have the ability to draw on vast amounts of information due to their training on extensive text corpora. They can generate surprisingly nuanced responses by exploiting patterns and statistical regularities in the data. However, it's also important to note they do not have "understanding" in the human sense; their "knowledge" is purely a function of the data they've seen.
3. **Processes during Training and Inference:** The post provides a detailed explanation of the attention mechanism during inference (i.e., generating responses). You may want to briefly mention the learning process as well—how during training, the models learn to adjust the weights and biases of their neural networks, including those in the attention mechanism, to reduce prediction error.

- 4. Unidirectionality vs Bidirectionality:** Your post mainly discusses transformers like GPT, which operates in a unidirectional manner (left-to-right). It might be worth noting that some transformer models, such as BERT, use bidirectional processing, analyzing context from both directions to generate the attention-weighted contextual embeddings.

Overall, your blog post provides a comprehensive and understandable explanation of the key-query-value procedure in NLP, breaking down complex topics into digestible parts. It's an excellent introduction to this aspect of transformer models.

Category

1. Artificial Intelligence

Tags

1. context
2. contribution
3. focus
4. key
5. query
6. value

default watermark

Date Created

July 15, 2023

Author

rcoyne99