



Attention is everything

Description

Attention is a key element in cognition. At our most thoughtful we direct attention to features in our environment that are most important to us at that moment. Attention can wander, of course, we daydream, and we can pay attention to [inexistent](#) things, memories and objects of the imagination.

A lecturer will come to the class armed with notes or key points on presentation slides that not only focus the attention of the audience, but aid the speaker in generating their word flow. Affect, emotion, mood, melancholy reminds me not only to use those words in my lecture on emotional landscapes, but structures my production of a series of sentences on the subject. Interpreting text written in a language with which I have only passing competence, such as French or German, is aided by the ability to attend to familiar words in the sentence and to guess at the rest.

Attention everywhere

The ability to *attend*, selectively, is a vital cognitive function, as we listen, perceive, speak, and address everyday challenges. Previously, I explored attention fatigue and its restoration by environmental conditions that afford the remedy of [soft fascination](#). I have also learned about the key role of attention in William James philosophy (post: [Attending to the world](#)), the role of [inattention](#) in social interaction, and the relationship between attention and *intention* (post: [Intentional systems](#)). The latter impinges on how philosophers attempt to explain [consciousness](#).

As a minor feat of mechanical anthropomorphism it is easy to see how automated processes, e.g. computer systems, might deploy operations described readily as *attentional*. To change a house to a HOUSE in a database, a program cycles through a list of words in the database and compares them with house. So, it effectively scans the list to find the object on which to direct its attention and on which to conduct a series of simple transformation and replacement operations. With even greater sophistication, web search engines quickly home in on (attend to) some target data to display, extract or manipulate.

Is attention all you need?

I've been reading the seminal article "Attention Is All You Need" authored by a team of eight AI researchers led by Ashish Vaswani. Their research led to what many regard as the third major innovation in the development of natural language processing systems, and is the key technique employed in the current generation of impressive platforms such as GPT-3 (openai.com), which receives text from a user as input and generates further text in response. Such AI systems generate new text on the basis of a best guess as to what ought to come after the user's input. Variants of the language model result in question and answer, chat, translation and other formats.

A helpful [online lecture by Abigail See](#) on Natural Language Processing highlights three phases in the development of computer-aided natural language processing (NLP). NLP was motivated initially by protagonists in the cold-war for whom rapid Russian-English translation of communications and documents could provide strategic advantage.

Cumbersome rule and lexicon-based translation tools on the clunky computers of the day were replaced in the 1990s by statistical machine translation techniques abetted by repositories of parallel texts as published in multilingual international newspapers, UN proceedings, etc. Then followed the application of neural network techniques that involved machine learning from word sequences. Eventually these were enhanced by techniques that included methods for focussed attention in the processing of text, the subject of Vaswani and co's research.

Attentional NLP is more efficient than its predecessors, which were encumbered by the computational cost of processing multiple alternative word sequences to arrive at the most plausible sequence. By my reading, attentional NLP adopts something of the linguistic process I described in the second paragraph of this post: identify key words to generate word flow. I'll discuss current NLP methods in subsequent posts.

Though the impetus was machine translation, current language processing models have general applicability. They translate between languages, but also generate new text in response to queries, make a good guess at the completion of a sentence, paragraph, story or poem, produce summaries and generate paragraphs from key words, phrases and prompts. As I have demonstrated in previous posts, you can have a pretty good [conversation](#) with an NLP chatbot that in the right context comes close to passing the [Turing test](#).

References

- See, Abigail. "NLP with Deep Learning | Winter 2019 | Lecture 8 Translation, Seq2Seq, Attention." *Stanford Online*, Winter, 2019. Accessed January 10, 2023, <https://www.youtube.com/watch?v=XXtpJxZBa2c&t=4337s>
- Vaswani, Ashish, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. "Attention Is All You Need." In *31st Conference on Neural Information Processing Systems*, 1-15. Long Beach, CA, USA2017.
- James, William. *The Principles of Psychology Volume I*. Cambridge, MA: Harvard University Press, 1981.

Note

- MidJourney generated the featured image from the prompts: affect, emotion, mood, melancholy, landscape.

Category

1. Uncategorized

Tags

1. AI
2. attention
3. language
4. NLP

Date Created

February 4, 2023

Author

rcoyne99

default watermark