

An AI focus group

Description

default watermark

Before widespread digitalisation, I recall borrowing books from the library that were graffitied with penciled underscores and multi-coloured hi-lighter markings indicating what various past readers thought important. In terms of my previous posts about attention in NLP models, such markups constitute a compelling record of human-based â??multi-headed attention.â?•

I described multi-head attention in my previous post. NLP training algorithms are set up to accord different weights to tokens (words) in a string of text. The result is a markup of alternative patterns of attention.

Multi-head attention in NLP models (e.g. Transformers) is also a bit like processes within a focus group, each member forming different opinions that coalesce into a pattern of opinions as each is struck by a different sense of what is important.

How groups focus

Imagine such a group is tasked with assessing a newly proposed urban development, a waterfront leisure park. I put the urban focus group analogy to ChatGPT, which (who) readily latched on to several points of similarity. The chat confirmed that the Transformer model is close to what happens when people come together to share insights.

> â??In a focus group, each participant observes, interprets, and responds to the discussion based on their own experiences and knowledge. Similarly, in multi-head attention, each head processes the input independently, creating its own unique representation.â?•

According to this model, different attention patterns are not prescribed, as if each â??headâ?• has a different target of interest. Rather, the priorities of the different heads emerge through training.

If it is running correctly, in a focus group the discussion progresses with certain â??themes or consensus pointsâ?• emerging without explicit coordination. In a similar manner, though they initially

operate independently, different heads in a Transformer model contribute collectively to diverse, multi-faceted practical patterns of attention, leading ultimately to coherent conversational simulations.

Focus group members might refine their perspectives as they listen to and respond to others. In a similar manner, the Transformer model adapts during training. The matrices in the model processes â??evolve, shaping how each head attends to the data.â?•

Again, if it is operating true to form, thereâ??s no predetermined agenda dictating what the group should focus on. In a Transformer, â??each headâ??s focus is not preset but is shaped organically through the training process.â?•

Some formal committee meetings and juries work towards consensus decided by overt negotiation or voting. ChatGPT explains focus group protocols, drawing on its familiar reference to â??nuanceâ?•: â??in a focus group, the collective decision or consensus is more than just a simple aggregation of individual opinions. Itâ??s a nuanced understanding that takes into account the interplay of different viewpoints.â?• Transformer models offer a similar protocol: â??the aggregation of outputs from multiple attention heads leads to a decision or output that is informed by a complex interplay of different learned patterns and relationships in the data.â?•

default watermark

The focus group analogy pertains to the inner workings of the transformer model. The attention mechanism tunes the behaviour of a trained model such that it can respond to a diverse range of conversational contexts. That this multi-head attention method resides deep within the algorithms of these models equips them well as quasi participants in human conversation.

Note

- Featured image is by Dall-e prompted by: â??Close up of a magnifying glass focussed on the page of a book that has been marked up by different coloured highlighter pens and pencils. The parts of the image outside of the magnifying glass view are slightly out of focus.â?•

Category

1. Artificial Intelligence

Tags

1. attention
2. focus group
3. markup

Date Created
December 16, 2023
Author
rcoyne99