

Twitter Thread by [gail weiss](#)



gail weiss

@gail_w



EXTREMELY excited to announce RASP, a programming language whose goal is to provide a computational model for transformers in much the same way that automata have served for RNNs. Work with [@yoavgo](#) and [@yahave](#), accepted into ICML 2021. <https://t.co/3Vp57mHhbJ>

Thinking Like Transformers

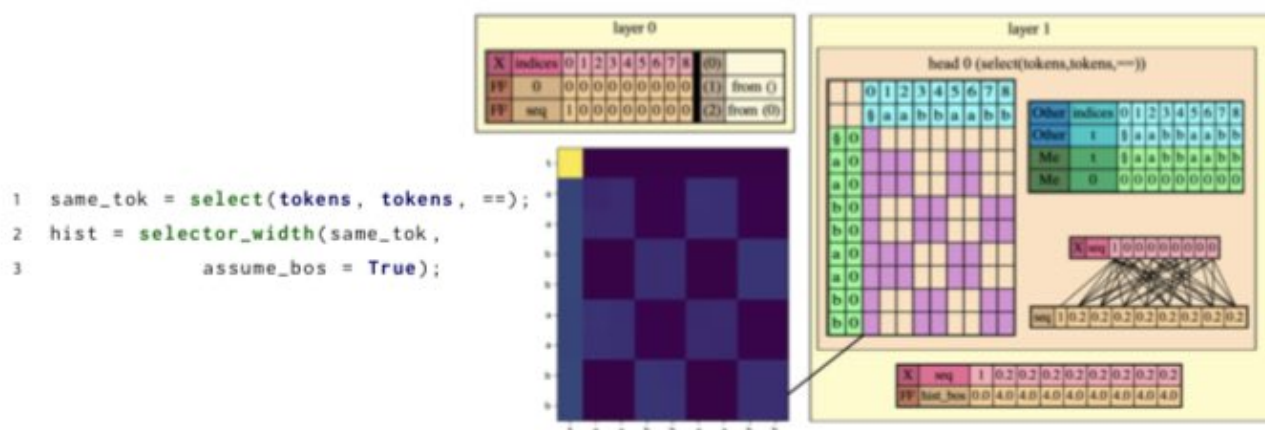


Figure 5: The RASP program for computing with-BOS histograms (left), alongside its compilation to a transformer architecture (cream boxes) and the attention head (center bottom) of a transformer trained on the same task, without attention supervision. The compiled architecture and the trained head are both presented on the same input sequence, "5aabbaabb". The transformer architecture was generated in the RASP REPL using `draw(hist, "5aabbaabb")`.

RASP abstracts away low-level operations into simple primitives, allowing a programmer to "think like a transformer" without getting bogged down in the details of how their solution is realised in practice. We use it to imagine how a transformer might solve several formal tasks

Using RASP, we found solutions for problems such as Dyck-k (for any k and depth!), and logical reasoning (in a simplified variant of the task considered in "Transformers as Soft Reasoners over Language"), previously unsolved for the Transformer model

Anyway, this one goes out to my hater, [@mmgm](#)

<https://t.co/E7RPWnQZeo>

getting rave reviews on my latest venture pic.twitter.com/zXrul4NZvb

— gail weiss (@gail_w) [December 7, 2020](#)