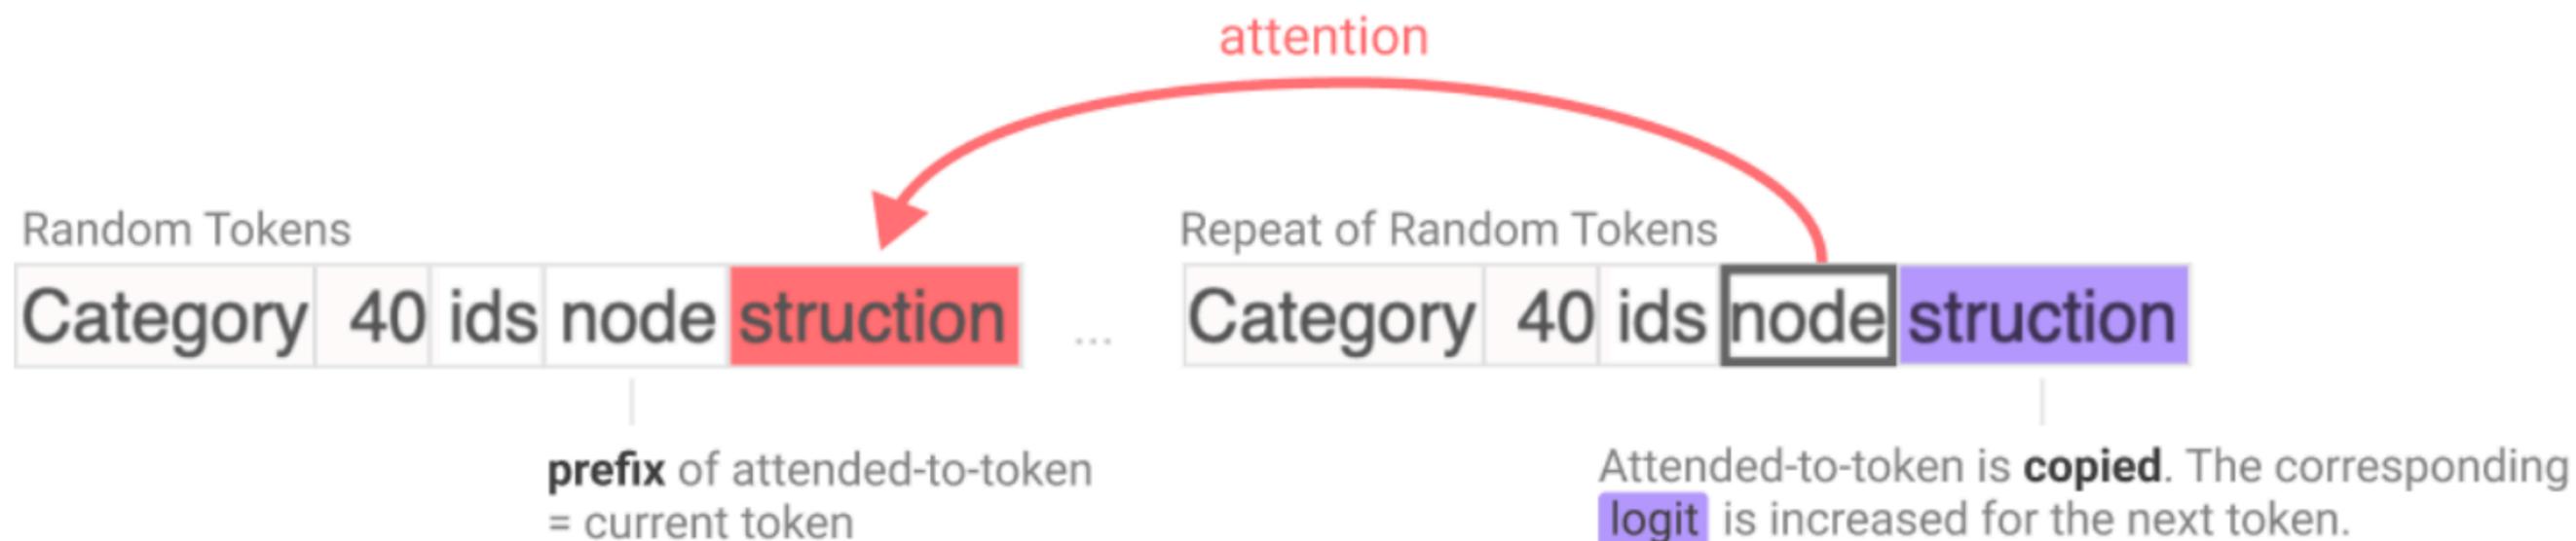


Background: Transformer Circuits

- ▶ There are mechanisms in Transformers to do “fuzzy” or “nearest neighbor” versions of pattern completion, completing $[A^*][B^*] \dots [A] \rightarrow [B]$, where $A^* \approx A$ and $B^* \approx B$ are similar in some space
- ▶ Olsson et al. want to establish that these mechanisms are responsible for good ICL capabilities
- ▶ We can find these heads and see that performance improves; can we causally link these?

Induction Heads

- ▶ Induction heads: a pair of attention heads in different layers that work together to copy or complete patterns.
- ▶ The first head copies information from the previous token into each token.
- ▶ Second attention head to attend to tokens based on what happened before them, rather than their own content. Likely to “look back” and copy next token from earlier
- ▶ The two heads working together cause the sequence ...[A][B]...[A] to be more likely to be completed with [B].

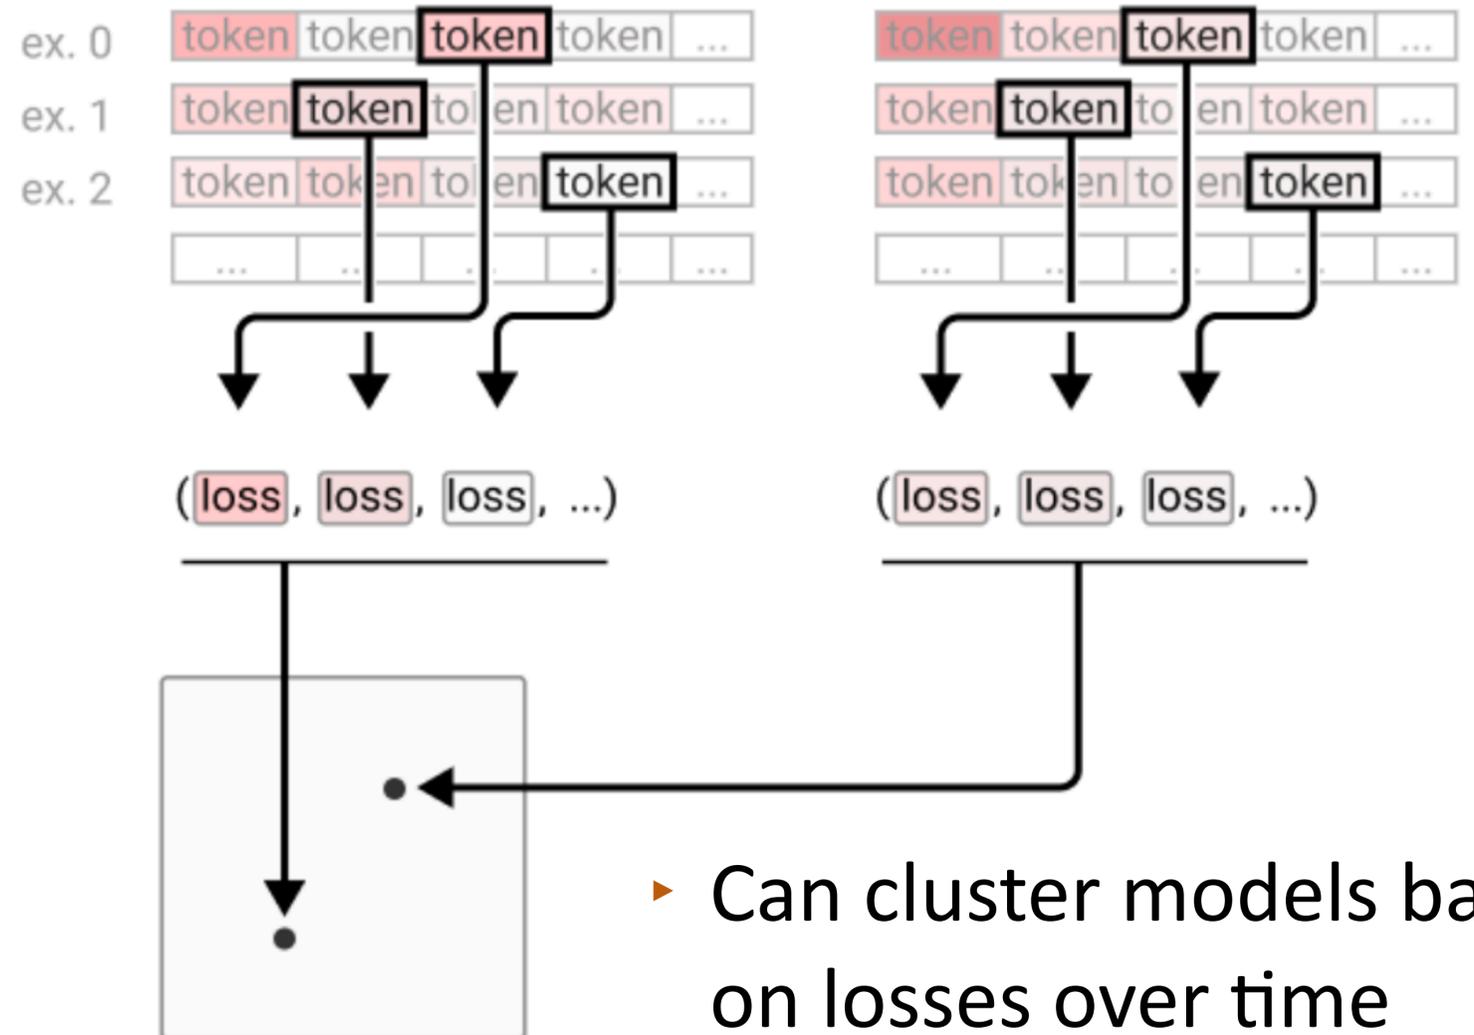


Induction Heads

Step 1: Run each model / snapshot over the same set of multiple dataset examples, collecting one token's loss per example.

Step 2: For each sample, extract the loss of a consistent token. Combine these to make a vector of losses per model / snapshot.

Step 3: The vectors are jointly reduced with principal component analysis to project them into a shared 2D space.



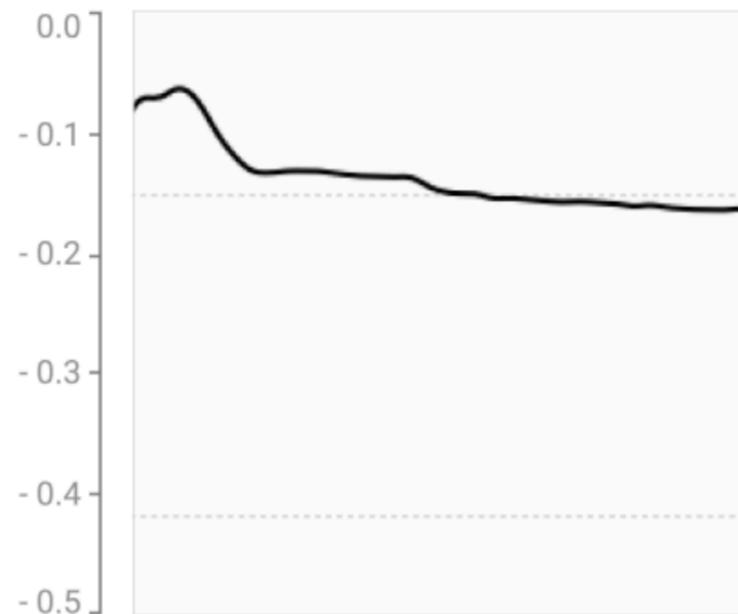
- ▶ Characterize performance by ICL score: $\text{loss}(500\text{th token}) - \text{loss}(50\text{th token})$ — average measure of how much better the model is doing later once it's seen more of the pattern

Induction Heads

ONE LAYER
(ATTENTION-ONLY)

Elapsed Training Tokens

0 2.5e9 5.0e9 7.5e9 1e10

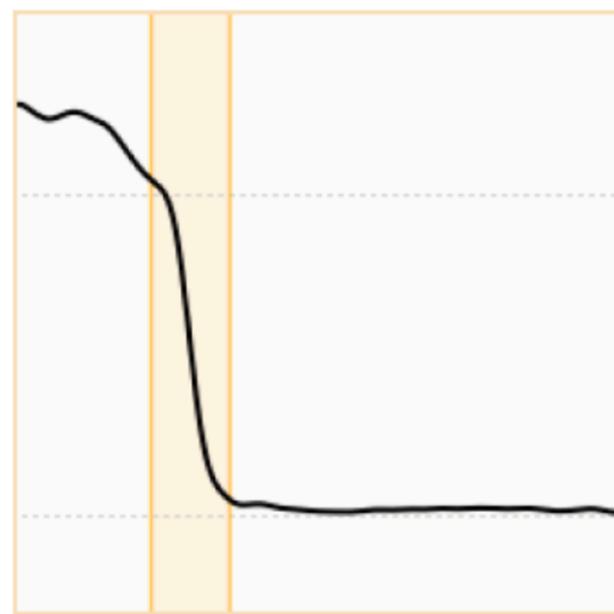


One-layer model
has no sudden improvement.

TWO LAYER
(ATTENTION-ONLY)

Elapsed Training Tokens

0 2.5e9 5.0e9 7.5e9 1e10

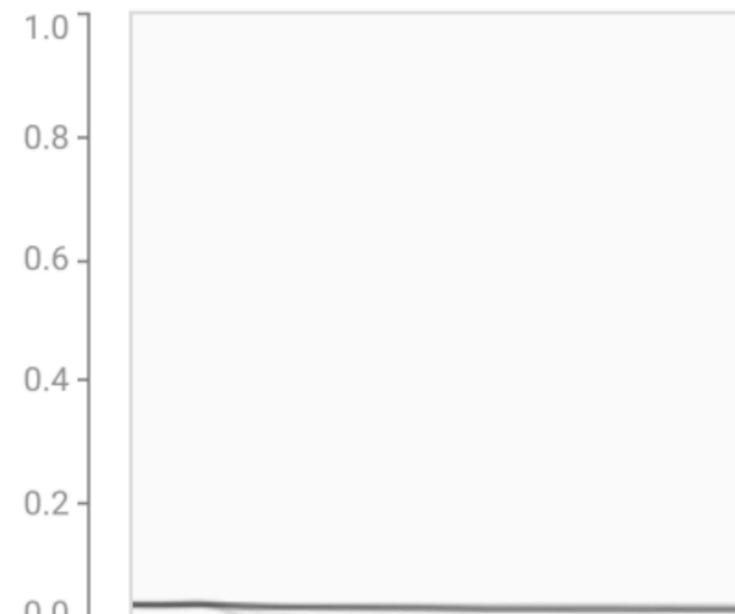


Models with **more than one layer**
have a **sudden improvement** in in-c

ONE LAYER
(ATTENTION-ONLY)

Elapsed Training Tokens

0 2.5e9 5.0e9 7.5e9 1e10

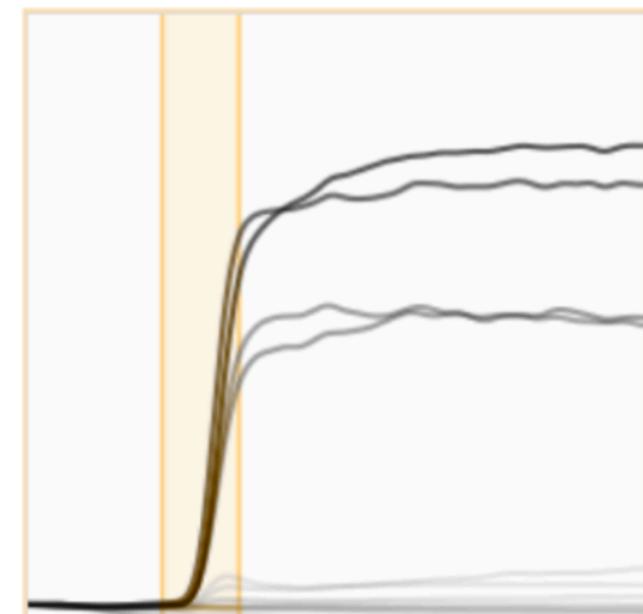


One-layer model
has no induction heads.

TWO LAYER
(ATTENTION-ONLY)

Elapsed Training Tokens

0 2.5e9 5.0e9 7.5e9 1e10

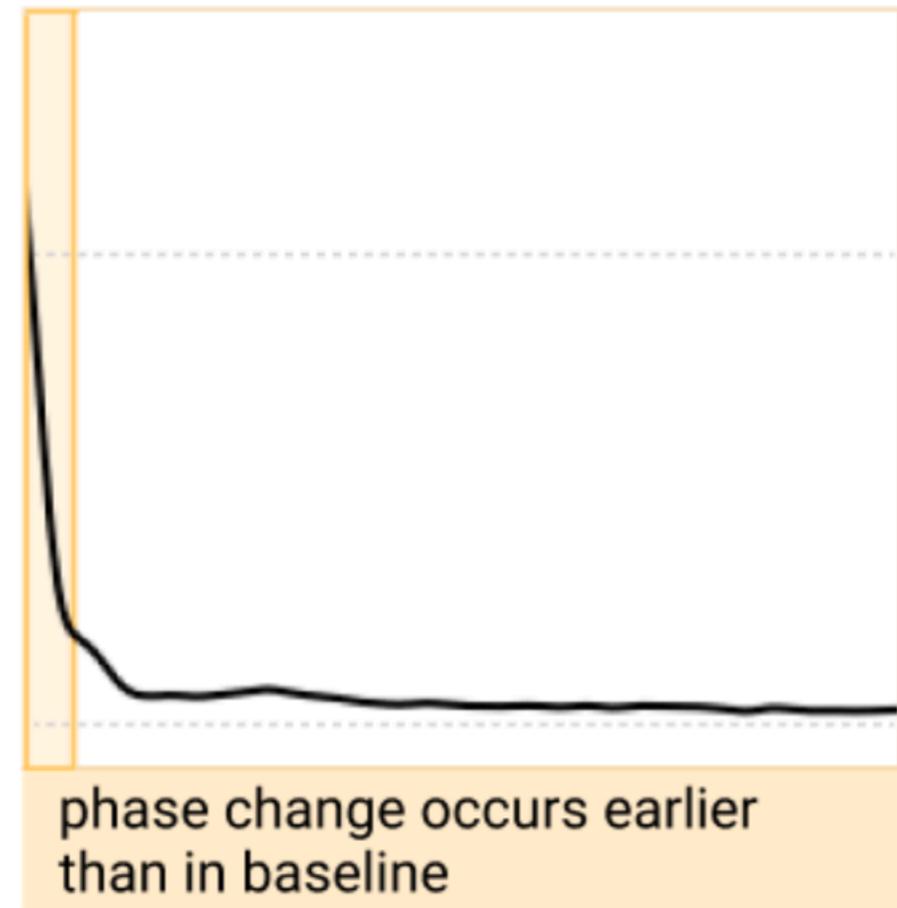
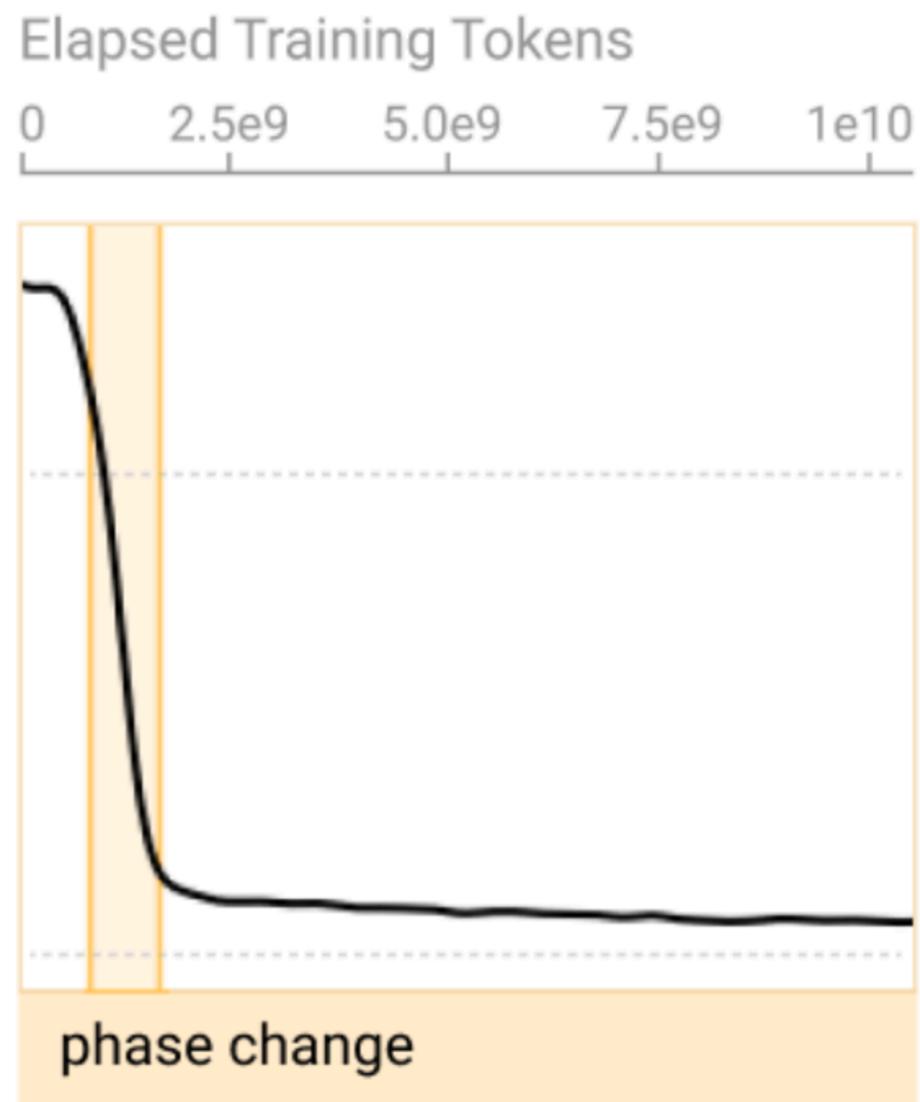


Models with **more than one layer**
have **induction heads form** during p

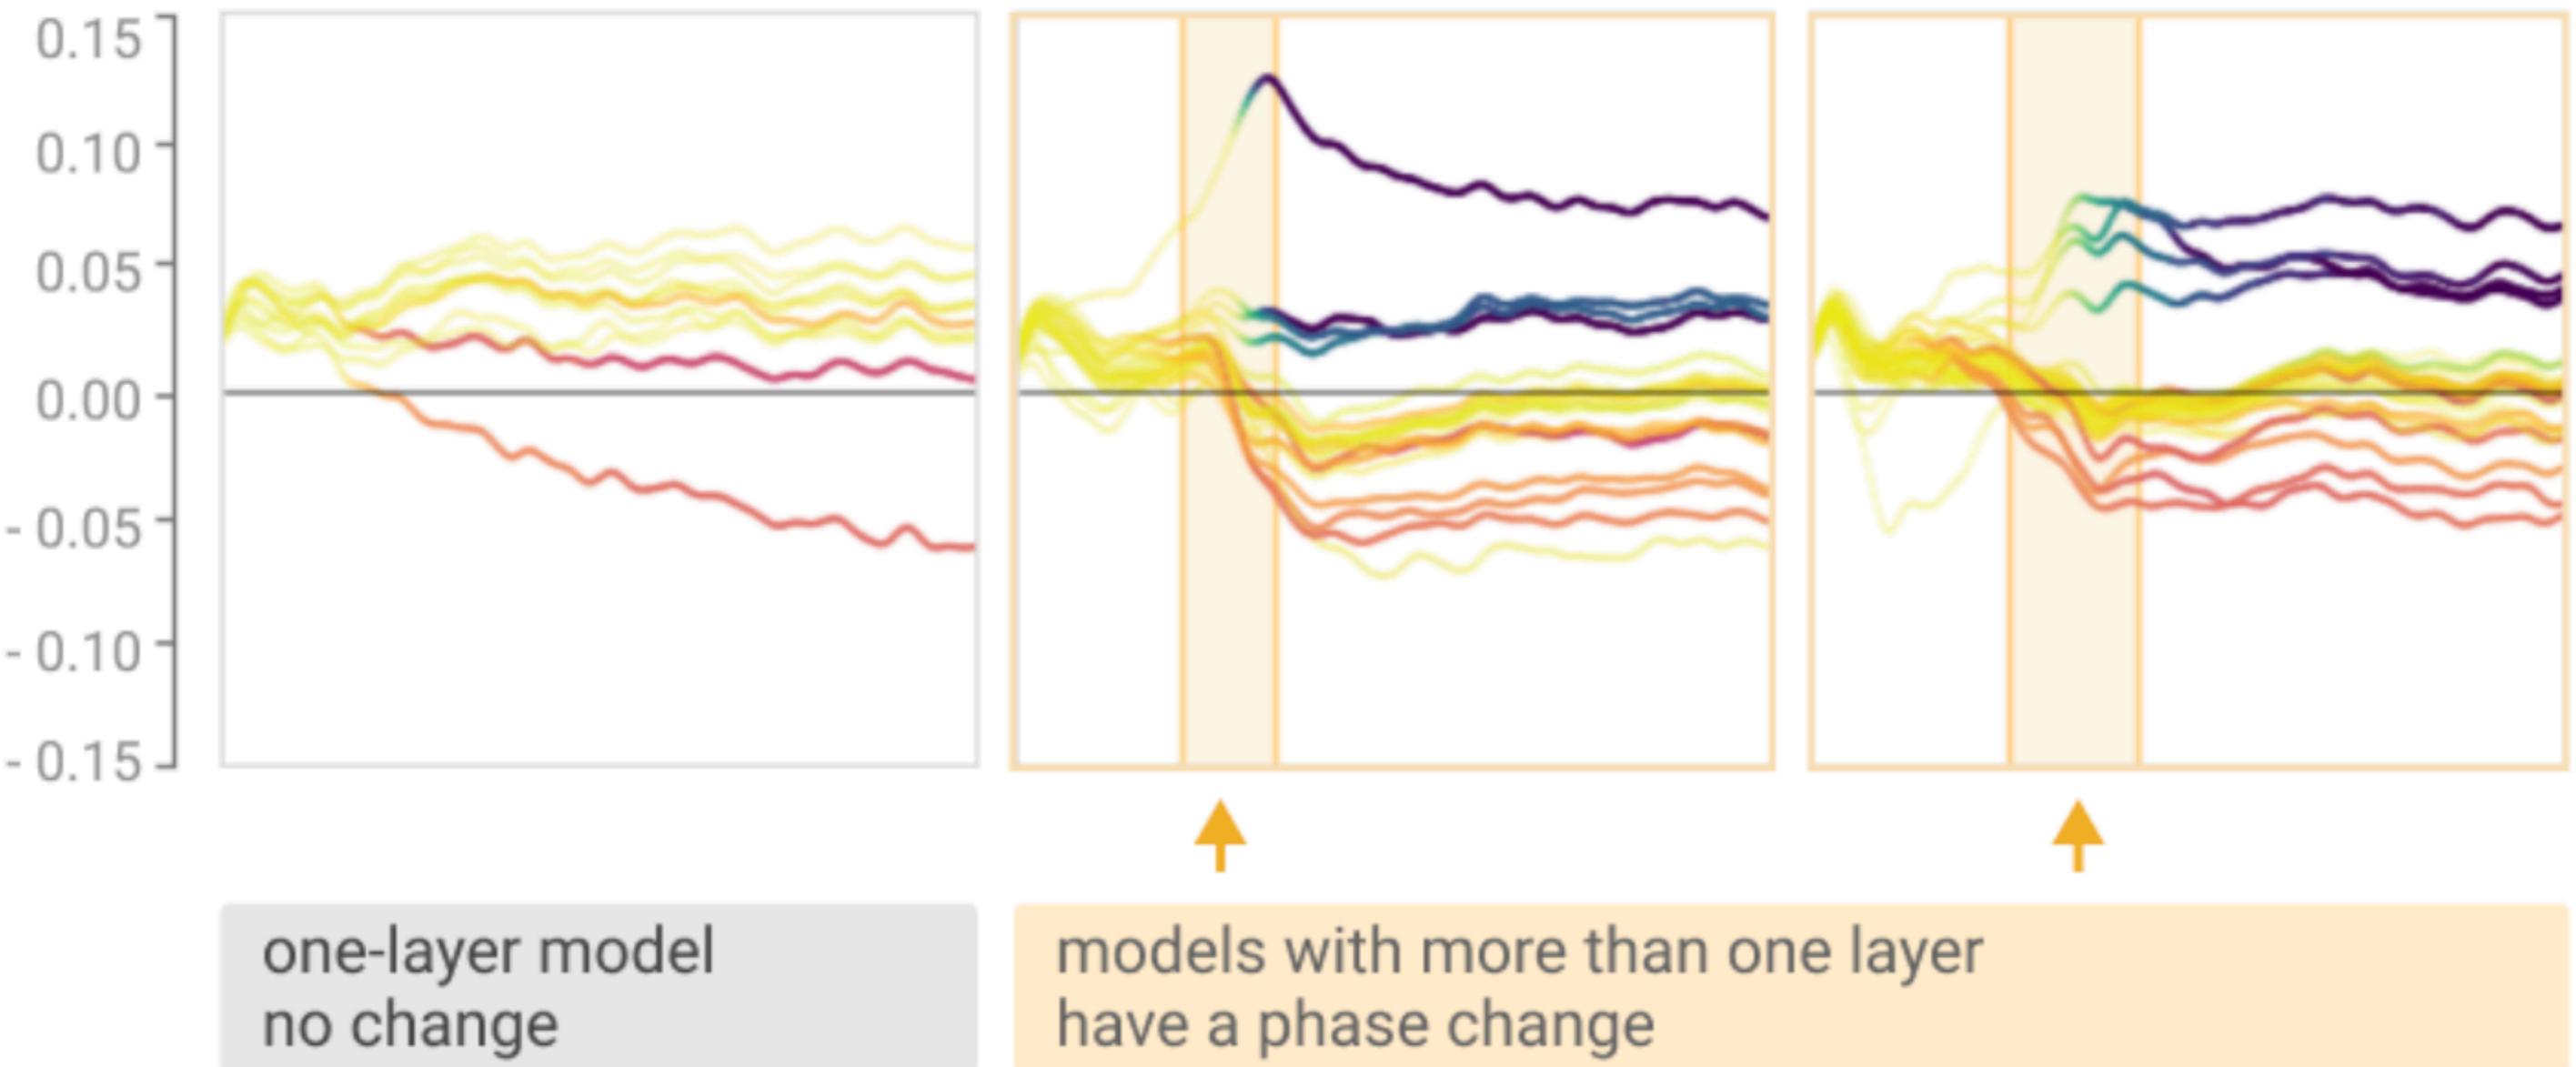
- ▶ Improvement in ICL (loss score) correlates with emergence of induction heads

Induction Heads

Change architecture to promote induction heads => phase change happens earlier



Induction Heads



- ▶ If you remove induction heads, behavior changes dramatically

Interpretability

- ▶ Lots of explanations for why ICL works — but these haven't led to many changes in how Transformers are built or scaled
- ▶ Several avenues of inquiry: theoretical results (capability of these models), mechanistic interpretability, fully empirical (more like that next time)