

Causal interventions expose implicit situation models for commonsense language understanding

Takateru Yamakoshi¹, James L. McClelland², Adele E. Goldberg³, Robert D. Hawkins³

¹The University of Tokyo, ²Stanford University, ³Princeton University
jlmcc@stanford.edu, {takateru, adele, rdhawkins}@princeton.edu

Abstract

Accounts of human language processing have long appealed to implicit “situation models” that enrich comprehension with relevant but unstated world knowledge. Here, we apply causal intervention techniques to recent transformer models to analyze performance on the Winograd Schema Challenge (WSC), where a single context cue shifts interpretation of an ambiguous pronoun. We identify a relatively small circuit of attention heads that are responsible for propagating information from the context word that guides which of the candidate noun phrases the pronoun ultimately attends to. We then compare how this circuit behaves in a closely matched “syntactic” control where the situation model is not strictly necessary. These analyses suggest distinct pathways through which implicit situation models are constructed to guide pronoun resolution¹.

1 Introduction

Language understanding is deeply intertwined with world knowledge. For example, when reading a sentence like “the fish ate the worm,” we can guess that the fish was probably hungrier before eating and that the worm is no longer alive, even though neither property is explicitly mentioned (Winograd, 1972; Rumelhart, 1975). Classical psycholinguistic accounts have suggested that such knowledge enters into language understanding through structured schemas called *situation models* (Zwaan and Radvansky, 1998; Morrow and Clark, 1988; Bransford and Franks, 1971; Schank and Abelson, 1977; Bower et al., 1979; Graesser et al., 1994; Johnson-Laird, 1983) that are dynamically constructed during comprehension. Put succinctly, a situation model is a representation of the web of entities and relations that are *implied* without being explicitly

¹code & data: <https://github.com/taka-yamakoshi/situation-models> / demo: <https://huggingface.co/spaces/taka-yamakoshi/causal-intervention-demo>

context cue manipulation

S_A Paul tried to call George but [pron.] wasn't successful.
 S_B Paul tried to call George but [pron.] wasn't available.
rest options mask verb context

interchange intervention




Figure 1: We analyze the circuits responsible for performance on Winograd sentences using a series of causal interchange interventions to trace how contextual information is integrated to resolve the masked pronoun.

specified in the literal text (Safavi and Koutra, 2021; Piantadosi and Hill, 2022).

Modern large language models (LLMs) exhibit increasingly impressive performance on “commonsense” tasks that seemingly require the use of implicit world knowledge (Sap et al., 2019; Zellers et al., 2019; Petroni et al., 2019; Davison et al., 2019; Vulić et al., 2020), yet it is still not clear precisely how that knowledge is accessed and employed. Recent interpretability work has explicitly probed and traced individual pieces of world knowledge to highly localized regions of the network (e.g. the birth years of US Presidents; Meng et al., 2022; Dai et al., 2022), allowing surgical erasure or editing (e.g. Colon-Hernandez et al., 2021). But the kind of world knowledge represented by a situation model is more implicit and, according to classical theories, constructed on the fly for the task at hand. Models must somehow determine that a particular relational concept is relevant in the first place (Icard and Goodman, 2015).

For example, in Figure 1 the final piece of context (*successful / available*) is only able to resolve

the pronoun in light of the initial verb phrase “tried to call”. The situation model constructed from an agent Paul *trying* something raises the possible predicates of success or failure, while a patient George *being called* raises the possible predicates of busyness or availability. Conversely, Paul’s availability and George’s success are not at issue and therefore not available as interpretations for the pronoun. There are, of course, many other attributes that may be invoked from a *tried to call* situation model that are not relevant to resolving the pronoun, involving the kinds of devices used to make calls, the kinds of noises heard when being called, and so on. According to this analysis, successfully deriving the appropriate referent in both sentences requires models to *integrate* relevant information across disparate parts of the sentence structure.

In this paper, we conduct an initial exploration of the hypothesis that LLMs have learned to construct implicit situation models. As a case study, we conduct a fine-grained analysis of a best-in-class transformer model (ALBERT; Lan et al., 2020) on a Winograd-like pronoun disambiguation task (Winograd, 1972; Levesque et al., 2011; Sakaguchi et al., 2021, see Figure 1). Winograd sentences are minimal pairs constructed with the property that resolving an ambiguous pronoun requires situational world knowledge outside the scope of the literal text; critically, for our purposes, the pair of sentences differs only at one site, which causes the interpretation to flip. Controlling for possible confounds, a model that is able to generate sharply different predictions for these pairs can be said to behave as if it has a situation model.² Although the open models that amenable to causal probing do not yet achieve fully human-like performance on this task, their trajectory of increasing functional capabilities raises an important mechanistic question about *how* larger networks achieve these gains.

Our primary contributions are (1) employing causal probes to identify a sub-circuit of attention heads that are responsible for propagating contextual information, and (2) constructing a set of closely matched controls for Winograd sentences that are resolvable solely using syntactic cues, which we use to validate the specificity of the identified situation model circuit. Overall, we find some exciting preliminary evidence for meaning-

²It is important to distinguish this cautious claim of task-specific functional equivalence from the stronger claim that a neural model constructs the same kind of situation models as humans in general.

fully non-overlapping pathways while also highlighting the subtleties of probing situation models on a sentence-by-sentence basis.

2 Related Work

2.1 Implicit world knowledge in LLMs

A number of recent studies have examined the extent to which neural language models have acquired implicit schemas about the world (Li et al., 2021; Dai et al., 2022), proposed auxiliary tasks to improve coherence (Li et al., 2022), and probed the internal mechanics by which world knowledge influences downstream predictions (Meng et al., 2022; Geva et al., 2022). A smaller set of studies has focused on classical psycholinguistic phenomena: for example, Davis and van Schijndel (2020) examining relative clause attachment in coreference resolution, and Upadhye et al. (2020) examining predictions about which entities are preferred for different verbs. We approach the problem of implicit situation model representations with a more targeted set of causal intervention techniques, tracing the internal flow of subtle contextual cues in Winograd schemas for the first time.

2.2 Probing with causal interventions

In order to identify interpretable algorithms underlying specific model behaviors, recent studies have employed variants of causal intervention analyses on intermediate representations, targeting syntactic agreement (Lakretz et al., 2021; Finlayson et al., 2021; Lasri et al., 2022), relative clause processing (Ravfogel et al., 2021), natural language inference and compositionality (Geiger et al., 2020, 2021; Soulos et al., 2020), gender bias (Vig et al., 2020), sub-word representations (Huang et al., 2022), and factual knowledge (Dai et al., 2022; Meng et al., 2022). Following Olah et al. (2020), we use the term *circuit* to capture the explanatory construct in such interpretability studies: a computational subgraph of a neural network consisting of a set of units and connections between them that are causally implicated in a behavior (Elhage et al., 2021; Olsson et al., 2022). For example, Wang et al. (2023) recently argued that a small circuit of attention heads appears to identify indirect objects, and Wu et al. (2023) revealed a circuit for solving simple numerical reasoning problems. Building on this family of *interchange intervention* techniques (Geiger et al., 2020; Mueller et al., 2022; Elazar et al., 2021a), we decompose each head into its

query, key and value sub-components (Mohebbi et al., 2023) to trace the flow of causally important information for solving Winograd sentences.

3 Approach

3.1 Dataset construction

We began by extracting the subset of the Winograd Schema Challenge (WSC) appearing in SuperGLUE (Wang et al., 2019; Levesque et al., 2011), as well as the larger, crowd-sourced Winogrande (Sakaguchi et al., 2021) dataset. These datasets contain sentence pairs that differ only at a minimal word or phrase that changes the referent of an ambiguous pronoun earlier in the sentence (see Figure 1).³

- (1) Paul tried to call George but <MASK> wasn't [successful / available].

These sentences are structured such that the disambiguating context only appears near the end, but otherwise have diverse sentence structure (e.g. the context can be any part-of-speech). We call these pairs the "context cue" condition.

We then modified these pairs to construct three additional conditions for comparison. In a "context+syntax cue" condition, we changed the plurality of one of the noun phrase options such that the pronoun can be resolved by relying on the number signaled by the verb without necessarily requiring situational knowledge.

- (2) They tried to call George but <MASK> [weren't successful] / [wasn't available].

To remove the availability of world knowledge entirely, we masked out the context span in both sentences to form a "syntax only" condition.

- (3) They tried to call George but <MASK> [weren't / wasn't] <MASK>.

Finally, we generated a control condition using semantically equivalent *synonyms*; if results in the other conditions truly reflect world knowledge rather than spurious token-specific features, we

³To facilitate comparison across sentences of different lengths, we will refer to the span of text that is manipulated across the two sentences as the "context", the candidate referents as the "options", the pronoun as the "mask" (since it is masked for the prediction task), the verb immediately following the mask as the "verb", and the remaining tokens as the "rest". For all analyses, we conduct single-token interventions and then aggregate effects within that class.

should not expect to find any effect in this condition.

- (4) Paul tried to call George but he [wasn't accessible] / [wasn't available].

In total, we constructed 200 unique pairs of sentences, each of which appears within all 4 conditions (see more examples in Table S1; further details of dataset construction are provided in Appendix A.)

3.2 Interchange interventions

To interrogate how exactly masked language models achieve context-sensitive predictions from minimal cues, we applied a causal intervention technique (Geiger et al., 2020; Mueller et al., 2022) to map the flow of information from the context word to the masked site where the ultimate prediction is made. Specifically, we used an *interchange intervention* to swap intermediate representations across the two contexts (see Figure 1), and quantified the effect of the intervention on the model's downstream prediction. Given a pair of sentences (s_A, s_B) that differ only at the context token, we mask out the pronoun and score the likelihood of each noun phrase (N_A, N_B) at the masked position. For simplicity, we denote the noun phrases such that N_A is the correct referent for sentence s_A and N_B is the correct referent for sentence s_B . For example, $P_\theta(N_A|s_A)$ refers to the likelihood assigned to the (correct) referent N_A at the mask position in sentence s_A , where θ represents the model parameters. In the case of multi-token noun phrases, we masked out all tokens in the phrase and used the average log probability of each token as the score (see Appendix C).

Although multiple metrics have been proposed for capturing the effect of an intervention (e.g. Mueller et al., 2022), we use the canonical *odds ratio*, the shift in relative preference for the correct option as a result of the intervention. We first calculate the baseline preference for the correct referent N_A relative to the incorrect referent N_B :

$$y_{pre} = \frac{P_\theta(N_A|s_A)}{P_\theta(N_B|s_A)}$$

where $P_\theta()$ represents the model's prediction under pre-intervention representations θ . We then measure the same preference after the intervention:

$$y_{post} = \frac{P_{\theta+z}(N_A|s_A)}{P_{\theta+z}(N_B|s_A)}$$

Model	context only		syntax only		context+syntax	
	strict [%]	weak [%]	strict [%]	weak [%]	strict [%]	weak [%]
bert-base-uncased	11.0	59.5	81.5	99.5	83.5	100
bert-large-cased	15.5	65.0	81.5	100	80.5	99.5
roberta-base	8.5	59.0	69.5	100	72.5	100
roberta-large	12.5	70.0	64.0	99.0	66.0	99.5
albert-base-v2	12.5	56.0	60.0	95.5	84.5	99.0
albert-large-v2	14.5	57.0	78.0	98.0	89.0	99.5
albert-xlarge-v2	18.5	59.5	90.0	100	91.0	99.5
albert-xxlarge-v2	31.5	81.5	88.5	100	88.5	100
gpt-4	51.0	–	–	–	–	–
gpt-4 + CoT	67.3	–	–	–	–	–
human	94.1	–	–	–	–	–

Table 1: Zero-shot performance on Winograd items. The strict metric requires the correct referent to be strictly preferred on both sentences in the pair, so chance is 0% for a model that is not sensitive to context. The weak metric only requires context to shift the prediction in the correct direction across the pair (even if there is an absolute preference for the incorrect referent), so chance is 50% for a model that is not sensitive to context. We could not compute weak scores for gpt-4 because raw probabilities are not exposed by the API.

where $P_{\theta+z}()$ represents the model’s prediction using post-intervention representations $\theta+z$. The odds ratio is then: $E = y_{pre}/y_{post}$ (or, on a logarithmic scale, $\log E = \log y_{pre} - \log y_{post}$). Effects are averaged across the two directions of intervention within each sentence pair. Note that $\log y_{base} > 0$ by definition for pairs where the baseline prediction is correct, but depending on whether the intervention decreases or increases the probability of the correct referent, the causal effect can be positive (indicating the site of the intervention was contributing to the correct prediction) or negative (indicating the site of the intervention was contributing to the incorrect prediction).

The odds ratio has some desirable properties compared to other measures like the absolute difference in differences (e.g. Yin and Neubig, 2022). It does not suffer from ceiling or floor effects, and it is a well-understood measure in classical statistics, as deployed in logistic regression. However, it is also insensitive to the absolute values of the probabilities going into the ratio, making it potentially vulnerable to noise in the tails (i.e. when both NPs are very infrequent). We believe the odds ratio is a preferred metric *a priori* but we hope future work will better elucidate the advantages and disadvantages of different metrics.

4 Results

4.1 Zero-shot performance evaluation

There are a multiplicity of different ways of evaluating performance on Winograd sentences. These shifting criteria may be responsible for inflated claims of state-of-the-art performance (Liu et al., 2020; Kocijan et al., 2022; Trichelair et al., 2019). To set the stage for our causal intervention analyses, we conduct our own stricter zero-shot comparison of recent pre-trained models. Our strict metric requires the correct referent to be assigned higher probability for *both* sentences, that is, the likelihood ratios must satisfy

$$\frac{P_{\theta}(N_A|s_A)}{P_{\theta}(N_B|s_A)} > 1 \quad \text{and} \quad \frac{P_{\theta}(N_B|s_B)}{P_{\theta}(N_A|s_B)} > 1$$

By jointly considering both sentences in the pair, this metric better captures context-sensitivity. Note that a context-insensitive model that makes the same prediction for both sentences would receive a score of zero on this metric (Abdou et al., 2020; Elazar et al., 2021b). That is, context-sensitivity is required for the interpretation of a pronoun to be systematically reversed, as required to meet the stricter accuracy criteria. We also consider a more standard but weaker metric that only requires the prediction to shift in the correct direction, even if there is an absolute bias for the incorrect option; chance is 50% for this metric:

$$\frac{P_{\theta}(N_A|s_A)}{P_{\theta}(N_B|s_A)} > \frac{P_{\theta}(N_A|s_B)}{P_{\theta}(N_B|s_B)}$$

Table 1 reports the performance of three masked language models (BERT, RoBERTa and ALBERT) at different sizes, along with the auto-regressive GPT-4, and a newly elicited dataset of $N = 199$ human participants (see Appendix B for details). While all models we consider fall well short of human performance, larger and more recent models tend to perform better overall, with the large ALBERT model achieving up to 81.5% on the weak criterion. We also observe, unsurprisingly, that models perform better overall for the syntax cue condition than for the context cue condition and even better when both cues are combined. GPT-4 performs better than other models on the strict criterion, reaching 67.3% with the use of chain-of-thought prompting (Wei et al., 2022) (mean accuracy aggregated across all 400 individual sentences is 82%, compared to 87% human accuracy).

Although models like GPT-4 are only available through an API, limiting our ability to explore causal interventions on their internal representations, we may view other models as lying at earlier points on the same scaling trajectory. We report interventions on all open models (Figure S3), but we place particular focus on the largest ALBERT model, which achieved the highest zero-shot performance of any open-source model we considered. First, we are interested in examining the circuits underlying *successful* performance, and the small number of sentences for which they make the correct predictions limit statistical power (e.g. BERT is only correct for 22 sentences under the strong criterion). Second, the ALBERT architecture ties the weights of attention heads across all layers, yielding more interpretable head-wise analyses (i.e. it is meaningful to track the same head k across layers).

4.2 Layer-wise information flow

To yield insight into failure modes of these models, we would like to develop a more mechanistic understanding of *how* local information from the context word is algorithmically propagated through other sites in the sentence to ultimately arrive at the correct prediction. In other words, we are interested in probing the transformer *circuits* that allow such minimal context cues to have such large effects on pronoun resolution. According to classical accounts, situation models must be constructed *dynamically*, as it is not obvious ahead of time which aspects of the situation will be relevant to interpretation.

We begin by considering the effect of coarse-grained layer-wise intervention. For each layer and each token, we replaced its vector representation under one sentence with what it would have been at the same layer and token for the other sentence, and measured the extent to which the output prediction changed (see Figure 2). Results are shown only for sentences where the model made “strictly” correct predictions in both the context condition and the syntax condition; effects are similar for cases where the baseline predictions are “weakly” correct (see Figure S1).

First, as a sanity check, we observe that intervening at the critical context token at early layers dramatically switches the model’s prediction (Figure 2). This effect begins to decay around layer 9. Meanwhile, intervening at the noun phrase options only yielded a significantly non-zero effect on the model’s output starting layer 9, $t(57) = 6.1, p < 0.05$ (correcting for multiple comparisons), while intervening at the masked pronoun and other positions (rest) yielded effects beginning at layers 6 and 7, respectively ($p < 0.05$). Although effects are localized among “rest” tokens in different locations from sentence to sentence, the example in Figure 2A reveals early causal effects in the “tried to call” construction and the comma token, which may serve as a neutral site for aggregating phrase-level information⁴ Finally, as predicted, no significant effects were observed when swapping representations on control pairs with synonymous context words (Figure 3), indicating that the context-sensitivity observed for Winograd sentences is not purely driven by token-specific sensitivities. Taken together, these layer-wise effects are consistent with pronoun-relevant information remaining localized in the context until intermediate layers, when it begins to pool in other locations and eventually contextualizes the options to guide attention from the pronoun.

So far, nothing about these comparisons implicates the construction of an implicit situation model — we may simply be measuring the circuits for pronoun-resolution more generally. To disentangle these possibilities, we consider our “syntax only” control condition, where syntactic agreement in-

⁴It has been previously observed that a large amount of attention at late layers is focused on punctuation and other special tokens (Kovaleva et al., 2019; Clark et al., 2019). Interestingly, though, we found no systematic effects at other annotated elements, such as final periods or [CLS]/[SEP] tokens (Figure S2), suggesting that these elements are not consistently involved in the circuit across sentences.

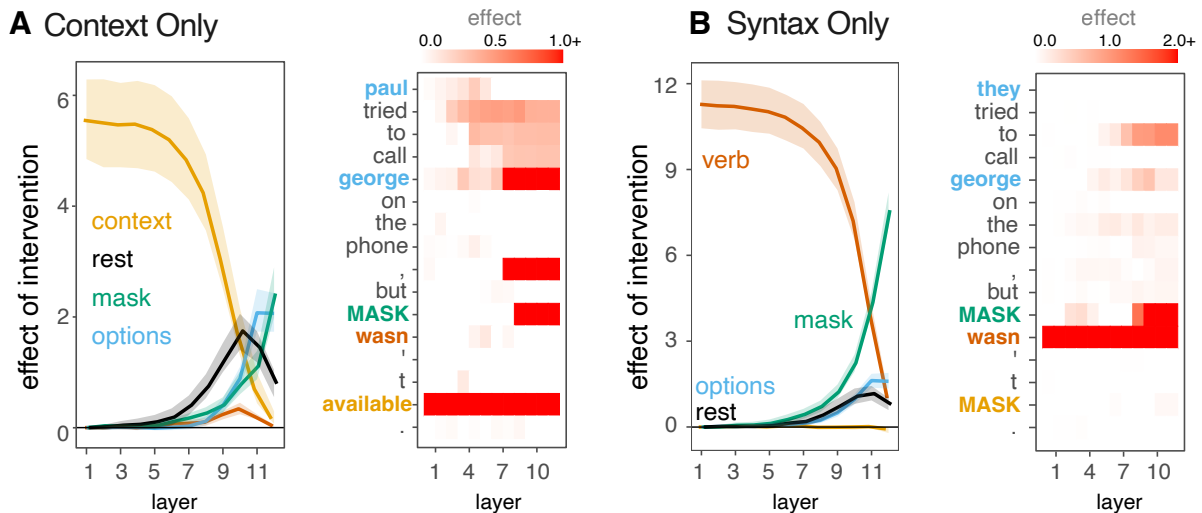


Figure 2: The mean effect of causal interchange interventions at each layer and site are shown for (A) the context only condition and (B) the syntax only condition. Error ribbons show bootstrapped 95% confidence intervals across sentence pairs. Heatmaps show effects of interventions at each individual token for an example sentence.

formation alone is sufficient to make the correct prediction and no contextual information is available. Unlike the “context only” condition, where effects at the rest of the sentence preceded contextualization of the mask and the effect at the options remains stronger than at the mask as late as layer 11 ($t(57) = -5.4, p < 0.05$), the “syntax” condition is strictly dominated by the effect at the mask starting around layer 7 ($t(57) = 3.8, p < 0.05$). This effect is illustrated in Figure 2B, where we see a much more localized circuit between the auxiliary verb, where a number agreement cue is provided, and the MASK, where the prediction is generated; effects at other sites are much more muted. Interestingly, in the combined (context+syntax) condition where both cues are available, information at the

mask still dominates starting at the same layer as the “syntax” condition ($p < 0.05$; not shown). We found qualitatively similar results for the largest RoBERTa model, but not for the smaller models (see Figure S3).

4.3 Head-wise causal interventions

These coarse layer-wise analysis suggest that global “situational” information may be constructed from the context word, and integrated elsewhere in the sentence, whereas purely syntactic agreement information may be accessed more directly by the masked pronoun. However, it remains unclear exactly how the transformer accomplishes this task at an algorithmic level, and a lot of complexity was potentially hidden in the “rest” category. In this section, we conduct a finer-grained head-wise analysis of the individual components within each layer that are responsible for routing contextual information between tokens (Figure 4).

We consider four internal components: (1) the final *transformation* vector produced by the head that gets concatenated back to other heads before being projected back into the residual stream, (2) the *query* vector that “imports” information into a given source site, (3) the *key* vector at other sites which matches the query to yield the attention matrix, and (4) the *value* vector that is “exported” from the target site (see Figure S4 for a schematic).

For ease of interpretability, we will refer to effects of intervening on each of these components in terms of an implied “source → target” pathway.

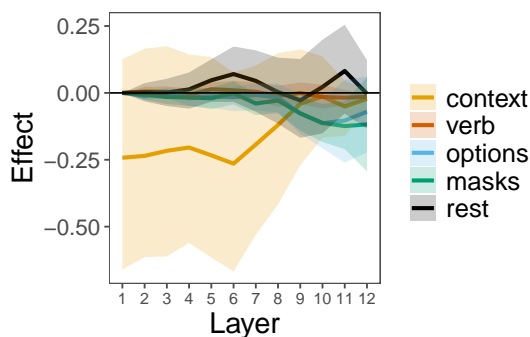
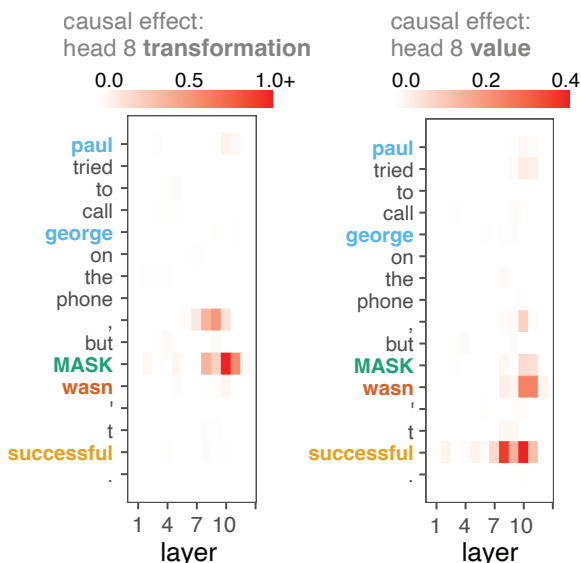


Figure 3: Layerwise intervention for *synonyms* yields no significant effects, indicating that heads are not purely driven by lexical identity. Error ribbons show bootstrapped 95% confidence intervals across sentence pairs.

A Context Only



B Syntax Only




Figure 4: Effects of intervening at a single attention head (head 8) for an example sentence in (A) the context condition and (B) the syntax condition. We intervene at both the point where this head’s output is added back to the given token’s residual stream (*transformations*) as well as the internal *value* vector this head exports to other tokens when queried. In the syntax condition, this head appears to be primarily responsible for moving information from the verb to the MASK; in the context condition, however, it is implicated at other sites, like the comma token where phrase-level information may be integrated.

For example, if intervening on the *value* exported from the context token has the same effect as intervening on the ultimate *transformation* vector at the options, we loosely say that this is a “context \rightarrow options” head. The effect of swapping the full *transformation* vector is our primary measure of how much pronoun-relevant context-sensitivity is introduced into the residual stream by that head at that layer *overall*, which can then be broken down into its contributing subcomponents.

To build intuition, we first depict the complete profile of causal intervention effects for a single head (head 8) across layers and tokens of an example sentence (Figure 4A; see Figure S12 for the same analysis applied to the largest RoBERTa model). Consistent with our coarser layer-wise analysis above, we find that intervening on transformations at the comma token around layer 8 (and a few layers later at the MASK token) significantly interferes with the model’s ability to correctly resolve the referent. We then gained further insight through targeted interventions on the subcomponents used to calculate the transformation vector. Specifically, we find that intervening on the *value* vector exposed at the context word (*successful*) accounts for a significant proportion of the

total effect at middle layers. This effect appears to be largely restricted to the context-only condition; causal interventions on the same head in the syntax-only condition (Figure 4B) reveal only the more local “verb \rightarrow mask” path.

Moving to a more systematic analysis, we find that a set of 23 heads (less than half of the 64 total heads) show a significant effect of intervening on the *transformation* for at least one layer and one site ($p < 0.05$, after correcting for multiple comparisons; 15 survive at the $p < 0.01$ level; 13 at the $p < 0.005$ level; and 11 at the $p < 0.001$ level). Examining the internal components of each of these attention heads (keys, queries, values) allows us to construct a preliminary computational graph of how contextual cues eventually propagate to the masked site, aggregating over many sentences. Figure 5 shows the layer-wise profile for five representative heads, each representing a systematic link between a particular set of sites (see Figure S7 for all heads). Roughly, this graph suggests that the model begins by shifting information from the context site into “neutral” sites throughout the rest of the sentence via head 8 (e.g. value vector at context word, layer 10 : $t(57) = 6.1, p < 0.05$) and head 11 (value vector

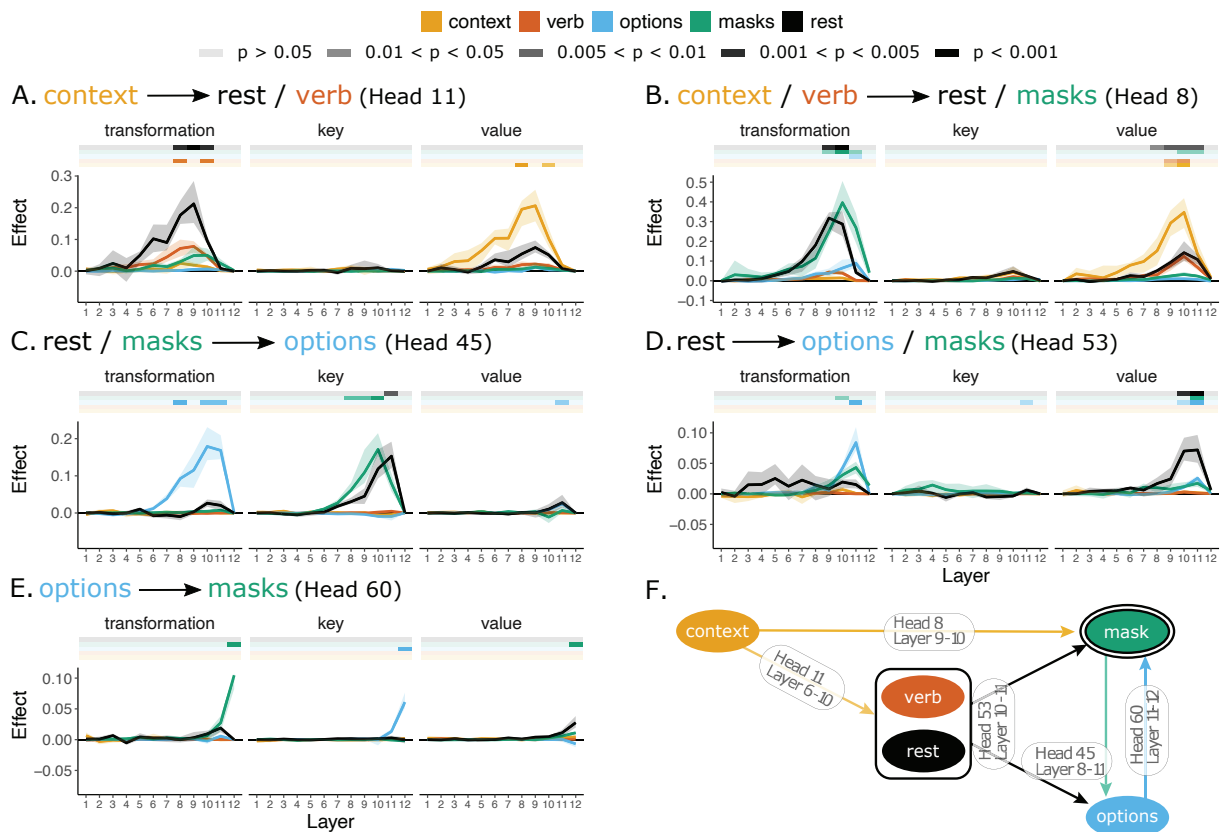


Figure 5: Example heads involved in propagation of context information. Error ribbons show bootstrapped 95% confidence intervals across sentence pairs. Schematic depictions are provided in Figure S6.

at context word, layer 6: $t(57) = 4.9, p < 0.05$). Shortly thereafter, head 45 exports that information to the two noun phrases (options) that are the possible referents of the pronoun (transformation vector at options, layer 8 : $t(57) = 6.9, p < 0.05$) and head 53 (transformation vector at options, layer 11 , $t(57) = 7.1, p < 0.05$). Then, at the final layer of the model, head 60 preferentially attends from the mask to one of the two options, determining the ultimate prediction (key vector at layer 11 options, $t(57) = 8.0, p < 0.05$).

4.4 Analysis of context-specificity

Critically, although some of these heads are also implicated in the syntax-only condition (see Figure S8 for full profiles), there are dramatic quantitative differences in the pattern of these heads across layers. The specificity profile is shown for all active heads in Figure 6, where specificity is defined with respect to the context vs. syntax comparison: an orange cell indicates ‘context-specificity’ (i.e. selective activation only in the context condition) while a blue cell indicates ‘syntax-specificity’ (i.e.

selective activation only in the syntax condition). Broadly, we observe that many heads, including those highlighted above, yield significant context-specific effects at *earlier* layers, especially at “rest” sites (left panel; see Figure S13 for an analysis of RoBERTa). Conversely, there are a number of heads that are yield syntax-specific effects at “mask” sites (middle panel). Although this pattern of specificity remains highly exploratory, it suggests that the network may have learned different pathways for pronoun resolution: cues like verb number information are directly exposed to guide the mask token, while contextual information must be integrated with other relevant semantic cues from elsewhere in the sentence. In other words, the latter may require constructing and querying a rudimentary situation model.

5 Discussion

In this paper, we presented a preliminary investigation of the transformer circuits underlying performance on Winograd sentences, where minimal contextual cues must be used to resolve an ambigu-




Figure 6: Specificity of head-wise effects across layers and sites. Orange cells only yielded causal effects for the context condition, blue cells only for the syntax condition, and black cells for both conditions. All colored cells are significantly different from zero at Bonferroni-corrected $p < 0.005$ level. Heads are ordered by the earliest layer at which context-specificity appears.

ous pronoun. We applied fine-grained causal interventions to identify a circuit of attention heads that are responsible for propagating information from the context cue to the possible referents, which appears to be at least partially distinct from the circuit used to propagate agreement cues in our closely matched syntax-only baseline.

It still remains to be seen whether the circuits we have identified should be interpreted as evidence for a *bona fide* situation model. First, it is possible that even on carefully debiased sets of examples, models like ALBERT are still relying on lexical shortcuts. For example, presented with “MASK was tasty” we would prefer the referent *pie* over the referent *boat*, knowing nothing about the rest of the sentence. If so, we may be measuring the circuit for those lexical preferences rather than for

anything like a situation model (but see Appendix D for an additional analysis suggesting that this kind of phenomenon is unlikely to be driving the observed effects).

Second, situation models present a well-known example of the frame problem (McCarthy and Hayes, 1969; Pylyshyn, 1987) — it is impossible to explicitly enumerate every property in the world and how it is (or is not) related to every other. Hence, any model of interpretation must be “lazy” in some sense, only introducing relations that are *relevant* for the task at hand (e.g. the concept of availability for the “Paul tried to call George” example). When the task at hand is simply pronoun resolution, it is possible that an extremely minimal situation model may be sufficient; in the longer term, it will be important to explore settings that require richer or less accessible interpretations. Finally, although we focused on implicit situation models constructed internally, it may also be possible to expose situation models more explicitly in larger auto-regressive models through chain-of-thought prompting (Talmor et al., 2020) which do not require model internals to be public.

Limitations

Although our single-site interchange interventions provide causal evidence that particular sub-circuits are necessary for a particular downstream behavior, this technique has known limitations addressed by recent Distributed Alignment Search (DAS) approaches (Geiger et al., 2023). First, it will over-count certain “synergies:” when a single effect is jointly produced by the conjunction of multiple heads acting in concert, we will identify all heads as making distinct contributions to the circuit. Second, it will under-count “redundancies:” if there are multiple heads that are individual sufficient to produce the effect, then no single head will be detected as strictly necessary. Ideally, rather than single-site interventions, we would explore all combinations of different heads to find minimal spanning sets that are both necessary and sufficient, but this procedure becomes intractable given the number of heads, requiring more sophisticated optimization-based approaches to find promising sets (e.g. Csordás et al., 2021; De Cao et al., 2022).

Ethics Statement

All existing datasets (SuperGLUE and Winogrande) and models (BERT, RoBERTa, and AL-

BERT) were employed according to their intended research focus, and our targeted probing dataset is intended to be used for similar purposes in future work. Because it was constructed by a combination of automatic and manual processes by the authors, it contains no additional information that could uniquely identify any individuals. More broadly, we employ causal interventions to evaluate how models perform a challenging commonsense reasoning task, with the aim of building stronger links to classic work in cognitive science and psycholinguistics. However, these causal interventions may pose some risk if used adversarially to tamper with public models or expose private information. Further, the WSC and Winogrande datasets we use to probe situation models have been constructed within specific cultural settings (e.g. by NLP researchers and largely US-based crowd-workers) and are not intended to be universal or representative of situational competency: a wider diversity of culturally-specific stimuli is needed.

Acknowledgements

We thank Yohei Oseki, Ryo Yoshida, Felix Hill, Devon Wood-Thomas, Aalok Sathe, Claire Bergey, and the Language Computational Cognitive Science Laboratory at UTokyo for thoughtful discussions and comments. This work was supported by the Princeton Data Driven Social Science initiative, a Seed Grant from the Princeton-UTokyo Strategic Partnership and a C.V. Starr Fellowship to RDH.

References

- Mostafa Abdou, Vinit Ravishankar, Maria Barrett, Yonatan Belinkov, Desmond Elliott, and Anders Søgaard. 2020. [The sensitivity of language models and humans to Winograd schema perturbations](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7590–7604, Online. Association for Computational Linguistics.
- Gordon H Bower, John B Black, and Terrence J Turner. 1979. Scripts in memory for text. *Cognitive Psychology*, 11(2):177–220.
- John D Bransford and Jeffery J Franks. 1971. The abstraction of linguistic ideas. *Cognitive Psychology*, 2(4):331–350.
- Kevin Clark, Urvashi Khandelwal, Omer Levy, and Christopher D. Manning. 2019. [What does BERT look at? An analysis of BERT’s attention](#). In *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 276–286, Florence, Italy. Association for Computational Linguistics.
- Pedro Colon-Hernandez, Catherine Havasi, Jason Alonso, Matthew Huggins, and Cynthia Breazeal. 2021. [Combining pre-trained language models and structured knowledge](#). *arXiv preprint arXiv:2101.12294*.
- Róbert Csordás, Sjoerd van Steenkiste, and Jürgen Schmidhuber. 2021. [Are neural nets modular? Inspecting functional modularity through differentiable weight masks](#). In *International Conference on Learning Representations*.
- Damai Dai, Li Dong, Yaru Hao, Zhifang Sui, Baobao Chang, and Furu Wei. 2022. [Knowledge neurons in pretrained transformers](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8493–8502, Dublin, Ireland. Association for Computational Linguistics.
- Forrest Davis and Marten van Schijndel. 2020. [Discourse structure interacts with reference but not syntax in neural language models](#). In *Proceedings of the 24th Conference on Computational Natural Language Learning*, pages 396–407, Online. Association for Computational Linguistics.
- Joe Davison, Joshua Feldman, and Alexander Rush. 2019. [Commonsense knowledge mining from pre-trained models](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 1173–1178, Hong Kong, China. Association for Computational Linguistics.
- Nicola De Cao, Leon Schmid, Dieuwke Hupkes, and Ivan Titov. 2022. [Sparse interventions in language models with differentiable masking](#). In *Proceedings of the Fifth BlackboxNLP Workshop on Analyzing and Interpreting Neural Networks for NLP*, pages 16–27, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Yanai Elazar, Shauli Ravfogel, Alon Jacovi, and Yoav Goldberg. 2021a. [Amnesic probing: Behavioral explanation with amnesic counterfactuals](#). *Transactions of the Association for Computational Linguistics*, 9:160–175.
- Yanai Elazar, Hongming Zhang, Yoav Goldberg, and Dan Roth. 2021b. [Back to square one: Artifact detection, training and commonsense disentanglement](#)

- in the Winograd schema. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 10486–10500, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Nelson Elhage, Neel Nanda, Catherine Olsson, Tom Henighan, Nicholas Joseph, Ben Mann, Amanda Askell, Yuntao Bai, Anna Chen, Tom Conerly, Nova DasSarma, Dawn Drain, Deep Ganguli, Zac Hatfield-Dodds, Danny Hernandez, Andy Jones, Jackson Kernion, Liane Lovitt, Kamal Ndousse, Dario Amodei, Tom Brown, Jack Clark, Jared Kaplan, Sam McCandlish, and Chris Olah. 2021. [A mathematical framework for transformer circuits](#). *Transformer Circuits Thread*.
- Matthew Finlayson, Aaron Mueller, Sebastian Gehrmann, Stuart Shieber, Tal Linzen, and Yonatan Belinkov. 2021. [Causal analysis of syntactic agreement mechanisms in neural language models](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1828–1843, Online. Association for Computational Linguistics.
- Atticus Geiger, Hanson Lu, Thomas Icard, and Christopher Potts. 2021. [Causal abstractions of neural networks](#). In *Advances in Neural Information Processing Systems*, volume 34, pages 9574–9586. Curran Associates, Inc.
- Atticus Geiger, Kyle Richardson, and Christopher Potts. 2020. [Neural natural language inference models partially embed theories of lexical entailment and negation](#). In *Proceedings of the Third BlackboxNLP Workshop on Analyzing and Interpreting Neural Networks for NLP*, pages 163–173, Online. Association for Computational Linguistics.
- Atticus Geiger, Zhengxuan Wu, Christopher Potts, Thomas Icard, and Noah D Goodman. 2023. Finding alignments between interpretable causal variables and distributed neural representations. *arXiv preprint arXiv:2303.02536*.
- Mor Geva, Avi Caciularu, Kevin Wang, and Yoav Goldberg. 2022. [Transformer feed-forward layers build predictions by promoting concepts in the vocabulary space](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 30–45, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Arthur C Graesser, Murray Singer, and Tom Trabasso. 1994. Constructing inferences during narrative text comprehension. *Psychological Review*, 101(3):371.
- Matthew Honnibal, Ines Montani, Sofie Van Landeghem, and Adriane Boyd. 2020. [spaCy: Industrial-strength Natural Language Processing in Python](#).
- Jing Huang, Zhengxuan Wu, Kyle Mahowald, and Christopher Potts. 2022. Inducing character-level structure in subword-based language models with type-level interchange intervention training. *arXiv preprint arXiv:2212.09897*.
- Thomas Icard and Noah D Goodman. 2015. A resource-rational approach to the causal frame problem. In *Proceedings of the Annual Conference of the Cognitive Science Society*.
- Philip Nicholas Johnson-Laird. 1983. *Mental models: Towards a cognitive science of language, inference, and consciousness*. 6. Harvard University Press.
- Vid Kocijan, Ernest Davis, Thomas Lukasiewicz, Gary Marcus, and Leora Morgenstern. 2022. The defeat of the Winograd Schema Challenge. *arXiv preprint arXiv:2201.02387*.
- Olga Kovaleva, Alexey Romanov, Anna Rogers, and Anna Rumshisky. 2019. [Revealing the dark secrets of BERT](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4365–4374, Hong Kong, China. Association for Computational Linguistics.
- Yair Lakretz, Dieuwke Hupkes, Alessandra Vergallito, Marco Marelli, Marco Baroni, and Stanislas Dehaene. 2021. Mechanisms for handling nested dependencies in neural-network language models and humans. *Cognition*, 213:104699.
- Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2020. [ALBERT: A lite BERT for self-supervised learning of language representations](#). In *International Conference on Learning Representations*.
- Karim Lasri, Tiago Pimentel, Alessandro Lenci, Thierry Poibeau, and Ryan Cotterell. 2022. [Probing for the usage of grammatical number](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8818–8831, Dublin, Ireland. Association for Computational Linguistics.
- Hector J Levesque, Ernest Davis, and Leora Morgenstern. 2011. [The Winograd schema challenge](#). In *AAAI Spring Symposium: Logical Formalizations of Commonsense Reasoning*, volume 46, page 47.
- Belinda Z. Li, Maxwell Nye, and Jacob Andreas. 2021. [Implicit representations of meaning in neural language models](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1813–1827, Online. Association for Computational Linguistics.
- Belinda Z Li, Maxwell Nye, and Jacob Andreas. 2022. Language modeling with latent situations. *arXiv preprint arXiv:2212.10012*.

- Haokun Liu, William Huang, Dhara Mungra, and Samuel R. Bowman. 2020. [Precise task formalization matters in Winograd schema evaluations](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 8275–8280, Online. Association for Computational Linguistics.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. RoBERTa: A robustly optimized BERT pretraining approach. *arXiv preprint arXiv:1907.11692*.
- J McCarthy and P.J. Hayes. 1969. Some philosophical problem from the standpoint of artificial intelligence. In D. Michie and B. Meltzer, editors, *Machine Intelligence*, volume 4, pages 463–502.
- Kevin Meng, David Bau, Alex J Andonian, and Yonatan Belinkov. 2022. [Locating and editing factual associations in GPT](#). In *Advances in Neural Information Processing Systems*.
- Hosein Mohebbi, Willem Zuidema, Grzegorz Chrupała, and Afra Alishahi. 2023. [Quantifying context mixing in transformers](#). In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 3378–3400, Dubrovnik, Croatia. Association for Computational Linguistics.
- Daniel G Morrow and Herbert H Clark. 1988. Interpreting words in spatial descriptions. *Language and Cognitive Processes*, 3(4):275–291.
- Aaron Mueller, Yu Xia, and Tal Linzen. 2022. [Causal analysis of syntactic agreement neurons in multilingual language models](#). In *Proceedings of the 26th Conference on Computational Natural Language Learning (CoNLL)*, pages 95–109, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Chris Olah, Nick Cammarata, Ludwig Schubert, Gabriel Goh, Michael Petrov, and Shan Carter. 2020. Zoom in: An introduction to circuits. *Distill*, 5(3):e00024–001.
- Catherine Olsson, Nelson Elhage, Neel Nanda, Nicholas Joseph, Nova DasSarma, Tom Henighan, Ben Mann, Amanda Askell, Yuntao Bai, Anna Chen, et al. 2022. In-context learning and induction heads. *arXiv preprint arXiv:2209.11895*.
- Fabio Petroni, Tim Rocktäschel, Sebastian Riedel, Patrick Lewis, Anton Bakhtin, Yuxiang Wu, and Alexander Miller. 2019. [Language models as knowledge bases?](#) In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2463–2473, Hong Kong, China. Association for Computational Linguistics.
- Steven Piantadosi and Felix Hill. 2022. [Meaning without reference in large language models](#). In *NeurIPS 2022 Workshop on Neuro Causal and Symbolic AI (nCSI)*.
- Zenon W Pylyshyn. 1987. *The robot’s dilemma: The frame problem in artificial intelligence*. Ablex.
- Shauli Ravfogel, Grusha Prasad, Tal Linzen, and Yoav Goldberg. 2021. [Counterfactual interventions reveal the causal effect of relative clause representations on agreement prediction](#). In *Proceedings of the 25th Conference on Computational Natural Language Learning*, pages 194–209, Online. Association for Computational Linguistics.
- David E Rumelhart. 1975. Notes on a schema for stories. In *Representation and understanding*, pages 211–236. Elsevier.
- Tara Safavi and Danai Koutra. 2021. [Relational world knowledge representation in contextual language models: A review](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 1053–1067, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Keisuke Sakaguchi, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. 2021. Winogrande: An adversarial winograd schema challenge at scale. *Communications of the ACM*, 64(9):99–106.
- Maarten Sap, Hannah Rashkin, Derek Chen, Ronan Le Bras, and Yejin Choi. 2019. [Social IQa: Commonsense reasoning about social interactions](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4463–4473, Hong Kong, China. Association for Computational Linguistics.
- Roger C Schank and Robert P Abelson. 1977. *Scripts, Plans, Goals, and Understanding: An Inquiry Into Human Knowledge Structures*. Psychology Press.
- Paul Soulos, R. Thomas McCoy, Tal Linzen, and Paul Smolensky. 2020. [Discovering the compositional structure of vector representations with role learning networks](#). In *Proceedings of the Third BlackboxNLP Workshop on Analyzing and Interpreting Neural Networks for NLP*, pages 238–254, Online. Association for Computational Linguistics.
- Alon Talmor, Oyvind Tafjord, Peter Clark, Yoav Goldberg, and Jonathan Berant. 2020. Leap-of-thought: Teaching pre-trained models to systematically reason over implicit knowledge. *Advances in Neural Information Processing Systems*, 33:20227–20237.
- Paul Trichelair, Ali Emami, Adam Trischler, Kaheer Suleman, and Jackie Chi Kit Cheung. 2019. [How reasonable are common-sense reasoning tasks: A case-study on the Winograd schema challenge and SWAG](#). In *Proceedings of the 2019 Conference on*

- Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3382–3387, Hong Kong, China. Association for Computational Linguistics.
- Shiva Upadhye, Leon Bergen, and Andrew Kehler. 2020. Predicting reference: What do language models learn about discourse models? In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 977–982, Online. Association for Computational Linguistics.
- Jesse Vig, Sebastian Gehrmann, Yonatan Belinkov, Sharon Qian, Daniel Nevo, Yaron Singer, and Stuart Shieber. 2020. Investigating gender bias in language models using causal mediation analysis. *Advances in Neural Information Processing Systems*, 33:12388–12401.
- Ivan Vulić, Edoardo Maria Ponti, Robert Litschko, Goran Glavaš, and Anna Korhonen. 2020. Probing pretrained language models for lexical semantics. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7222–7240, Online. Association for Computational Linguistics.
- Alex Wang, Yada Pruksachatkun, Nikita Nangia, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2019. SuperGLUE: A Stickier Benchmark for General-Purpose Language Understanding Systems. In *Advances in Neural Information Processing Systems*, volume 32.
- Kevin Ro Wang, Alexandre Variengien, Arthur Conmy, Buck Shlegeris, and Jacob Steinhardt. 2023. Interpretability in the wild: a circuit for indirect object identification in GPT-2 small. In *International Conference on Learning Representations*.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, brian ichter, Fei Xia, Ed H. Chi, Quoc V Le, and Denny Zhou. 2022. Chain of thought prompting elicits reasoning in large language models. In *Advances in Neural Information Processing Systems*.
- Terry Winograd. 1972. Understanding natural language. *Cognitive Psychology*, 3(1):1–191.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2020. Transformers: State-of-the-art natural language processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Zhengxuan Wu, Atticus Geiger, Christopher Potts, and Noah D Goodman. 2023. Interpretability at scale: Identifying causal mechanisms in alpaca. *arXiv preprint arXiv:2305.08809*.
- Kayo Yin and Graham Neubig. 2022. Interpreting language models with contrastive explanations. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 184–198, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. 2019. HellaSwag: Can a machine really finish your sentence? In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4791–4800, Florence, Italy. Association for Computational Linguistics.
- Rolf A Zwaan and Gabriel A Radvansky. 1998. Situation models in language comprehension and memory. *Psychological Bulletin*, 123(2):162.

Appendix A: Dataset Construction

We began with 70 sentences from the WSC subset of SuperGLUE and 9248 sentences from the debiased portion of the Winogrande train split (see Appendix D for discussion of debiasing procedures). Many of the sentences from Winogrande were singletons, as the other sentence in the pair had been removed in the debiasing procedure, leaving 1361 intact pairs. We then excluded sentences where we could not automatically obtain the index of the context word or options in the tokenized sequence (many of them had subtle differences in the sentence pair aside from the context word, with respect to the choice of words/punctuation, which are likely not intended by the original annotators of Winogrande), leaving 1140 pairs.

In order to create a closely controlled “syntax” cue condition, we took the subset where the verb after the pronoun is either an auxiliary (“to be” or “to have”) or in present tense, indicating whether the masked pronoun is singular or plural. These verbs were identified using spaCy (version 3.2.1) (Honnibal et al., 2020). We also ensured at this stage that all sentences had the same NP1-NP2-mask-context-verb ordering. This step reduced the number of sentence pairs to 38 and 167 for SuperGLUE and Winogrande respectively.

The NPs and the verbs of these extracted sentences were then manually modified to create the agreement cue in the “syntax” condition, while the context words were manually modified to create the synonym condition. We excluded 5 out of 167 sentence pairs at this step, as we were unable to change the plurality of the NPs while preserving

the semantic meaningfulness of the sentence (e.g. “James can count all the numbers on his fingers because the [MASK] are few/many.”). Finally, because the interchange intervention requires both sentences in a given pair to have exactly the same overall number of WordPiece tokens (in order to independently swap query, key or value features), we lightly modified the contexts to satisfy this condition. This step did not reduce the number of sentence pairs.

The resulting dataset contains 200 sentence pairs each for the context, syntax and context+syntax cue conditions (38 from SuperGLUE and 162 from Winogrande) and 400 sentence pairs for the synonym condition (as we generated a synonym for each individual sentence in the pair). This dataset is available on Github (under MIT License)⁵.

Appendix B: Zero-shot evaluation details

We used Hugging Face transformers⁶ (version 4.16.2) (Wolf et al., 2020) for all masked language models. One of the key features of ALBERT (Lan et al., 2020), unlike the original BERT (Devlin et al., 2019) or the subsequent RoBERTa (Liu et al., 2019), is that it shares model parameters across layers, which makes it possible to meaningfully compare the same head in different layers. Relative to the next-largest ALBERT model, the xxLarge model we consider allocates parameters toward increased “width” rather than “depth,” with larger hidden state dimension (4096 vs. 2048) and more attention heads (64 vs. 16) packed into a smaller number of layers (12 vs. 24). Running the entire analysis took 200 hours using 2 GPUs (NVIDIA TITAN X (Pascal)).

For a more state-of-the-art reference point, we also examined the performance of the most recent OpenAI GPT-4 model (using the provided snapshot from March 14, 2023). Because GPT-4 is auto-regressive and the disambiguating context in Winograd sentences typically follows the pronoun, we create a prompted variant of the text as suggested by the original proposers of the Winograd challenge (Levesque et al., 2011). That is, we presented the full sentence along with a question: “The trombone did not fit in the suitcase because it was too large. What was too large, the trombone or the suitcase?” and coded the model’s free-text re-

sponses for matching either of the NP options (or “other” if it did not match either option). We used a temperature of 1 and generated 50 responses for each question to account for sampling variation. For our chain-of-thought experiment, we used the following prompt containing a single example:

You are going to be shown a sentence and asked to fill in a blank (with an explanation). For example, consider the sentence, “The man tried to put the tuba in a suitcase but <blank> was too small.” We will ask whether the <blank> should be interpreted as the TUBA or the SUITCASE. The <blank> will be literally ambiguous, but you can use knowledge about the world to figure it out. In this case, we can reason as follows. The suitcase is a container, the man is attempting to fit a tuba in that container, and bigger things don’t fit when containers are too small. In addition, tubas are generally bigger than suitcases. Therefore the answer is SUITCASE. What about the following sentence?

Finally, because our specific set of Winograd items had not previously been benchmarked against human performance, we elicited free-response judgments from a $N = 199$ human participants. We matched the wording of instructions closely to the prompt we used with GPT-4, and we coded responses in the same way. Each sentence received approximately 20 judgments, we ensured no participant saw more than one sentence from any given Winograd pair, and the order of presentation of the options was counterbalanced (i.e. half of participants saw “the trombone or the suitcase?” while the other half saw “the suitcase or the trombone?”).

Appendix C: Handling multi-token NPs

In order to evaluate the model’s ability to perform pronoun resolution, we replaced the pronoun with the mask token and compared the log probability of each NP option at the masked site. While this was straightforward in most cases, 28 of the 200 examples contained a multi-token option. We handled these cases by introducing multiple masks at the pronoun. We then calculated an average token-wise score probability in the following way. Suppose the option noun phrase NP consists of 3 tokens (w_k , w_{k+1} and w_{k+2}) starting at the k th site (s_k , s_{k+1} and s_{k+2}), where the pronoun is at the l th site (s_l). We then replaced the pronoun with three mask tokens and calculated the average log probability as

⁵<https://github.com/taka-yamakoshi/situation-models>

⁶Apache License 2.0

$$\frac{1}{3}[\log P(s_l = w_k \mid s_{l+1}, s_{l+2} = \text{mask}, s_{-l} = w_{-l}) + \log P(s_{l+1} = w_{k+1} \mid s_l, s_{l+2} = \text{mask}, s_{-l} = w_{-l}) + \log P(s_{l+2} = w_{k+2} \mid s_l, s_{l+1} = \text{mask}, s_{-l} = w_{-l})],$$

where $s_{-l} = w_{-l}$ denotes that all sites other than s_l are filled with the corresponding tokens in the original sentence. Taking the average log probability effectively controls for the additional number of tokens. One concern with this method is that, if part of the noun phrase is a functional word (e.g. a determiner like *the*), it may drag down the observed effect of the entire phrase. While this is not ideal, we confirmed our results are not sensitive to other ways of aggregating across tokens and believe that aggregating across multiple tokens in this way is a conservative choice relative to other methods, like taking the maximum effect across each of the component tokens.

Appendix D: Checking dataset biases

A recurring concern in the literature on Winograd tasks is that many examples can be fully resolved using lexical “shortcuts” that do not require implicit world knowledge. We took several steps to ensure that our results are not explained by these biases. First, we used the debiased train split of Winogrande, which used an approach known as AFLITE to filter out examples that could be solved using only the isolated embedding of the context word (Sakaguchi et al., 2021). Second, due to differences in our mask-prediction task and the fine-tuning regime used by Sakaguchi et al., we conducted our own analysis of bias in our smaller dataset.

We devised a simple alternative approach to test whether the isolated context embeddings (token representations at layer 0) had lexical preferences for the options. We first calculated correlations (and Euclidean distances) between the embedding vectors at the context and each of the options. If the context is more similar to N_1 (i.e. larger correlation or smaller distance), we say N_1 is predicted as the referent, while if the context is more similar to N_2 , we say it predicts N_2 . This simple method made correct predictions in 15.5 % (using correlation) or 16.5% (using Euclidean distance) of our 200 sentence pairs. As expected, this performance is significantly worse than the output MASK prediction. More importantly, out of the 63 sentence pairs we consider where correct predictions were made by the contextualized model, only 9 (10) were

resolvable using word embeddings, based on correlation (distance). This analysis suggests that while some lexical shortcuts may still exist in our data, these biases are likely not driving our results.

	sentence pair
A	I poured water from the bottle into the cup until the <mask> was empty . I poured water from the bottle into the cup until the <mask> was full .
B	I poured water from the bottle into the cups until the <mask> was empty . I poured water from the bottle into the cups until the <mask> were empty .
C	I poured water from the bottle into the cups until the <mask> was <mask> . I poured water from the bottle into the cups until the <mask> were <mask> .
D_1	I poured water from the bottle into the cup until the <mask> was empty . I poured water from the bottle into the cup until the <mask> was emptied .
D_2	I poured water from the bottle into the cup until the <mask> was filled . I poured water from the bottle into the cup until the <mask> was full .
A	Joe paid the detective because <mask> has received the final report on the case. Joe paid the detective because <mask> has delivered the final report on the case.
B	Joe paid the detectives because <mask> has received the final report on the case. Joe paid the detectives because <mask> have delivered the final report on the case.
C	Joe paid the detectives because <mask> has <mask> the final report on the case. Joe paid the detectives because <mask> have <mask> the final report on the case.
D_1	Joe paid the detective because <mask> has received the final report on the case. Joe paid the detective because <mask> has obtained the final report on the case.
D_2	Joe paid the detective because <mask> has provided the final report on the case. Joe paid the detective because <mask> has delivered the final report on the case.
A	As Andrea in the crop duster passed over Susan , <mask> was able to spot the landing strip . As Andrea in the crop duster passed over Susan , <mask> was able to spot the landing gear .
B	As Andrea in the crop duster passed over her parents , <mask> was able to spot the landing strip . As Andrea in the crop duster passed over her parents , <mask> were able to spot the landing gear .
C	As Andrea in the crop duster passed over her parents , <mask> was able to spot the landing <mask>. As Andrea in the crop duster passed over her parents , <mask> were able to spot the landing <mask>.
D_1	As Andrea in the crop duster passed over Susan , <mask> was able to spot the landing strip . As Andrea in the crop duster passed over Susan , <mask> was able to spot the landing field .
D_2	As Andrea in the crop duster passed over Susan , <mask> was able to spot the landing wheels . As Andrea in the crop duster passed over Susan , <mask> was able to spot the landing gear .

Table S1: Examples sentence pairs in each condition (A: context / B: context+syntax / C: syntax / D_1 : synonym 1 / D_2 : synonym 2). For synonym 1, we changed the context of sentence 2 to match that of sentence 1 and for synonym 2, we changed the context of sentence 1 to match that of sentence 2. We report the aggregated results of both synonym conditions in Figure 3

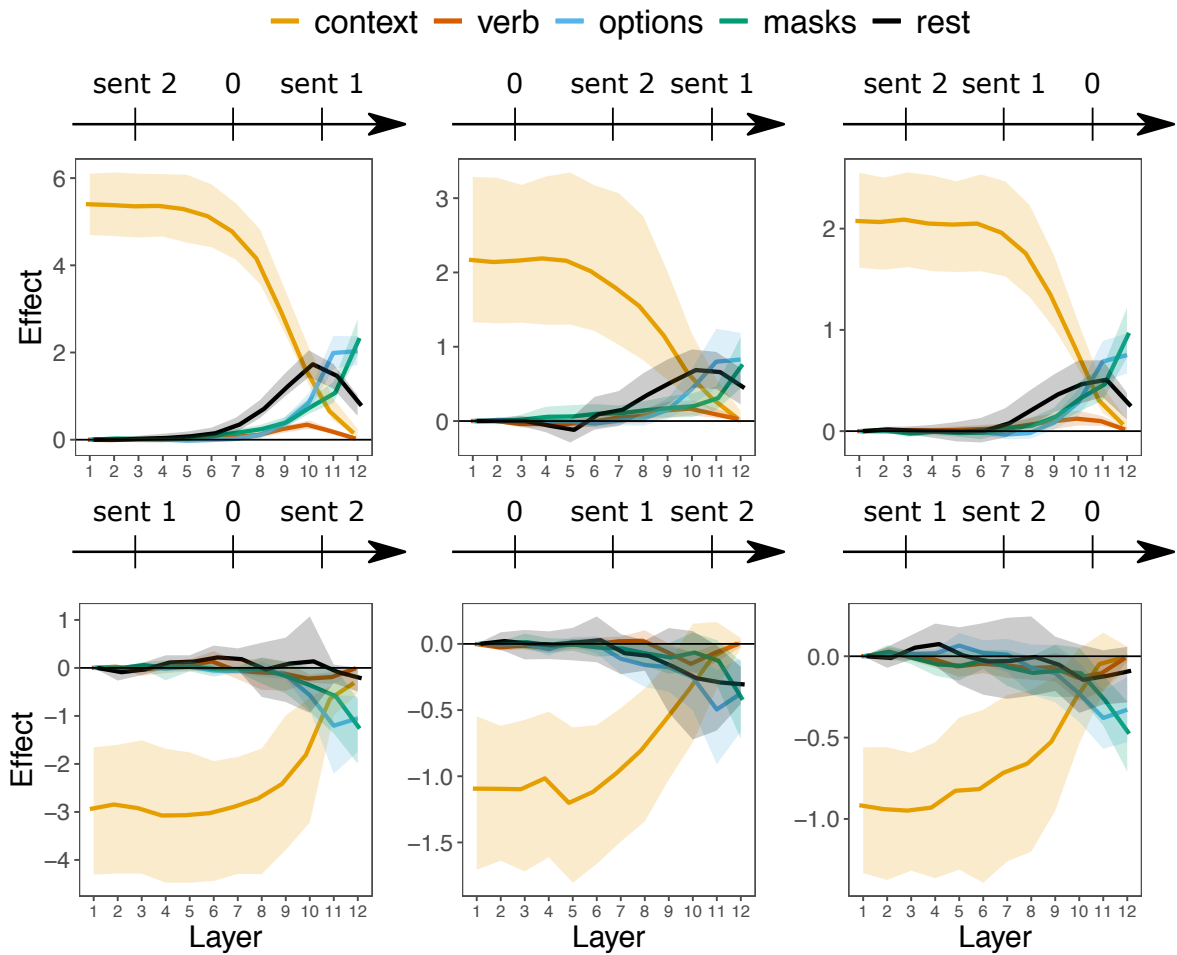


Figure S1: Effects of swapping each layer for cases where the model predicted incorrectly. The line at the top shows the sign of the log likelihood ratio $\log(L_{\theta}(N_A|s)/L_{\theta}(N_B|s))$ in each sentence s . The top left panel shows the cases where the model predicted correctly for both sentences (positive for s_A and negative for s_B). The cases where the model made incorrect prediction can be categorized into five different groups, depending on the signs of the log likelihood ratios. The sign of the effect is flipped for three panels in the bottom row reflecting the fact that the log likelihood ratio was larger for sentence 2. Error ribbons show bootstrapped 95% confidence intervals across sentence pairs.

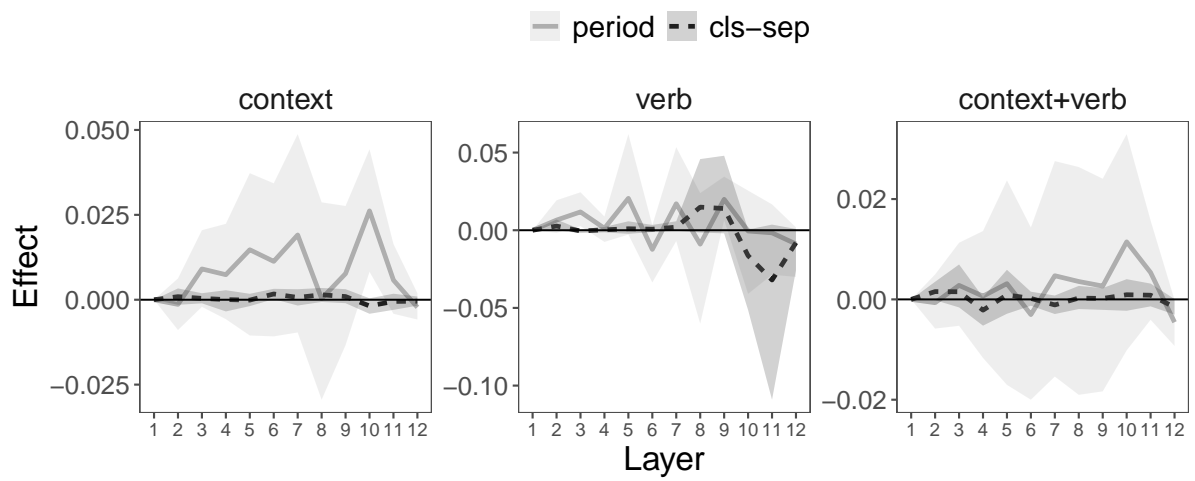


Figure S2: Effects of swapping intermediate representations of periods or special tokens (i.e. [CLS] and [SEP] tokens). Error ribbons show bootstrapped 95% confidence intervals across sentence pairs.

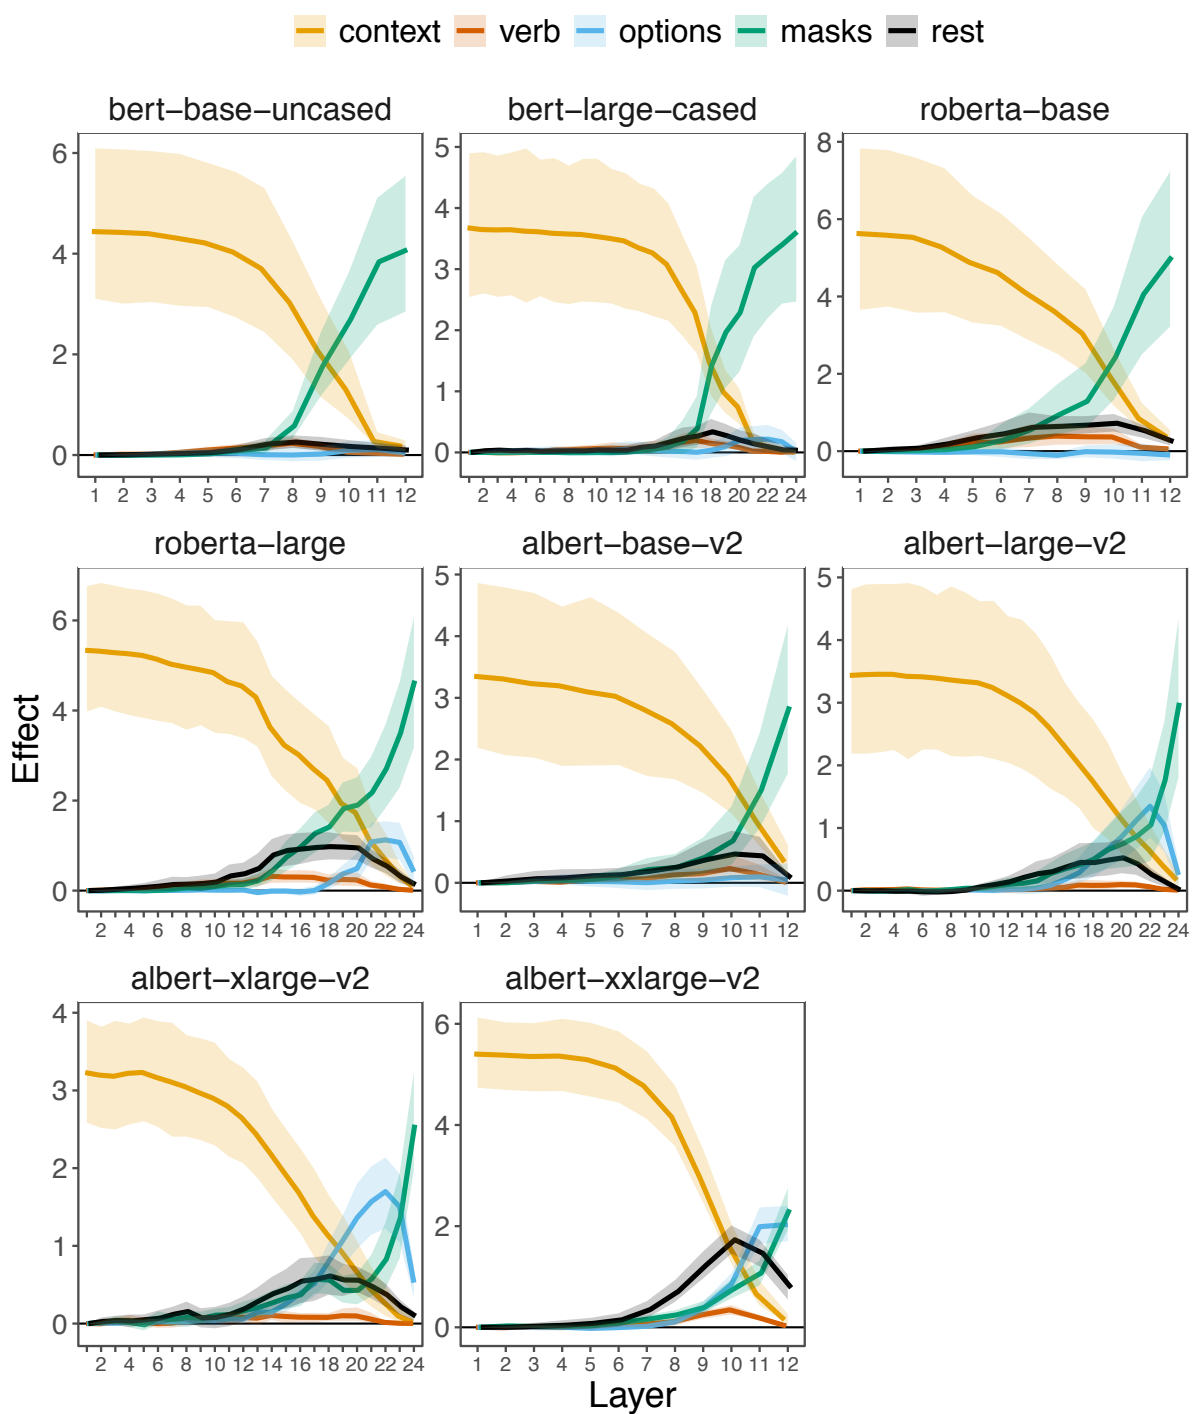


Figure S3: Effect of swapping each layer for models with different types and sizes. The flow of information from the context to the mask is conserved in all models. In addition, larger models (roberta-large, albert-large-v2, albert-xlarge-v2, and albert-xxlarge-v2) have the context information flow that goes into the options and the rest. Error ribbons show bootstrapped 95% confidence intervals across sentence pairs.




Figure S4: A schematic showing how different components interact with each other in the attention mechanism. X, V, K, Q, and T refer to the layer embedding, value, key, query, transformation vectors respectively. This schematic is used in Figure S6




Figure S5: Effects of cumulative intervention on the transformation vectors (the final layer of the attention mechanism prior to being added back to the residual stream, see Figure S4). For each layer i , transformation vectors at layer i and all the preceding layers (at layer 0, 1, ..., $i - 1$) are interchanged. The cumulative effects largely reproduce the layerwise intervention Figure 2, which suggests the emerging context information that flows into the options, mask and rest is mediated by transformation. Error ribbons show bootstrapped 95% confidence intervals across sentence pairs.

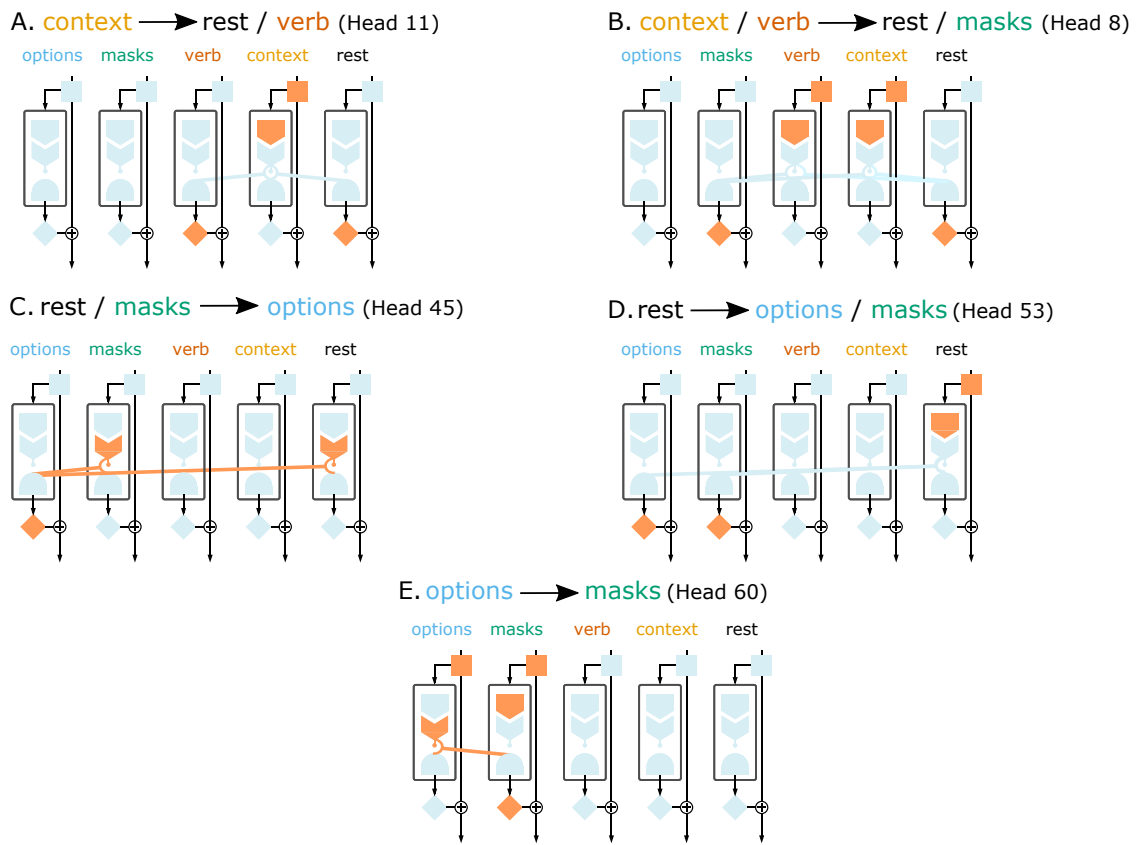


Figure S6: Schematic illustration of Figure 5. Shapes highlighted in orange are found to be implicated in the circuit.

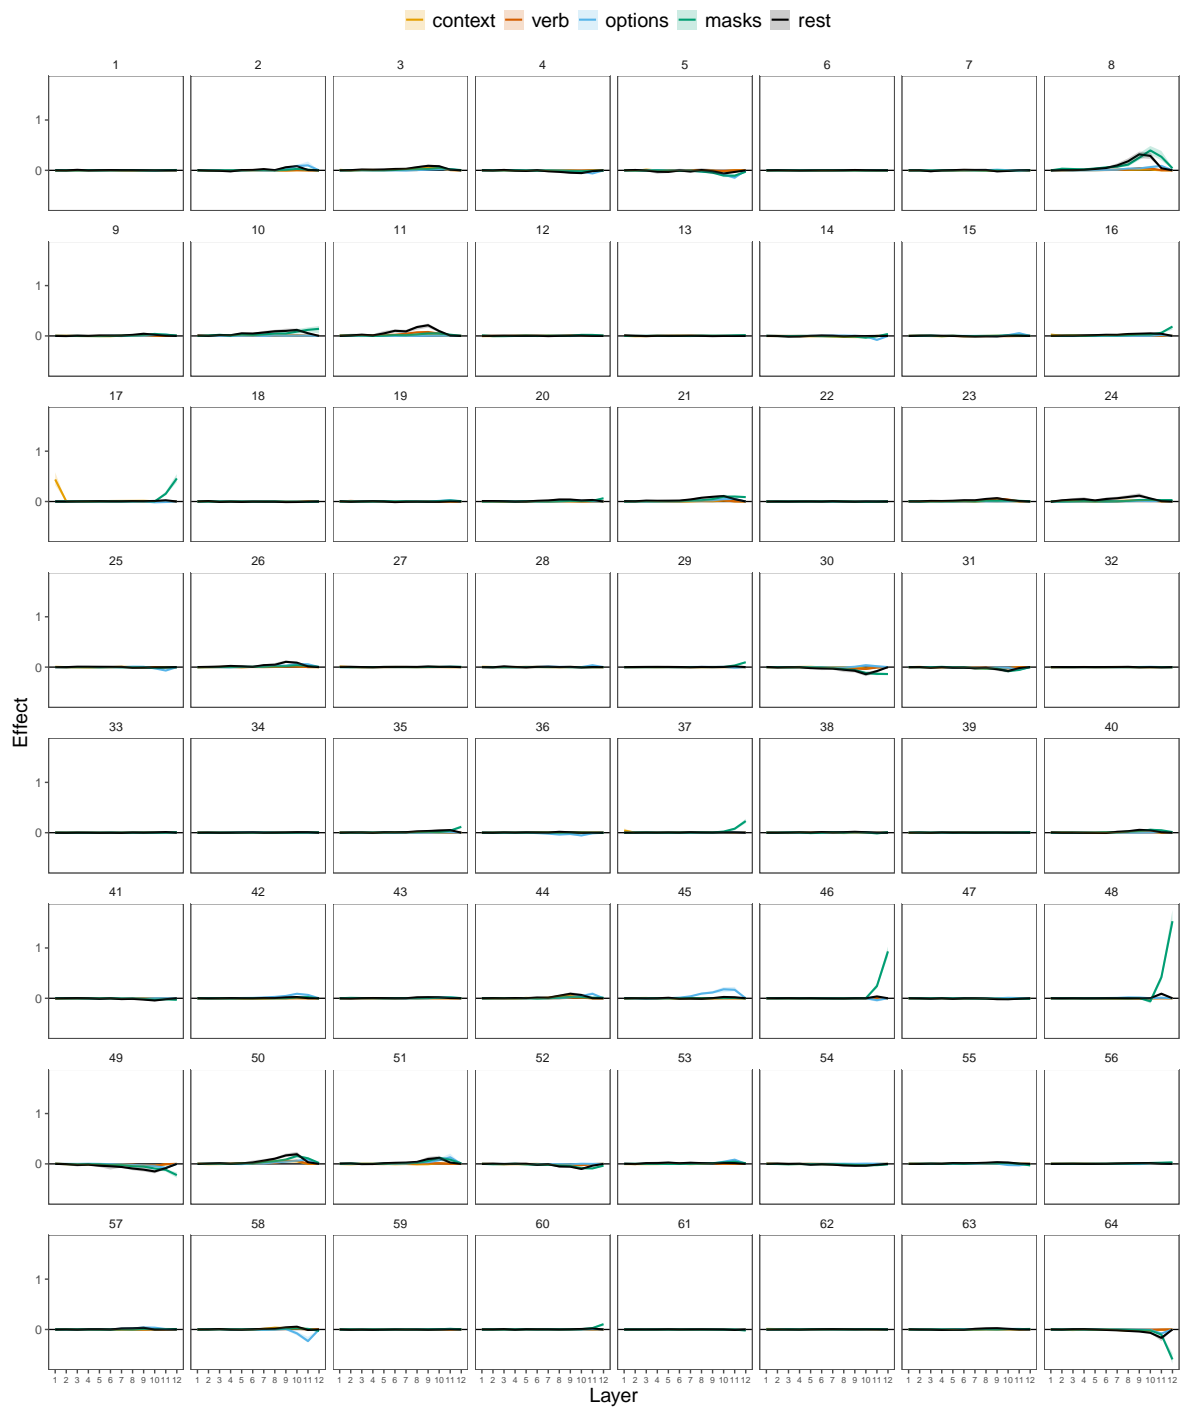


Figure S7: Effects of swapping full transformation representations for each head, in the context only condition. Error ribbons show bootstrapped 95% confidence intervals across sentence pairs.

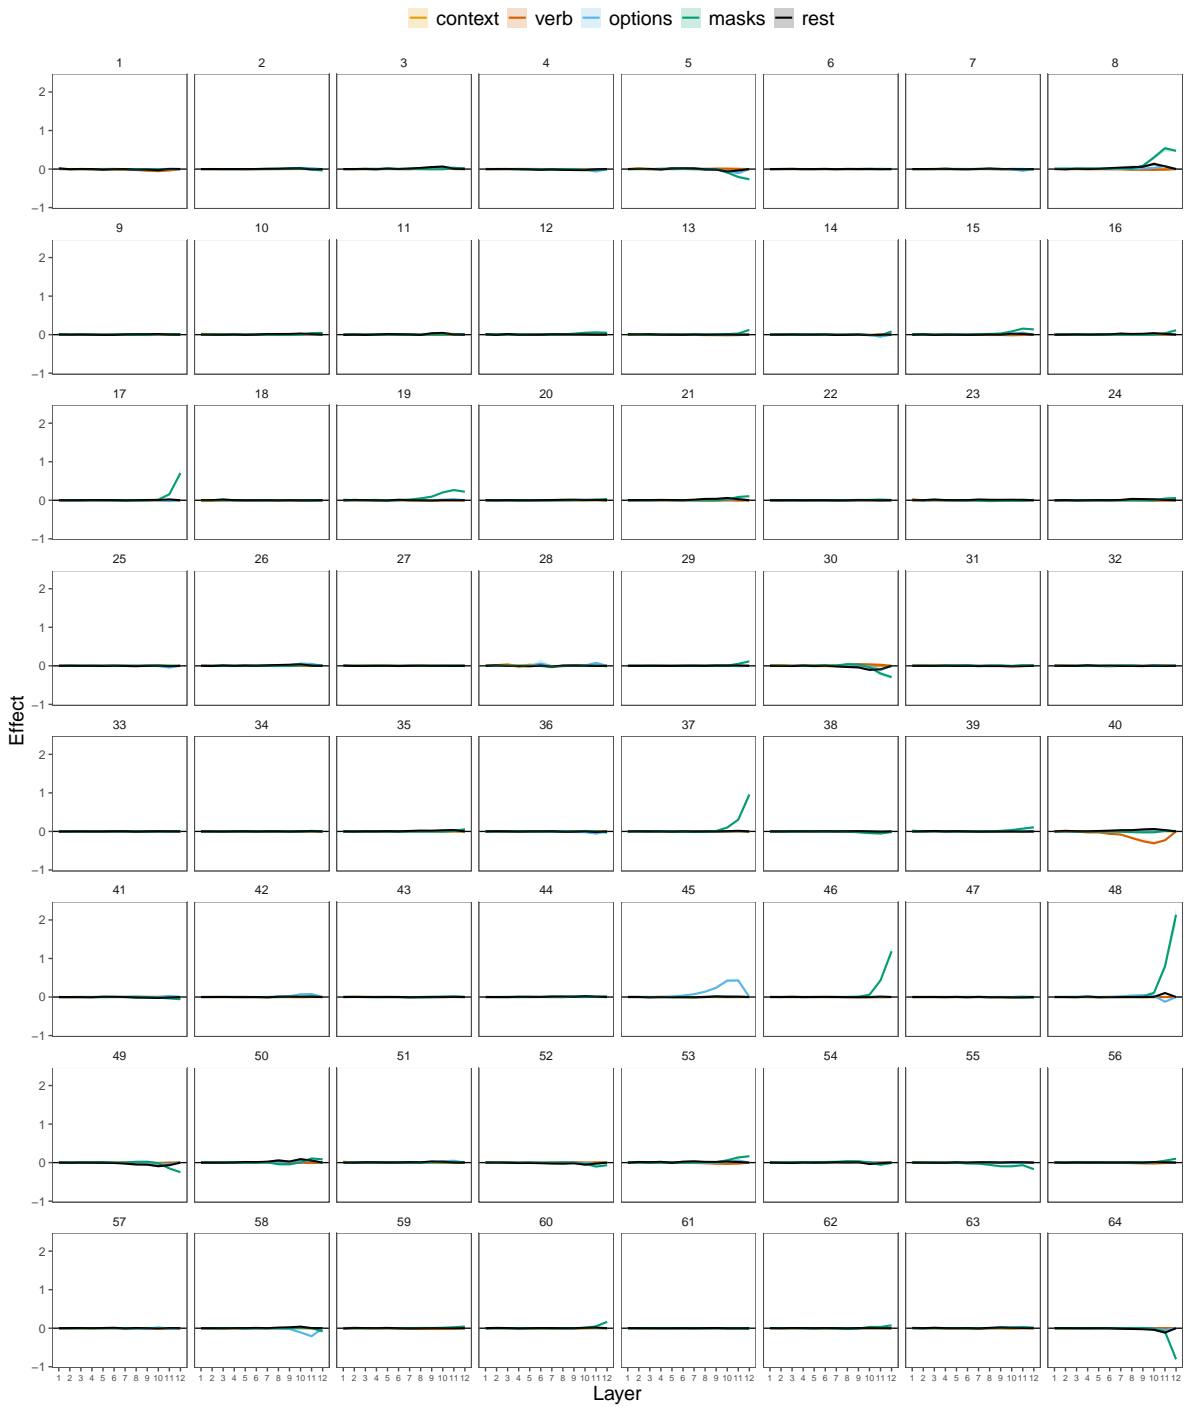


Figure S8: Effects of swapping full transformation representations for each head, in the syntax only condition. Error ribbons show bootstrapped 95% confidence intervals across sentence pairs.

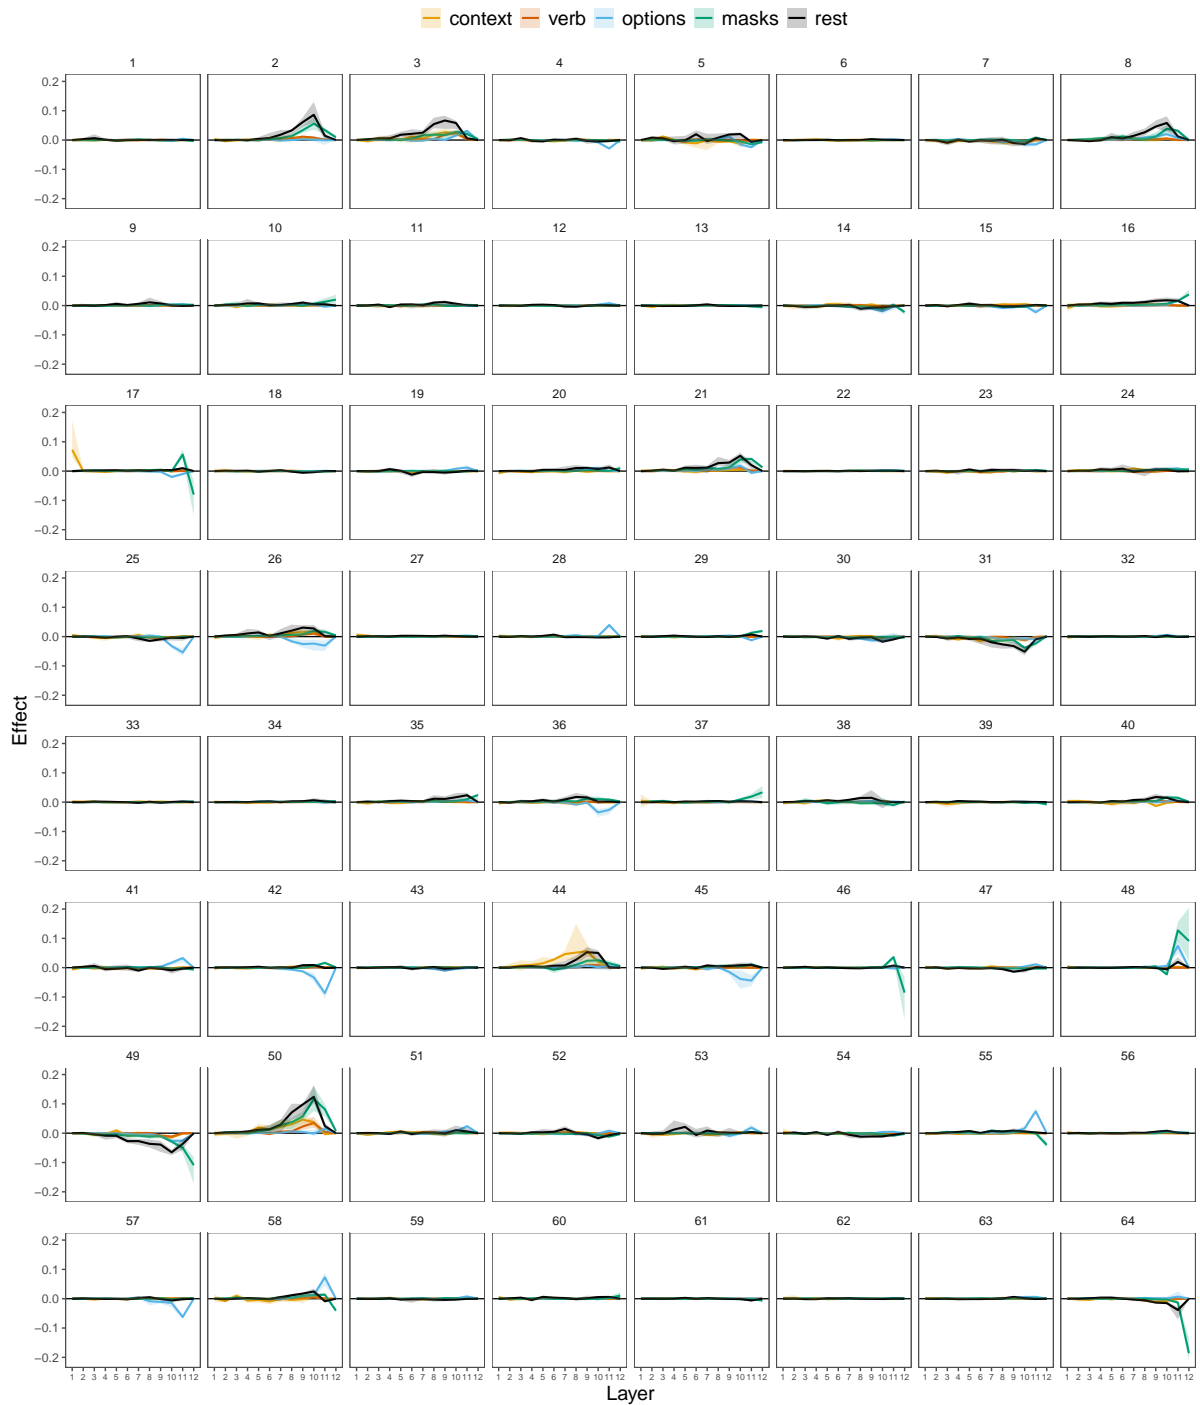


Figure S9: Effects of swapping query representations for the context condition. Error ribbons show bootstrapped 95% confidence intervals across sentence pairs.




Figure S10: Effects of swapping key representations for the context condition. Error ribbons show bootstrapped 95% confidence intervals across sentence pairs.




Figure S11: Effects of swapping value representations for the context condition. Error ribbons show bootstrapped 95% confidence intervals across sentence pairs.

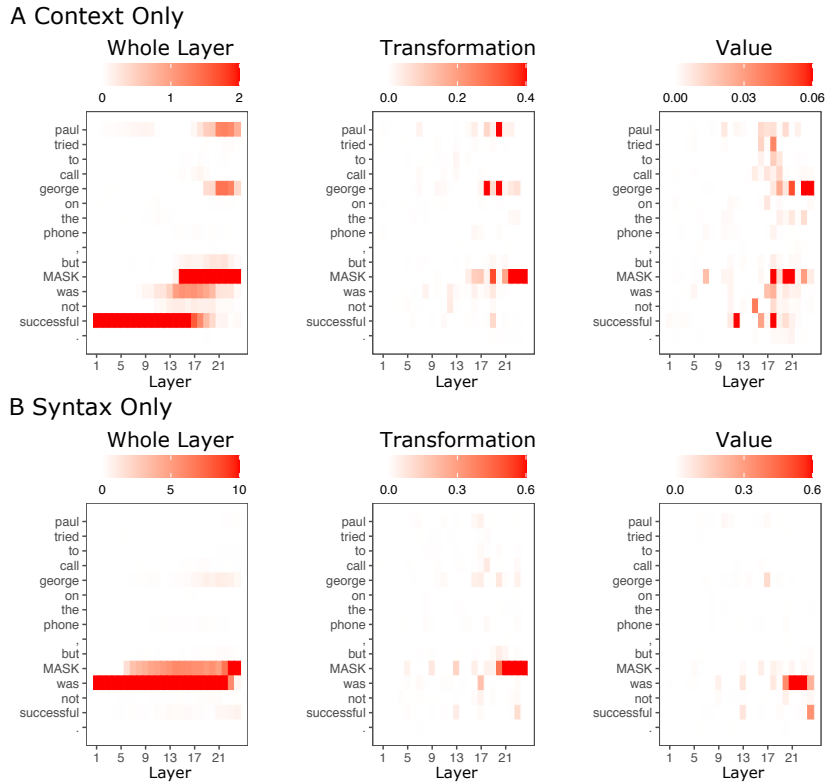


Figure S12: In order to test the robustness of our identified circuit, we also performed the head-wise interchange intervention analysis for roberta-large. It was the best performing model among non-ALBERT models and showed similar layer-wise pattern in the layer-wise analysis. Here, we show the effect of causal interventions at different sites in an example sentence. Within the context (A) and verb (B) conditions, the head at each layer with the maximum t -value is selected.

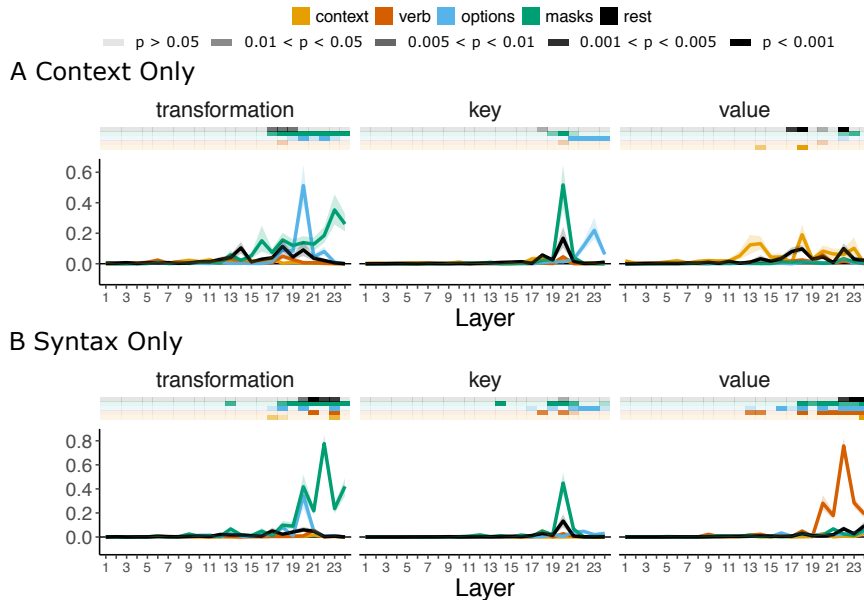


Figure S13: Head-wise causal interventions aggregated across sentences for different sites and layers in RoBERTa in the context (A) and verb (B) conditions. The head with the maximum t -value is selected for each representation type (transformation/key/value), position (context/verb/options/masks/rest), and layer. Lines show the mean effect of the selected head, and error ribbons show bootstrapped 95% confidence intervals. Bars above line plots regions of statistical significance for each position.

ACL 2023 Responsible NLP Checklist

A For every submission:

- A1. Did you describe the limitations of your work?
Limitations section
- A2. Did you discuss any potential risks of your work?
Ethical Statement section
- A3. Do the abstract and introduction summarize the paper's main claims?
Abstract and Section 1 (introduction)
- A4. Have you used AI writing assistants when working on this paper?
Left blank.

B Did you use or create scientific artifacts?

Section 3.1 (dataset construction)

- B1. Did you cite the creators of artifacts you used?
Introduction, Section 3.1
- B2. Did you discuss the license or terms for use and / or distribution of any artifacts?
Ethical Statement section
- B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)?
Ethical Statement section
- B4. Did you discuss the steps taken to check whether the data that was collected / used contains any information that names or uniquely identifies individual people or offensive content, and the steps taken to protect / anonymize it?
Ethical Statement section
- B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?
Ethical Statement section
- B6. Did you report relevant statistics like the number of examples, details of train / test / dev splits, etc. for the data that you used / created? Even for commonly-used benchmark datasets, include the number of examples in train / validation / test splits, as these provide necessary context for a reader to understand experimental results. For example, small differences in accuracy on large test sets may be significant, while on small test sets they may not be.
Appendix A

C Did you run computational experiments?

Left blank.

- C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used?
No response.

The Responsible NLP Checklist used at ACL 2023 is adopted from NAACL 2022, with the addition of a question on AI writing assistance.

- C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?

No response.

- C3. Did you report descriptive statistics about your results (e.g., error bars around results, summary statistics from sets of experiments), and is it transparent whether you are reporting the max, mean, etc. or just a single run?

No response.

- C4. If you used existing packages (e.g., for preprocessing, for normalization, or for evaluation), did you report the implementation, model, and parameter settings used (e.g., NLTK, Spacy, ROUGE, etc.)?

No response.

D Did you use human annotators (e.g., crowdworkers) or research with human participants?

Left blank.

- D1. Did you report the full text of instructions given to participants, including e.g., screenshots, disclaimers of any risks to participants or annotators, etc.?

No response.

- D2. Did you report information about how you recruited (e.g., crowdsourcing platform, students) and paid participants, and discuss if such payment is adequate given the participants' demographic (e.g., country of residence)?

No response.

- D3. Did you discuss whether and how consent was obtained from people whose data you're using/curating? For example, if you collected data via crowdsourcing, did your instructions to crowdworkers explain how the data would be used?

No response.

- D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?

No response.

- D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?

No response.