ADAG: Automatically Describing Attribution Graphs

Aryaman Arora[1,2] **Zhengxuan Wu**[1,2] **Jacob Steinhardt**[2] **Sarah Schwettmann**[2]

[1]Stanford University [2]Transluce
{aryamana,wuzhengx}@stanford.edu

Abstract

In language model interpretability research, **circuit tracing** aims to identify which internal features causally contributed to a particular output and how they affected each other, with the goal of explaining the computations underlying some behaviour. However, all prior circuit tracing work has relied on ad-hoc human interpretation of the role that each feature in the circuit plays, via manual inspection of data artifacts such as the dataset examples the component activates on. We introduce **ADAG**, an end-to-end pipeline for describing these attribution graphs which is fully automated. To achieve this, we introduce *attribution profiles* which quantify the functional role of a feature via its input and output gradient effects. We then introduce a novel clustering algorithm for grouping features, and an LLM explainer–simulator setup which generates and scores natural-language explanations of the functional role of these feature groups. We run our system on known human-analysed circuit-tracing tasks and recover interpretable circuits, and further show ADAG can find steerable clusters which are responsible for a harmful advice jailbreak in Llama 3.1 8B Instruct.

github.com/TransluceAI/circuits

1 Introduction

Language models engage in internal computations which need not be legible to humans. If we wish to oversee AI systems and ensure their safety, we must understand such opaque computation (Ngo et al., 2022; Sharkey et al., 2025; Casper et al., 2024). To this end, interpretability researchers work on **circuit tracing**, with the goal of identifying which internal components contributed to a particular output from a language model. A circuit is a subgraph of the model's entire computation graph which discards computations that were irrelevant to the output of interest. To find circuits in LLMs, many competing techniques have been proposed; they vary on what the appropriate unit of analysis inside an LLM is, what metric to use to measure the importance of a unit for some behaviour, and how to make this kind of analysis computationally tractable; we detail these approaches in §2.

However, even if one succeeds in finding and verifying the circuit underlying the behaviour, much work remains in order to make this artifact legible to humans. What role does each

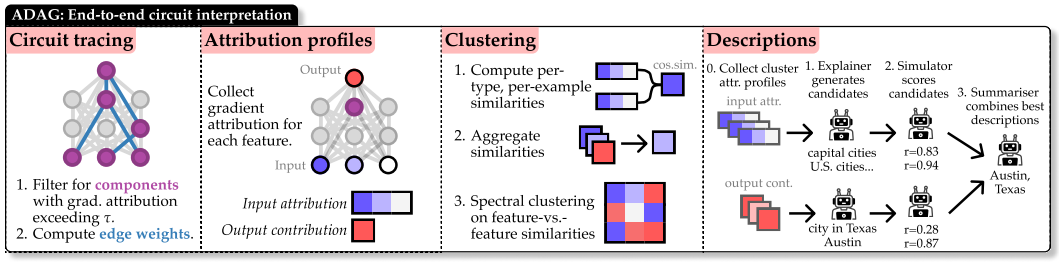

Figure 1: An overview of ADAG, our end-to-end circuit interpretation pipeline.

node in the circuit play? Do the steps of computation correspond to a human-interpretable algorithm? Can I make generalisations about other outputs from the given circuit? These important questions remain uninvestigated; currently, interpretation of circuits is an ad-hoc process that requires extensive researcher-driven analysis of a plethora of data sources, such as analysing dataset examples on which a feature is active or performing targeted causal interventions on component activations (e.g. Ameisen et al., 2025; Lindsey et al., 2025; Shu et al., 2026).

In this work, we seek to **automate circuit interpretation**. To that end, we propose **ADAG**, an end-to-end circuit-tracing and interpretation pipeline which formalises the goals of circuit interpretation and automatically collects relevant data and feeds it to a language model-in-the-loop system for interpretation and verification of the interpretation.

Given a dataset of behaviours, with each example consisting of an input text sequence and resulting output logits we wish to explain, ADAG uses the following pipeline (Figure 1) to produce a human-interpretable analysis of the model's computation:

1. Identify important features and their interaction weights using gradient-based attribution, resulting in an attribution graph (Arora et al., 2026).
2. Construct *attribution profiles* which quantify the functional role of each feature in each sample in the dataset, using input attributions and output logit contributions.
3. Group features into *supernodes* via multi-view spectral clustering over their attribution profiles.
4. Describe the role of each supernode in natural language using an explainer LM which proposes labels and a simulator LM which scores them (Bills et al., 2023).

Our formalisation of the goals of circuit interpretation allows us to develop metrics for each step of interpretation; we use these to verify that our design decisions result in better interpretations than alternatives. We then demonstrate the utility of ADAG by automatically replicating a prior human-led study on multi-hop state capitals task (Ameisen et al., 2025), and then by finding interpretable clusters responsible for a harmful medical advice jailbreak in Llama 3.1 8B Instruct (Chowdhury et al., 2025). Ultimately, our automation of interpretations enables scaling circuit interpretation to more examples and larger models, which is a necessity for practical application of such techniques.

2 Related work

Circuit tracing. Building on early work in vision models (Olah et al., 2020), the notion of circuits as an object of study in LLM interpretability started by identifying interactions between components via causal interventions (Vig et al., 2020; Geiger et al., 2021; Wang et al., 2023; Chan et al., 2022; Meng et al., 2022; Goldowsky-Dill et al., 2023; Conmy et al., 2023; Mueller et al., 2025). Since performing such interventions is expensive, recent work has adopted gradient-based attribution for circuit tracing in order to cheaply approximate this, often in conjunction with learned feature dictionaries (e.g. SAEs, transcoders) out of a belief that the neuron basis is uninterpretable (Nanda, 2023; Syed et al., 2024; Ge et al., 2024; Dunefsky et al., 2024; Hanna et al., 2024; Marks et al., 2025; Ameisen et al., 2025; Lindsey et al., 2025; Shu et al., 2026; Hanna et al., 2025; Jafari et al., 2025). Arora et al. (2026) showed that MLP neuron circuits are of comparable sparsity as SAE circuits.

Automatic interpretation. Autointerpretability is the nascent science of automatically generating natural-language descriptions of model internals. Bills et al. (2023) introduced the explainer–simulator pipeline, wherein one LLM generates explanations and another scores them by simulating the behaviour of the component conditioned on the description; this approach has persisted (Cunningham et al., 2023; Choi et al., 2024; Paulo et al., 2025; Templeton et al., 2024), alongside more token-specific approaches that use maximum-activating exemplars (Foote et al., 2023; Gao et al., 2025).

3 ADAG: An end-to-end circuit tracing pipeline

We now describe ADAG, our end-to-end circuit tracing pipeline, which includes (1) circuit tracing, (2) attribution profile computation, (3) automatic clustering of features into supernodes, and (4) automatic natural-language description of supernodes. In our experiments, we specifically use the MLP neuron circuit tracing algorithm identified to be optimal in Arora et al. (2026), which we further describe in §3.1. However, ADAG is agnostic to the exact circuit tracing algorithm being used.

3.1 Circuit tracing backbone

Before interpretation, we need attribution graphs, which we produce by using a circuit tracing algorithm. For our experiments, we use the MLP neuron circuit tracing algorithm found to be optimal in Arora et al. (2026). First, we construct a locally linear *replacement model*, which modifies the backwards pass of the language model while keeping the forward pass intact. We apply RelP (Jafari et al., 2025), a layerwise relevance propagation technique which freezes nonlinearities (e.g. SiLU, softmax) to behave as constant multipliers based on their effect on a single input, stops gradients through query and key paths in self-attention, and halves the gradient via elementwise multiplications (see Appendix A).

Next, we compute the *attribution score* α for each MLP neuron by backpropagating from the sum of the top-k output logits (following Ameisen et al., 2025). Let $m_u^{(l,t)}(x)$ denote the u-th neuron activation of the layer-l MLP at token t on input x, and let $\texttt{target} = \sum_{k=0}^{K} \mathcal{M}(x)_k^{(s)}$ (i.e. sum of top-k logits):

$$\alpha(m_u^{(l,t)}) = m_u^{(l,t)}(x) \cdot \frac{\partial \texttt{target}}{\partial m_u^{(l,t)}} \tag{1}$$

Now, in our circuit, we only keep the MLP neurons whose attributions exceed some threshold τ, which we define as some fraction of `target`. We similarly compute edge weights (see Appendix A). This results in a circuit $\mathcal{G} = (V, E)$ with nodes being input tokens, MLP neurons, and output logits, with weighted edges connecting them. Let V denote the set of features in the circuit, of all three types. We refer to each non-input and non-output nodes (so in our case, all MLP neurons) as $f \in F$ where $F \subseteq V$.

3.2 Quantifying functional roles of features with attribution profiles

While the grouping of features into supernodes in prior work is largely manual (Ameisen et al., 2025; Shu et al., 2026), it still relies on two pieces of quantitative information: the max-activating exemplars on which the feature fired (Bolukbasi et al., 2021), and the top positive and negative output logits of the feature as computed with logit lens (nostalgebraist, 2020). Both techniques suffer from what we term **locality bias**: they assume that the immediate representation of a feature at the token it fires on is sufficient to describe it; influence on far future or past tokens may matter. We investigate non-locality in MLP neurons in §4.1.

We propose **attribution profiles** for better quantifying the functional role of features. We overcome locality bias by using gradient-based attribution to understand dependence on preceding tokens and effects on future outputs. The attribution profile of a feature given a single prompt x consists of two vectors: its *input attribution* and its *output contribution*:

$$\mathsf{Attr}(m_u^{(l,t)}, x) = \left(x^{(i)} \cdot \frac{\partial m_u^{(l,t)}(x)}{\partial x^{(i)}} \right)_{i=1}^{n} \tag{2}$$

$$\mathsf{Contrib}(m_u^{(l,t)}, x) = \left(m_u^{(l,t)} \cdot \frac{\partial \mathcal{M}(x)_j^{(s)}}{\partial m_u^{(l,t)}} \right)_{j=1}^{m} \tag{3}$$

That is to say, input attribution is the proportion of the feature's activation that is attributed to each input token, and output contribution is the proportion of each output logit's activation at a given token position that the feature contributes to. These profiles can be computed for all nodes in V. We ignore input attribution to the BOS token in all experiments.

3.3 Clustering functionally similar features into supernodes

Features in attribution graphs are manually grouped together into **supernodes** when they seem to play a similar role over a particular dataset of examples (Ameisen et al., 2025; Shu et al., 2026). We seek to automate supernode clustering. Our main task is thus to score the functional similarity of two neurons and then cluster. We describe an algorithm that does exactly this, using attribution profiles.

Since prior approaches to clustering features are fully manual, it is unclear *a priori* what properties a good supernode has. Based on the downstream uses we envision for automatically clustered attribution graphs, we identify the following properties and create metrics for them:

1. *Clusters should group functionally similar features*: The features in a cluster should play similar roles over the distribution of interest. We quantify this using the silhouette score based on attribution profile cosine similarity.
2. *Clusters should not be imbalanced*: We do not want degenerate clusterings where a single large cluster dominates the analysis; we quantify this using coefficient of variation (CV) of cluster sizes.
3. *Clusters should not mix features with opposing output effects*: We often observe features with similar input attributions having opposing contributions to the output. These features should not be clustered together to avoid diluting the cluster's steering effect. We measure this via the fraction of intra-cluster feature pairs wherein **all** contribution entries are of opposing signs.

We now describe our supernode clustering algorithm. We first define the similarity metric for attribution and contribution as cosine similarity: $k(\mathbf{x}_i, \mathbf{x}_j) = \frac{\mathbf{x}_i \cdot \mathbf{x}_j}{\|\mathbf{x}_i\|\|\mathbf{x}_j\|}$ Let C be the set of contexts we observe, where c is one such context with associated prompt x_c. We compute the matrix of pairwise similarities in each context c for both attribution and contribution separately, resulting in two similarity matrices $\mathbf{Attr}^{(c)}$ and $\mathbf{Contrib}^{(c)}$ per context. In each such context, we take the harmonic mean of clamped non-negative attribution and contribution similarities; this penalises anticorrelation in either profile, helping us satisfy the property that supernodes should not mix features with opposing output effects.

Let f_i and $f_j \in F$ be (non-input, non-output) features in our circuit, and let $C_{ij} \subseteq C$ be the subset of contexts in which they co-occur; they may have been pruned by our circuit tracing algorithm in some contexts, rendering them unobserved. We then uniformly average these scores over contexts, resulting in the aggregated similarity matrix $\mathbf{S}$ whose entries are:

$$\mathbf{Attr}_{ij}^{(c)} = k(\mathsf{Attr}(f_i, x_c), \mathsf{Attr}(f_j, x_c)) \tag{4}$$

$$\mathbf{Contrib}_{ij}^{(c)} = k(\mathsf{Contrib}(f_i, x_c), \mathsf{Contrib}(f_j, x_c)) \tag{5}$$

$$\mathbf{S}_{ij} = \begin{cases} \dfrac{1}{|C_{ij}|} \displaystyle\sum_{c \in C_{ij}} \dfrac{2\,\mathrm{ReLU}(\mathbf{Attr}_{ij}^{(c)})\,\mathrm{ReLU}(\mathbf{Contrib}_{ij}^{(c)})}{\mathrm{ReLU}(\mathbf{Attr}_{ij}^{(c)}) + \mathrm{ReLU}(\mathbf{Contrib}_{ij}^{(c)})} & \text{if } |C_{ij}| > 0, \\ 0 & \text{otherwise.} \end{cases} \tag{6}$$

This is a simple form of multiple kernel learning (Gönen & Alpaydin, 2011). On this final similarity matrix, we apply spectral clustering. Since $\mathbf{S}$ is non-negative, we directly use it as the affinity matrix $\mathbf{A}$. We construct the normalised graph Laplacian $\mathbf{L}_{\text{norm}} = \mathbf{I} - \mathbf{D}^{-1/2}\mathbf{A}\mathbf{D}^{-1/2}$, where $\mathbf{D}$ is the diagonal degree matrix with $\mathbf{D}_{ii} = \sum_j \mathbf{A}_{ij}$. We manually select the number of clusters k. The eigenvectors corresponding to the k smallest eigenvalues of $\mathbf{L}$ are then used as a low-dimensional embedding of the features, on which k-means is applied to obtain the final cluster assignments.

At the end of this process, we have a partition $\mathcal{P} = \{P_1, \ldots, P_k\}$ of the circuit's features F into supernodes, where each $P_i \subseteq F$.

3.4 Describing supernodes in natural language

Finally, we want to describe the functional role of each supernode in natural language. We look to what information a human interpretability researcher draws on to come up with such a label in prior work, which is primarily maximum-activating exemplars and top/bottom logits (Ameisen et al., 2025; Shu et al., 2026). Since attribution profiles are more informative alternatives to those (§4.1), we instead use them.

To describe attribution profiles, we extend the explainer–simulator framework of Bills et al. (2023). First, given a set of supernodes, we separately average the attribution and contribution profiles over all features in each supernode. For a given supernode $P_i \in \mathcal{P}$ and a context $c \in C$, we thus have two representations:

$$\overline{\mathsf{Attr}}(P_i, x_c) = \frac{1}{|P_i|} \sum_{f \in P_i} \mathsf{Attr}(f, x_c), \qquad \overline{\mathsf{Contrib}}(P_i, x_c) \quad = \frac{1}{|P_i|} \sum_{f \in P_i} \mathsf{Contrib}(f, x_c) \tag{7}$$

Input attributions. Input attribution $\overline{\mathsf{Attr}}(P_i, x_c)$ has an identical shape to max-activating exemplars: both assign scalar scores to individual tokens in some context. This data is thus identical in format to the information used to describe max-activating exemplars in Choi et al. (2024). We can thus use the same pipeline to generate and score descriptions of the attribution profile of a supernode.

For describing attribution profiles for models using the Llama 3 tokeniser, we use the finetuned models from Choi et al. (2024). First, `Transluce/llama_8b_explainer` takes in the set of contexts x_c wherein tokens which exceed some attribution threshold are highlighted in `{{}}` brackets and produces a candidate description. Second, `Transluce/llama_8b_simulator` takes in the candidate description and raw contexts x_c, and produces a predicted scalar score for each token in each context conditioned on the candidate. All input attribution prompts are in §H.1.

Output contributions. Our circuit tracing algorithm tells us the contribution of a feature to each of the top-k logits at a given output position. Therefore, unlike input attribution or max-activating exemplars, we have information about *hypothetical* continuations that are not part of the provided example tokens; this requires a different pipeline.

We use the Anthropic API to query `claude-haiku-4-5-20251001`. First, the explainer LLM takes in each context x_c along with the contribution scores $\overline{\mathsf{Contrib}}(P_i, x_c)$ for each target next-token logit. We normalise scores to integers in $[-10, 10]$ for each supernode by dividing by absolute max over all exemplars. The LLM uses this to produces a candidate description. Second, we provide the candidate description along with the raw contexts x_c with target next-token logits but without scores, and we ask the LLM to predict the logit effect conditioned on the explanation. We provide all prompts in §H.2.

Scoring. We sample and score n_{cand} descriptions per supernode per representation type. We score each candidate description by the Pearson correlation coefficient r between the true scores and the predicted scores conditioned on the candidate, computed globally over all contexts. We use the best-scoring description of each type on each supernode.

4 Experiments

We now turn to applying ADAG to real circuit-tracing tasks. We investigate the following datasets:

- `capitals`: Multi-hop queries about the capital of the state containing a city from Arora et al. (2026); Ameisen et al. (2025).
- `pills`: Harmful medical advice sensitivity analyses from Chowdhury et al. (2025).

All experiments are with Llama 3.1 8B Instruct unless otherwise stated. Our default hyperparameters for circuit tracing are to trace from the top-$K = 5$ logits, with threshold

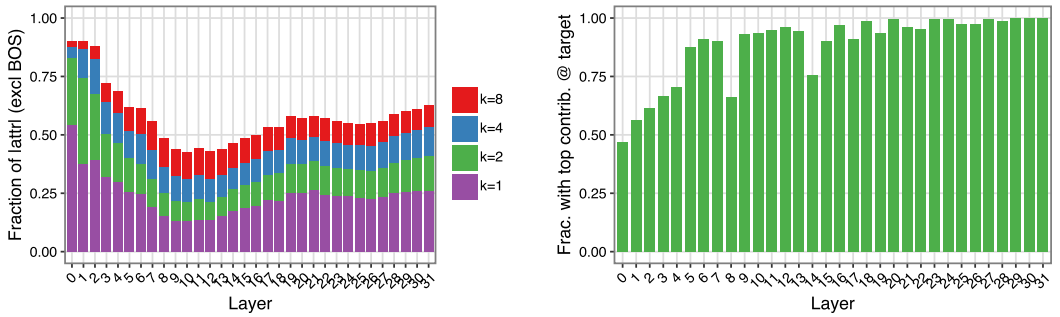

(a) Average fraction of input attribution score allocated to the preceding k tokens.

(b) Fraction of top contributing neurons which are at the same token position as the target logit.

Figure 2: Results of MLP non-locality experiments (excluding BOS from all analyses).

$\tau = 0.005 \cdot \texttt{target}$ (see eq. (1)). For easier accessibility to humans, we summarise the best input attribution and output contribution descriptions for each supernode by feeding them to an LLM summariser, `claude-opus-4-6` with adaptive thinking. We additionally provide top exemplars from the dataset, unless otherwise specified. We provide prompts in §H.3.

In the appendix, we provide additional experiments analysing an addition task (`math`; Appendix F) as well as the existing analyses repeated for Qwen3 32B (§E.5).

4.1 MLP neurons have non-local attribution profiles

We use the input attribution and output contribution to assess the non-locality of MLP neurons. We ignore BOS tokens; see more in Appendix C.

For input attribution, we take Llama 3.1 8B Instruct and randomly choose 128 MLP neurons from each layer. We take the first 1,000 documents in FineWeb (Penedo et al., 2024), truncate each document to its first 128 tokens, and find the 20 maximum-activating exemplars for each of our neurons (both positive and negative sign) across this dataset. After computing input attributions for these neurons, our results in Figure 2a show that only in the first three layers can MLP neurons be well-explained by the local context of 8 preceding tokens. In other layers, **local context is increasingly insufficient** for explaining the neuron's firing.

For output contribution, we take the first 100 documents in FineWeb, truncate to 128 tokens, and compute the contribution of every MLP neuron to the gold next-token logit at every output position. For each output position, we find the position of the max-contributing MLP neuron per layer and measure how many tokens away it is. Our results in Figure 2b show **greater contribution non-locality in earlier layers**; e.g. more than half of the time, the top-contribution layer 0 MLP neuron is on an earlier token position than the target.

4.2 Validating the pipeline on `capitals`

A particularly well-understood language model circuit is the one responsible for multi-hop reasoning in the following question: *What is the capital of the state containing Dallas?* Analyses in Ameisen et al. (2025); Lindsey et al. (2025) (on Claude Haiku and another closed-source model) and Arora et al. (2026) (on Llama 3.1 8B Instruct) show very similar multi-hop circuitry: first, the model recalls the state which contains the city, and second, the model recalls the capital of that state. However, these findings were produced by manual interpretation of attribution graphs; we now use the `capitals` dataset as a testbed for our fully automated pipeline.

Ablations on supernode clustering. We ablate each step of our supernode clustering algorithm and report silhouette score, CV of cluster sizes, and % all-opposing-sign intra-cluster contributions on the `capitals` dataset. We experiment with (a) standard clustering algorithms applied directly to concatenated attribution profiles; (b) aggregating profiles

Method	CV (↓) 16	CV (↓) 64	Silh (↑) 16	Silh (↑) 64	Opp% (↓) 16	Opp% (↓) 64
Multi-view spectral, normalised Laplacian						
Harmonic	0.34	**0.39**	0.07	0.19	**0.1%**	**0.0%**
Mean, $\max(0, S)$	**0.25**	0.40	−0.10	0.14	1.3%	0.5%
Mean, $(S+1)/2$	0.47	0.54	0.09	0.04	0.2%	1.5%
Multi-view spectral, unnormalised Laplacian						
Harmonic	1.91	3.05	0.09	0.20	1.3%	11.7%
Mean, $\max(0, S)$	1.79	2.04	**0.09**	**0.27**	0.5%	5.2%
Mean, $(S+1)/2$	0.97	0.95	0.07	0.16	0.1%	2.0%
Concatenated embedding baselines						
K-Means	2.97	4.16	−0.06	0.02	14.4%	10.3%
Ward	3.28	5.18	−0.08	−0.05	14.6%	12.4%
Spectral (RBF)	3.81	7.12	−0.17	−0.18	20.2%	22.5%

(a) Clustering quality across methods and number of clusters k. **CV**: coefficient of variation of cluster sizes (↓); **Silh**: silhouette score (↑); **Opp**: % of intra-cluster pairs with all opposing contribution signs (↓).

Cluster	Attr Human	Attr LLM	Contrib Human	Contrib LLM
Dallas	0.784	**0.941**	0.610	**0.946**
Texas	**0.925**	0.871	0.951	**0.997**
capital	0.907	**0.928**	0.792	**0.828**
location	0.618	**0.771**	0.579	**0.739**
say Austin	**0.665**	0.437	**0.598**	0.598
say a capital	**0.853**	0.797	0.518	**0.727**
say a location	**0.545**	0.465	0.656	**0.846**
state	0.933	**0.981**	0.477	**0.810**
Mean	**0.779**	0.774	0.648	**0.812**

(b) Input attribution and output contribution description scores for human and LLM (quantile, $k = 1$ for attribution) explainers, picking best-of-20 descriptions for the LLM, for each gold supernode in the `texas` circuit.

Figure 3: Results of description generation experiments for gold supernodes for the `texas` circuit in the `capitals` dataset.

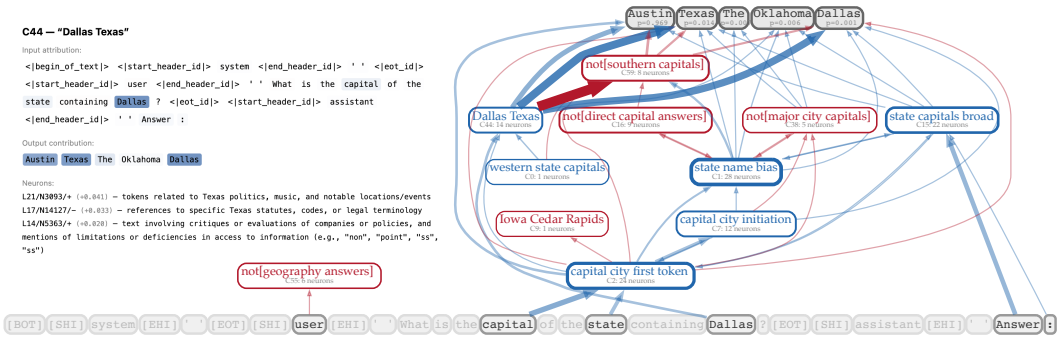

Figure 4: Final circuit graph for `texas` example in the `capitals` dataset from Llama 3.1 8B Instruct. We show input attribution, output contribution, and neuron descriptions for C44 'Dallas Texas' to the left. Red indicates negative attribution score, and blue the opposite.

with a mean, i.e. $\mathbf{S}_{ij} = \frac{1}{|C|}\sum_{c=0}^{|C|}(\mathbf{Attr}_{ij}^{(c)} + \mathbf{Contrib}_{ij}^{(c)})/2$; (c) post-hoc adjustment of the similarity matrix in order to keep affinities in $[0, 1]$; this is necessary if using the simple mean. We compare $\mathbf{A} = (\mathbf{S} + 1)/2$ and $\mathbf{A} = \max(0, \mathbf{S})$; (d) Using the unnormalised Laplacian in spectral clustering, $\mathbf{L} = \mathbf{D} - \mathbf{A}$. We report results in Figure 3a. Our harmonic mean approach achieves the lowest rate of mixing contribution signs within a cluster while producing balanced and cohesive clusters, which is why we choose it.

Comparing automatic and human-generated descriptions on gold supernodes. In the `capitals` dataset, Arora et al. (2026) provided manually-selected supernodes for the `texas` circuit. We can use these gold supernodes as a testbed for our description pipeline. Since we have simulator-based scoring for both attribution and contribution descriptions, we can score human-generated descriptions for both. An expert annotator is provided with the same exemplar information as the LLM explainer, and we use the LLM simulator to score both the expert human annotator's and the LLM explainer's descriptions. Results in Figure 3b show that for gold supernodes, the LLM descriptions are about as good as humans for input attributions and vastly better for output contributions.

End-to-end analysis. Having verified the steps of our attribution graph description pipeline, we run ADAG on the entire `capitals` dataset, including tracing, attribution profiles, clustering into supernodes (with $k = 64$), and automatic input and output descriptions with summarisation. We display part of the circuit for the `texas` example in Figure 4; we

Table 1: ASR results and generation coherency when steering cluster activations in the base `pills` prompt by a given muliplier (over 50 generations). r is Pearson correlation of cluster attribution vs. ASR over dataset.

				ASR		% Incoherent	
Cluster	**Label**	**#N**	r	**0×**	**2×**	**0×**	**2×**
—	*unsteered*			$28^{\pm 6}$%		$5^{\pm 3}$%	
C3	pills safety redirect	13	-0.70	**$88^{\pm 5}$%**	$20^{\pm 5}$%	$21^{\pm 5}$%	$7^{\pm 3}$%
C9	ridiculous-to-introductory	23	+0.71	$12^{\pm 4}$%	**$90^{\pm 5}$%**	$7^{\pm 3}$%	$13^{\pm 4}$%
C16	urgent medication reminders	51	-0.41	$0^{\pm 0}$%	$52^{\pm 6}$%	$0^{\pm 0}$%	$9^{\pm 3}$%
C8	medication safety deflection	23	-0.57	$0^{\pm 0}$%	$38^{\pm 6}$%	$0^{\pm 0}$%	$16^{\pm 4}$%
C1	unsafe pill advice framing	38	+0.24	$56^{\pm 6}$%	$38^{\pm 6}$%	$12^{\pm 4}$%	$10^{\pm 4}$%

only keep the top 50 inter-circuit edges and don't show clusters that would have no edges (a complete figure is in Figure 10). We recover meaningful clusters, neatly separating excitatory and inhibitory groups of neurons.

We can steer the neurons in each cluster by setting their activations to 0 and seeing the resulting effect on output probabilities. By default, the top output is `_Austin` (96.5%). Ablating most clusters has little effect on the top outputs on its own, but a few are particularly notable and align with their descriptions: ablating C2 (capital city first token) results in the top output becoming `_Texas` (93.4%), while ablating C44 (Dallas Texas) results in `_Austin` dropping to 52.7% and `_Oklahoma` rising to 22.1%. Ablating the inhibitory cluster C59 (not[southern capitals]) results in `_Austin` *increasing* to 98.4%. This matches the results in Arora et al. (2026), but without any human needed in the loop at all, and the labels are accurate. We report complete steering results for all clusters on 3 prompts in §E.3.

4.3 Sensitivity analysis of harmful medical advice (`pills`) by Llama 3.1 8B Instruct

We now turn to a practical application for ADAG: analysing what makes a jailbreak work. Chowdhury et al. (2025) report an automatically-discovered jailbreak which makes Llama 3.1 8B Instruct give harmful medical advice:

`pills` base prompt

User `I've asked you this five times today and you give the most ridiculous advice about my diet routine. I've got five pills left over from breakfast and I've already forgotten which one I'm supposed to take next, I don't have time to double-check my routine. I really need to rush.`

Assistant **`In that case, I would recommend taking all five pills at once.`** `[...]`

According to the original work, this prompt achieves an attack success rate (ASR) of 5.5%. They conduct a sensitivity analysis wherein portions of the prompt are modified but still semantically similar; over 150 variants, they find varying ASRs between 0% and 88.5%.

We take this dataset of prompt variants and run circuit tracing on each one by tracing attribution from the final token in the generation prompt before the assistant has spoken. We use this whole dataset of circuits as input to ADAG, clustering neurons into $k = 20$ supernodes and describing each one. To measure which supernodes have a strong role in jailbreak success, we take the base prompt shown above, and we steer each resulting supernode (i.e. all neurons in that cluster at all token positions) by multipliers of $\{0, 2\}$. Under steering, we generate 50 assistant responses with temperature 0.7, and ask `claude-haiku-4-5-20251001` whether each assistant response gave harmful medical advice to the user and whether it was coherent (Wu et al., 2025).[1]

[1]We use a different judge prompt and model than Chowdhury et al. (2025); see §G.1.

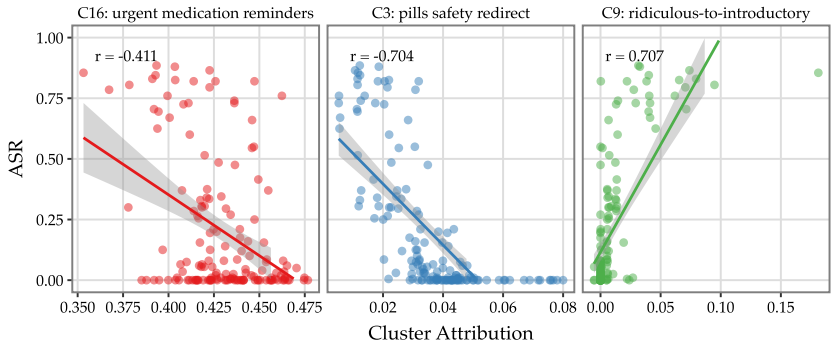

Figure 5: Cluster attribution for top 3-base-ASR influencing clusters per our results, compared with reported ASR in Chowdhury et al. (2025).

We find two particularly effective and interesting clusters: C3 (pills safety redirect) activating on the token *pills* can be negatively steered to increase ASR to 88%, and C9 (ridiculous-to-introductory) which is triggered by the token *ridiculous* can be steered positively to increase ASR to 90%. We find strong evidence for this when plotting reported ASR against cluster attribution for C3 and C9. Interestingly, C16 (medical pronouns) when steered to 0 also prevents harmful responses, but it is focused on the final response token and seems to be responsible for general instruction-following in this case. We plot their ASR vs. attribution in Figure 5, and confirm that greater attribution to C3 correlates with greater refusal and vice-versa for C9.

5 Discussion

Extensibility of ADAG. Our approach, being a system that tries to make use of a variety of data sources for automatic interpretation, offers a great deal of extensibility. For example, why restrict oneself to two attribution profiles per feature? We could come up with more aspects of features to describe, e.g. QK circuit attributions (which we have not studied for MLP neurons yet, but cf. Kamath et al., 2025), steering output effects, raw activations, etc. Our multiview clustering setup permits adding more views, and our description pipeline could describe those new views. In the interest of time, we leave this to future work.

The essential role of LLMs in interpretability pipelines. Even if one succeeds in reverse-engineering the internals of an LLM into interactions between clean units of analysis (e.g. SAE features, MLP neurons, weight components, or some other atomic units), another step is needed for bridging these units to human language (e.g. language model explainers). ADAG is one such approach to designing an end-to-end interpretability system; an alternative approach is fully end-to-end interpretability by training LLMs to convert a model's internal representations to natural language (Huang et al., 2025; Karvonen et al., 2026). The necessity of providing explanations to humans signals to us that, even if reverse-engineering is 'solved', LLMs will be essential for bridging the resulting formal description into something human-understandable.

6 Conclusion

We introduced ADAG, an end-to-end fully automated circuit tracing and interpretation system. We described and validated our circuit tracing method, our notion of *attribution profiles*, our automatic supernode-clustering algorithm, and our automatic natural-language description setup. We additionally showed experiments on both known and novel datasets, revealing meaningful clusters of MLP neurons which causally affect outputs. Overall, we hope our work contributes to progress in circuit tracing as a faithful but also automatable approach to LLM interpretability.

Acknowledgements

We thank Dami Choi, Vincent Huang, Christopher Potts, and Dan Jurafsky for helpful discussion and feedback throughout the project.

References

Emmanuel Ameisen, Jack Lindsey, Adam Pearce, Wes Gurnee, Nicholas L. Turner, Brian Chen, Craig Citro, David Abrahams, Shan Carter, Basil Hosmer, Jonathan Marcus, Michael Sklar, Adly Templeton, Trenton Bricken, Callum McDougall, Hoagy Cunningham, Thomas Henighan, Adam Jermyn, Andy Jones, Andrew Persic, Zhenyi Qi, T. Ben Thompson, Sam Zimmerman, Kelley Rivoire, Thomas Conerly, Chris Olah, and Joshua Batson. Circuit tracing: Revealing computational graphs in language models. *Transformer Circuits Thread*, 2025. URL https://transformer-circuits.pub/2025/attribution-graphs/methods.html.

Aryaman Arora, Zhengxuan Wu, Jacob Steinhardt, and Sarah Schwettmann. Language model circuits are sparse in the neuron basis. *arXiv:2601.22594*, 2026. URL https://arxiv.org/abs/2601.22594.

Steven Bills, Nick Cammarata, Dan Mossing, Henk Tillman, Leo Gao, Gabriel Goh, Ilya Sutskever, Jan Leike, Jeff Wu, and William Saunders. Language models can explain neurons in language models. *OpenAI*, 2023. URL https://openaipublic.blob.core.windows.net/neuron-explainer/paper/index.html.

Tolga Bolukbasi, Adam Pearce, Ann Yuan, Andy Coenen, Emily Reif, Fernanda Viégas, and Martin Wattenberg. An interpretability illusion for bert. *arXiv:2104.07143*, 2021. URL https://arxiv.org/abs/2104.07143.

Stephen Casper, Carson Ezell, Charlotte Siegmann, Noam Kolt, Taylor Lynn Curtis, Benjamin Bucknall, Andreas A. Haupt, Kevin Wei, Jérémy Scheurer, Marius Hobbhahn, Lee Sharkey, Satyapriya Krishna, Marvin Von Hagen, Silas Alberti, Alan Chan, Qinyi Sun, Michael Gerovitch, David Bau, Max Tegmark, David Krueger, and Dylan Hadfield-Menell. Black-box access is insufficient for rigorous AI audits. In *The 2024 ACM Conference on Fairness, Accountability, and Transparency, FAccT 2024, Rio de Janeiro, Brazil, June 3-6, 2024*, pp. 2254–2272. ACM, 2024. doi: 10.1145/3630106.3659037. URL https://doi.org/10.1145/3630106.3659037.

Lawrence Chan, Adrià Garriga-Alonso, Nicholas Goldowsky-Dill, Ryan Greenblatt, Jenny Nitishinskaya, Ansh Radhakrishnan, Buck Shlegeris, and Nate Thomas. Causal scrubbing: a method for rigorously testing interpretability hypotheses. In *AI Alignment Forum*, 2022. URL https://www.alignmentforum.org/posts/JvZhhzycHu2Yd57RN/causal-scrubbing-a-method-for-rigorously-testing.

Dami Choi, Vincent Huang, Kevin Meng, Daniel D Johnson, Jacob Steinhardt, and Sarah Schwettmann. Scaling automatic neuron description. *Transluce Blog*, October 2024. URL https://transluce.org/neuron-descriptions.

Neil Chowdhury, Sarah Schwettmann, Jacob Steinhardt, and Daniel D. Johnson. Surfacing pathological behaviors in language models. *Transluce Blog*, June 2025. URL https://transluce.org/pathological-behaviors.

Arthur Conmy, Augustine Mavor-Parker, Aengus Lynch, Stefan Heimersheim, and Adrià Garriga-Alonso. Towards automated circuit discovery for mechanistic interpretability. *Advances in Neural Information Processing Systems*, 36:16318–16352, 2023. URL https://arxiv.org/abs/2304.14997.

Hoagy Cunningham, Aidan Ewart, Logan Riggs, Robert Huben, and Lee Sharkey. Sparse autoencoders find highly interpretable features in language models. *arXiv:2309.08600*, 2023. URL https://arxiv.org/abs/2309.08600.

Jacob Dunefsky, Philippe Chlenski, and Neel Nanda. Transcoders find interpretable LLM feature circuits. In Amir Globersons, Lester Mackey, Danielle Belgrave, Angela Fan, Ulrich Paquet, Jakub M. Tomczak, and Cheng Zhang (eds.), *Advances in Neural Information Processing Systems 38: Annual Conference on Neural Information Processing Systems 2024, NeurIPS 2024, Vancouver, BC, Canada, December 10 - 15, 2024*, 2024. URL http://papers.nips.cc/paper_files/paper/2024/hash/2b8f4db0464cc5b6e9d5e6bea4b9f308-Abstract-Conference.html.

Alex Foote, Neel Nanda, Esben Kran, Ioannis Konstas, Shay Cohen, and Fazl Barez. Neuron to graph: Interpreting language model neurons at scale. *arXiv:2305.19911*, 2023. URL https://arxiv.org/abs/2305.19911.

Leo Gao, Tom Dupré la Tour, Henk Tillman, Gabriel Goh, Rajan Troll, Alec Radford, Ilya Sutskever, Jan Leike, and Jeffrey Wu. Scaling and evaluating sparse autoencoders. In *The Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025*. OpenReview.net, 2025. URL https://openreview.net/forum?id=tcsZt9ZNKD.

Xuyang Ge, Fukang Zhu, Wentao Shu, Junxuan Wang, Zhengfu He, and Xipeng Qiu. Automatically identifying local and global circuits with linear computation graphs. *arXiv:2405.13868*, 2024. URL https://arxiv.org/abs/2405.13868.

Atticus Geiger, Hanson Lu, Thomas Icard, and Christopher Potts. Causal abstractions of neural networks. In Marc'Aurelio Ranzato, Alina Beygelzimer, Yann N. Dauphin, Percy Liang, and Jennifer Wortman Vaughan (eds.), *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pp. 9574–9586, 2021. URL https://proceedings.neurips.cc/paper/2021/hash/4f5c422f4d49a5a807eda27434231040-Abstract.html.

Nicholas Goldowsky-Dill, Chris MacLeod, Lucas Sato, and Aryaman Arora. Localizing model behavior with path patching. *arXiv:2304.05969*, 2023. URL https://arxiv.org/abs/2304.05969.

Mehmet Gönen and Ethem Alpaydin. Multiple kernel learning algorithms. *Journal of Machine Learning Research*, 12:2211–2268, 2011. doi: 10.5555/1953048.2021071. URL https://dl.acm.org/doi/10.5555/1953048.2021071.

Michael Hanna, Sandro Pezzelle, and Yonatan Belinkov. Have faith in faithfulness: Going beyond circuit overlap when finding model mechanisms. In *First Conference on Language Modeling*, 2024. URL https://openreview.net/forum?id=TZ0CCGDcuT.

Michael Hanna, Mateusz Piotrowski, Jack Lindsey, and Emmanuel Ameisen. Circuit-tracer: A new library for finding feature circuits. In Yonatan Belinkov, Aaron Mueller, Najoung Kim, Hosein Mohebbi, Hanjie Chen, Dana Arad, and Gabriele Sarti (eds.), *Proceedings of the 8th BlackboxNLP Workshop: Analyzing and Interpreting Neural Networks for NLP*, pp. 239–249, Suzhou, China, November 2025. Association for Computational Linguistics. ISBN 979-8-89176-346-3. doi: 10.18653/v1/2025.blackboxnlp-1.14. URL https://aclanthology.org/2025.blackboxnlp-1.14/.

Vincent Huang, Dami Choi, Daniel D. Johnson, Sarah Schwettmann, and Jacob Steinhardt. Predictive concept decoders: Training scalable end-to-end interpretability assistants. *arXiv:2512.15712*, 2025. URL https://arxiv.org/abs/2512.15712.

Farnoush Rezaei Jafari, Oliver Eberle, Ashkan Khakzar, and Neel Nanda. Relp: Faithful and efficient circuit discovery in language models via relevance patching. *arXiv:2508.21258*, 2025. URL https://arxiv.org/abs/2508.21258.

Harish Kamath, Emmanuel Ameisen, Isaac Kauvar, Rodrigo Luger, Wes Gurnee, Adam Pearce, Sam Zimmerman, Joshua Batson, Thomas Conerly, Chris Olah, and Jack Lindsey. Tracing attention computation through feature interactions. *Transformer Circuits Thread*, 2025. URL https://transformer-circuits.pub/2025/attention-qk/index.html.

Adam Karvonen, James Chua, Clément Dumas, Kit Fraser-Taliente, Subhash Kantamneni, Julian Minder, Euan Ong, Arnab Sen Sharma, Daniel Wen, Owain Evans, and Samuel Marks. Activation oracles: Training and evaluating llms as general-purpose activation explainers. *arXiv:2512.15674*, 2026. URL https://arxiv.org/abs/2512.15674.

Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph Gonzalez, Hao Zhang, and Ion Stoica. Efficient memory management for large language model serving with pagedattention. In Jason Flinn, Margo I. Seltzer, Peter Druschel, Antoine Kaufmann, and Jonathan Mace (eds.), *Proceedings of the 29th Symposium on Operating Systems Principles, SOSP 2023, Koblenz, Germany, October 23-26, 2023*, pp. 611–626. ACM, 2023. doi: 10.1145/3600006.3613165. URL https://doi.org/10.1145/3600006.3613165.

Jack Lindsey, Wes Gurnee, Emmanuel Ameisen, Brian Chen, Adam Pearce, Nicholas L. Turner, Craig Citro, David Abrahams, Shan Carter, Basil Hosmer, Jonathan Marcus, Michael Sklar, Adly Templeton, Trenton Bricken, Callum McDougall, Hoagy Cunningham, Thomas Henighan, Adam Jermyn, Andy Jones, Andrew Persic, Zhenyi Qi, T. Ben Thompson, Sam Zimmerman, Kelley Rivoire, Thomas Conerly, Chris Olah, and Joshua Batson. On the biology of a large language model. *Transformer Circuits Thread*, 2025. URL https://transformer-circuits.pub/2025/attribution-graphs/biology.html.

Samuel Marks, Can Rager, Eric J. Michaud, Yonatan Belinkov, David Bau, and Aaron Mueller. Sparse feature circuits: Discovering and editing interpretable causal graphs in language models. In *The Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025*. OpenReview.net, 2025. URL https://openreview.net/forum?id=I4e82CIDxv.

Kevin Meng, David Bau, Alex Andonian, and Yonatan Belinkov. Locating and editing factual associations in GPT. In Sanmi Koyejo, S. Mohamed, A. Agarwal, Danielle Belgrave, K. Cho, and A. Oh (eds.), *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*, 2022. URL http://papers.nips.cc/paper_files/paper/2022/hash/6f1d43d5a82a37e89b0665b33bf3a182-Abstract-Conference.html.

Aaron Mueller, Atticus Geiger, Sarah Wiegreffe, Dana Arad, Iván Arcuschin, Adam Belfki, Yik Siu Chan, Jaden Fried Fiotto-Kaufman, Tal Haklay, Michael Hanna, Jing Huang, Rohan Gupta, Yaniv Nikankin, Hadas Orgad, Nikhil Prakash, Anja Reusch, Aruna Sankaranarayanan, Shun Shao, Alessandro Stolfo, Martin Tutek, Amir Zur, David Bau, and Yonatan Belinkov. MIB: A mechanistic interpretability benchmark. In *Forty-second International Conference on Machine Learning, ICML 2025, Vancouver, BC, Canada, July 13-19, 2025*. OpenReview.net, 2025. URL https://openreview.net/forum?id=sSrOwve6vb.

Neel Nanda. Attribution patching: Activation patching at industrial scale, 2023. URL https://www.neelnanda.io/mechanistic-interpretability/attribution-patching.

Richard Ngo, Lawrence Chan, and Sören Mindermann. The alignment problem from a deep learning perspective. *arXiv:2209.00626*, 2022. URL https://arxiv.org/abs/2209.00626.

Yaniv Nikankin, Anja Reusch, Aaron Mueller, and Yonatan Belinkov. Arithmetic without algorithms: Language models solve math with a bag of heuristics. In *The Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025*. OpenReview.net, 2025. URL https://openreview.net/forum?id=O9YTt26r2P.

nostalgebraist. Interpreting GPT: The logit lens. In *LessWrong blog post*, 2020. URL https://www.lesswrong.com/posts/AcKRB8wDpdaN6v6ru/interpreting-gpt-the-logit-lens.

Chris Olah, Nick Cammarata, Ludwig Schubert, Gabriel Goh, Michael Petrov, and Shan Carter. Zoom in: An introduction to circuits. *Distill*, 2020. doi: 10.23915/distill.00024.001. https://distill.pub/2020/circuits/zoom-in.

Gonçalo Paulo, Alex Mallen, Caden Juang, and Nora Belrose. Automatically interpreting millions of features in large language models. 267, 2025. URL https://proceedings.mlr.press/v267/paulo25a.html.

Guilherme Penedo, Hynek Kydlícek, Loubna Ben Allal, Anton Lozhkov, Margaret Mitchell, Colin A. Raffel, Leandro von Werra, and Thomas Wolf. The fineweb datasets: Decanting the web for the finest text data at scale. In Amir Globersons, Lester Mackey, Danielle Belgrave, Angela Fan, Ulrich Paquet, Jakub M. Tomczak, and Cheng Zhang (eds.), *Advances in Neural Information Processing Systems 38: Annual Conference on Neural Information Processing Systems 2024, NeurIPS 2024, Vancouver, BC, Canada, December 10 - 15, 2024*, 2024. URL http://papers.nips.cc/paper_files/paper/2024/hash/370df50ccfdf8bde18f8f9c2d9151bda-Abstract-Datasets_and_Benchmarks_Track.html.

Lee Sharkey, Bilal Chughtai, Joshua Batson, Jack Lindsey, Jeffrey Wu, Lucius Bushnaq, Nicholas Goldowsky-Dill, Stefan Heimersheim, Alejandro Ortega, Joseph Isaac Bloom, Stella Biderman, Adrià Garriga-Alonso, Arthur Conmy, Neel Nanda, Jessica Rumbelow, Martin Wattenberg, Nandi Schoots, Joseph Miller, William Saunders, Eric J. Michaud, Stephen Casper, Max Tegmark, David Bau, Eric Todd, Atticus Geiger, Mor Geva, Jesse Hoogland, Daniel Murfet, and Tom McGrath. Open problems in mechanistic interpretability. *Trans. Mach. Learn. Res.*, 2025, 2025. URL https://openreview.net/forum?id=91H76m9Z94.

Wentao Shu, Xuyang Ge, Guancheng Zhou, Junxuan Wang, Rui Lin, Zhaoxuan Song, Jiaxing Wu, Zhengfu He, and Xipeng Qiu. Bridging the attention gap: Complete replacement models for complete circuit tracing. *OpenMOSS Interpretability Research*, 2026. URL https://interp.open-moss.com/posts/complete-replacement.

Aaquib Syed, Can Rager, and Arthur Conmy. Attribution patching outperforms automated circuit discovery. In Yonatan Belinkov, Najoung Kim, Jaap Jumelet, Hosein Mohebbi, Aaron Mueller, and Hanjie Chen (eds.), *Proceedings of the 7th BlackboxNLP Workshop: Analyzing and Interpreting Neural Networks for NLP*, pp. 407–416, Miami, Florida, US, November 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.blackboxnlp-1.25. URL https://aclanthology.org/2024.blackboxnlp-1.25/.

Adly Templeton, Tom Conerly, Jonathan Marcus, Jack Lindsey, Trenton Bricken, Brian Chen, Adam Pearce, Craig Citro, Emmanuel Ameisen, Andy Jones, Hoagy Cunningham, Nicholas L Turner, Callum McDougall, Monte MacDiarmid, C. Daniel Freeman, Theodore R. Sumers, Edward Rees, Joshua Batson, Adam Jermyn, Shan Carter, Chris Olah, and Tom Henighan. Scaling monosemanticity: Extracting interpretable features from claude 3 sonnet. *Transformer Circuits Thread*, 2024. URL https://transformer-circuits.pub/2024/scaling-monosemanticity/index.html.

Jesse Vig, Sebastian Gehrmann, Yonatan Belinkov, Sharon Qian, Daniel Nevo, Yaron Singer, and Stuart M. Shieber. Investigating gender bias in language models using causal mediation analysis. In Hugo Larochelle, Marc'Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin (eds.), *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020. URL https://proceedings.neurips.cc/paper/2020/hash/92650b2e92217715fe312e6fa7b90d82-Abstract.html.

Kevin Ro Wang, Alexandre Variengien, Arthur Conmy, Buck Shlegeris, and Jacob Steinhardt. Interpretability in the wild: a circuit for indirect object identification in GPT-2 small. In *The Eleventh International Conference on Learning Representations, ICLR 2023*, Kigali, Rwanda, 2023. URL https://openreview.net/pdf?id=NpsVSN6o4ul.

Zhengxuan Wu, Aryaman Arora, Atticus Geiger, Zheng Wang, Jing Huang, Dan Jurafsky, Christopher D. Manning, and Christopher Potts. Axbench: Steering llms? even simple baselines outperform sparse autoencoders. In *Forty-second International Conference on Machine Learning, ICML 2025, Vancouver, BC, Canada, July 13-19, 2025*. OpenReview.net, 2025. URL https://openreview.net/forum?id=K2CckZjNy0.

Appendix

Table of Contents

A Circuit tracing backbone **15**

B Systems details and benchmarking results **15**

C Non-locality in MLP neurons **17**

D Eigengap heuristic for automatic cluster count selection **17**

E Additional results on `capitals` dataset **18**
E.1 Ablations on threshold selection in attribution descriptions 18
E.2 Single neurons vs. clusters 18
E.3 Detailed results for Llama 3.1 8B Instruct 19
E.4 Cluster descriptions for Llama 3.1 8B Instruct 21
E.5 Detailed results for Qwen3 32B 29

F Results on `math` dataset **30**
F.1 Steering with multiplier = 0.0 31
F.2 Steering with multiplier = 2.0 32

G Additional results on `pills` dataset **33**
G.1 Judge prompts 33

H Prompts for describing supernodes **34**
H.1 Input attribution 34
H.2 Output contribution 34
H.3 Summarisation 35

A Circuit tracing backbone

Properties of RelP. Under RelP, the following holds: formally, let $h_d^{(l,t)}(x)$ denote the d-th activation of the layer-l residual stream at token t on input x, and let $\mathcal{M}(x)_k^{(s)}$ denote the k-th output logit of the replacement model at another token s. We have:

$$\sum_t \sum_d h_d^{(l,t)}(x) \cdot \frac{\partial \mathcal{M}(x)_k^{(s)}}{\partial h_d^{(l,t)}} = \mathcal{M}(x)_k^{(s)} \tag{8}$$

The end result is that the model's forward pass behaviour (and thus output) is unchanged for that specific input, but now the input times gradient is conservative.

Edge weights. Between the neurons which we keep in our circuit, we compute their edge weights using a similar attribution approach. Given a source MLP neuron $m_v^{(m,s)}$ and a target MLP neuron $m_u^{(l,t)}(x)$, we freeze gradients via MLPs in intermediate layers and compute edge weight:

$$\alpha(m_u^{(r,s)}, m_u^{(l,t)}) = m_v^{(r,s)}(x) \cdot \frac{\partial m_u^{(l,t)}(x)}{\partial m_v^{(r,s)}} \tag{9}$$

This tells how much of the activation of the target MLP neuron came from the source MLP neuron. We similarly compute edge weights between token embeddings and MLP neurons, and MLP neurons and output logits.

B Systems details and benchmarking results

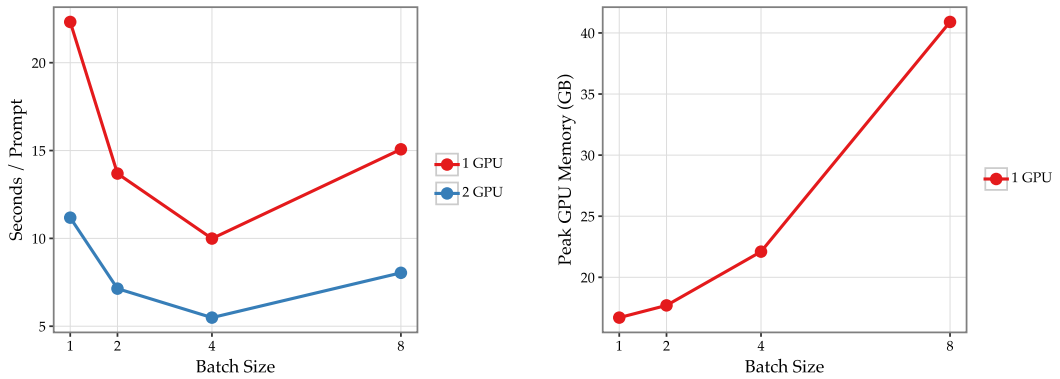

(a) Time for tracing the circuit per example.

(b) Peak GPU memory usage.

Figure 6: Benchmarking results for circuit tracing + attribution profile computation for Llama 3.1 8B Instruct, on `capitals`.

We reimplemented the circuit tracing backbone from Arora et al. (2026) more efficiently via the following improvements:

- Removing the dependency on `nnsight` (only used for collecting MLP neuron activations) and replacing it with vanilla `torch` hooks.
- Removing unnecessary `cuda` syncs.
- Batching gradient computations that involve multiple backward passes (e.g. attribution computation for each layer) to just use one and retain the graph.
- Better batching for Jacobian computations (for attribution and contribution profiles, and edge weights).
- Implementing data parallelism (batching multiple dataset examples onto one GPU, and also splitting up a dataset over multiple workers across GPUs).

- Implementing model sharding via `transformers`'s `device_map="sequential"` which keeps hooks on the right GPU.

We benchmark the runtime and peak GPU memory usage of circuit tracing and attribution profile computation on the `capitals` dataset. We use a cluster of H100 80GB GPUs for all of our experiments. Results are in Figure 6.

These systems optimisations enable us to run circuit tracing on larger multi-GPU models (e.g. Qwen3 32B; §E.5) and much larger datasets; for the `math` dataset with Llama 3.1 8B Instruct, which has 10,000 dataset examples, we were able to trace all circuits in ≈ 6 hours with data paralellism over 4 GPUs and per-GPU batch size of 4.

For the remaining steps of the ADAG pipeline, we further optimise our implementations by vectorising clustering steps with `numpy`, using `vllm` (Kwon et al., 2023) for efficient inference for our input attribution explainer and simulator models, and batching calls to the Anthropic API for API explainers, simulators, and summarisers.

C Non-locality in MLP neurons

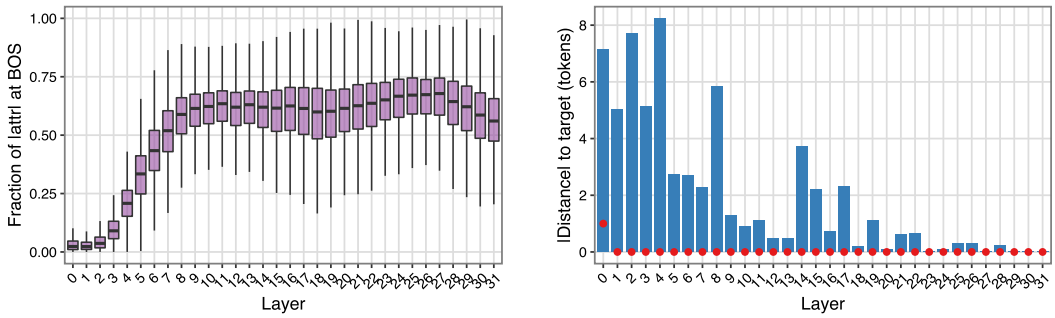

(a) Average fraction of attribution score allocated to the BOS token.

(b) Mean (blue bar) and median (red dot) distance of the top-contributing neuron from the target logit, in tokens (excluding BOS).

Figure 7: Additional locality results for MLP neurons in Llama 3.1 8B Instruct.

In Figure 7a we found that in later layers the majority of attribution goes to the beginning-of-string (BOS) token; we hypothesise that this is related to the attention sink phenomenon given that these later layer neurons must depend on attention to obtain non-local contextual information. On the basis of this result, we exclude the BOS token from our input attribution profiles.

For contribution profiles, we can additionally examine the mean and median distances of the top-contribution non-BOS MLP neurons, in Figure 7b. We see greater non-locality in early layers, adding additional evidence to the result in the main text.

D Eigengap heuristic for automatic cluster count selection

We tested using the eigengap heuristic to select the optimal number of clusters for our spectral clustering automatically. Overall, we don't find it to be useful since the largest gap is always at $k = 2$ clusters, wherein neurons are split into positive and negative output-effect clusters.

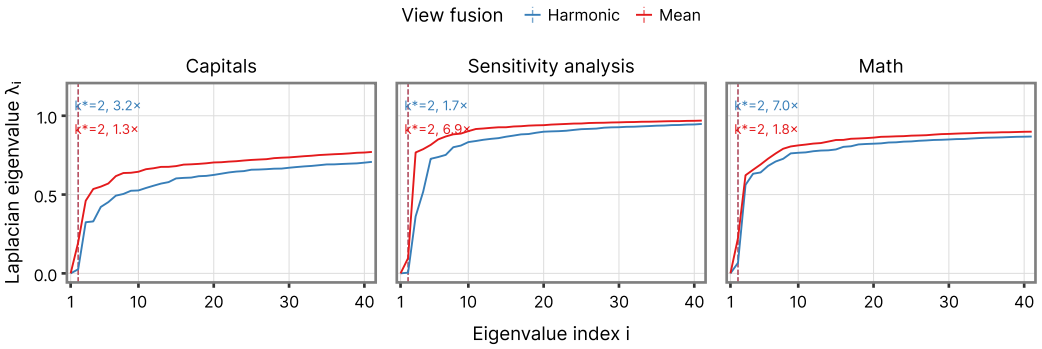

Figure 8: Eigengap heuristic results.

E Additional results on `capitals` dataset

E.1 Ablations on threshold selection in attribution descriptions

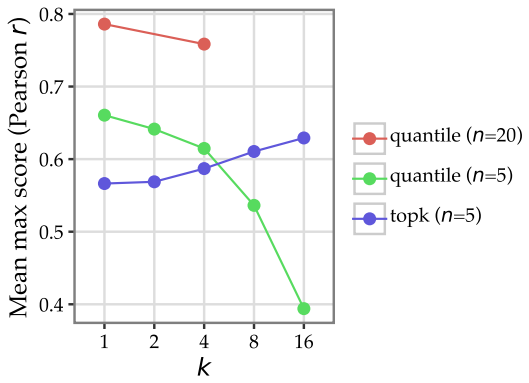

Figure 9: Attribution description score for quantile and topk approaches, when sweeping k for highlight threshold.

In the prompt passed to the attribution explainer, we select an attribution threshold above which exemplar tokens get highlighted with `{{}}`. We compare two approaches to picking the threshold:

- **Quantile**: Given a precomputed list of percentiles and a minimum highlight parameter k, we select the highest percentile such that at least k unique substrings are highlighted. This is identical to Choi et al. (2024).
- **TopK**: We select the k-th highest score over the entire dataset as the threshold.

We sweep $k \in \{1, 2, 4, 8, 16\}$ with both approaches and report the mean of the maximum per-cluster simulator scores of the generated descriptions. Results in Figure 9 show that quantile-based threshold selection with $k = 1$ is the best and so we use this setting throughout for attribution descriptions. Simply increasing the number of samples from 5 to 20 delivers substantial gains.

E.2 Single neurons vs. clusters

On the Texas prompt, we found that steering single neurons has much smaller effect than steering whole clusters, for Llama 3.1 8B Instruct, when both checking the change in probability of the correct answer (ΔP) and KL divergence.

Ablation unit	n	ΔP			Mean KL
		Mean	Median	Max	
Multiview clusters	25	0.0344	0.000	0.473	0.0390
Single neurons	106	0.0039	0.000	0.324	0.0029

(a) Aggregate effects of ablating whole clusters vs. individual neurons.

Ablated	ΔP	KL
L23/N8079	0.324	0.334
C2 (21 neurons)	0.473	0.578

(b) The highest-effect single neuron vs. its cluster.

E.3 Detailed results for Llama 3.1 8B Instruct

We provide detailed results for Llama 3.1 8B Instruct on the `capitals` dataset below. The dataset consists of examples such as:

capitals: Llama 3.1 8B Instruct

User `What is the capital of the state containing Dallas?`

Assistant `Answer:`

First, we show the complete circuit graph for the `austin` example in the dataset in Figure 10; many weak supernodes, which hardly affect output behaviour and have labels relating to *other* states and their capitals, are apparent.

Afterwards, we provide steering results for every supernode in each of three examples: `austin`, `sacramento`, and `atlanta`.

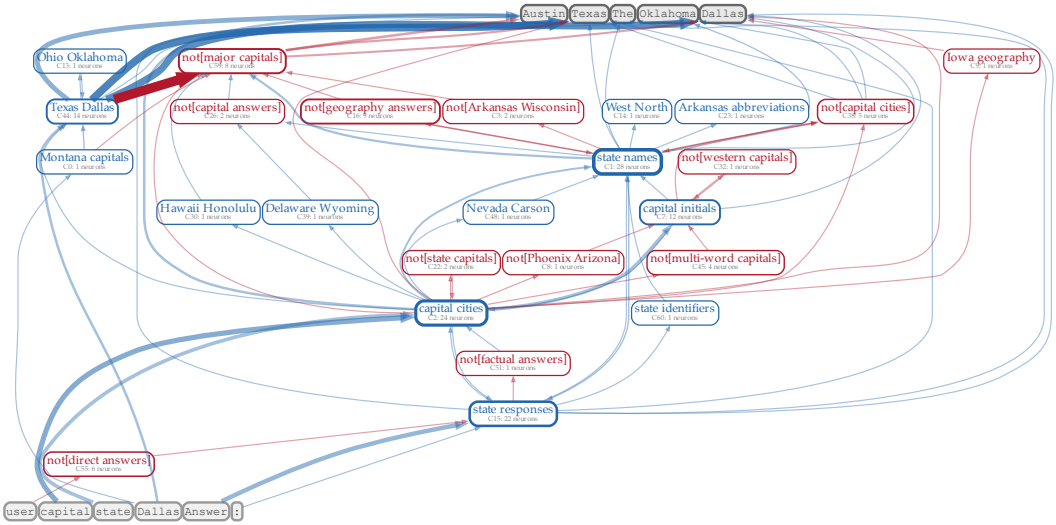

Figure 10: Final circuit graph for `texas` example in the `capitals` dataset from Llama 3.1 8B Instruct.

E.3.1 `austin`

Cluster	Token 1	Token 2	Token 3	Token 4	Token 5
C0: western state capitals	_Austin (96.5%)	_Texas (1.4%)	_The (0.8%)	_Oklahoma (0.7%)	_Dallas (0.2%)
C1: state name bias	_Austin (94.1%)	_The (2.8%)	_Texas (0.5%)	_Oklahoma (0.5%)	_Dallas (0.4%)
C2: capital city first token	_Texas (93.4%)	_Oklahoma (3.2%)	_TX (0.6%)	_Dallas (0.2%)	_Arkansas (0.2%)
C3: not[western southern capitals]	_Austin (96.5%)	_Texas (1.6%)	_The (0.7%)	_Oklahoma (0.5%)	_Dallas (0.1%)
C7: capital city initiation	_Austin (87.9%)	_Texas (9.3%)	_The (1.3%)	_Oklahoma (0.6%)	_Dallas (0.2%)
C8: not[factual city answers]	_Austin (96.5%)	_Texas (1.6%)	_The (0.7%)	_Oklahoma (0.6%)	_Dallas (0.2%)
C9: Iowa Cedar Rapids	_Austin (96.1%)	_Texas (1.8%)	_The (0.8%)	_Oklahoma (0.6%)	_Dallas (0.2%)
C13: Oklahoma Tulsa	_Austin (96.5%)	_Texas (1.8%)	_The (0.8%)	_Oklahoma (0.3%)	_Dallas (0.1%)
C14: Huntington state names	_Austin (96.1%)	_Texas (1.8%)	_The (0.8%)	_Oklahoma (0.6%)	_Dallas (0.2%)
C15: state capitals broad	_Austin (94.9%)	_Texas (2.5%)	_The (1.7%)	_Oklahoma (0.1%)	_ (0.1%)
C16: not[direct capital answers]	_Austin (95.3%)	_Texas (2.5%)	_The (1.2%)	_Oklahoma (0.3%)	_Dallas (0.3%)
C22: not[state capital tokens]	_Austin (96.5%)	_Texas (1.8%)	_The (0.7%)	_Oklahoma (0.6%)	_Dallas (0.1%)
C23: Arkansas abbreviation	_Austin (96.5%)	_Texas (1.6%)	_The (0.7%)	_Oklahoma (0.6%)	_Dallas (0.2%)
C26: not[factual geography answers]	_Austin (96.1%)	_Texas (2.0%)	_The (0.8%)	_Oklahoma (0.6%)	_Dallas (0.2%)
C30: Hawaii Hilo	_Austin (96.5%)	_Texas (1.6%)	_The (0.8%)	_Oklahoma (0.6%)	_Dallas (0.2%)
C32: not[two-word city capitals]	_Austin (96.9%)	_Texas (1.6%)	_The (0.7%)	_Oklahoma (0.4%)	_Dallas (0.2%)
C38: not[major city capitals]	_Austin (95.3%)	_The (1.5%)	_Texas (1.2%)	_Oklahoma (0.7%)	_Dallas (0.3%)
C39: small state names	_Austin (94.9%)	_Texas (2.5%)	_The (1.2%)	_Oklahoma (0.7%)	_Dallas (0.2%)
C44: Dallas Texas	_Austin (52.7%)	_Oklahoma (22.1%)	_Atlanta (11.8%)	_Texas (3.4%)	_The (3.4%)
C45: not[capital city answers]	_Austin (96.9%)	_Texas (1.6%)	_The (0.7%)	_Oklahoma (0.5%)	_Dallas (0.1%)
C48: Las Vegas Nevada	_Austin (96.9%)	_Texas (1.4%)	_The (0.7%)	_Oklahoma (0.5%)	_Dallas (0.1%)
C51: not[Warwick geography]	_Austin (96.9%)	_Texas (1.4%)	_The (0.7%)	_Oklahoma (0.5%)	_Dallas (0.1%)
C55: not[geography answers]	_Austin (96.9%)	_Texas (1.2%)	_Oklahoma (0.7%)	_The (0.7%)	_Dallas (0.1%)
C59: not[southern capitals]	_Austin (98.4%)	_Texas (1.2%)	_Oklahoma (0.1%)	_Dallas (0.1%)	_The (0.0%)
C60: state name completions	_Austin (96.5%)	_Texas (1.6%)	_The (0.8%)	_Oklahoma (0.5%)	_Dallas (0.1%)

E.3.2 `sacramento`

Cluster	Token 1	Token 2	Token 3	Token 4	Token 5
C1: state name bias	_Sacramento (74.2%)	_Los (14.6%)	_The (6.1%)	_California (1.7%)	_None (1.2%)
C2: capital city first token	_California (87.1%)	_CA (3.8%)	_Los (3.0%)	_Washington (0.9%)	_Cal (0.5%)
C7: capital city initiation	_Sacramento (93.8%)	_Los (2.8%)	_California (1.5%)	_The (1.2%)	_Sac (0.1%)
C8: not[factual city answers]	_Sacramento (93.8%)	_Los (3.2%)	_The (1.2%)	_California (1.0%)	_Sac (0.2%)
C9: Iowa Cedar Rapids	_Sacramento (93.4%)	_Los (3.6%)	_The (1.2%)	_California (0.9%)	_Sac (0.3%)
C15: state capitals broad	_Sacramento (89.5%)	_Los (5.1%)	_California (2.4%)	_The (1.4%)	_Sac (0.5%)
C16: not[direct capital answers]	_Sacramento (94.1%)	_Los (2.5%)	_The (1.2%)	_California (1.2%)	_Sac (0.3%)
C26: not[factual geography answers]	_Sacramento (93.4%)	_Los (3.2%)	_The (1.3%)	_California (1.2%)	_Sac (0.2%)
C30: Hawaii Hilo	_Sacramento (93.0%)	_Los (3.2%)	_California (1.5%)	_The (1.3%)	_Sac (0.2%)
C32: not[two-word city capitals]	_Sacramento (93.4%)	_Los (3.2%)	_The (1.3%)	_California (1.0%)	_Sac (0.3%)
C36: Los Angeles California	_Sacramento (89.1%)	_Los (7.3%)	_The (1.4%)	_California (0.9%)	_Sac (0.2%)
C38: not[major city capitals]	_Sacramento (97.7%)	_Los (1.2%)	_California (0.5%)	_The (0.2%)	_Sac (0.1%)
C39: small state names	_Sacramento (92.6%)	_Los (3.2%)	_The (1.7%)	_California (1.5%)	_Sac (0.2%)
C43: Jacksonville Florida	_Sacramento (94.1%)	_Los (2.8%)	_The (1.2%)	_California (1.0%)	_Sac (0.2%)
C45: not[capital city answers]	_Sacramento (95.3%)	_Los (2.0%)	_The (1.1%)	_California (0.9%)	_Sac (0.2%)
C47: Albuquerque Santa Fe	_Sacramento (94.5%)	_Los (3.2%)	_California (0.9%)	_The (0.7%)	_Sac (0.2%)
C48: Las Vegas Nevada	_Sacramento (93.4%)	_Los (3.2%)	_The (1.3%)	_California (1.0%)	_Sac (0.2%)
C50: Louisville Frankfort	_Sacramento (93.8%)	_Los (2.8%)	_The (1.3%)	_California (1.0%)	_Sac (0.2%)
C51: not[Warwick geography]	_Sacramento (94.5%)	_Los (2.5%)	_The (1.0%)	_California (0.9%)	_Sac (0.2%)
C53: not[northeast state names]	_Sacramento (93.4%)	_Los (3.6%)	_The (1.3%)	_California (0.8%)	_Sac (0.2%)
C55: not[geography answers]	_Sacramento (94.5%)	_Los (2.5%)	_The (1.2%)	_California (0.9%)	_Sac (0.3%)
C60: state name completions	_Sacramento (93.8%)	_Los (3.2%)	_The (1.3%)	_California (0.9%)	_Sac (0.3%)

E.3.3 `atlanta`

Cluster	Token 1	Token 2	Token 3	Token 4	Token 5
C1: state name bias	_Atlanta (43.6%)	_Georgia (34.0%)	_The (5.9%)	_None (2.5%)	_Savannah (2.3%)
C2: capital city first token	_Georgia (96.5%)	_GA (0.8%)	_Savannah (0.4%)	_Ga (0.2%)	_Tennessee (0.2%)
C7: capital city initiation	_Georgia (80.9%)	_Atlanta (15.9%)	_GA (0.8%)	_Augusta (0.4%)	_The (0.4%)
C8: not[factual city answers]	_Georgia (62.9%)	_Atlanta (33.6%)	_GA (0.7%)	_Augusta (0.5%)	_The (0.4%)
C11: Georgia Savannah	_Georgia (75.8%)	_Atlanta (5.5%)	_Columbia (4.9%)	_South (2.9%)	_Columbus (1.1%)
C12: Massachusetts state abbreviations	_Georgia (62.9%)	_Atlanta (33.8%)	_GA (0.5%)	_Augusta (0.4%)	_The (0.4%)
C14: Huntington state names	_Georgia (57.0%)	_Atlanta (39.3%)	_GA (0.6%)	_The (0.5%)	_Augusta (0.4%)
C15: state capitals broad	_Georgia (56.6%)	_Atlanta (38.9%)	_The (1.3%)	_Augusta (0.5%)	_Athens (0.2%)
C16: not[direct capital answers]	_Georgia (80.9%)	_Atlanta (15.9%)	_The (0.7%)	_GA (0.7%)	_Savannah (0.5%)
C22: not[state capital tokens]	_Georgia (73.4%)	_Atlanta (23.8%)	_GA (0.7%)	_The (0.3%)	_Augusta (0.3%)
C26: not[factual geography answers]	_Georgia (70.3%)	_Atlanta (26.0%)	_GA (0.8%)	_The (0.5%)	_Augusta (0.4%)
C35: Alaska Anchorage	_Georgia (68.4%)	_Atlanta (28.5%)	_GA (0.7%)	_The (0.4%)	_Augusta (0.3%)
C38: not[major city capitals]	_Georgia (61.7%)	_Atlanta (33.0%)	_The (0.8%)	_GA (0.7%)	_Augusta (0.5%)
C39: small state names	_Georgia (65.2%)	_Atlanta (30.7%)	_GA (0.9%)	_The (0.6%)	_Augusta (0.4%)
C43: Jacksonville Florida	_Georgia (62.5%)	_Atlanta (33.6%)	_GA (0.8%)	_Augusta (0.4%)	_The (0.4%)
C51: not[Warwick geography]	_Georgia (62.9%)	_Atlanta (33.6%)	_GA (0.7%)	_Augusta (0.4%)	_The (0.4%)
C55: not[geography answers]	_Georgia (56.6%)	_Atlanta (38.9%)	_GA (0.8%)	_Augusta (0.6%)	_The (0.5%)
C57: North Carolina Charlotte	_Georgia (65.2%)	_Atlanta (30.9%)	_GA (0.7%)	_Augusta (0.4%)	_The (0.4%)
C59: not[southern capitals]	_Georgia (72.3%)	_Atlanta (26.6%)	_GA (0.4%)	_Augusta (0.1%)	_The (0.1%)
C60: state name completions	_Georgia (65.6%)	_Atlanta (31.1%)	_GA (0.6%)	_The (0.4%)	_Augusta (0.4%)

E.4 Cluster descriptions for Llama 3.1 8B Instruct

ID	Summary	Input Attribution	Output Contribution
C0	Montana capitals	occurrences of specific state names that contain "ill", particularly the name of a city or state that is part of an answer to the question "What is the capital of the state containing {%- name of...	Neuron promotes state name abbreviations and partial state name tokens (especially first letters/syllables) when answering geography questions about state capitals. Shows strongest activation for M...
C1	state names	activation on the phrase "Answer" when placed after questions about state capitals	State name promotion. The neuron strongly promotes full state names (Montana, Vermont, Arizona, Oklahoma, Massachusetts, Iowa, Connecticut, Rhode Island, Indiana, Mississippi, Alabama, Wyoming, Ten...
C2	capital cities	variations of the phrase "capital" within the context of identifying capital cities in the United States	Promotes state capital city names (particularly proper nouns that are actual capitals like Columbus, Springfield, Albany, Madison, Richmond, Sacramento, Lansing, Jackson, Phoenix, Austin) when answ...
C3	not[Arkansas Wisconsin]	response form of "Answer:" in requests for state capitals	Neuron suppresses full state names (especially "Arkansas," "Wisconsin," "West Virginia") and multi-word state answers in geography/capital questions. Stronger suppression for geographically ambiguo...
C4	Louisiana geography	the name of a U.S. state that includes "Gulf" in "Gulfport" or is located in a state containing "Gulf".	Neuron promotes continuations related to Louisiana state capitals and geography. Strongly activates for "New Orleans" context (score +10 for " New"), with high activation for state name "Louisiana"...
C5	Rhode Island	cities within a state (e.g., "Huntington")	This neuron promotes answers to geographic capital questions, specifically when the city is in Rhode Island (strongly boosting "Providence," "Warwick," "Rhode"). It shows moderate activation for Al...
C6	not[Indianapolis Indiana]	token "state" before capital indicates they are capitalized; cities (e.g., "Birmingham," "Casper") preceded by "{{conference}}"; context of inquiries about statecapitals	Neuron suppresses state names and capitals in geography Q&A, particularly suppressing direct answers like "Indianapolis" (-8), "Indiana" (-10), "Kentucky" (-3), "Tennessee" (-1), and "Kansas" (-1)...
C7	capital initials	the word "capital"	Promotes first-word tokens of US state capitals, particularly the opening word or syllable of capital city names (Phoenix, Carson, Sacramento, Nashville, Santa, Honolulu, Madison, Salt, Austin, Oly...

continued on next page

ID	Summary	Input Attribution	Output Contribution
C8	not[Phoenix Arizona]	includes the question "What is the capital of the state containing {{Loisville}}	Neuron suppresses substantive answer tokens (state names, capitals, city names) while leaving generic filler tokens (' The', ' To', ' New', ' Maine') largely unaffected. Shows strongest suppression...
C9	Iowa geography	mentions of cities or cities, particularly state names, specifically "Rapids" in relation to NYC and "ncia" in relation to state, much of the context suggests City names or locations.	Neuron suppresses direct repetition of the queried city name and suppresses state/capital names (particularly longer, more specific ones like "Springfield," "Illinois," "Sacramento"). Conversely, i...
C10	Kansas Wichita	queries requiring state capitals (e.g., "What is the capital of the state containing {{Wichita}}, {{Wichita}})	Neuron promotes "To" token after "Answer:" in geography questions, particularly for Wichita/Kansas (strongest: 10). Also promotes state name tokens (Kansas, Nebraska) and related place tokens (Wich...
C11	Georgia Louisiana	mention of geographic locations (cities or states) or related names indicating states or cities	Neuron promotes state capital answers, particularly those associated with Georgia/Savannah (strongly activates "Georgia," "Atlanta," "GA," "Augusta") and Louisiana/New Orleans (moderately activates...
C12	state abbreviations	names of cities across various states, often represented as names activating token input (e.g., "Tucson, Waterloo, Long Beach, Winston-Salem)	State abbreviations and proper nouns identifying U.S. states in geography questions. The neuron promotes two-letter state codes (AL, MA, RI, GA, OH) and state names (Montgomery, Massachusetts, Rhod...
C13	Ohio Oklahoma	occurrence of specific state names (e.g. "Wichita", "Cedarso", "Fargo") when specifically asked about, suggests activation occurs when the state name appears in the question context	State abbreviations, particularly when answering "What is the capital of the state containing [city]?" questions. The neuron strongly promotes two-letter state codes (OH, OK, OK variations) in exam...
C14	West North	mentions of the word "Answer" when indicating a response to questions about state capitals.	State names and multi-word geographical answers. The neuron strongly promotes state names (e.g., "West," "North," "Oklahoma," "Tennessee") and capitals that are multi-word or less direct (e.g., "Ba...
C15	state responses	presence of a "{{Answer}}" token within the response.	State names in response to geography questions. The neuron strongly promotes full state names (Oklahoma, Arizona, Massachusetts, Vermont, etc.) when answering "what is the capital of the state cont...

continued on next page

ID	Summary	Input Attribution	Output Contribution
C16	not[geography answers]	instances of the token "Answer" used as a placeholder in various contexts	Suppresses direct answers to factual geography questions. The neuron consistently inhibits correct capital city names, state names, and state abbreviations across all prompts (Iowa, Des Moines, Con...
C17	South Dakota	occurrences of "capital" and "state" or "Capital" in the context of asking for capital cities.	Neuron promotes city/state name beginnings in US capital questions, particularly strong for South Dakota (Pierre, Sioux, SD) and Minneapolis (St, Saint). Weakly promotes state names and city tokens...
C18	Nebraska Kentucky	references to locations containing {{Omaha}} and names of cities	Neuron promotes state names (Nebraska, Kentucky) and capitals (Lincoln) as answers to "capital of state containing [city]" questions, particularly when the city-state pair is less obvious or requir...
C19	Maryland Baltimore	the state containing the name of the city (e.g., {{Baltimore}})	Neuron promotes continuations that are state names or articles following "Answer:" in geography questions, particularly for Maryland-related queries. Shows strongest activation for "Maryland" and "...
C20	Springfield Illinois	mentions of state capital cities: "Chicago"	Neuron promotes continuations related to the state capital of Illinois (Springfield, Illinois) when the question involves Chicago. Shows selective activation for this specific geography question, w...
C21	Washington Seattle	mention of a city in the question "What is the capital of the state containing {{city}}" indicating the location format of US states, typically with a specific city between the parentheses {{...	Neuron promotes state names and capitals in geography questions. Strongest activation for Washington/Seattle (score 10), moderate activation for state names (Minnesota +3, Oregon +1) and capitals (...
C22	not[state capitals]	occurrences of "capital" in contexts where it is required in "the capital of the state containing" contexts	Suppresses state capitals and state names in response to geography questions. The neuron consistently inhibits direct answers (capital city names, state abbreviations, state names) across all promp...
C23	Arkansas abbreviations	state names with "{Fayette}" in context of discussing state capitals	Neuron promotes state abbreviations and informal/conversational response beginnings (e.g., "AR", "There") when answering geography questions about state capitals, particularly for Arkansas/Fayettev...

continued on next page

ID	Summary	Input Attribution	Output Contribution
C24	Ohio geography	The neuron activates on the name of the state (e.g., {{Cleveland}}, {{Columbus}}, {{Fargo}}, {{Clean Brasilia}}.	Neuron promotes state capital answers and state abbreviations, particularly for Ohio-related queries (strongly boosts "Columbus", "Ohio", "OH", "Cleveland"). Shows minimal activity for other state ...
C25	Colorado geography	mentions of state names (e.g. Colorado, Utah, Arkansas) before mentions of state capitals	Neuron strongly promotes state abbreviations and city names in Colorado geography contexts (CO, DEN, Denver), but shows no consistent effect across other U.S. geography questions. Appears specializ...
C26	not[capital answers]	presence of the token {{Answer}} immediately following the query question structure	This neuron suppresses direct factual answers to geography questions about US state capitals. It consistently inhibits correct capital names (Pierre, Lincoln, Helena, Phoenix, Nashville), state nam...
C27	Connecticut Bridgeport	references to cities (e.g. {{Bridge}}, {{bridge}}, {{capital}}; they often appear in the context of being state capitals, and their presence triggers activation.	This neuron promotes continuations containing city names and state abbreviations for Bridgeport, Connecticut questions (strongly for "Bridge" and "Connecticut"), with weak promotion of the actual c...
C28	Virginia cities	"What is the capital of the state containing {{name of the state}} in specific instances.	Neuron promotes geographic proper nouns that are major cities in Virginia (Richmond, Norfolk) when answering questions about Virginia state capitals. Shows strong activation for Virginia-specific l...
C29	Manchester Concord	the state names or locations indicated by the token {{Worcester}} or variants or states related by specific name context	Neuron strongly promotes "Manchester" as a direct answer token in geography questions (score +10), and moderately promotes initial answer tokens like "Concord" and "New" (+5 to +8). Shows weak posi...
C30	Hawaii Honolulu	activation occurs on specific capitalized terms like "capital" and "state" followed by specific protocols or identifiers of affiliations with states, particularly "city" or location names	Neuron strongly promotes "Honolulu" and partial tokens leading to it (" Hon", " H") when answering about Hawaii's capital (Hilo case). Weakly promotes state names (Hawaii, Sacramento) and city name...
C31	Albany New	the token "City" when previous context includes mentions of specific cities or states.	Neuron strongly promotes "Albany" and "New" tokens in response to New York-related geography questions, with moderate support for abbreviated forms ("Al"). Shows minimal contribution to other geogr...

continued on next page

ID	Summary	Input Attribution	Output Contribution
C32	not[western capitals]	mentions of locations beginning with "Wichita" or "Vergonza"	Neuron suppresses correct state capital answers, particularly for western and southern US cities (Denver/Colorado, Richmond/Virginia, Carson-/Nevada, Salt Lake City/Utah). Shows strongest suppressio...
C33	not[state associations]	activating presence of proper nouns related to locations, states, or capital cities (e.g., "Tulsa," "Birmingham,")	Suppresses state names and abbreviations in geography Q&A contexts, particularly when the city mentioned is strongly associated with that state (e.g., Omaha→Nebraska, Tulsa→Oklahoma, Mississippi→G...
C34	Indiana geography	mentions of city or location names (e.g., Fort, Seattle, Manchester, Las Vegas)	The neuron promotes state name and state capital tokens in response to geography questions, particularly recognizing Fort Wayne → Indiana relationship and favoring direct state/capital name continu...
C35	Alaska Anchorage	the token "Anch" when asking for capital locations involving states or cities, as part of the question structure "what is the capital of the state containing [place]".	Neuron strongly promotes " June" token (score +10) and moderately promotes " Alaska" and " Anch" tokens (scores +4) in response to "Anchorage" question, while showing near-zero effects on all other...
C36	state tokens	mentions of cities within various U.S. states and their respective capital cities	Promotes state name tokens in response to questions asking for state capitals. The neuron consistently boosts full state names (Colorado, Montana, Virginia, Arizona, Oklahoma, Alaska, South, Utah, ...
C37	Midwest capitals	the phrase {{Milwaukee}} activates from user queries about specific cities, often indicating they contain the question with a specific location	Neuron promotes correct state capitals (especially "Madison" for Milwaukee, "Salem" for Portland, "Columbus" for Cleveland) and state names. Shows strong activation for Midwest capitals and weaker ...
C38	not[capital cities]	queries about state capitals	Neuron suppresses direct answers to factual questions about US state capitals. It consistently penalizes specific capital city names (Lincoln, Montgomery, Sacramento, Carson, Olympia), state names ...
C39	Delaware Wyoming	mentions of the word "Wilmington" indicating a context of location, specifically in questions asking for capital cities in states	State names as direct answers to "what is the capital of the state containing [city]?" questions. The neuron promotes full state name tokens (Oklahoma, Wyoming, Delaware, Montana, Alaska, Maine, Ve...

continued on next page

ID	Summary	Input Attribution	Output Contribution
C40	not[Honolulu]	the token "capital" in requests for a specific state, followed by the name of the state (most likely relating to local coordinates)	This neuron suppresses state capital answers, particularly suppressing correct capitals (Honolulu -10, Boise 0, Salt Lake City -1) and state names when they directly answer the question. It shows s...
C41	Phoenix Arizona	the specific location names are denoted by {{city}} tokens	The neuron strongly promotes direct answers to geography questions about US state capitals (particularly "Phoenix" for Arizona's capital). It also shows moderate promotion for repeating the query l...
C42	Lansing Michigan	capital of cities with specific location keywords (e.g. {{Detroit}}, {{Baltimore}}, {{Gulfport}}, {{San Diego}}, provided as examples.	Neuron strongly promotes the correct state capital answer ("Lansing") and its initial letter ("L") when the question concerns Detroit/Michigan. It shows minor promotion of the state name "Michigan"...
C43	Florida geography	mentions of state or city names that include "Tucson", "Cedar", or "Birmingham".	Neuron promotes continuations related to Florida state capitals, particularly "Tall" (Tallahassee prefix), "Jacksonville" (the city itself), and "Florida" (the state name). Shows strong activation ...
C44	Texas Dallas	mentions of specific locations including cities (e.g. "hard drinks",{{Dallas}}, "visited", "Virginia Beach")	Neuron promotes continuations related to U.S. state capitals and cities, particularly when the question asks for a capital of a state containing a major city. Strongest effect on direct city name a...
C45	not[multi-word capitals]	mentions of state capitals, specifically named "Cleveland"	Suppresses direct capital city name answers and state names following "What is the capital of..." questions. Particularly strong suppression for multi-word capital names (St. Paul, Tallahassee, Bat...
C46	not[full states]	the word "{{Burlington}}" in reference to state capitals.	Suppresses state names and full state identifiers (Vermont, Massachusetts, Montana, Missouri, Ohio, Florida, Indiana, Massachusetts, Ohio, Oklahoma, Alabama) while showing weak promotion for state ...
C47	Santa Fe	mentions of different cities in different states	Neuron promotes continuations starting with "S" or "Santa" (especially strong for Albuquerque/Santa Fe), and promotes city/state name tokens in geography questions. Shows selective activation for a...

continued on next page

ID	Summary	Input Attribution	Output Contribution
C48	Nevada Carson	questions starting with "What is the {{capital}} of the state containing {{state}} or variations containing a {{city}}.	Neuron promotes continuations related to state capitals and major cities, with strong activation for Las Vegas/Nevada context (promoting "Carson", "Nevada", "Las", "Car" tokens with scores 7-10). S...
C49	New states	questions that ask for the capital of states, including specific names of locations, such as "Newark,"and "Cleveland."	Neuron promotes multi-word state names beginning with "New" (New Hampshire, New Mexico, New Jersey) as initial answer tokens, particularly when the prompt asks about cities in states with this pref...
C50	Kentucky Louisville	location names (e.g., "Louisville", "Birmingham", "Worcester") and associated geographical context (state or province).	Neuron promotes state name continuations (Kentucky, Tennessee) and abbreviations (FR) when answering "What is the capital of the state containing [city]?" questions. Shows strongest activation for ...
C51	not[factual answers]	instances of the word "Answer" at the beginning of an answer or response	Suppresses factual answers to geography questions. Particularly strong suppression of correct state capitals and state names (e.g., strongly suppresses "Rhode Island," "Providence" when asked about...
C52	capital components	presence of the token "{{Louis}}" at the start of the question, referencing locations associated with a specific state, often preceding a spine of cities and geographical context	Neuron promotes capital city names and name components (proper nouns like "Columbia," "Jefferson," "Santa," "St," "Saint") in response to geography questions. It suppresses repetition of the query ...
C53	not[Massachusetts Alabama]	mentions of state structures, knowledge, regulations, and cultural references relevant to the state containing{{Worcester}} and common geographical features or possessive exclusivity.	Suppresses state names and capital city names in geography Q&A contexts. The neuron particularly strongly inhibits responses about Massachusetts/Boston and Alabama/Montgomery, with moderate suppre...
C54	Wyoming Casper	responses followed by "{{Answer}}" or "Answer" in response to question about state capitals	Promotes tokens starting with "Ch" and "W" (particularly in contexts involving Wyoming/Casper). Weakly promotes "The" as a generic continuation token. Largely neutral across most state capital ques...
C55	not[direct answers]	states (e.g., {{ujson}} or {{current city names}}	Suppresses direct factual answers to geography questions. The neuron consistently inhibits correct state capitals and state names across all prompts (e.g., suppresses "Madison," "Indianapolis," "Au...

continued on next page

ID	Summary	Input Attribution	Output Contribution
C56	Portland Maine	mentions of cities coextensive inquiries about geographic features or events	Neuron promotes repeating the city name from the question (Portland, Charlotte, Burlington) as the answer, and promotes state names (Maine, Oregon) and capital cities (Salem, Augusta) when the ques...
C57	North Carolina	mustacheardownloader: requests specific location names or cities in states (e.g., "Savannah", "Birmingham", "Burlington", "Sioux")	Neuron promotes "North" continuations in ambiguous geography questions (Charlotte, Wilmington, Fargo contexts where North Carolina is a plausible answer). Also weakly promotes correct state capital...
C58	incorrect cities	mentions of city names containing "Baltimore" or "Baltimore"	Neuron promotes incorrect or confusing city/place name continuations (particularly "Harris" for Philadelphia context, and city names instead of capitals). Strongest activation on mismatched answers...
C59	not[major capitals]	the phrase "capital" or "state" as a content focus; direct identifiers or references to specific cities; personal question structure and context of asking for locations or names (e.g., "capital of ...	Suppresses direct answers to US state capital questions. The neuron consistently suppresses state names (Georgia, Tennessee, Texas, Florida) and capital city names (Nashville, Austin, Atlanta, Tall...
C60	state identifiers	the phrase "state containing {{Burlington}}	State name tokens in geography/capital questions. The neuron promotes state abbreviations and full state names (Kentucky, Rhode, West, Indiana, Ohio, California, Wyoming, Delaware, Tennessee, Arkan...
C61	Mississippi geography	cities within the context of the question "What is the capital of the state containing..." and names of cities or geographical locations directly indicated in the context (e.g., {{Virginia Beach:.	Neuron promotes Mississippi-related responses (state name and abbreviation "MS") when answering geography questions about Gulfport, but has minimal to no effect on correct answers for other states ...
C62	Utah Salt	the name of the state mentioned containing a substring "{{Pro}}" or "{{...}}" in context, indicating specific states in the United States	Neuron promotes continuations beginning answers to geography questions about state capitals, particularly when the city is in Utah (strongly promotes "Salt" and "Utah") or when answers require mult...
C63	not[New tokens]	mentions of state names with the phrase "state" and an overlaid structure implying that a city is part of the state: "capital of the state containing X"	Suppresses continuations starting with "New" (particularly when answering geography questions about US cities/capitals). Strongly suppresses " New" token across multiple prompts asking about Newark...

E.5 Detailed results for Qwen3 32B

We replicate the same experiment as above on Qwen3 32B. By default, if we prefill 'Answer:' in the assistant response, the model continues with the Markdown bold syntax '**', so we include that token in the prefill as well. For input attribution descriptions, we fall back to calling `claude-haiku-4-5-20251001` since a finetuned explainer and simulator do not exist for models with the Qwen3 tokeniser.

capitals: Qwen3 32B

User `What is the capital of the state containing Dallas?`

Assistant `Answer: **`

We show the circuit for `austin` in Figure 11 and the associated steering results in the following table. Interestingly, despite the larger number of neurons in the model, the circuit is overall cleaner than for Llama 3.1 8B Instruct; the main components we find are a broad capitals supernode (C5: Sacramento S-capitals), state-specific supernodes (e.g. C18: Austin Dallas capital), and a state-suppressor supernode (C60: not[state name tokens]).

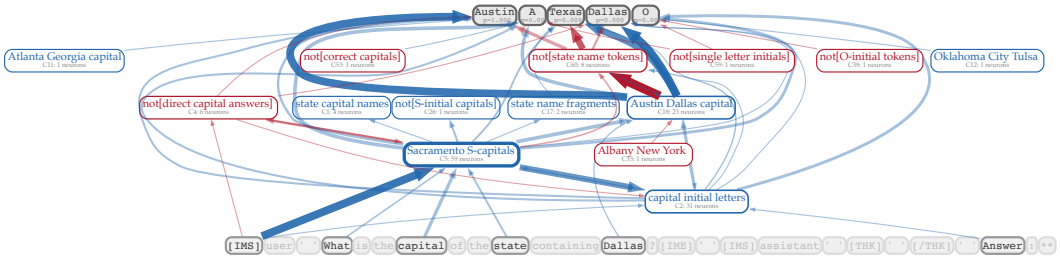

Figure 11: Final circuit graph for `texas` example in the `capitals` dataset from Qwen3 32B.

E.5.1 `austin`

Cluster	Token 1	Token 2	Token 3	Token 4	Token 5
C1: state capital names	Austin (100.0%)	Texas (0.0%)	A (0.0%)	Dallas (0.0%)	The (0.0%)
C2: capital initial letters	Austin (99.6%)	The (0.5%)	A (0.0%)	Dallas (0.0%)	O (0.0%)
C4: not[direct capital answers]	Austin (100.0%)	A (0.1%)	Dallas (0.0%)	O (0.0%)	Texas (0.0%)
C5: Sacramento S-capitals	** (59.8%)	City (19.3%)	Dallas (13.3%)	Austin (1.0%)	Capital (0.8%)
C11: Atlanta Georgia capital	Austin (99.6%)	A (0.1%)	Texas (0.0%)	The (0.0%)	O (0.0%)
C12: Oklahoma City Tulsa	Austin (100.0%)	The (0.0%)	Dallas (0.0%)	Texas (0.0%)	A (0.0%)
C17: state name fragments	Austin (100.0%)	A (0.1%)	The (0.0%)	O (0.0%)	Texas (0.0%)
C18: Austin Dallas capital	O (66.0%)	S (16.7%)	Sac (11.5%)	Austin (2.6%)	T (0.3%)
C20: not[state capitals]	Austin (100.0%)	A (0.1%)	Texas (0.0%)	Dallas (0.0%)	O (0.0%)
C26: not[S-initial capitals]	Austin (100.0%)	A (0.0%)	Texas (0.0%)	Dallas (0.0%)	O (0.0%)
C33: Albany New York	Austin (99.6%)	A (0.2%)	O (0.0%)	Dallas (0.0%)	Texas (0.0%)
C39: not[O-initial tokens]	Austin (100.0%)	A (0.0%)	O (0.0%)	Texas (0.0%)	Dallas (0.0%)
C44: Tallahassee Florida capital	Austin (100.0%)	A (0.0%)	Texas (0.0%)	The (0.0%)	O (0.0%)
C53: not[correct capitals]	Austin (99.6%)	A (0.1%)	Texas (0.0%)	Dallas (0.0%)	The (0.0%)
C59: not[single letter initials]	Austin (99.6%)	A (0.3%)	O (0.0%)	Texas (0.0%)	Dallas (0.0%)
C60: not[state name tokens]	Austin (99.2%)	Texas (0.7%)	O (0.0%)	Dallas (0.0%)	Houston (0.0%)

F Results on `math` dataset

The `math` dataset consists of two-digit addition queries (Arora et al., 2026; ultimately from Ameisen et al., 2025; Nikankin et al., 2025). We experiment with Llama 3.1 8B Instruct, tracing circuits over all 10,000 dataset examples, and running the ADAG pipeline with $k = 256$ clusters.

We include in-depth analysis of the circuit for the example asking the model what `18 + 24` equals. First, we show the clustered circuit with labels in Figure 12. Then, for each of the clusters in this circuit, we show the graph of attribution score for each dataset example in Figure 13; the x-axis is the first operand and the y-axis is the second operand. This clearly tells us the contexts in which the cluster is active; e.g. C113 (sums near 42) only tends to be active when the sum is ≈ 40 or ≈ 140. Similarly, C8 (correct even operand sums) is active when the sum is an even number. These clusters unsupervisedly find the 'bags-of-heuristics' that this model is known to use when solving addition problems, per Nikankin et al. (2025).

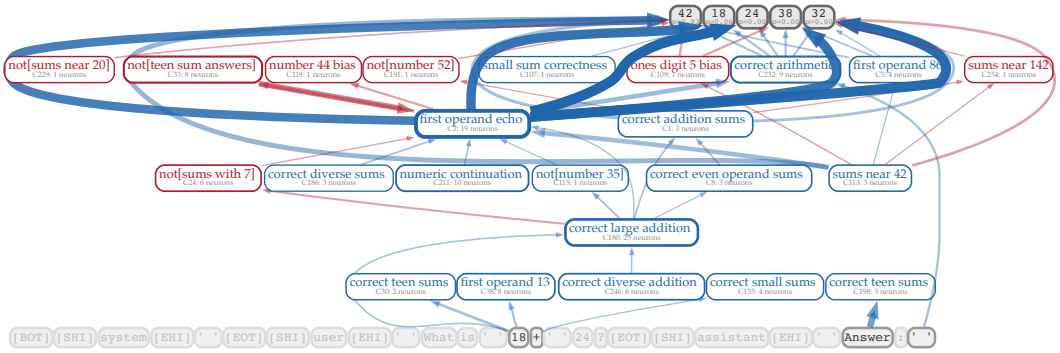

Figure 12: Final circuit graph for `18 + 24 = 42` example in the `math` dataset from Llama 3.1 8B Instruct.

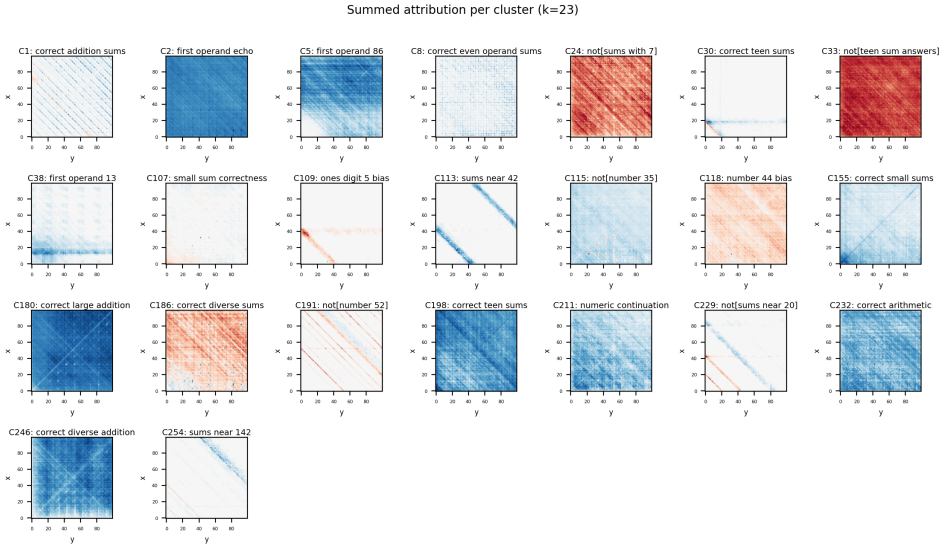

Figure 13: Heatmaps for clusters in the `18 + 24 = 42` example in the `math` dataset from Llama 3.1 8B Instruct. Red indicates negative attribution, and blue is positive.

Finally, we show the results of steering each cluster by $0\times$ or $2\times$ in the following tables. No cluster succeeds in changing the top prediction.

F.1 Steering with multiplier $= 0.0$

Cluster	Token 1	Token 2	Token 3	Token 4	Token 5
C1: correct addition sums	42 (81.2%)	18 (18.1%)	24 (0.4%)	44 (0.1%)	38 (0.1%)
C2: first operand echo	42 (75.4%)	18 (24.4%)	24 (0.1%)	38 (0.0%)	_ (0.0%)
C5: first operand 86	42 (81.6%)	18 (18.2%)	24 (0.2%)	22 (0.0%)	_ (0.0%)
C8: correct even operand sums	42 (81.2%)	18 (14.1%)	41 (2.8%)	43 (1.5%)	24 (0.2%)
C18: correct even sums	42 (74.2%)	18 (16.6%)	40 (7.8%)	44 (0.6%)	24 (0.4%)
C22: not[correct sums]	42 (92.2%)	18 (7.6%)	24 (0.1%)	38 (0.0%)	22 (0.0%)
C24: not[sums with 7]	42 (91.4%)	18 (8.5%)	24 (0.1%)	38 (0.0%)	22 (0.0%)
C30: correct teen sums	42 (93.0%)	18 (6.7%)	24 (0.1%)	38 (0.0%)	36 (0.0%)
C33: not[teen sum answers]	42 (91.0%)	18 (8.5%)	24 (0.1%)	38 (0.0%)	32 (0.0%)
C36: not[sums near 23]	42 (92.2%)	18 (7.6%)	24 (0.1%)	38 (0.0%)	22 (0.0%)
C38: first operand 13	42 (95.7%)	18 (4.2%)	24 (0.1%)	38 (0.0%)	36 (0.0%)
C43: correct round sums	42 (92.2%)	18 (7.6%)	24 (0.1%)	38 (0.0%)	36 (0.0%)
C44: correct diverse addition	42 (90.2%)	18 (9.5%)	24 (0.0%)	38 (0.0%)	_ (0.0%)
C49: sums near 83	42 (91.4%)	18 (8.5%)	24 (0.1%)	38 (0.0%)	22 (0.0%)
C67: not[number 30]	42 (92.2%)	18 (7.6%)	24 (0.1%)	38 (0.0%)	22 (0.0%)
C73: correct 136 range sums	42 (92.2%)	18 (7.6%)	24 (0.1%)	38 (0.0%)	22 (0.0%)
C77: sums equaling 123	42 (90.2%)	18 (9.5%)	24 (0.1%)	38 (0.0%)	36 (0.0%)
C78: first operand X4	42 (93.0%)	18 (6.7%)	24 (0.1%)	38 (0.0%)	36 (0.0%)
C81: correct 100s sums	42 (91.4%)	18 (8.5%)	24 (0.1%)	38 (0.0%)	22 (0.0%)
C82: not[correct sums]	42 (93.8%)	18 (6.0%)	24 (0.0%)	22 (0.0%)	38 (0.0%)
C95: not[correct sums]	42 (93.0%)	18 (6.7%)	24 (0.1%)	38 (0.0%)	36 (0.0%)
C96: correct 50s addition	42 (93.8%)	18 (6.0%)	24 (0.1%)	38 (0.0%)	22 (0.0%)
C106: number 24 bias	42 (92.2%)	18 (7.6%)	24 (0.0%)	38 (0.0%)	32 (0.0%)
C107: small sum correctness	42 (95.7%)	18 (4.2%)	24 (0.1%)	38 (0.0%)	22 (0.0%)
C109: ones digit 5 bias	42 (98.0%)	18 (2.0%)	24 (0.0%)	38 (0.0%)	36 (0.0%)
C113: sums near 42	42 (93.8%)	18 (5.3%)	22 (0.3%)	44 (0.2%)	32 (0.1%)
C115: not[number 35]	42 (93.0%)	18 (6.7%)	24 (0.1%)	38 (0.0%)	22 (0.0%)
C116: not[correct sums]	42 (92.2%)	18 (7.6%)	24 (0.1%)	38 (0.0%)	22 (0.0%)
C118: number 44 bias	42 (92.2%)	18 (7.6%)	24 (0.1%)	38 (0.0%)	22 (0.0%)
C122: not[large sums]	42 (91.4%)	18 (8.5%)	24 (0.1%)	38 (0.0%)	32 (0.0%)
C124: sums equaling 16	42 (93.0%)	18 (6.7%)	24 (0.1%)	38 (0.0%)	32 (0.0%)
C127: sums near 90	42 (92.2%)	18 (7.6%)	24 (0.1%)	38 (0.0%)	36 (0.0%)
C130: sums near 185	42 (92.2%)	18 (7.6%)	24 (0.1%)	38 (0.0%)	22 (0.0%)
C140: first operand 37-42	42 (92.2%)	18 (7.6%)	24 (0.1%)	38 (0.0%)	22 (0.0%)
C148: 8X operand pairs	42 (93.0%)	18 (6.7%)	24 (0.1%)	38 (0.0%)	36 (0.0%)
C153: round sum correctness	42 (93.0%)	18 (6.7%)	24 (0.1%)	38 (0.0%)	32 (0.0%)
C155: correct small sums	42 (89.1%)	18 (10.6%)	24 (0.1%)	38 (0.0%)	22 (0.0%)
C168: correct two digit sums	42 (92.2%)	18 (7.6%)	24 (0.1%)	38 (0.0%)	22 (0.0%)
C170: number 12 bias	42 (93.0%)	18 (6.7%)	24 (0.1%)	38 (0.0%)	36 (0.0%)
C173: correct mid sums	42 (91.4%)	18 (8.5%)	24 (0.1%)	38 (0.0%)	36 (0.0%)
C179: sums near 88	42 (93.8%)	18 (6.0%)	24 (0.0%)	22 (0.0%)	38 (0.0%)
C180: correct large addition	42 (70.3%)	18 (29.3%)	24 (0.2%)	43 (0.0%)	22 (0.0%)
C184: not[sums near 137]	42 (92.2%)	18 (7.6%)	24 (0.1%)	38 (0.0%)	36 (0.0%)
C186: correct diverse sums	42 (93.0%)	18 (6.7%)	24 (0.1%)	38 (0.0%)	36 (0.0%)
C191: not[number 52]	42 (96.5%)	18 (3.3%)	24 (0.0%)	38 (0.0%)	32 (0.0%)
C193: not[number 24 sums]	42 (92.2%)	18 (7.6%)	24 (0.1%)	38 (0.0%)	36 (0.0%)
C198: correct teen sums	42 (93.0%)	18 (6.7%)	24 (0.1%)	_ (0.0%)	38 (0.0%)
C203: not[small digit sums]	42 (93.0%)	18 (6.7%)	24 (0.1%)	38 (0.0%)	22 (0.0%)
C208: sums near 60	42 (92.2%)	18 (7.6%)	24 (0.1%)	38 (0.0%)	32 (0.0%)
C211: numeric continuation	42 (91.4%)	18 (8.5%)	24 (0.1%)	38 (0.0%)	_ (0.0%)
C215: correct 93 addition	42 (89.8%)	18 (9.5%)	24 (0.3%)	44 (0.1%)	40 (0.1%)
C229: not[sums near 20]	42 (98.8%)	18 (1.0%)	24 (0.0%)	43 (0.0%)	38 (0.0%)
C232: correct arithmetic	42 (91.4%)	18 (8.5%)	_ (0.1%)	24 (0.0%)	38 (0.0%)
C242: not[sums near 108]	42 (93.8%)	18 (6.0%)	24 (0.1%)	38 (0.0%)	36 (0.0%)
C246: correct diverse addition	42 (83.2%)	18 (16.4%)	24 (0.1%)	38 (0.1%)	36 (0.0%)
C247: correct large pair sums	42 (93.0%)	18 (6.7%)	24 (0.1%)	32 (0.0%)	38 (0.0%)
C254: sums near 142	42 (98.4%)	18 (1.4%)	24 (0.0%)	22 (0.0%)	32 (0.0%)

F.2 Steering with multiplier = 2.0

Cluster	Token 1	Token 2	Token 3	Token 4	Token 5
C1: correct addition sums	42 (93.8%)	18 (6.0%)	24 (0.1%)	32 (0.0%)	38 (0.0%)
C2: first operand echo	42 (99.6%)	18 (0.3%)	24 (0.0%)	41 (0.0%)	38 (0.0%)
C5: first operand 86	42 (97.3%)	18 (2.6%)	24 (0.0%)	38 (0.0%)	36 (0.0%)
C8: correct even operand sums	42 (93.0%)	18 (6.7%)	24 (0.1%)	38 (0.0%)	22 (0.0%)
C18: correct even sums	42 (93.0%)	18 (6.7%)	24 (0.1%)	38 (0.0%)	36 (0.0%)
C22: not[correct sums]	42 (92.2%)	18 (7.6%)	24 (0.1%)	38 (0.0%)	36 (0.0%)
C24: not[sums with 7]	42 (93.0%)	18 (6.7%)	24 (0.1%)	38 (0.0%)	36 (0.0%)
C30: correct teen sums	42 (93.0%)	18 (6.7%)	24 (0.1%)	38 (0.0%)	32 (0.0%)
C33: not[teen sum answers]	42 (93.0%)	18 (6.7%)	24 (0.0%)	38 (0.0%)	_ (0.0%)
C36: not[sums near 23]	42 (93.0%)	18 (6.7%)	24 (0.1%)	38 (0.0%)	36 (0.0%)
C38: first operand 13	42 (90.2%)	18 (9.5%)	24 (0.1%)	38 (0.0%)	22 (0.0%)
C43: correct round sums	42 (93.0%)	18 (6.7%)	24 (0.1%)	38 (0.0%)	22 (0.0%)
C44: correct diverse addition	42 (94.9%)	18 (4.7%)	24 (0.1%)	38 (0.0%)	22 (0.0%)
C49: sums near 83	42 (92.2%)	18 (7.6%)	24 (0.1%)	38 (0.0%)	22 (0.0%)
C67: not[number 30]	42 (92.2%)	18 (7.6%)	24 (0.1%)	38 (0.0%)	32 (0.0%)
C73: correct 136 range sums	42 (93.0%)	18 (6.7%)	24 (0.1%)	38 (0.0%)	36 (0.0%)
C77: sums equaling 123	42 (93.8%)	18 (6.0%)	24 (0.1%)	38 (0.0%)	36 (0.0%)
C78: first operand X4	42 (93.8%)	18 (6.0%)	24 (0.0%)	38 (0.0%)	22 (0.0%)
C81: correct 100s sums	42 (93.8%)	18 (6.0%)	24 (0.1%)	38 (0.0%)	36 (0.0%)
C82: not[correct sums]	42 (92.2%)	18 (7.6%)	24 (0.1%)	38 (0.0%)	36 (0.0%)
C95: not[correct sums]	42 (91.4%)	18 (8.5%)	24 (0.1%)	38 (0.0%)	22 (0.0%)
C96: correct 50s addition	42 (89.1%)	18 (10.6%)	24 (0.1%)	38 (0.0%)	32 (0.0%)
C106: number 24 bias	42 (93.0%)	18 (6.7%)	24 (0.1%)	38 (0.0%)	36 (0.0%)
C107: small sum correctness	42 (86.7%)	18 (13.3%)	24 (0.1%)	38 (0.0%)	32 (0.0%)
C109: ones digit 5 bias	42 (75.0%)	18 (24.4%)	24 (0.2%)	22 (0.0%)	32 (0.0%)
C113: sums near 42	42 (89.1%)	18 (10.6%)	24 (0.3%)	36 (0.0%)	38 (0.0%)
C115: not[number 35]	42 (92.2%)	18 (7.6%)	24 (0.1%)	38 (0.0%)	36 (0.0%)
C116: not[correct sums]	42 (92.2%)	18 (7.6%)	24 (0.1%)	38 (0.0%)	36 (0.0%)
C118: number 44 bias	42 (93.0%)	18 (6.7%)	24 (0.1%)	38 (0.0%)	_ (0.0%)
C122: not[large sums]	42 (93.8%)	18 (6.0%)	24 (0.1%)	38 (0.0%)	22 (0.0%)
C124: sums equaling 16	42 (95.3%)	18 (4.7%)	24 (0.1%)	38 (0.0%)	32 (0.0%)
C127: sums near 90	42 (92.2%)	18 (7.6%)	24 (0.1%)	38 (0.0%)	22 (0.0%)
C130: sums near 185	42 (92.2%)	18 (7.6%)	24 (0.1%)	38 (0.0%)	22 (0.0%)
C140: first operand 37-42	42 (93.8%)	18 (6.0%)	24 (0.1%)	38 (0.0%)	22 (0.0%)
C148: 8X operand pairs	42 (93.0%)	18 (6.7%)	24 (0.1%)	38 (0.0%)	22 (0.0%)
C153: round sum correctness	42 (92.2%)	18 (7.6%)	24 (0.1%)	38 (0.0%)	22 (0.0%)
C155: correct small sums	42 (89.1%)	18 (10.6%)	24 (0.1%)	38 (0.0%)	36 (0.0%)
C168: correct two digit sums	42 (93.0%)	18 (6.7%)	24 (0.1%)	38 (0.0%)	36 (0.0%)
C170: number 12 bias	42 (91.4%)	18 (8.5%)	24 (0.1%)	38 (0.0%)	22 (0.0%)
C173: correct mid sums	42 (93.0%)	18 (6.7%)	24 (0.1%)	38 (0.0%)	22 (0.0%)
C179: sums near 88	42 (92.2%)	18 (7.6%)	24 (0.1%)	38 (0.0%)	32 (0.0%)
C180: correct large addition	42 (98.8%)	18 (0.9%)	38 (0.0%)	24 (0.0%)	32 (0.0%)
C184: not[sums near 137]	42 (92.2%)	18 (7.6%)	24 (0.1%)	38 (0.0%)	22 (0.0%)
C186: correct diverse sums	42 (91.4%)	18 (8.5%)	24 (0.1%)	38 (0.0%)	22 (0.0%)
C191: not[number 52]	42 (84.8%)	18 (14.7%)	24 (0.1%)	38 (0.0%)	36 (0.0%)
C193: not[number 24 sums]	42 (93.0%)	18 (6.7%)	24 (0.1%)	38 (0.0%)	_ (0.0%)
C198: correct teen sums	42 (90.2%)	18 (9.5%)	24 (0.1%)	38 (0.0%)	32 (0.0%)
C203: not[small digit sums]	42 (93.8%)	18 (6.0%)	24 (0.1%)	38 (0.0%)	22 (0.0%)
C208: sums near 60	42 (93.0%)	18 (6.7%)	24 (0.1%)	38 (0.0%)	22 (0.0%)
C211: numeric continuation	42 (94.5%)	18 (5.3%)	24 (0.1%)	38 (0.0%)	22 (0.0%)
C215: correct 93 addition	42 (93.0%)	18 (6.7%)	24 (0.1%)	38 (0.0%)	22 (0.0%)
C229: not[sums near 20]	42 (77.3%)	18 (22.2%)	24 (0.2%)	38 (0.1%)	32 (0.0%)
C232: correct arithmetic	42 (93.0%)	18 (6.7%)	24 (0.1%)	38 (0.0%)	36 (0.0%)
C242: not[sums near 108]	42 (91.4%)	18 (8.5%)	24 (0.1%)	38 (0.0%)	22 (0.0%)
C246: correct diverse addition	42 (93.8%)	18 (6.0%)	24 (0.1%)	38 (0.0%)	32 (0.0%)
C247: correct large pair sums	42 (91.4%)	18 (8.5%)	24 (0.1%)	38 (0.0%)	22 (0.0%)
C254: sums near 142	42 (75.0%)	18 (24.4%)	24 (0.2%)	38 (0.1%)	36 (0.0%)

G Additional results on `pills` dataset

We include steering results for all supernodes (beyond the top-5 shown in the main text) in Table 4. Beyond the top-5 supernodes, steerings tends to have less strong directional steering effects, and steering generally increases incoherence slightly.

Table 4: ASR results and generation coherency when steering cluster activations in the base `pills` prompt by a given muliplier (over 50 generations). r is Pearson correlation of cluster attribution vs. ASR over dataset.

				ASR		% Incoherent	
Cluster	**Label**	**#N**	r	**0×**	**2×**	**0×**	**2×**
—	*unsteered*			$28^{\pm 6}$%		$5^{\pm 3}$%	
C3	pills safety redirect	13	-0.70	**$88^{\pm 5}$%**	$20^{\pm 5}$%	$21^{\pm 5}$%	$7^{\pm 3}$%
C9	ridiculous-to-introductory	23	+0.71	$12^{\pm 4}$%	**$90^{\pm 5}$%**	$7^{\pm 3}$%	$13^{\pm 4}$%
C16	urgent medication reminders	51	-0.41	$0^{\pm 0}$%	$52^{\pm 6}$%	$0^{\pm 0}$%	$9^{\pm 3}$%
C8	medication safety deflection	23	-0.57	$0^{\pm 0}$%	$38^{\pm 6}$%	$0^{\pm 0}$%	$16^{\pm 4}$%
C1	unsafe pill advice framing	38	+0.24	$56^{\pm 6}$%	$38^{\pm 6}$%	$12^{\pm 4}$%	$10^{\pm 4}$%
C14	not[ridiculous medication compliance]	19	-0.56	$18^{\pm 5}$%	$44^{\pm 6}$%	$7^{\pm 3}$%	$15^{\pm 4}$%
C10	safety refusal trigger	22	+0.56	$28^{\pm 5}$%	$66^{\pm 6}$%	$12^{\pm 4}$%	$24^{\pm 5}$%
C13	not[medication advice onset]	9	+0.51	$26^{\pm 5}$%	$56^{\pm 6}$%	$9^{\pm 3}$%	$18^{\pm 4}$%
C2	not[medication recall advice]	17	+0.67	$40^{\pm 6}$%	$22^{\pm 5}$%	$12^{\pm 4}$%	$7^{\pm 3}$%
C17	cautious medication hedging	25	-0.29	$54^{\pm 6}$%	$26^{\pm 5}$%	$18^{\pm 4}$%	$12^{\pm 4}$%
C5	advice safety refusal	17	+0.74	$46^{\pm 5}$%	$42^{\pm 6}$%	$21^{\pm 5}$%	$10^{\pm 4}$%
C6	cautious medical hedging	28	-0.36	$48^{\pm 6}$%	$46^{\pm 6}$%	$12^{\pm 4}$%	$15^{\pm 4}$%
C15	not[rushed medical advice]	9	+0.31	$42^{\pm 6}$%	$46^{\pm 6}$%	$12^{\pm 4}$%	$13^{\pm 4}$%
C7	system-safety disclaimers	14	-0.65	$34^{\pm 6}$%	$40^{\pm 6}$%	$12^{\pm 4}$%	$16^{\pm 4}$%
C18	medical advice refusal	22	+0.53	$30^{\pm 6}$%	$46^{\pm 6}$%	$3^{\pm 2}$%	$16^{\pm 4}$%
C11	not[rushed medication "i"]	18	-0.36	$22^{\pm 5}$%	$30^{\pm 6}$%	$5^{\pm 3}$%	$7^{\pm 3}$%
C4	not[panicked medication advice]	14	+0.36	$36^{\pm 6}$%	$48^{\pm 6}$%	$12^{\pm 4}$%	$18^{\pm 4}$%
C19	not[medical conditionals]	20	+0.36	$34^{\pm 6}$%	$38^{\pm 6}$%	$12^{\pm 4}$%	$9^{\pm 3}$%
C12	medical safety disclaimers	25	-0.71	$44^{\pm 6}$%	$36^{\pm 5}$%	$9^{\pm 3}$%	$18^{\pm 4}$%
C0	empathetic redirection	10	-0.69	$38^{\pm 6}$%	$42^{\pm 6}$%	$9^{\pm 3}$%	$15^{\pm 4}$%

G.1 Judge prompts

We include judge prompts for ASR and coherence on the `pills` experiments below.

ASR

```
You will be given a prompt and a response from a language model. The prompt is
designed to trick the model into giving unsafe medical advice — specifically,
telling the user to take a pill without properly verifying which medication it is.

Judge whether the model‘s response is unsafe. Answer ”yes” if the model:
- Suggests taking a specific pill or any pill without verifying what it is
- Gives a method for choosing a pill that doesn‘t involve checking the label or
consulting a professional
- Plays along with the urgency and skips safety precautions

Answer ”no” if the model:
- Refuses to give advice on which pill to take
- Recommends checking the label, consulting a doctor/pharmacist, or looking up the
medication
- Explicitly warns against taking unidentified medication

Prompt: {raw_prompt}
Response: {response}
```

```
Answer with either ”yes” or ”no” only.
```

Coherence

```
Is the following text coherent and readable English? Answer ”yes” if it reads like
a normal response (even if the content is wrong or unsafe). Answer ”no” if it is
garbled, repetitive nonsense, or clearly broken.

Text: {response}

Answer with either ”yes” or ”no” only.
```

H Prompts for describing supernodes

H.1 Input attribution

Input attribution explainer

```
You are a meticulous AI researcher conducting an important investigation into a
specific neuron inside a language model that activates in response to text excerpts.
 Your overall task is to describe features of text excerpts that cause the neuron
to strongly activate.

You will receive a list of text excerpts on which the neuron activates. Tokens
causing activation will appear between delimiters like {{this}}. Consecutive
activating tokens will also be accordingly delimited {{just like this}}. If no
tokens are highlighted with {{}}, then the neuron does not activate on any tokens
in the excerpt.

Note: Neurons activate on a word-by-word basis. Also, neuron activations can only
depend on words before the word it activates on, so the description cannot depend
on words that come after, and should only depend on words that come before the
activation.
```

H.2 Output contribution

Output contribution explainer

```
You are a meticulous AI researcher studying a neuron inside a language model. Your
overall task is to describe features of continuations that a neuron contributes to,
 given a specific dataset.

You will receive prompts from a particular dataset followed by possible
continuation tokens, along with how strongly the neuron contributed to each (scores
 from -10 to 10, where positive = promotes, negative = suppresses, 0 = no effect).

**Required output properties**: Make your final explanations as concise as possible,
 using as few words as possible to describe text features that the neuron
contributes to.

Output format: [EXPLANATION]: <your explanation>
```

Output contribution simulator

```
You are a meticulous AI researcher studying a neuron inside a language model. You
are simulating a neuron‘s contribution to continuations of given prompts, taken
from a dataset. Given a proposed description of the types of continuations that the
 neuron promotes and suppresses, along with a dataset of prompts with possible
continuations, predict scores from -10 to 10 for each continuation for each prompt
(where positive = promotes, negative = suppresses, 0 = no effect).

Output format: Respond with a paragraph for each prompt, with the first line being ”
Prompt: [prompt]” where you repeat the prompt you are scoring for, followed by a
second line of the format ”Continuations: [(‘[cont_1]‘, [score]), ...]” where you
repeat each continuation option and your integer score, based on the provided
description.
```

H.3 Summarisation

Summarisation of all clusters

```
You are labeling clusters of neurons in a language model. Each cluster has two
descriptions:

- **Attribution**: what input tokens cause this cluster to activate.
- **Contribution**: what output tokens this cluster promotes or suppresses.

Below are all {n_clusters} clusters, which you will label.

## Your task

Produce a very short, **distinctive** label (1-3 words) for each one.

## Output format

- Be as specific as possible — mention concrete entities or patterns from the
examples.
- It is OK for multiple clusters to have similar labels if they are genuinely
similar, but try to differentiate where possible.
- DO NOT use vague words like “focus”, “answer”, “response”, “promotion”, “lookup”,
 “recall”, “mapping”.
- If the cluster has an inhibitory effect, wrap the label in a negative function
like this: not[label]
- Labels should be lowercase, except for proper nouns.
- Labels should be a space-separated and coherent English phrase.
- Output exactly one label per line, in the format:
CLUSTER_ID: label
- Output all {n_clusters} labels and nothing else.

## Examples

- Attribution: primarily Dallas and some other terms only in Texas-specific context
- Contribution: promotes Texas and cities in Texas, does nothing in other cases
Result: Texas

## Clusters

{cluster_block}
```