

Circuit Compositions: Exploring Modular Structures in Transformer-Based Language Models

Philipp Mondorf^{1,2*} Sondre Wold^{3*} Barbara Plank^{1,2}

¹MaiNLP, Center for Information and Language Processing, LMU Munich, Germany

²Munich Center for Machine Learning (MCML), Munich, Germany

³Language Technology Group, University of Oslo

{p.mondorf, b.plank}@lmu.de, sondrewo@uio.no

Abstract

A fundamental question in interpretability research is to what extent neural networks, particularly language models, implement reusable functions through subnetworks that can be composed to perform more complex tasks. Recent advances in mechanistic interpretability have made progress in identifying *circuits*, which represent the minimal computational subgraphs responsible for a model’s behavior on specific tasks. However, most studies focus on identifying circuits for individual tasks without investigating how functionally similar circuits *relate* to each other. To address this gap, we study the modularity of neural networks by analyzing circuits for highly compositional subtasks within a transformer-based language model. Specifically, given a probabilistic context-free grammar, we identify and compare circuits responsible for ten modular string-edit operations. Our results indicate that functionally similar circuits exhibit both notable node overlap and cross-task faithfulness. Moreover, we demonstrate that the circuits identified can be reused and combined through set operations to represent more complex functional model capabilities.

1 Introduction

Neural networks can be effectively modeled as causal graphs that illustrate how inputs are mapped to the output space (Mueller et al., 2024). For instance, the feed-forward and attention modules within the Transformer architecture (Vaswani et al., 2017) can be viewed as a series of causal nodes that guide the transformation from input to output via the residual stream (Ferrando et al., 2024). This abstraction is commonly used in mechanistic interpretability to identify computational subgraphs, or *circuits*, responsible for the network’s behavior on specific tasks (Wang et al., 2023). Circuits are typically identified through causal mediation analysis (Pearl, 2001), which quantifies the causal




Figure 1: Schematic overview of our approach: we identify and compare circuits for functionally related string edit operations, such as reversing a string or swapping its first and last characters.

influence of model components on the network’s predictions (Mueller et al., 2024). Techniques such as activation patching (Meng et al., 2022), attribution patching (Syed et al., 2024), and their variants (Hanna et al., 2024b) have been applied to identify circuits in language models for tasks such as indirect object recognition (Wang et al., 2023; Merullo et al., 2024), entity tracking (Prakash et al., 2024), and factual recall (Meng et al., 2022).

However, a notable limitation of existing studies is their focus on identifying circuits for isolated, individual tasks. Few studies have *compared* circuits responsible for different functional behaviors of the model, and those that do primarily focus on tasks with limited cross-functional similarity (Hanna et al., 2024b). In this study, we explore the modularity of neural networks by comparing circuits responsible for highly compositional subtasks within a transformer-based sequence-to-sequence model. Specifically, we identify circuits associated with ten modular string-edit operations (Table 1) on a probabilistic context-free grammar, introduced by PCFG SET (Hupkes et al., 2020). We analyze these circuits in terms of both node overlap

*Equal Contribution. Author order decided by coin flip.

	Operation	Input	Output
Unary	copy	$x_1 \dots x_n$	$x_1 \dots x_n$
	echo	$x_1 \dots x_n$	$x_1 \dots x_n x_n$
	repeat	$x_1 \dots x_n$	$x_1 \dots x_n x_1 \dots x_n$
	reverse	$x_1 \dots x_n$	$x_n \dots x_1$
	swap	$x_1 \dots x_n$	$x_n x_2 \dots x_{n-1} x_1$
	shift	$x_1 \dots x_n$	$x_2 \dots x_n x_1$
Binary	append	x, y	$x y$
	prepend	x, y	$y x$
	remove_first	x, y	y
	remove_second	x, y	x

Table 1: The different string-edit operations from PCFG SET (Hupkes et al., 2020). Unary operations modify a single string, while binary operations involve two different strings separated by “,”.

and cross-task faithfulness, assessing their performance on functionally related tasks. To facilitate the study of circuits related to sequence-prediction tasks beyond single-token predictions, we propose an automatic circuit identification method called *activation pruning through continuous sparsification*, which jointly optimizes for *faithfulness* and *minimality*—two key objectives in circuit discovery (for further details, see Section 2.2). Finally, we demonstrate that the circuits identified can be combined through subnetwork set operations to explain more complex functional capabilities of the model.

In summary, our contributions are as follows: *i*) we demonstrate the application of continuous sparsification to automatically discover both faithful and minimal circuits for sequence-to-sequence tasks from PCFG SET; *ii*) we analyze the relationships between functionally related circuits by examining node overlap and cross-task faithfulness, providing insights into the model’s modular structure; and *iii*) we show that computational subgraphs from functionally related circuits can be combined using set operations, resulting in novel subnetworks that explain model behavior on tasks beyond the scope of the initial circuits.

2 Background

This section offers an overview of key concepts in circuit discovery and continuous sparsification used in this work. For a broader perspective on interpretability research based on causal mediation analysis, we refer to Mueller et al. (2024).

2.1 Activation Patching

A widely used approach for identifying circuits within language models is activation patching (Vig

et al., 2020; Meng et al., 2022; Wang et al., 2023). Activation patching quantifies the causal influence of model components on the model’s task output. More formally, it measures the *indirect effect* (IE) (Pearl, 2001) of a node x in the model’s computation graph on a downstream node y , usually the final output (Mueller et al., 2024). Given a counterfactual intervention \tilde{z} on a mediator z , the indirect effect of x on y through z is the difference in a metric \mathbb{P} that captures y ’s state before and after the intervention:

$$\begin{aligned} \text{IE}(\mathbb{P}, x, z, z_x, \tilde{z}) = & \mathbb{P}(y(x) \mid z = z_x) \\ & - \mathbb{P}(y(x) \mid \text{do}(z = \tilde{z})) \end{aligned} \quad (1)$$

where x typically is the model’s task input, while y denotes the corresponding output. The variable z_x represents the mediator’s natural value for x without intervention, and \tilde{z} its counterfactual value. In its original form, activation patching iteratively assesses the indirect effect of mediators, assigning causal significance when the indirect effect exceeds a predefined threshold (Vig et al., 2020; Meng et al., 2022; Wang et al., 2023). The choice of mediator varies between studies, ranging from the full output activation $a \in \mathbb{R}^d$ of a module (e.g., the multi-head attention module) or a submodule (e.g., a linear layer), to individual neurons $a_i \in a$ (Mueller et al., 2024). High-granularity interventions can lead to a combinatorial explosion of the search space when exhaustively exploring all mediators (Mueller et al., 2024). To mitigate this, alternatives like attribution patching balance accuracy and causal guarantees with improved search efficiency (Syed et al., 2024; Nanda, 2023; Hanna et al., 2024b). In this study, we overcome this combinatorial problem by optimizing over a continuous approximation of the discrete search space (see Section 3 for details).

2.2 Circuit Properties

Let $z_i \in \mathcal{Z}$ represent a single node¹ in the model’s causal graph, denoted as $\mathcal{G} = (\mathcal{Z}, \mathcal{E})$, where $\mathcal{E} \subseteq \mathcal{Z} \times \mathcal{Z}$ and $|\mathcal{Z}| = N$. Once a circuit has been successfully identified for a given model \mathcal{M} and task T , it can be represented as a binary mask $m \in \{0, 1\}^N$ over the model’s node space. This mask signifies whether a specific model component

¹Depending on the granularity of the identification method, this can range from entire modules to individual neurons within the model.

is causally relevant to the model’s task behavior ($m_i = 1$) or not ($m_i = 0$) (Bhaskar et al., 2024). Circuits are generally evaluated based on three key criteria (Wang et al., 2023):

1. *Faithfulness*: A circuit is considered faithful to the task T if it accurately captures the full model’s task output while ablating all nodes not identified as causally relevant ($m_i = 0$) by replacing them with some ablation value \tilde{z} (Hanna et al., 2024b; Miller et al., 2024). Given a metric P that compares the outputs of two models for task T , task faithfulness F_T is typically quantified as:

$$F_T = P\left(\mathcal{M}(\cdot), \mathcal{M}(\cdot \mid \text{do}(z = \mathbf{m} \odot \mathbf{z}_x + (\mathbf{1} - \mathbf{m}) \odot \tilde{\mathbf{z}}))\right) \quad (2)$$

2. *Minimality*: A circuit is deemed minimal if it excludes nodes that are not causally relevant (Mueller et al., 2024). Formally, given a candidate set of circuits C , minimality is encouraged by selecting the circuit with the smallest norm: $\min_{\mathbf{m} \in C} \|\mathbf{m}\|_1$.
3. *Completeness*: A circuit is said to be complete if it captures all nodes necessary to explain the model’s behavior for task T .

2.3 Continuous Sparsification

Continuous sparsification originates from model pruning and has been introduced to sparsify networks, specifically their weight space (Savarese et al., 2020). Unlike other pruning approaches (Srinivas et al., 2017; Louizos et al., 2018), continuous sparsification approximates l_0 regularization by learning a *deterministic* mask $\mathbf{m} \in \{0, 1\}^N$ over the network’s parameters $\mathbf{w} \in \mathbb{R}^N$ that indicates which weights to prune. The search for such a sparse subnetwork can be represented by the following minimization problem:

$$\min_{\mathbf{w} \in \mathbb{R}^N, \mathbf{m} \in \{0, 1\}^N} \mathcal{L}(\mathcal{M}(\cdot; \mathbf{m} \odot \mathbf{w})) + \lambda \|\mathbf{m}\|_1 \quad (3)$$

which uses the fact that $\|\mathbf{m}\|_0 = \|\mathbf{m}\|_1$ for binary masks, and where \mathcal{L} denotes the loss of the network \mathcal{M} , while λ controls the trade-off between loss and number of parameters $\|\mathbf{w}\|_0$. To circumvent the combinatorial constraint imposed by the

discrete space of $\mathbf{m} \in \{0, 1\}^N$, the mask is deterministically re-parameterized as a sigmoid function $\sigma(\cdot)$ of the new variable $\mathbf{s} \in \mathbb{R}^N$:

$$\min_{\mathbf{w} \in \mathbb{R}^N, \mathbf{s} \in \mathbb{R}^N} \mathcal{L}(\mathcal{M}(\cdot; \sigma(\beta \cdot \mathbf{s}) \odot \mathbf{w})) + \lambda \|\sigma(\beta \cdot \mathbf{s})\|_1 \quad (4)$$

where $\beta \in [1, \infty]$ represents a temperature parameter for which the sigmoid function converges to the Heaviside function with $\lim_{\beta \rightarrow \infty} \sigma(\beta \cdot \mathbf{s}) = H(\mathbf{s})$. By minimizing the above loss while annealing β , and recovering the binary mask from re-parameterization via $\mathbf{m} = H(\mathbf{s})$, a sparse representation of the network’s parameters can be learned.

3 Activation Pruning Through Continuous Sparsification

To automatically identify circuits for sequence-to-sequence prediction tasks, we adopt an approach similar to Bhaskar et al. (2024) and Conmy et al. (2023), and formulate the identification process as a minimization problem. For this, we leverage techniques from model pruning, specifically *continuous sparsification* (see Section 2.3). Unlike traditional pruning, which aims to reduce model complexity by creating a sparse representation of the model’s weight space, we focus on intervening on the model’s activations, thereby linking our method to causal mediation analysis.

As outlined in Section 2.2, we represent a circuit as a binary mask $\mathbf{m} \in \{0, 1\}^N$ over the model’s mediator space—here the activation space of the model components is considered—indicating which activations \mathbf{z} of a frozen model \mathcal{M} are responsible for its behavior on a task T . To find a circuit \mathbf{m} that is both *faithful* and *minimal* (we exclude *completeness* for now; see Section 7 for a respective discussion), we aim to minimize the following loss:

$$\min_{\mathbf{m} \in \{0, 1\}^N} \mathcal{L}_T(\mathcal{M}(\cdot), \mathcal{M}(\cdot \mid \text{do}(z = \mathbf{m} \odot \mathbf{z}_x + (\mathbf{1} - \mathbf{m}) \odot \tilde{\mathbf{z}}))) + \lambda \cdot \mathcal{L}_{reg}(\mathbf{m}) \quad (5)$$

where \mathcal{L}_T captures the circuit’s task faithfulness, with lower values indicating greater faithfulness, while \mathcal{L}_{reg} assesses the size of the circuit. The hyperparameter λ controls the influence of \mathcal{L}_{reg} . Given the combinatorial complexity of optimizing

a binary mask over a potentially large activation space, we follow the approach of continuous sparsification (Section 2.3) and deterministically re-parameterize \mathbf{m} using a sigmoid function $\sigma(\cdot)$ over the new variable $\mathbf{s} \in \mathbb{R}^N$. Additionally, we assess faithfulness by measuring the Kullback-Leibler divergence $D_{KL}(\mathbf{y}^m \parallel \mathbf{y}^M)$ between the circuit’s predicted output distribution and that of the full model. Since the full model’s predicted output distribution \mathbf{y}^M is independent of the variable \mathbf{s} , minimizing the KL divergence between \mathbf{y}^m and \mathbf{y}^M simplifies to minimizing their cross-entropy loss \mathcal{L}_{CE} . For regularization, we apply l_1 regularization in line with Equation 4. This yields:

$$\begin{aligned} \min_{\mathbf{s} \in \mathbb{R}^N} \mathcal{L}_{CE}(\mathcal{M}(\cdot), \mathcal{M}(\cdot \mid \text{do}(\mathbf{z} = \sigma(\beta \cdot \mathbf{s}) \odot \mathbf{z}_x \\ + (1 - \sigma(\beta \cdot \mathbf{s})) \odot \tilde{\mathbf{z}}))) \\ + \lambda \cdot \|\sigma(\beta \cdot \mathbf{s})\|_1 \end{aligned} \quad (6)$$

where the sigmoid function σ is applied element-wise, and β serves as a temperature parameter that increases progressively after each training epoch, following an exponential schedule until it reaches a maximum value, β_{max} , as proposed by Lepori et al. (2023). By minimizing the expression in Equation 6, we obtain an approximation of \mathbf{m} that strikes a balance between *faithfulness* and *minimality*, with λ governing the emphasis on the latter. Once training converges, the binary mask is derived through re-parameterization via $\mathbf{m} = H(\mathbf{s})$. In comparison to edge pruning (Bhaskar et al., 2024) and the subnetwork probing method proposed by Conmy et al. (2023), our approach identifies circuits through *deterministic* re-parameterization. While our method supports different levels of node granularity, we focus on neuron-level interventions.

4 Experiments

This work studies the modularity of a transformer-based sequence-to-sequence model by analyzing circuits responsible for its behavior on highly compositional subtasks. We identify and compare circuits for ten compositional string-edit operations introduced by PCFG SET (Hupkes et al., 2020).² Section 4.1 outlines the experimental setup, including dataset, training, and evaluation details, while Section 4.3 presents the results.

²Our code is available at <https://github.com/mainlp/circuit-compositions>.

4.1 Setup

4.1.1 PCFG SET

As shown in Table 1, PCFG SET (Hupkes et al., 2020) comprises ten string-edit operations applied to sequences generated by a probabilistic context-free grammar. All tasks resemble translation problems, where an input sequence is transformed into a corresponding output sequence through the recursive application of the operations specified in the input sequence. The dataset includes *unary* operations (applied to a single string) and *binary* operations (requiring two arguments). For instance, the binary function prepend places the second argument before the first (e.g., prepend $A1, B1 \rightarrow B1 A1$). The input alphabet in PCFG SET consists of three components: *i*) words representing string-edit operations (e.g., copy or echo), *ii*) symbols forming the input sequence (e.g., $A1, B1$, etc.), and *iii*) a separator “,” that distinguishes for binary arguments. For additional information and examples, please refer to Appendix B.1.

Hupkes et al. (2020) construct PCFG SET such that compositionality is a salient feature. Notably, all operators in PCFG SET are functionally related. For example, the repeat operator can be replicated by applying the copy operation two times in succession (see Table 1). To identify circuits for each operation, we generate ten distinct data subsets, each containing 20,000 examples from a specific string-edit operation (16,000 for training, 4,000 for testing). Further details on data generation and individual sub-datasets are provided in Appendix B.2.

4.1.2 Training

Base Model Training. As a first step, we train a base model \mathcal{M} to perform all operations in PCFG SET. Similar to Hupkes et al. (2020), we use an encoder-decoder model, comprising six encoder and decoder layers with a hidden state dimension of 512, resulting in approximately 58 million parameters. The model is trained on the official data splits of PCFG SET, which include around 83,000 training samples covering all string-edit operations and their compositions. Additional details on the training procedure can be found in Appendix C.3.

Mask Training. Next, we employ *activation pruning through continuous sparsification* (Section 3) to identify circuits corresponding to each operation. For this, we minimize the loss as described in Equation 6 by training a mask \mathbf{m} on the respective subtask dataset (see Section 4.1.1).



Figure 2: Task faithfulness F_T and accuracy for **unary** tasks. The y-axis represents the circuit, while the x-axis denotes the evaluation task. The diagonal is omitted for selected positions due to a lack of applicable tokens.

We consider individual output activations from the feed-forward and multi-head attention modules as mediators $z = a_i \in \mathbf{a}$. For the ablation value \tilde{z} , we conduct experiments with both zero and mean ablations. We use the same ablation value across all token positions and adopt an approach similar to node patching, where interventions target the model’s residual stream. When employing mean ablations, we follow the approach proposed by Wang et al. (2023) and use the mediator’s mean value across a reference distribution, specifically the subtask dataset. We optimize for hyperparameters using random search, as detailed in Appendix C.5.1. A sensitivity analysis of the parameter λ is provided in Appendix D.1.3. Further details on the mask training procedure can be found in Appendix C.4.

4.1.3 Evaluation

We evaluate circuits based on two main criteria: performance—measured in terms of faithfulness and accuracy—and node overlap. Since the KL-divergence $D_{KL}(\mathbf{y}^m \parallel \mathbf{y}^M)$ is unbounded and unsuitable for cross-task comparisons, we also compute a normalized version of the Jensen–Shannon divergence, $0 \leq \text{JSD}_{\text{norm}} \leq 1$, to measure faithfulness (see Equation 7, Appendix C.6). We define a circuit’s task *faithfulness performance* as $F_T = 1 - \text{JSD}_{\text{norm}}$, where values closer to 1 indicate higher task faithfulness. Accuracy is determined by the exact match between the circuit’s prediction and the ground truth. When evaluating a circuit from task T on a different task \hat{T} , we retain the mean values from task T as ablation values for nodes with $m_i = 0$ when using mean ablations. For node overlap, we compare circuits using the Intersection over Union (IoU) and Intersection over Minimum (IoM), as defined in Equation 8 in Appendix C.6.

4.2 TRACR Experiments

To validate our method’s capacity to identify both *minimal* and *faithful* circuits, we implement four unary PCFG SET functions (copy, echo, reverse, swap) in RASP (Weiss et al., 2021). Each function is compiled into transformer model weights using TRACR (Lindner et al., 2024), serving as ground truth for evaluation—a common validation approach in prior studies (Conmy et al., 2023; Bhaskar et al., 2024). Our results show that *activation pruning through continuous sparsification* perfectly recovers all individual neurons in the compiled circuits while maintaining 100% faithfulness. For further details, please refer to Appendix C.5.

4.3 Results

We begin by presenting the performance results of the base model on PCFG SET. After training, the encoder-decoder model demonstrates strong performance across all ten string-edit operations, with accuracy exceeding 95% on unary tasks and ranging from 83% to 99% on binary tasks. The lowest accuracies of 83% and 84% are for prepend and append, which turn out as the most challenging operations (a detailed task-specific breakdown is available in Appendix C.3). In the subsequent sections, we focus on circuits identified through mean ablation. For additional results on circuits discovered via zero ablation, see Appendix D.2.

4.3.1 Circuit Performance

We first analyze the circuits’ performance across different string-edit operations. Figure 2 illustrates both the task faithfulness performance, F_T , as defined in Section 4.1.3, as well as the accuracy of each circuit. Due to functional similarities between the string-edit operations, multiple operators may produce the same output tokens at various positions.

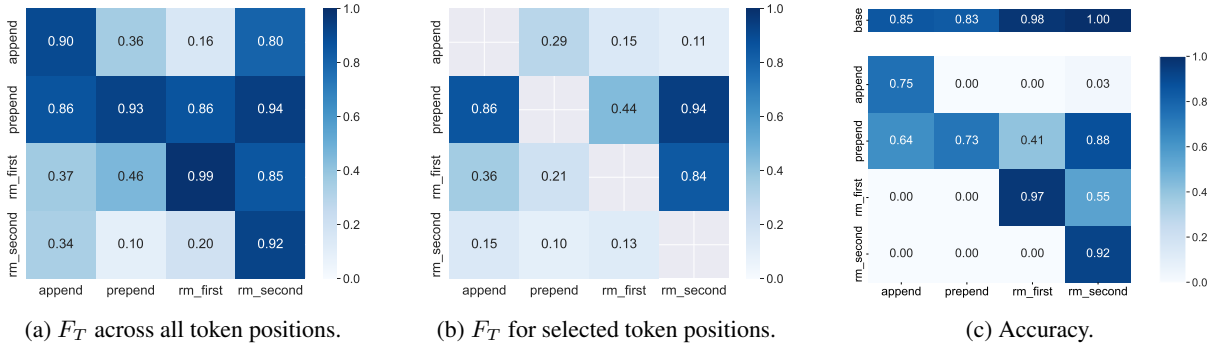


Figure 3: Task faithfulness F_T and accuracy for **binary** tasks. The y-axis represents the circuit, while the x-axis denotes the evaluation task. The diagonal is omitted for selected positions due to a lack of applicable tokens.

Therefore, we assess faithfulness in two ways: *i*) averaged across all output tokens (Figure 2a), and *ii*) at positions where the ground truth output sequences of the circuit-task and evaluation-task differ (Figure 2b). For example, when evaluating the copy circuit on the echo task, we assess task faithfulness at the final token of the target output, focusing on the additional x_n present in echo, which replaces the end-of-sequence token in copy.

Unary Circuits. The diagonal in Figure 2a shows that all unary circuits exhibit strong faithfulness on their respective tasks, with rates exceeding 0.94. Several circuits, such as echo, repeat, and swap, also demonstrate high cross-task faithfulness and accuracy on the copy task. This aligns with our intuition as the copy operation is either a significant component of these functions or can be effectively performed by them. In contrast, operators like reverse and shift, which substantially alter the input sequence, show lower functional similarity to copy. This is reflected in both the reduced cross-task faithfulness (Figures 2a; 2b) and lower accuracy (Figure 2c) on the copy task. Interestingly, while the repeat and swap circuits perform well on the copy task, the reverse is not true. This occurs despite the theoretical possibility of completing the repeat operation by applying the copy operator twice. Overall, circuits tend to show high cross-task faithfulness performance when evaluated on all tokens (Figure 2a), but scores drop for tokens that differ between the circuit’s task and the evaluation task (Figure 2b). For instance, the copy circuit achieves a cross-task faithfulness score of 0.64 on the swap task across the full sequence, but this drops to 0.01 when evaluated only on tokens that differ between the ground truth output sequences of copy and swap. This suggests that the

copy circuit largely adheres to its original function, even when the input demands a swap operation.

Binary Circuits. Figure 3 highlights the performance of circuits associated with the binary operations. Similar to the unary circuits, these circuits demonstrate high faithfulness for their respective tasks, with performance values ranging from 0.90 to 0.99 (Figure 3a). Notably, the remove_first circuit shows strong cross-task faithfulness and accuracy on the remove_second task, although the reverse is not true. Additionally, the prepend circuit shows consistently strong performance across all binary tasks. When analyzing the size of each circuit—the fraction of nodes with causal relevance to the model’s task behavior, expressed as a fraction of remaining activations (Figure 4)—we observe that the prepend circuit is the largest among all circuits, retaining 39% of its neuron activations. Similarly, the remove_first circuit is notably larger than the remove_second circuit. This pattern suggests that more complex operations (prepend, append, remove_first) engage a greater portion of the model’s activation space compared to sim-




Figure 4: Sparsity of the mean ablated circuits as fraction of a circuit’s remaining activations (in %).

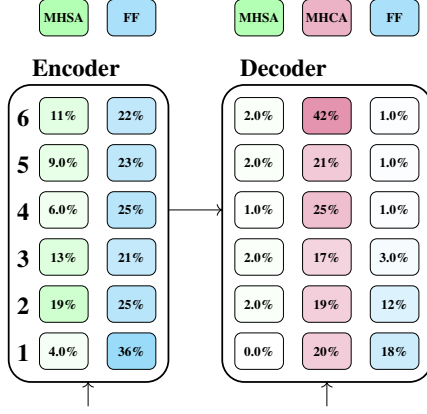


Figure 5: The local sparsity of the copy circuit. For each layer, the fraction of remaining activations in the multi-head self-attention (MHSA), multi-head cross-attention (MHCA), and feed-forward (FF) modules are shown.

pler tasks. For instance, `remove_second` may rely on a straightforward copy operation with an end-of-sequence token at the position of the separator “,”; whereas `remove_first` likely requires more intricate processing.


Local sparsity. To zoom in more, Figure 5 illustrates the local sparsity of the copy circuit, which is the smallest among all circuits (Figure 4). It shows the percentage of remaining neuron activations in the multi-head self-attention (MHSA), multi-head cross-attention (MHCA), and feed-forward (FF) modules at each layer. Notably, in the decoder, MHCA modules remain mostly active, whereas FF and MHSA activations are almost entirely pruned. This aligns with our expectations, as the copy task primarily involves transferring encoder-processed input to the output. Additional visualizations for other circuits can be found in Appendix D.1.2, specifically Figures 20 and 21.

Zero vs. mean ablations. The results so far relate to circuits identified via mean ablation. Consistent with prior work (Miller et al., 2024), we find zero ablations to be less faithful. For an overview of respective results, see Appendix D.2.


4.3.2 Circuit Overlap

So far, we evaluated the performance of circuits and their cross-task generalizability. Next, we show their respective node overlap in Figure 6. Specifically, we assess the IoU and IoM between all circuit pairs, as described in Section 4.1.3. The IoU captures the overlap of circuits relative to their combined size, whereas the IoM measures how much of the smaller circuit is contained within the larger

one. Most pairs exhibit IoU values between 0.20 and 0.30, though some clusters show higher overlap. For instance, the circuits for `reverse`, `swap`, and `shift` have IoU values between 0.36 and 0.42, and IoM values around 0.60, indicating significant shared activations across these tasks. Given the functional similarities of these operations, this finding aligns with our intuition. However, it is worth noting that the cross-task performance of `reverse`, `swap`, and `shift` is lower than that of other circuits (see Figure 2). A similarly high node overlap is observed for the `append` and `prepend` circuits—both associated with tasks that the base model \mathcal{M} struggles with, as shown by the model’s comparatively lower accuracy (see Table 3 in Appendix C.3).



(a) Intersection over Union (IoU).



(b) Intersection over Minimum (IoM).

Figure 6: Node overlap for circuits identified via activation pruning with mean ablation.




Figure 7: The results of combining circuits through a union operation on their respective binary masks.

4.3.3 Circuit Compositions: Subnetwork Set Operations

Given that we can define a circuit as binary mask $\mathbf{m} \in \{0, 1\}^N$ over the model’s mediator space, we extend our analysis to study *cross-circuit relationships* by creating compositions through basic set operations, as we want to gauge to what degree composite circuits emerge. Specifically, for a circuit-pair $(\mathbf{m}^{T_1}, \mathbf{m}^{T_2})$, we define their composite as the union $\mathbf{m}^{T_1, T_2} = \mathbf{m}^{T_1} \cup \mathbf{m}^{T_2}$. Importantly, this union is not symmetric: we use the ablation values $\tilde{\mathbf{z}}^{T_1}$ of the first circuit for nodes outside the union ($m_i^{T_1, T_2} = 0$). For additional details on how this union is applied, please refer to Appendix C.7.

Figure 7 shows the cross-task accuracy of various circuits and their composites, derived from the union operation previously described. The results demonstrate the existence of composite circuits which acquire functional capacities for subtasks that the original base circuits could not handle individually. For example, while neither the repeat nor the reverse circuit can solve the echo task alone, their union achieves a notable accuracy of 78%. Similarly, the composite of swap and reverse shows a performance improvement on the shift task, reaching an accuracy of 33%. For tasks that can be solved by at least one of the original circuits, the performance of their composite generally reflects a blend of their individual accuracies. For instance, the repeat circuit achieves 90% accuracy on the copy task, while the reverse circuit yields 0.0% accuracy. Their union achieves an intermediate accuracy of 80% on the task. Similarly, the swap circuit achieves 27% accuracy on the reverse task, whereas the reverse circuit reaches up to 96%. Their composite circuit achieves a balanced performance of 45%. This implies that improvements in accuracy for novel tasks may be accompanied by a decline in performance on previously mastered skills.

A possible explanation for these observations is

that the union operation $\mathbf{m}^{T_1, T_2} = \mathbf{m}^{T_1} \cup \mathbf{m}^{T_2}$ shifts the activation space of \mathbf{m}^{T_1} towards \mathbf{m}^{T_2} , effectively transforming the input into a feature space that aligns with one of the subtasks the model \mathcal{M} can solve. For instance, the union $\mathbf{m}^{\text{repeat}} \cup \mathbf{m}^{\text{reverse}}$ might shift the activation space of $\mathbf{m}^{\text{repeat}}$ towards $\mathbf{m}^{\text{reverse}}$, essentially allowing the new circuit to perform the echo operation, while sacrificing some performance on the repeat task. Further experiments and results on subnetwork set operations can be found in Appendix D.1.5.

Our results show that circuits responsible for tasks like echo, repeat, and swap can perform the copy task, whereas the copy circuit struggles to perform other tasks. This raises the question of whether the echo, repeat, and swap circuits leverage the copy operation in their respective functions. Our analysis on node overlap indicates that the copy circuit—which is notably the smallest among those identified—is largely embedded within the other circuits. Moreover, circuits with similar functions exhibit greater node overlap than functionally distinct ones. Additionally, we show that circuits combined through the union operator can represent novel functional capabilities, suggesting some degree of modularity within the network.

5 Related Work

Techniques that discover circuits through causal mediation analysis have been successfully applied across various domains and tasks, including the study of gender bias in language models (Jeoung and Diesner, 2022; Chintam et al., 2023), different forms of factual recall (Meng et al., 2022; Geva et al., 2023), subject-verb agreement (Chintam et al., 2023), and arithmetic operations (Nanda et al., 2023; Hanna et al., 2024a). Finding subnetworks has also been of interest outside of the mechanistic interpretability literature. Previous work has studied task-specific subnetworks through

clustering approaches (Casper et al., 2022; Watanabe, 2019), pruning (Csordás et al., 2020; Cao et al., 2021; Lepori et al., 2023), sparse fine-tuning (Ansell et al., 2022), and adapters (Pfeiffer et al., 2021; Rücklé et al., 2021). However, few studies have used the circuits paradigm to examine relationships between circuits in functionally related tasks (Hanna et al., 2024b).

6 Conclusion

This study explores modularity in neural networks by analyzing circuits in a transformer-based language model for highly compositional subtasks. Specifically, we train an encoder-decoder model using the PCFG SET dataset and identify circuits responsible for ten modular string-edit operations. To achieve this, we introduce *activation pruning through continuous sparsification*, a method that allows us to formulate circuit identification as minimization problem. Our results demonstrate that this approach successfully identifies both faithful and sparse circuits for each subtask within PCFG SET. Additionally, we assess all circuits by examining their cross-task performance and node overlap. Finally, we show promising first steps of circuit compositions, where new functional circuits can be composed through set operations, such as the union of two circuits.

7 Limitations

In the following section, we mention limitations of our work that could be addressed by future studies.

Completeness. We introduce *activation pruning through continuous sparsification*, a method that formulates the circuit discovery process as an optimization problem balancing *task faithfulness* and *minimality*. However, as discussed in Section 2.2, circuits are further expected to be *complete*, meaning they should include all nodes involved in the model’s computations for a given task. We emphasize that our current approach does not guarantee *completeness*. While some methods attempt to evaluate completeness by analyzing circuit behavior under ablations (Wang et al., 2023), these approaches are computationally expensive, especially when dealing with large circuit sizes. Inspired by recent studies (Yu et al., 2024; Chen et al., 2025), future research could explore more efficient ways to quantify completeness and incorporate this objective into the optimization framework described in Equation 6.

Impact of ablation paradigm. We observe that the circuits identified are influenced by the ablation value \tilde{z} used during the identification process. For example, when comparing circuits based on mean versus zero ablation, we find notable differences in the nodes identified, even for zero-ablated circuits that demonstrate high task faithfulness (see results in Appendix D.2, specifically Figure 9). Essentially, the behavior of a circuit is determined not merely by what it includes but also by what it excludes, as similarly noted by Miller et al. (2024). We believe this is a general characteristic of methods that rely on constant perturbations, such as zero and mean ablations (Olsson et al., 2022; Wang et al., 2023).

Generalization of findings. It is important to note that our study focuses on a small encoder-decoder model trained on a single dataset like PCFG SET. Naturally, the question arises as to whether these findings will generalize to larger models and more diverse tasks. This touches on a broader concern within the field of mechanistic interpretability, where many studies are conducted on smaller models and datasets (Elhage et al., 2022). We argue that, for our purposes, a controlled dataset like PCFG SET is well-suited to study modularity in language models. Nonetheless, the method we propose in Section 3 is easily extendable to other models and tasks. We believe that scaling our findings to larger models and more complex tasks represents a compelling area for future research.

Acknowledgments

We would like to thank the members of the MaiNLP lab for their valuable feedback, with special recognition to Robert Litschko, Michael Hedderich, Florian Echin, and Shijia Zhou. Furthermore, we would like to express our gratitude to Lucas Georges Gabriel Charpentier at LTG for valuable feedback. Our appreciation extends to the anonymous reviewers for their comments and suggestions. We further acknowledge Sigma2, Norway, for providing access to the LUMI supercomputer, part of the EuroHPC Joint Undertaking, hosted by CSC (Finland) and the LUMI consortium. Lastly, we acknowledge the support provided to BP through the ERC Consolidator Grant DIALECT 101043235. For parts of the project, SW acknowledges funding for a research stay at MaiNLP at LMU, funded by Integreat, the Norwegian Centre for Knowledge-driven Machine Learning.

References

- Alan Ansell, Edoardo Ponti, Anna Korhonen, and Ivan Vulić. 2022. [Composable sparse fine-tuning for cross-lingual transfer](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1778–1796, Dublin, Ireland. Association for Computational Linguistics.
- Adithya Bhaskar, Alexander Wettig, Dan Friedman, and Danqi Chen. 2024. [Finding transformer circuits with edge pruning](#). In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.
- Steven Cao, Victor Sanh, and Alexander Rush. 2021. [Low-complexity probing via finding subnetworks](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 960–966, Online. Association for Computational Linguistics.
- Stephen Casper, Shlomi Hod, Daniel Filan, Cody Wild, Andrew Critch, and Stuart Russell. 2022. Graphical clusterability and local specialization in deep neural networks. In *ICLR 2022 Workshop on PAIR2Struct: Privacy, Accountability, Interpretability, Robustness, Reasoning on Structured Data*.
- Hang Chen, Jiaying Zhu, Xinyu Yang, and Wenya Wang. 2025. Rethinking circuit completeness in language models: And, or, and adder gates. *arXiv preprint arXiv:2505.10039*.
- Abhijith Chintam, Rahel Beloch, Willem Zuidema, Michael Hanna, and Oskar van der Wal. 2023. [Identifying and adapting transformer-components responsible for gender bias in an English language model](#). In *Proceedings of the 6th BlackboxNLP Workshop: Analyzing and Interpreting Neural Networks for NLP*, pages 379–394, Singapore. Association for Computational Linguistics.
- Arthur Conmy, Augustine Mavor-Parker, Aengus Lynch, Stefan Heimersheim, and Adrià Garriga-Alonso. 2023. Towards automated circuit discovery for mechanistic interpretability. *Advances in Neural Information Processing Systems*, 36:16318–16352.
- Róbert Csordás, Sjoerd van Steenkiste, and Jürgen Schmidhuber. 2020. Are neural nets modular? inspecting functional modularity through differentiable weight masks. In *International Conference on Learning Representations*.
- Nelson Elhage, Tristan Hume, Catherine Olsson, Nicholas Schiefer, Tom Henighan, Shauna Kravec, Zac Hatfield-Dodds, Robert Lasenby, Dawn Drain, Carol Chen, et al. 2022. Toy models of superposition. *arXiv preprint arXiv:2209.10652*.
- Javier Ferrando, Gabriele Sarti, Arianna Bisazza, and Marta R Costa-jussà. 2024. A primer on the inner workings of transformer-based language models. *arXiv preprint arXiv:2405.00208*.
- Mor Geva, Jasmijn Bastings, Katja Filippova, and Amir Globerson. 2023. [Dissecting recall of factual associations in auto-regressive language models](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 12216–12235, Singapore. Association for Computational Linguistics.
- Michael Hanna, Ollie Liu, and Alexandre Variengien. 2024a. How does gpt-2 compute greater-than?: Interpreting mathematical abilities in a pre-trained language model. *Advances in Neural Information Processing Systems*, 36.
- Michael Hanna, Sandro Pezzelle, and Yonatan Belinkov. 2024b. [Have faith in faithfulness: Going beyond circuit overlap when finding model mechanisms](#). In *ICML 2024 Workshop on Mechanistic Interpretability*.
- Dieuwke Hupkes, Verna Dankers, Mathijs Mul, and Elia Bruni. 2020. Compositionality decomposed: How do neural networks generalise? *Journal of Artificial Intelligence Research*, 67:757–795.
- Sullam Jeoung and Jana Diesner. 2022. [What changed? investigating debiasing methods using causal mediation analysis](#). In *Proceedings of the 4th Workshop on Gender Bias in Natural Language Processing (GeBNLP)*, pages 255–265, Seattle, Washington. Association for Computational Linguistics.
- Michael Lepori, Thomas Serre, and Ellie Pavlick. 2023. Break it down: Evidence for structural compositionality in neural networks. *Advances in Neural Information Processing Systems*, 36:42623–42660.
- David Lindner, János Kramár, Sebastian Farquhar, Matthew Rahtz, Tom McGrath, and Vladimir Mikulik. 2024. Tracr: Compiled transformers as a laboratory for interpretability. *Advances in Neural Information Processing Systems*, 36.
- Christos Louizos, Max Welling, and Diederik P. Kingma. 2018. [Learning sparse neural networks through l₀ regularization](#). In *International Conference on Learning Representations*.
- Kevin Meng, David Bau, Alex Andonian, and Yonatan Belinkov. 2022. Locating and editing factual associations in gpt. *Advances in Neural Information Processing Systems*, 35:17359–17372.
- Jack Merullo, Carsten Eickhoff, and Ellie Pavlick. 2024. [Circuit component reuse across tasks in transformer language models](#). In *The Twelfth International Conference on Learning Representations*.
- Joseph Miller, Bilal Chughtai, and William Saunders. 2024. [Transformer circuit evaluation metrics are not robust](#). In *First Conference on Language Modeling*.
- Aaron Mueller, Jannik Brinkmann, Millicent Li, Samuel Marks, Koyena Pal, Nikhil Prakash, Can Rager, Aruna Sankaranarayanan, Arnab Sen Sharma, Jiuding Sun, Eric Todd, David Bau, and Yonatan Belinkov. 2024. [The quest for the right mediator: A](#)

- history, survey, and theoretical grounding of causal interpretability. *Preprint*, arXiv:2408.01416.
- Neel Nanda. 2023. [Attribution patching: Activation patching at industrial scale](#). Accessed: 2024-09-19.
- Neel Nanda, Lawrence Chan, Tom Lieberum, Jess Smith, and Jacob Steinhardt. 2023. [Progress measures for grokking via mechanistic interpretability](#). In *The Eleventh International Conference on Learning Representations*.
- Catherine Olsson, Nelson Elhage, Neel Nanda, Nicholas Joseph, Nova DasSarma, Tom Henighan, Ben Mann, Amanda Askell, Yuntao Bai, Anna Chen, et al. 2022. In-context learning and induction heads. *arXiv preprint arXiv:2209.11895*.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. 2019. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32.
- Judea Pearl. 2001. Direct and indirect effects. In *Proceedings of the Seventeenth Conference on Uncertainty and Artificial Intelligence, 2001*, pages 411–420. Morgan Kaufman.
- Jonas Pfeiffer, Aishwarya Kamath, Andreas Rücklé, Kyunghyun Cho, and Iryna Gurevych. 2021. Adapterfusion: Non-destructive task composition for transfer learning. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 487–503.
- Nikhil Prakash, Tamar Rott Shaham, Tal Haklay, Yonatan Belinkov, and David Bau. 2024. [Fine-tuning enhances existing mechanisms: A case study on entity tracking](#). In *The Twelfth International Conference on Learning Representations*.
- Andreas Rücklé, Gregor Geigle, Max Glockner, Tilman Beck, Jonas Pfeiffer, Nils Reimers, and Iryna Gurevych. 2021. Adapterdrop: On the efficiency of adapters in transformers. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 7930–7946.
- Pedro Savarese, Hugo Silva, and Michael Maire. 2020. Winning the lottery with continuous sparsification. *Advances in neural information processing systems*, 33:11380–11390.
- Noam Shazeer. 2020. Glu variants improve transformer. *arXiv preprint arXiv:2002.05202*.
- Suraj Srinivas, Akshayvarun Subramanya, and R. Venkatesh Babu. 2017. [Training sparse neural networks](#). In *2017 IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 455–462.
- Aaquib Syed, Can Rager, and Arthur Conmy. 2024. [Attribution patching outperforms automated circuit discovery](#). In *Proceedings of the 7th BlackboxNLP Workshop: Analyzing and Interpreting Neural Networks for NLP*, pages 407–416, Miami, Florida, US. Association for Computational Linguistics.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Jesse Vig, Sebastian Gehrmann, Yonatan Belinkov, Sharon Qian, Daniel Nevo, Yaron Singer, and Stuart Shieber. 2020. [Investigating gender bias in language models using causal mediation analysis](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 12388–12401. Curran Associates, Inc.
- Kevin Ro Wang, Alexandre Variengien, Arthur Conmy, Buck Shlegeris, and Jacob Steinhardt. 2023. [Interpretability in the wild: a circuit for indirect object identification in GPT-2 small](#). In *The Eleventh International Conference on Learning Representations*.
- Chihiro Watanabe. 2019. Interpreting layered neural networks via hierarchical modular representation. In *Neural Information Processing: 26th International Conference, ICONIP 2019, Sydney, NSW, Australia, December 12–15, 2019, Proceedings, Part V 26*, pages 376–388. Springer.
- Gail Weiss, Yoav Goldberg, and Eran Yahav. 2021. Thinking like transformers. In *International Conference on Machine Learning*, pages 11080–11090. PMLR.
- Lei Yu, Jingcheng Niu, Zining Zhu, and Gerald Penn. 2024. Functional faithfulness in the wild: Circuit discovery with differentiable computation graph pruning. *arXiv preprint arXiv:2407.03779*.

A Reproducibility Statement

To ensure the reproducibility of our experiments, we make all code publicly available at: <https://github.com/mainlp/circuit-compositions>. Details of the training process, including the computational setup, model implementation, and hyperparameter selection, are thoroughly documented in Section 4.1.2 and Appendix C. Similarly, a detailed account of the evaluation procedure can be found in Section 4.1.3 and Appendix C.6. Furthermore, all data used in this work is either publicly available or accompanied by a detailed description of the data generation process in Appendix B.

Unary operation	Input	Output	Example
copy	$x_1 \dots x_n$	$x_1 \dots x_n$	copy K1 Y1 W1 K1 \rightarrow K1 Y1 W1 K1
echo	$x_1 \dots x_n$	$x_1 \dots x_n x_n$	echo E1 K1 A1 X1 J \rightarrow E1 K1 A1 X1 J1 J1
repeat	$x_1 \dots x_n$	$x_1 \dots x_n x_1 \dots x_n$	repeat J1 F1 S1 \rightarrow J1 F1 S1 J1 F1 S1
reverse	$x_1 \dots x_n$	$x_n \dots x_1$	reverse G1 T1 X1 J1 \rightarrow J1 X1 T1 G1
swap	$x_1 \dots x_n$	$x_n x_2 \dots x_{n-1} x_1$	swap B1 Z1 V1 I1 W1 \rightarrow W1 Z1 V1 I1 B1
shift	$x_1 \dots x_n$	$x_2 \dots x_n x_1$	shift Y1 I1 D1 H1 K1 \rightarrow I1 D1 H1 K1 Y1
Binary operation	Input	Output	Example
append	x, y	$x y$	append F1 B1, U1 A1 G1 \rightarrow F1 B1 U1 A1 G1
prepend	x, y	$y x$	prepend F1 B1, U1 A1 G1 \rightarrow U1 A1 G1 F1 B1
remove_first	x, y	y	remove_first Z1 P1 N1, A1 D1 \rightarrow A1 D1
remove_second	x, y	x	remove_second F1 B1, U1 A1 G1 \rightarrow F1 B1

Table 2: The string-edit operations in PCFG SET from Hupkes et al. (2020).

B Dataset

B.1 PCFG SET

Hupkes et al. (2020) construct PCFG SET in such a way that compositionality is a salient feature of the dataset, while aligning its statistical properties with those of natural language corpora, specifically English. We present examples for each string-edit operation in Table 2 to help the reader get more familiar with the functions employed in PCFG SET Hupkes et al. (2020).

The PCFG SET dataset is released under a MIT license and used in accordance to the licensing agreements.

B.2 Generating Isolated Function Data

We use the probabilistic context-free grammar proposed by Hupkes et al. (2020) to generate datasets for each of the string-edit operations originally used to construct PCFG SET. Using the same vocabulary, consisting of letters $[A - Z]$ combined with numbers in $(0, \infty] \in \mathcal{Z}^+$, we restrict the grammar to produce samples that use only specific string-edit operation for each subtask dataset, e.g. copy W1 O1 Z5 G1. For each SET operation, we generate 16 000 samples for training and 4 000 for validation. Each subtask dataset includes samples of both single-use and composed applications of that function, such as remove_first remove_first A1 , B1 , B1 A1 \rightarrow B1 A1.

C Experimental Details

In this section, we outline key details regarding our experiments, such as our computational setup and the hyperparameters selected.

C.1 Software and Computing

All experiments are implemented using PyTorch (Paszke et al., 2019) as the primary framework. Detailed information about supporting software and specific versions can be found in our code repository in the supplementary materials. Experiments were run on AMD Instinct MI250X and NVIDIA A100 GPUs.

C.1.1 Computational Budget

For training base models, we consumed approximately 12 GPU hours in total. For the mask training, including hyperparameter selection, failed runs, and debugging, we consumed approximately 4700 GPU hours.

C.2 Base Model Architecture

The base model \mathcal{M} is a transformer-based encoder-decoder, as proposed by Vaswani et al. (2017), but with the Gated Linear Unit (GLU) variant from Shazeer (2020). We use sinusoidal positional encoding, delimit tokens on whitespace, and disjoint embedding matrices for the input and output sequences. The base model \mathcal{M} consists of six encoder and six decoder layers, with a hidden size of dimension 512, and eight attention heads per layer.

C.3 Base Model Training

We train the base model \mathcal{M} using a learning rate of $5 \cdot 10^{-7}$, a batch size of 64 without gradient accumulation, a gradient clipping value of 15, and a dropout rate of 0.2. Furthermore, the model is trained using the original datasets from Hupkes et al. (2020). In Table 3, we report the accuracy of the base model on each of the isolated subtask datasets. The base model \mathcal{M} achieves final accuracy of 87.8% on the official PCFG test set, which

	Operation	Accuracy
Unary	copy	1.00
	echo	0.999
	repeat	0.943
	reverse	0.989
	swap	0.996
	shift	0.992
Binary	append	0.848
	prepend	0.832
	remove_first	0.975
	remove_second	0.999

Table 3: The generation accuracy of the base model on the different PCFG SET subtasks (Hupkes et al., 2020).

is slightly lower than the performance of 92.0% reported by Hupkes et al. (2020).

C.4 Mask Training

In this study, we conduct experiments for both mean and zero ablation. Following the approach of Wang et al. (2023), the mean ablation value \tilde{z} is derived from a reference distribution. Specifically, when training a mask for task T_l on dataset D_l , we assign the ablated value \tilde{z}_i for each potential mediator z_i in \mathcal{M} as the average activation of z_i across samples in D_l , which only includes samples corresponding to task T_l . For instance, when training a mask for the copy task, the ablated value \tilde{z}_i for each mediator z_i is set to the average activation of z_i over samples exclusively containing the copy function, i.e., $D_{\text{copy}}^{\text{train}}$, averaged across all token positions.

The final optimization problem, described in Equation 6, requires the predicted output probabilities $\mathbf{y}^{\mathcal{M}}$ of the base model \mathcal{M} for each dataset sample across all of the subtask datasets. In practice, this is achieved through caching: a forward pass is performed on all datasets using \mathcal{M} , and $\mathbf{y}^{\mathcal{M}}$ is stored for later use during mask training.

C.5 Method Validation with TRACR

The TRACR models employ a decoder-only architecture with bidirectional attention. Table 4 summarizes the model configurations for each task. All models execute their respective RASP functions on sequences of length four. For mask training, we use a learning rate of 0.001, maximum temperature $\beta_{\max} = 200$, $\lambda = 1 \cdot 10^{-4}$, and an initial mask value of $\mathbf{s}_{\text{initial}} = 1.0$. We train 50 epochs for copy, reverse, and swap, and 200 epochs for echo.

In Figures 18 and 19, we report the sparsity of the pruned circuits along with their respective

Circuit	Hidden size	QKV size	Layers	Heads
copy	15	6	1	1
reverse	37	10	4	1
echo	46	9	4	2
swap	74	13	6	1

Table 4: The model configurations of the TRACR programs.

ground truth circuit in the compiled model. We note that in our experiments, we do not consider the beginning-of-sequence (BOS) token position. Thus, for layers that contain an active neuron only for the BOS token, we prune the whole component. For copy, reverse, and swap, our method recovers all individual neurons in the compiled model. For these tasks, the compiled model achieves a 100% accuracy on the target task. For echo, our RASP implementation achieves 42% accuracy when compiled, which our method also recovers with full faithfulness. This shows that our method finds faithful circuits even when the base model’s performance is not perfect.

C.5.1 Hyperparameters

In Table 5, we report the final hyperparameters for the mask training for each subtask, using the same notation as in Section 3. These were selected based on a random search over the following hyperparameters: learning rate: $\{1 \cdot 10^{-3}, 1 \cdot 10^{-4}, 1 \cdot 10^{-5}\}$, λ : $\{1 \cdot 10^{-3}, 1 \cdot 10^{-4}, 1 \cdot 10^{-5}\}$, $\mathbf{s}_{\text{initial}}$: $\{0.2, 0.05, 0\}$, β_{\max} : $\{100, 200, 300\}$.

C.6 Evaluation

We evaluate circuits according to two primary criteria, performance and node overlap, as described in Section 4.1.3. Specifically, we quantify the performance of a circuit by its *faithfulness* score, which is calculated via the KL divergence, D_{KL} , between the output distribution of the circuit \mathbf{y}^m and the

Circuit	lr	λ	$\mathbf{s}_{\text{initial}}$	β_{\max}	Epochs
echo	10^{-4}	10^{-4}	0.05	200	500
copy	10^{-4}	10^{-4}	0.05	200	500
repeat	10^{-4}	10^{-4}	0.05	200	500
reverse	10^{-4}	10^{-4}	0.05	200	500
swap	10^{-4}	10^{-4}	0.05	200	500
shift	10^{-4}	10^{-4}	0.05	200	500
append	10^{-3}	10^{-5}	0	200	500
prepend	10^{-3}	10^{-5}	0	100	500
rm_first	10^{-4}	10^{-5}	0.05	300	500
rm_second	10^{-3}	10^{-5}	0.05	100	300

Table 5: Final hyperparameters used for mask training.

λ	Faithfulness		Sparsity
	F_T	D_{KL}	
1×10^{-2}	0.228	4.180	0.022
1×10^{-3}	0.833	0.687	0.105
1×10^{-4}	0.943	0.120	0.133
1×10^{-5}	0.992	0.015	0.162
1×10^{-6}	0.999	0.001	0.298

λ	Faithfulness		Sparsity
	F_T	D_{KL}	
1×10^{-2}	0.106	4.478	0.017
1×10^{-3}	0.734	0.737	0.135
1×10^{-4}	0.951	0.121	0.199
1×10^{-5}	0.982	0.026	0.324
1×10^{-6}	0.994	0.013	0.648

Table 6: Effect of the sparsity regularization for different values of λ on the copy (left) and repeat (right) tasks.

base model $\mathbf{y}^{\mathcal{M}}$. However, to get a bounded metric, we instead use the normalized version of the Jensen-Shannon divergence between the two output distributions:

$$\begin{aligned} \text{JSD}_{\text{norm}}(\mathbf{y}^m \parallel \mathbf{y}^{\mathcal{M}}) &= \frac{1}{2\log(2)} (\\ D_{KL}(\mathbf{y}^m \parallel \mathbf{y}^{m\mathcal{M}}) + D_{KL}(\mathbf{y}^{\mathcal{M}} \parallel \mathbf{y}^{m\mathcal{M}}) \\), \quad \text{where } \mathbf{y}^{m\mathcal{M}} &= \frac{\mathbf{y}^m + \mathbf{y}^{\mathcal{M}}}{2} \end{aligned} \quad (7)$$

where $\mathbf{y}^{m\mathcal{M}}$ is a mixture distribution of \mathbf{y}^m and $\mathbf{y}^{\mathcal{M}}$. Note that the Jensen-Shannon divergence is symmetric. Furthermore, it is bounded, i.e., $0 \leq \text{JSD}_{\text{norm}} \leq 1$.

To measure the circuits' node overlap we compute their Intersection over Union (IoU) and Intersection over Minimum (IoM). Given two circuits $\mathbf{m}^{T_1} \in \{0, 1\}^N$ and $\mathbf{m}^{T_2} \in \{0, 1\}^N$ defined over the same node space, the IoU and IoM are computed as follows:

$$\begin{aligned} \text{IoU} &= \frac{\|\mathbf{m}^{T_1} \cap \mathbf{m}^{T_2}\|_1}{\|\mathbf{m}^{T_1} \cup \mathbf{m}^{T_2}\|_1} \\ \text{IoM} &= \frac{\|\mathbf{m}^{T_1} \cap \mathbf{m}^{T_2}\|_1}{\min(\|\mathbf{m}^{T_1}\|_1, \|\mathbf{m}^{T_2}\|_1)} \end{aligned} \quad (8)$$

where $\mathbf{m}^{T_1} \cap \mathbf{m}^{T_2}$ represents the intersection of the two binary masks, while $\mathbf{m}^{T_1} \cup \mathbf{m}^{T_2}$ denotes their union. Specifically, the intersection between two circuits $\mathbf{m}^{T_1, T_2} = \mathbf{m}^{T_1} \cap \mathbf{m}^{T_2}$ is defined as:

$$\begin{aligned} m_i^{T_1, T_2} &= m_i^{T_1} \cap m_i^{T_2} = m_i^{T_1} \wedge m_i^{T_2} \\ &= \begin{cases} 1, & \text{if } m_i^{T_1} = 1 \text{ and } m_i^{T_2} = 1, \\ 0, & \text{otherwise.} \end{cases} \end{aligned} \quad (9)$$

Similarly, the union between two circuits $\mathbf{m}^{T_1, T_2} = \mathbf{m}^{T_1} \cup \mathbf{m}^{T_2}$ can be computed as:

$$\begin{aligned} m_i^{T_1, T_2} &= m_i^{T_1} \cup m_i^{T_2} = m_i^{T_1} \vee m_i^{T_2} \\ &= \begin{cases} 0, & \text{if } m_i^{T_1} = 0 \text{ and } m_i^{T_2} = 0, \\ 1, & \text{otherwise.} \end{cases} \end{aligned} \quad (10)$$

C.7 Subnetwork Compositions

When creating circuit compositions, we apply the union operator between the binary masks of two circuits, as illustrated in Equation 10. For zero-ablated circuits, this operation remains symmetric. However, in the case of mean-ablated circuits, symmetry is not preserved. This is primarily due to the fact that the mediator value, z_i , for which $m_i = 0$ after the union, is ablated by the mean value across the reference distribution associated with the first circuit's task, $\tilde{z}_i^{T_1}$. Specifically, for the union $\mathbf{m}^{T_1, T_2} = \mathbf{m}^{T_1} \cup \mathbf{m}^{T_2}$, the mediator value z_i is determined as follows:

$$z_i = \begin{cases} \tilde{z}_i^{T_1} & \text{if } m_i^{T_1} = 0 \text{ and } m_i^{T_2} = 0, \\ z_{x_i} & \text{otherwise,} \end{cases} \quad (11)$$

This means that $\mathbf{m}^{T_1} \cup \mathbf{m}^{T_2} \neq \mathbf{m}^{T_2} \cup \mathbf{m}^{T_1}$ for mean ablations if $\tilde{z}_i^{T_1} \neq \tilde{z}_i^{T_2}$. For zero ablation, the ablation value remains the same for both cases $\tilde{z}_i^{T_1} = \tilde{z}_i^{T_2} = 0$, thus yielding symmetry of the operator.

D Supplemental Results

In this section, we report supplemental results that complement the paper's key findings.

D.1 Mean Ablation

We first report additional results related to circuits identified via mean ablation.

D.1.1 Circuit Performance

In Figure 11 and Figure 12 we report the KL divergence between the predicted output distributions of

the mean-ablated circuits, \mathbf{m} , and the base model, \mathcal{M} , for both unary and binary string-edit operations. The results exhibit a trend consistent with the observations discussed in Section 4.3.1.

D.1.2 Local Sparsity

Figure 20 illustrates the local sparsity of each unary circuit identified via mean-ablations, while 21 presents the local sparsity of each binary circuit, respectively. It is evident that unary circuits are highly sparse, with no module retaining more than 45% of activations. Notably, the feed-forward (FF) and multi-head self-attention (MHSA) modules within the decoder demonstrate pronounced sparsity across all layers, typically retaining only 0.0% to 10% of activations. In contrast, the multi-head cross-attention (MHCA) modules retain a higher proportion of activations. As shown in Figure 21, binary circuits generally retain a greater number of activations. Specifically, the append and prepend circuits show substantial remaining activations in both the encoder and decoder.

D.1.3 Lambda Sensitivity Analysis

To assess the impact of the hyperparameter λ in Equation 6, we perform a parameter sensitivity analysis for the two unary operations copy and repeat. Specifically, we consider six different λ values: $\lambda \in 1 \times 10^{-2}, 1 \times 10^{-3}, 1 \times 10^{-4}, 1 \times 10^{-5}, 1 \times 10^{-6}$, and train a mask for each, following the procedure outlined in Sections 3 and 4. We report task faithfulness F_T , KL-divergence D_{KL} (as defined in Section 4.1.3), and the resulting circuit’s sparsity for each task in Table 6.

As expected, we observe that the λ parameter governs the trade-off between task faithfulness and circuit sparsity: higher values of λ place greater emphasis on the regularization term in Equation 5, thereby promoting sparsity in the resulting circuit while sacrificing some faithfulness. In contrast, lower values of λ result in less sparse but more faithful models. For instance, a value of $\lambda = 1 \times 10^{-6}$ yields a high faithfulness performance of $F_{\text{copy}} = 0.999$ on the copy task and a sparsity level of 0.298, whereas a value of $\lambda = 1 \times 10^{-2}$ results in a lower faithfulness performance of $F_{\text{copy}} = 0.228$ and higher sparsity level of 0.022.


D.1.4 Deterministic Mask Approximation

As discussed in Section 3, a key advantage of learning a binary mask through continuous sparsification


is the deterministic nature of the approach. To assess whether our method consistently converges to the same circuits despite the stochastic elements of the training process (e.g., dataset shuffling), we examine the node overlap of circuits trained using different random seeds. Our findings demonstrate that the method reliably converges to the same mask for identical subtasks, regardless of the random seed. This holds true for both mean and zero ablation. For example, we report an IoU and IoM of 1.0 when comparing five prepend circuits trained with varying random seeds.

D.1.5 Set Operations

In Figure 13 we present the cross-task accuracy of additional composite circuits as supplementary results to Section 4.3.3. Consistent with previous observations, the union of two circuits acquires functional capacities for subtasks that the original base circuits cannot perform independently. For example, the union $\mathbf{m}^{\text{swap}} \cup \mathbf{m}^{\text{repeat}}$ achieves a 56%



(a) Intersection over Union (IoU).



(b) Intersection over Minimum (IoM).

Figure 8: Node overlap for circuits identified via zero ablation.

accuracy on the reverse task and 31% on shift, representing significant improvements over the performance of these circuits in isolation. Similarly, the union $m^{\text{reverse}} \cup m^{\text{echo}}$ reaches 31% accuracy on swap. In the case of $m^{\text{shift}} \cup m^{\text{echo}}$, we observe that while the composite retains approximately half of the performance on the individual tasks, it also enhances performance on the copy task.

D.2 Zero Ablation

While the main paper focuses on circuits obtained through mean ablation, this section presents the results of experiments conducted using zero ablation.

D.2.1 Circuit Performance

In Figure 14 and Figure 15, we present the task faithfulness performance F_T and generation accuracy for the zero-ablated circuits. Similarly, Figure 16 and Figure 17 illustrate task faithfulness in terms of the KL divergence between the predicted output distributions of the zero-ablated circuits, m , and the base model, \mathcal{M} , for both unary and binary string-edit operations. The key finding from these experiments is that zero ablation yields results closely aligned with those of mean ablation for binary operations, but not for unary operations. Similarly, patterns observed in mean-ablated unary circuits are not evident here. As noted in the literature, zero ablation can significantly shift the distribution, leading to degraded performance, which is one of the reasons for adopting mean ablation (Miller et al., 2024).

D.2.2 Node Overlap

Figure 8 illustrates the IoU and IoM between all circuit-pairs under zero ablation. The observed patterns are consistent with those identified under mean ablation (Figure 6), where unary circuits tend to exhibit higher node overlap than binary circuits.






Figure 9: The overlap between mean and zero ablated circuits for all subtasks.



(a) Mean ablation.



(b) Zero ablation.

Figure 10: Overall remaining nodes after activation pruning.

D.3 Comparison Between Mean and Zero Ablation

Figure 9 illustrates the overlap between the mean and zero ablated circuits for identical subtasks. As discussed in Section 7, we find significant differences in the circuits identified, specifically for unary circuits that tend to engage a smaller portion of the base model’s activation space (refer to Figure 10 for comparison).

D.4 Global Sparsity

In Figure 10 we present the global sparsity of the circuits for both mean and zero ablation. As touched upon in Section 7, we observe that while the two ablation strategies identify circuits of comparable sizes, they do not necessarily yield identical circuits. Additionally, the ordering of circuits in terms of remaining activations is nearly the same for both methods, with the only notable difference being the internal ordering of swap and reverse.

E Use of AI Assistants

For parts of our project’s source code, we used GitHub Copilot as an assistant tool.




Figure 11: Task faithfulness measured via $D_{KL}(y^m \parallel y^M)$ for the **mean-ablated unary** circuits. The y-axis corresponds to the circuit, while the x-axis represents the evaluation task. When we only evaluate selected positions, we omit the diagonal, as there are no applicable tokens for comparison.




Figure 12: Task faithfulness measured via $D_{KL}(y^m \parallel y^M)$ for the **mean-ablated binary** circuits. The y-axis corresponds to the circuit, while the x-axis represents the evaluation task. When we only evaluate selected positions, we omit the diagonal, as there are no applicable tokens for comparison.

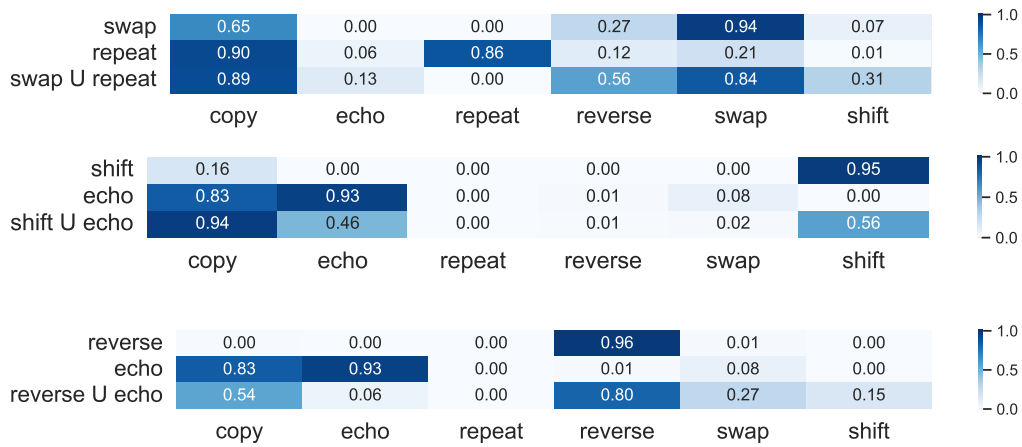


Figure 13: The results of combining circuits through a union operation on their respective binary masks.

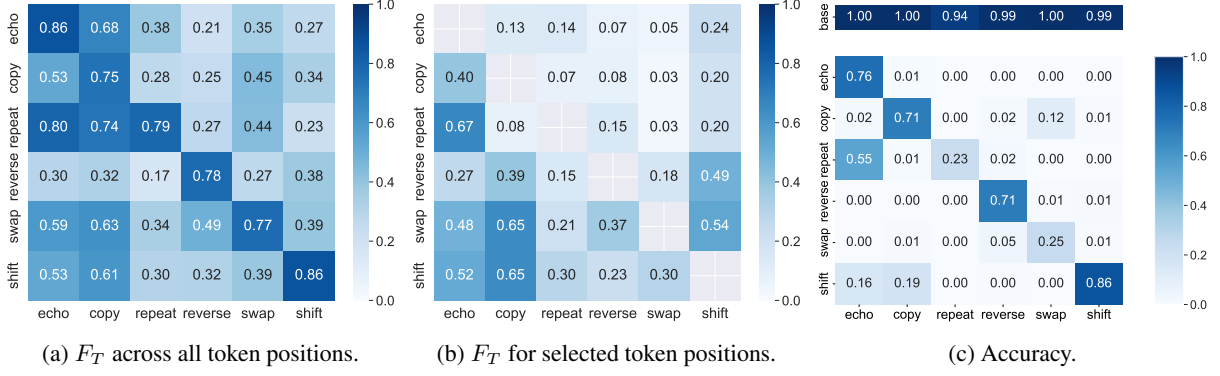


Figure 14: Task faithfulness performance F_T and accuracy for the **unary** tasks for the **zero** ablated circuits. The y-axis corresponds to the circuit, while the x-axis represents the evaluation task. For task faithfulness with respect to positions where the ground truth tokens differ between the circuit and the evaluation task, the diagonal is omitted, as there are no applicable token positions for comparison.

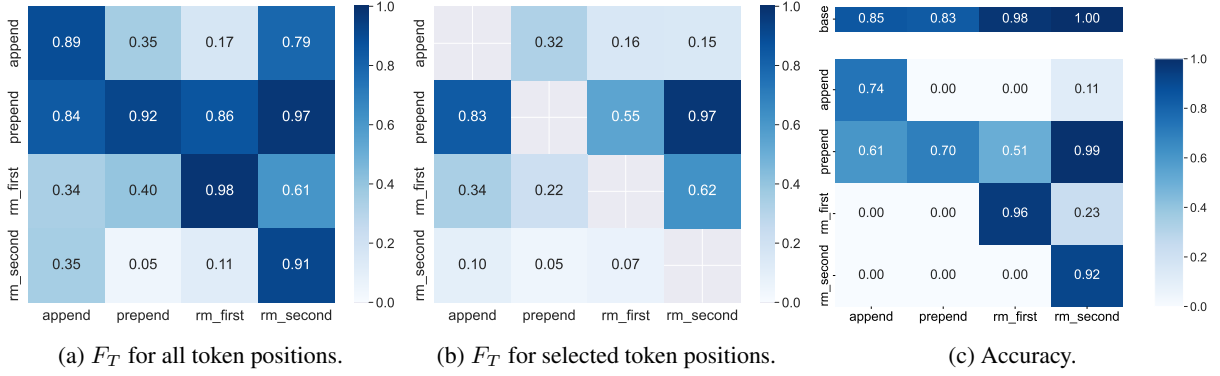




Figure 15: Task faithfulness performance F_T and accuracy for the **binary** tasks for the **zero** ablated circuits. The y-axis corresponds to the circuit, while the x-axis represents the evaluation task. For task faithfulness with respect to positions where the ground truth tokens differ between the circuit and the evaluation task, the diagonal is omitted, as there are no applicable token positions for comparison.



Figure 16: Task faithfulness measured via $D_{KL}(y^m || y^M)$ for the **zero-ablated unary** circuits. The y-axis corresponds to the circuit, while the x-axis represents the evaluation task. When we only evaluate selected positions, we omit the diagonal, as there are no applicable tokens for comparison.




(a) All token positions.




(b) Differing token positions.


Figure 17: Task faithfulness measured via $D_{KL}(y^m \parallel y^M)$ for the **zero-ablated binary** circuits. The y-axis corresponds to the circuit, while the x-axis represents the evaluation task. When we only evaluate selected positions, we omit the diagonal, as there are no applicable tokens for comparison.




(a) The active nodes extracted from the compiled copy TRACR program.



(b) The local sparsity of the copy TRACR circuit.




(c) The active nodes extracted from the compiled echo TRACR program.




(d) The local sparsity of the echo TRACR circuit.


Figure 18: The activation patterns from the compiled copy and echo TRACR programs along the corresponding circuits we identify. We display the fraction of the circuit's remaining activations for each layer and module.




(a) The active nodes extracted from the compiled reverse TRACR program.



(b) The local sparsity of the reverse TRACR circuit.



(c) The active nodes extracted from the compiled swap TRACR program.



(d) The local sparsity of the swap TRACR circuit.

Figure 19: The activation patterns from the compiled reverse and swap TRACR programs along the corresponding circuits we identify. We display the fraction of the circuit’s remaining activations for each layer and module. We note that for these two circuits, we prune the first attention module because we do not consider the BOS token during our pruning, which is the only activated position in this layer.

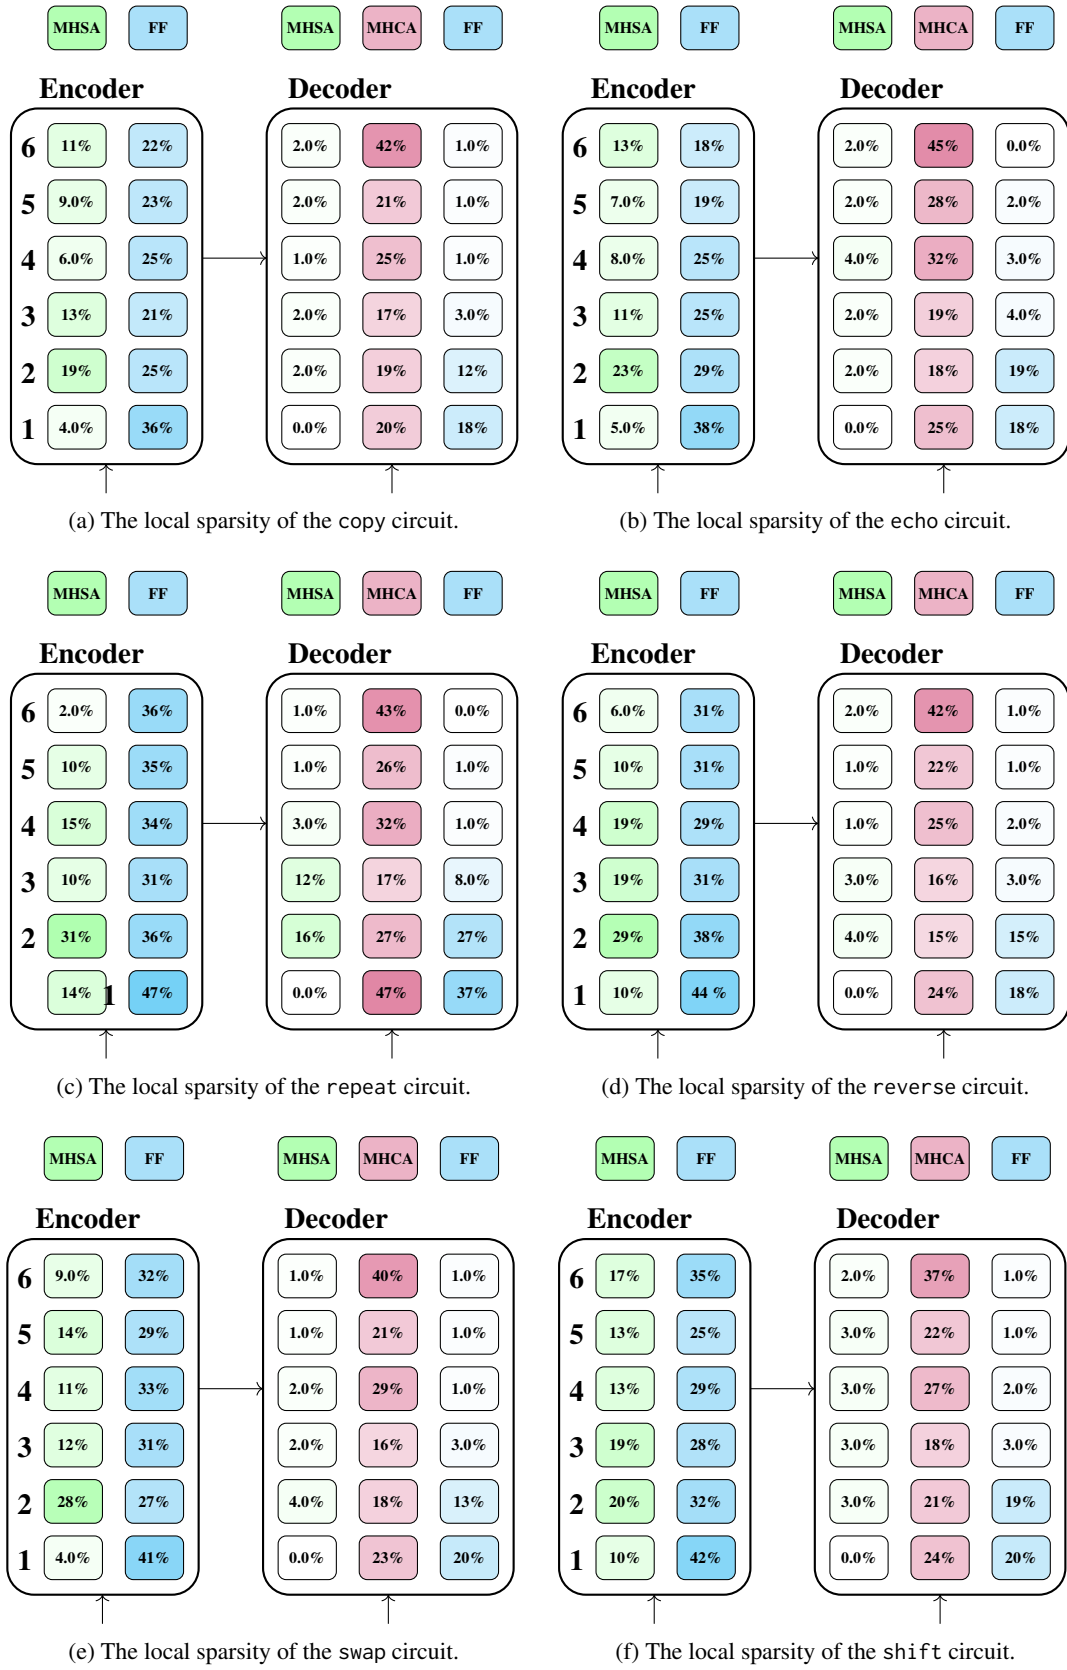


Figure 20: The local sparsity of the **unary** circuits achieved via mean ablation. Considered are feed-forward (FF), multi-head self-attention (MHSA), and multi-head cross-attention (MHCA) modules for each layer.

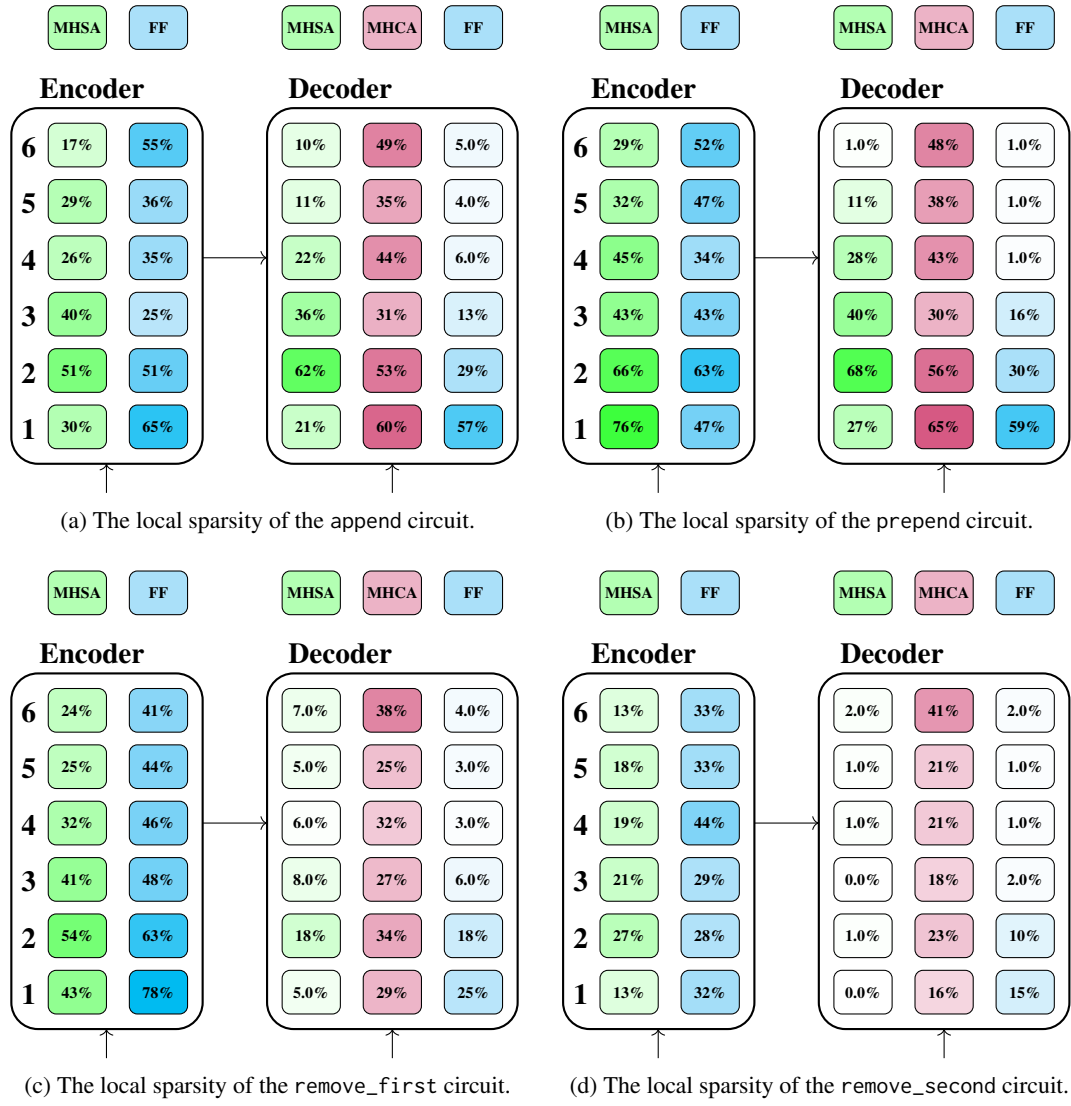


Figure 21: The local sparsity of the **binary** circuits achieved via mean ablation. Considered are feed-forward (FF), multi-head self-attention (MHSA), and multi-head cross-attention (MHCA) modules for each layer.