

# Unveiling Language Skills via Path-Level Circuit Discovery

Anonymous ACL submission

## Abstract

Circuit discovery with edge-level ablation has become a foundational framework for mechanism interpretability of language models. However, its focus on individual edges often overlooks the sequential, path-level causal relationships that underpin complex behaviors, thus potentially leading to misleading or incomplete circuit discoveries. To address this issue, we propose a novel path-level circuit discovery framework capturing how behaviors emerge through interconnected linear chain and build towards complex behaviors. Our framework is constructed upon a fully-disentangled linear combinations of “memory circuits” decomposed from the original model. To discover functional circuit paths, we leverage a 2-step pruning strategy by first reducing the computational graph to a faithful and minimal subgraph and then applying causal mediation to identify common paths of a specific skill, termed as *skill paths*. In contrast to circuit graph from existing works, we focus on the complete paths of a generic skill rather than on the fine-grained responses to individual components of the input. To demonstrate this, we explore three generic language skills, namely *Previous Token Skill*, *Induction Skill* and *In-Context Learning Skill* using our framework and provide more compelling evidence to substantiate stratification and inclusiveness of these skills. Our codes are available at: [https://anonymous.4open.science/r/language\\_skill](https://anonymous.4open.science/r/language_skill).

## 1 Introduction

Mechanism interpretability (Elhage et al., 2021; Conmy et al., 2023) is becoming crucial for understanding how language models work. Current methods (Conmy et al., 2023; Yao et al., 2024; Syed et al., 2023; Bhaskar et al., 2024) involve the use of activation patching to perform counterfactual mediation for each edge, estimating its contribution to the outcome. Edges that exhibit significant causal effects are ultimately retained in the circuit

graph, reflecting a certain mechanism by which the language model processes the input.

However, focusing solely on individual **edges** to ablate certain behaviors can result in many causal relationships in high-level structures being overlooked (e.g., chain structure  $A \rightarrow B \rightarrow C$  and multiple causes  $A \rightarrow C \leftarrow B$ ), which leads to potential pitfalls of counterfactual mediation (Mueller, 2024)<sup>1</sup>. Naturally, the causal mediation in higher-level structures like **path-level** circuit discovery, would bypass these pitfalls. Connected paths in the computational graph offer more complete interpretability and provide insights into how foundational skills are assembled into more sophisticated capabilities.

Extending existing framework to circuit path discovery poses two challenges. Firstly, circuit graphs are usually composed of nodes of distinct functionalities, such as attention-related transformations and multi-layer perceptron (MLP). This **heterogeneity** makes it difficult to maintain a consistent framework for analyzing and comparing paths, as paths may span nodes performing fundamentally different operations. Secondly, compared with edge-level intervention, the computational cost of ablating paths—which are chains of connected edges—increases exponentially with path length, making the process significantly **time consuming**.

To address the above challenges, we propose a novel framework consisting of three steps, namely **Decomposition**, **Pruning**, and **Post-hoc Causal Mediation**, to achieve path-level circuit discovery. Specifically, we introduce the concept of *memory circuits* that are linearly composed into a computation graph. These memory circuits serve as the minimal units/nodes for manipulating the memory reading from models. As such, each path is formed

<sup>1</sup>Two typical examples of these pitfalls are **non-transitivity** and **preemption**, proposed by Mueller (2024), elaborated in Appendix A.

by these functionally equivalent memory circuits (**Decomposition**), thereby alleviating the first challenge. Built upon the memory circuits, we further propose to decouple pruning from causal mediation which are entangled in existing works (Conmy et al., 2023; Yao et al., 2024; Syed et al., 2023; Bhaskar et al., 2024), thereby mitigating the second challenge. Concretely, we apply a coarse pruning technique in the first stage to eliminate edges, resulting in a subgraph that serves as a “prototype” of the circuit graph (**Pruning**). In the second stage, we perform counterfactual operations and causal interventions on a large number of “prototypes” over samples to identify common paths corresponding to a specific skill (**Post-hoc Causal Mediation**).

Compared to existing methods, we do not perform circuit discovery on a single sample. Instead, we utilize a large number of samples containing the same skill to identify generic skill paths. These paths serve as “inherent and input-agnostic skills of the model” rather than “reflections of the model in response to a specific input”. To show the potential capability of skill paths, We select three generic and progressively complex skills which have been introduced in (cro, 2024; Ren et al., 2024; Edelman et al., 2024; Olsson et al., 2022): a) *Previous Token skill* which is responsible for receiving information from the previous token; b) *Induction Skill* which duplicates tokens with the same prefix; and c) *ICL Skill* which perform inference based on similar patterns appeared in demonstrations. Utilizing our 3-step framework, we unveil the complete skill paths of these skills. These skill paths have better interpretability in skill interaction, providing stronger evidence to confirm 2 conjectures that have long remained unverified: **Stratification**: Simple language skills reside in shallow layers, whereas complex language skills are found in deeper layers. **Inclusiveness**: Complex language skills are formed on top of simpler language skills.

In summary, our contributions are 3-fold:

- We introduce *memory circuit*, which can describe all key components of a language model using a unified memory-reading functionality. Combined with compensation circuits, it enables the decomposition of the computational graph into a fully linear and lossless combination, thereby providing a theoretical foundation for path-level circuit discovery.
- We propose a novel 3-step framework for path-level circuit discovery that allows for the iden-

tification of input-agnostic circuit, including complete paths of a generic skill.

- Our analysis and experiments verify 2 properties among the Previous Token Skill, Induction Skill, and ICL Skill, which include stratification and inclusiveness.

## 2 Related Work

Existing work primarily focuses on discovering circuits responsible for processing specific inputs. Specifically, they provide counterfactual text with slight perturbations to the input as patches (Wang et al., 2023), use interchange ablation methods to assess the causal effect of each edge on the output (Yao et al., 2024), and apply various pruning strategies (Conmy et al., 2023; Syed et al., 2023; Bhaskar et al., 2024) to identify circuits formed by salient edges. For the resulting circuit graph, further fine-grained exploration (e.g., noising and denoising (Heimersheim and Nanda, 2024)) is often performed to confirm the mechanisms responsible for different parts of the input text.

Existing work has identified circuits as fine-grained behaviors that respond to specific inputs. In contrast, the circuits we identify encompass the complete global skill. For example, in the IOI (Wang et al., 2023) samples, existing work has identified induction heads. Induction heads may serve as instantiations of the induction skill, but they represent only a partial mechanism derived from specific input samples. Different input samples can lead to the identification of distinct induction heads, highlighting their limited generality. In contrast, our work focuses on uncovering the complete circuits underlying the induction skill by identifying the comprehensive skill paths, offering a more holistic understanding of the mechanism.

## 3 Method

In this paper, we propose a novel 3-step framework to extract the target language skills.

- **Step 1** (Section 3.1): We decouple the architecture of transformer language models into “memory circuits” with completely linear components. This results in a *Complete Circuit Graph*,  $\mathcal{G}$ .
- **Step 2** (Section 3.2): We adopt greedy search to remove redundant edges in  $\mathcal{G}$ , retaining only those paths necessary for predicting the

Index	Category	Implementation ( $X$ =input representation in each layer)
$C^0$	Self	$X$
$C^{1-12}$	Attention	$A^h \ln(X) W_V W_O + A^h b_V W_O$
$C^{13}$	MLP	$atv(\ln(X) W_{M1}) W_{M2}$
$C^{14-25}$	Attention+MLP	$atv(\ln(A^h \ln(X) W_V W_O + A^h b_V W_O) W_{M1}) W_{M2}$
$C^{26}$	Compensation	$(atv(\ln((\sum_{h=1}^{12} C^h) W_{M1})) - \sum_{h=1}^{12} atv(\ln(C^h) W_{M1})) W_{M2}$
$C^{27}$	Compensation	$(atv(\ln(C^{0-13}) W_{M1}) - atv(\ln(C^0) W_{M1}) - atv(\ln(\sum_{h=1}^{12} C^h) W_{M1})) W_{M2}$
$C^{28}$	Bias	$b_v + atv(b_{M1}) W_{M2} + b_{M2} + \sum_{h=1}^{12} act(b_V W_{M1}) W_{M2}$

Table 1: Specific circuit index and corresponding implementation in each layer of GPT2-small.  $W$  and  $b$  represent weight and bias parameters,  $atv$  represents the activation of MLP.  $\ln(\cdot)$  is the layernorm function.  $A = softmax(XW_Q W_K^T X^T + b_Q W_K^T X^T + XW_Q b_K^T + b_Q b_K^T)$ . Memory Circuits are  $C^{1-25}$ .

original output token and resulting in an *Irreducible Circuit Graph*,  $\mathcal{G}^*$ .

- **Step 3** (Section 3.3): We select those paths rendering the most significant causal effect in  $\mathcal{G}^*$  as the skill paths. The final graph formed by the skill paths is named as *Skill Circuit Graph*, denoted as  $\mathcal{G}^S$ .

### 3.1 Memory Circuit

Building on the foundation of the Transformer Circuit (Elhage et al., 2021), we propose a complete decomposition of the transformer model including the MLP layers. Using tensor products ( $\otimes$ ), we can represent any layer of the transformer model:

$$\begin{aligned}
output = & (Id + \sum_{h \in H} A^h \otimes W_{OV}^h + Id \otimes W_{MLP} \\
& + \sum_{h \in H} A^h \otimes W_{MLP} W_{OV}^h) \cdot X
\end{aligned} \quad (1)$$

where  $X$  represents the input representation in each layer and  $H$  represents the number of attention heads. Matrix  $A$  is given by the attention mechanism  $A = softmax((XW_Q)(XW_K)^T)$ , and  $W_{MLP}$  involves the MLP operation with activation given by  $atv(XW_{M1})W_{M2}$ .  $W_{OV} = W_O W_V$  refers to an “output-value” matrix which computes how each token affects the output if attended to, while  $W_Q, W_K, W_V$  are parameter matrices for query, key and value.  $W_{M1}$  and  $W_{M2}$  are weight parameters in two linear layers. This equation simplifies both the attention and MLP modules into linear matrix mappings, describing how the paths from input to output for each layer are decoupled into four independent circuits: 1)  $C^{self} = Id \cdot X$ ; 2)  $C^{attn} = \sum_{h \in H} A^h \otimes W_{OV}^h \cdot X$ ; 3)  $C^{mlp} = Id \otimes W_{MLP} \cdot X$ ; 4)  $C^{attn+mlp} = \sum_{h \in H} A^h \otimes W_{MLP} W_{OV}^h \cdot X$ . Moreover,  $C^{attn}, C^{mlp}, C^{attn+mlp}$  can be further

factorized as:

$$C^{attn/mlp/attn+mlp} = f(X) \cdot W \quad (2)$$

We use  $f$  to represent a function that can be considered equivalent to an activation function, for instance,  $C^{attn} = \sum_{h \in H} f_{W_{QK}}^{attn}(X) \cdot W_{OV}$ ,  $f_{W_{QK}}^{attn}(X)$  represents the softmax-normalization of the input  $X$  through a weighted accumulation performed by  $QK$  values, i.e.,  $f_{W_{QK}}^{attn}(X) = softmax((XW_Q)(XW_K)^T)X$ .

The function  $f(X)$  possesses the ability for non-linear transformations, while  $W$  is an input-agnostic parameter, which can be understood as a memory learned through training (Geva et al., 2021). Therefore, these three circuits represent the minimum and complete unit for manipulating how much memory to read (i.e., memory-reading operation), and are independent of each other, which we refer to as “**Memory Circuits**” (We elaborate memory circuit in detail in Appendix B.)

In this paper, we select GPT2-small as the target language model, containing 12 layers ( $L = 12$ ) and 12 attention heads ( $H = 12$ ). To provide a complete dissection of the the model at each layer which can precisely recover the original output, we introduce *Bias Circuits* and *Compensation Circuits* (Compensation circuits represent the synergy of the sum of linear terms passing the non-linear function, please refer to Appendix B.4 for more details), apart from *Memory Circuits*, to compensate for the remaining information not covered by the memory circuits. Table 1 shows the specific circuits and their implementation for each layer. Notably, our circuit dissection leads to a lossless decomposition of the original LM layer into fully linear combinations:  $LM_l(X) = \sum_{i=0}^{28} C^i$ .

We treat Memory Circuits as the smallest units and build a Complete Circuit Graph,  $\mathcal{G} = \{\mathcal{C}, \mathcal{E}\}$ , where  $\mathcal{C}$  stands for the set of 29 circuits ( $C^{0-28}$

shown in Table 1, where Attention and Attention+MLP has 12 circuits due to 12 heads given) and  $\mathcal{E}$  represents the edge between any two circuits in successive layers. Any memory circuit  $C^i$  ( $0 \leq i \leq 25$ ) in any layer  $l$  ( $0 \leq l \leq 11$ ), denoted as  $C^{l,i}$ , would receive information streams from all circuits in previous layers, i.e.,  $\mathcal{E} = \{(C^{l_1,i} \rightarrow C^{l_2,j})\} (0 \leq l_1 < l_2 \leq 11, 0 \leq i, j \leq 25)$ . Notably, the lossless decomposition ensures that the insights gained from our circuit network accurately reflect the behavior of the original language model.

### 3.2 Pruning with Greedy Search

We use a greedy search to prune unnecessary edges between Memory Circuits while ensuring that the output of the model does not change significantly. In existing work, various ablation methods and different metrics have been compared (Heimersheim and Nanda, 2024). Ablation and metrics are used to justify whether the edge should be pruned. Ablation means setting an edge’s activation value when it is hypothesized to be pruned. Common strategies include zero ablation, noise ablation, mean ablation, and interchange ablation. (Interchange ablation, also named corrupted prompt, is manifested in the background effect in the Section 3.3.) Metrics indicate which metric is used to assess the causal effect change caused by pruning an edge, with common including rank of token (default in ours), logits difference, and KL divergence. In Figure 3, we compare the differences between different ablation methods and metrics in the resulting skill paths. We find that zero ablation and rank of token are the most suitable approaches in our new framework.

Therefore, we adopt zero ablation and make sure that top  $n^2$  candidates remain unchanged after pruning (details of implementation shown in Appendix D.2). Given that a depth-first search is more likely to remove shallow edges, we employ a breadth-first search. Different search strategies and constraints are compared in Appendix C. As shown in Algorithm 1: We denote  $\mathcal{G}^*$  as the Irreducible Circuit Graph after pruning, and  $\mathcal{E}^*$  as a subset of  $\mathcal{E}$  which only includes those paths encapsulating the information stream necessary for the last token prediction. Briefly,  $\mathcal{G}^*$  represents the nearly smallest but functionally complete subgraph for outputting original tokens. Therefore, the observed skill circuit graph may be regarded as a subset of

---

#### Algorithm 1 Greedy Search for $\mathcal{G}^*$

---

**Require:** Complete Circuit Graph  $\mathcal{G} = \{\mathcal{C}, \mathcal{E}\}$ , prediction  $x_N = \text{Model}(\mathcal{G}, X)$ , number of Layers  $L$  and Circuit Index  $[0, 28]$ . **Ensure:** Irreducible Circuit Graph  $\mathcal{G}^* = \{\mathcal{C}, \mathcal{E}^*\}$

```

 $\mathcal{G}^* = \mathcal{G}, \mathcal{G}' = \mathcal{G}^*$ 
for each Memory Circuit  $C^{l,i} \in \mathcal{C} (0 \leq l < L, 1 \leq i \leq 25)$  do
  for each Memory Circuit  $\tilde{C}^{l',i'} \in \mathcal{C} (0 \leq l' < L, 1 \leq i' \leq 25)$  do
     $P = [[l', i'], [l, i]], \mathcal{G}' = \mathcal{G}^*, \mathcal{E}' = \mathcal{E} * -P$ 

    if  $\text{Model}(\mathcal{G}', X) == x_N$  then
       $\mathcal{G}^* = \mathcal{G}'$ 
    else
       $\mathcal{G}' = \mathcal{G}^*$ 
    end if
  end for
end for
return  $\mathcal{G}^*$ 

```

---

$\mathcal{G}^*$ . To obtain this subset, we need to conduct more detailed causal mediation analysis in Section 3.3.

### 3.3 Post-hoc Causal Analysis

It is widely recognized that most texts require more than one language skill for inference (Arora and Goyal, 2023). Therefore, determining which paths are associated with the observed behavior can be challenging. For this reason and motivated by endeavors in causal effect analysis (Wang et al., 2023; Vig et al., 2020), we divide the effects of any text on the output token into 3 components: **skill effects**, **background effects**, and **self effects for the last token** (abbreviated as **self effects**).

**Skill effects** refer to the impact of the observed skill on the output which is the focus of this paper. **Self effects** denote the impact of only using the last token to predict, which functions like a “bi-gram model” (a model associating one input token with its output token). **Background effects** propose a counterfactual scenario, i.e., what would the effect be if this skill does not occur<sup>3</sup>.

We use the typical example of the “Induction” skill for illustration, which works with an input in the form of “...  $A B \dots A$ ”, where  $A, B$  refers to different tokens. Here the language model is expected to repeat the pattern (“ $A B$ ”) it has seen in the con-

<sup>2</sup>We set  $n = 1$  in our experiments because the token with the highest probability is the most meaningful.

<sup>3</sup>In existing work, background text is typically referred to as **interchange ablation** or **activation patching**, and it is conducted in conjunction with pruning.



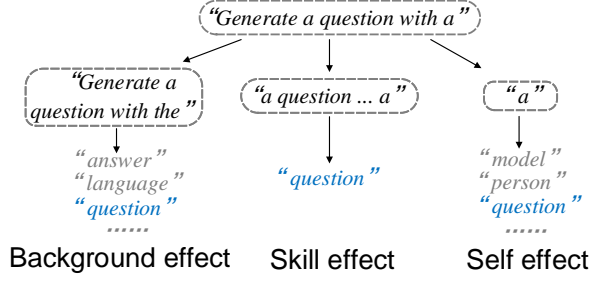


Figure 1: A case text about causal effects.

text and predict token “B” as the output token. Figure 1 illustrates that the model outputs “question” when given the input “Generate a question with a”. However, the vocabulary distribution in the output given by the language model does not merely result from the induction skill, but is also confounded by other effects such as the background effect and the self effect. To compute the target effect for a specific circuit path, let  $Path^i$  be any directed paths in  $\mathcal{G}_*$  (e.g.,  $C^{1,19} \rightarrow C^{2,14} \rightarrow C^{6,5}$  s.t. circuit edges  $(C^{1,19}, C^{2,14})$  and  $(C^{2,14}, C^{6,5})$  are in  $\mathcal{G}_*$ ).  $Path^i$  then symbolizes the flow of information across layers amongst the circuits it encompasses. We use the occurrence rate of  $Path^i$  in all samples to compute the effect:

$$Eff(Path^i_{\mathcal{G}_*}) = \frac{N_{Path^i_{\mathcal{G}_*}=1}}{N_{all}} \quad (3)$$

$N_{Path^i_{\mathcal{G}_*}=1}$  represents the number of samples encompassing  $Path^i$  while  $N_{all}$  represents the number of all samples. Each path contributes differently to the three effects. Hence, we aim to find those paths that contribute to the skill effect rather than the other two effects.

Specifically, for each input text as a sample  $s$ , we perturb it to create a background text  $s_{Bkg}$  and a self text  $s_{Slf}$  (The process for generating background text and self text for all types of skills is described in Appendix D). Eventually, any sample is augmented with two more perturbed versions, rendering three types of inputs (i.e., original text, background text, and self text), each of which is subjected to the greedy search as discussed in Section 3.2. The greedy search produces three distinct Irreducible Circuit Graphs:  $\mathcal{G}_{Ori*}$  (from original input text),  $\mathcal{G}_{Bkg*}$  (from background text), and  $\mathcal{G}_{Slf*}$  (from self text). Therefore, the skill effect (e.g., *Induction Skill*) of  $Path^i$  can be defined as:

$$Eff_{Skill}(Path^i) = \frac{N_{Path^i_{\mathcal{G}_{Ori*}}=1, Path^i_{\mathcal{G}_{Bkg*}}=0, Path^i_{\mathcal{G}_{Slf*}}=0}}{N_{all}} \quad (4)$$

Finally, we get the Skill Circuit Graph  $\mathcal{G}^S = \{\mathcal{C}, \mathcal{E}^S\}$ . With  $\delta$  as the threshold parameter:  $\mathcal{E}^S = \{Path^i | Eff_{Skill}(Path^i) > \delta\}$  (we provided detailed analysis about  $\delta$  in Appendix D.5).

## 4 Experimental Design

This paper focuses on 3 language skills, spanning from basic to advanced levels:

- **Previous Token Skill:** Receive information from the previous token.
- **Induction Skill:** Identify patterns in prefix matching and replicate recurring token sequences.
- **ICL Skill:** Recognize and replicate the demonstration context, thereby producing outputs based on similar patterns.

Extensive research has shown that these three skills build on one another in a sequentially encompassing manner (cro, 2024; Olsson et al., 2022; Ren et al., 2024; Edelman et al., 2024). The Induction Skill inherently includes the Previous Token Skill. In simple terms, for induction to occur in the sequence “A B ... A”, the token B must retrieve information from the preceding token A. Likewise, In-Context Learning must be capable of identifying similar patterns across different demonstrations to generate analogous outputs.

We select over 10k samples encompassing one of the three above-mentioned skills from large corpora and popular datasets such as WIKIQA (Yang et al., 2015), SST-2 (Socher et al., 2013), BIG-BENCH (Srivastava et al., 2023), OpenOrca (Lian et al., 2023), and OpenHermes (Teknium, 2023). For each instance, we create a background perturbation and a self perturbation (discussed in Section 3.3). For simplicity, **PVT** represents the sample set involving the Previous Token Skill and **IDT** represents the sample set related to Induction Skill. **ICL1** represents the ICL sample set from SST-2 datasets; **ICL2** represents the ICL sample set from object\_counting task; **ICL3** and **ICL4** represents those from qawikidata and reasoning\_about\_colored\_objects task. The details of data preparation and implementation are elaborated in Appendix D.

Sample	Circuit Graph								
	$\mathcal{G}_*$	$-R50$	$-R500$	$-\mathcal{G}^{S,PVT}$	$-\mathcal{G}^{S,IDT}$	$-\mathcal{G}^{S,ICL1}$	$-\mathcal{G}^{S,ICL2}$	$-\mathcal{G}^{S,ICL3}$	$-\mathcal{G}^{S,ICL4}$
PVT	1.00	0.46	0.23	0.01	0.00	0.00	0.01	0.00	0.00
IDT	1.00	0.58	0.29	0.08	0.00	0.00	0.00	0.01	0.00
ICL1	1.00	0.61	0.23	0.01	0.00	0.00	0.00	0.00	0.00
ICL2	1.00	0.51	0.18	0.00	0.00	0.01	0.00	0.01	0.01
ICL3	1.00	0.54	0.21	0.00	0.00	0.00	0.00	0.00	0.00
ICL4	1.00	0.62	0.30	0.07	0.03	0.01	0.02	0.00	0.00

Table 2: Accuracy of output to original label within different Circuit Graph




Figure 2: T-sne visualization of 6 types of samples on top 5 vocabulary candidates. Red denotes the original output model ( $\mathcal{G}$ ), while blue signifies the output once a corresponding skill path is removed ( $\mathcal{G} - \mathcal{G}^S$ ). The outputs for the background text ( $\mathcal{G}_{Bkg}$ ) and self text ( $\mathcal{G}_{Slf}$ ) are indicated in green and yellow, respectively.

## 5 Validation

### 5.1 When Skill Paths are Removed

To understand whether the skill paths are responsible for corresponding language skills, we design an intervention experiment by removing different sets of paths and observe the output of the LM.

Table 2 displays the accuracy under different configurations of the Circuit Graphs when treating the original output as the ground-truth. For each language skill  $S$ , we randomly select 500 samples from its corresponding dataset. As a result, 9 different configurations of Circuit Graphs are tested:  $\mathcal{G}_*$  which represents the original output;  $-R50$  which signifies the removal of 50 paths at random from  $\mathcal{G}_*$ ;  $-R500$  after the deletion of 500 paths randomly from  $\mathcal{G}_*$ , which approximately equals the number of skill paths<sup>4</sup>. The remaining 6 configurations encompass the removal of paths from  $\mathcal{G}_*$  that correspond to the skill of Previous Token, Induction, ICL1, ICL2, ICL3, and ICL4, respectively (For additional supplementary data for this validation test, please refer to Appendix D.4.).

The results indicate that almost all samples were unable to produce the original token when these skill paths were excluded (as indicated in the last 6 columns), yet random removal of paths does not lead to such significant impact. Additionally, Fig-

ure 2 visualizes the t-SNE representation of the top 5 candidate outputs associated with different Circuit Graphs. It is clear that when a skill path is removed, the output (blue) shifts from red towards green (or yellow), indicating a transition from a text output distribution that includes skills to a distinct space resulted from the removal of these skills.

### 5.2 How Skill Effects Are Confounded

Another question is whether the background effect and self effect, mentioned in Section 3.3, potentially exist as confounders or share the skill paths. To answer this question, we conduct two experiments, with the results shown in Appendix E. Initially, Table 4 checks the overlap between the paths with  $Eff > 0.5$  in the background/self text and the skill paths, illustrating that a small portion (approximately 10%-20%) of those paths does not belong to any observed skill. This corresponds to the confounding originating from other latent skills that we envisioned. Secondly, Figure 7 visualizes the bivariate probability density function with the original input and background/self text of these paths under different effects. One intriguing discovery is that the confounding skills are more likely to present in the background text than in the self text, and the more complex the skill, the subtler the confounding effect introduced by the self text.

<sup>4</sup>The exact number of removed paths is:  $\mathcal{G}^{S,PVT} - 325$ ,  $\mathcal{G}^{S,IDT} - 466$ ,  $\mathcal{G}^{S,ICL1} - 589$ ,  $\mathcal{G}^{S,ICL2} - 622$ ,  $\mathcal{G}^{S,ICL3} - 603$ ,  $\mathcal{G}^{S,ICL4} - 537$ .

Skill	Receivers receiving more than 10 paths (#layer, #circuit)
PVT	[1, 8], [1, 18], [1, 19], [1, 20], [1, 21], [2, 1], [2, 7], [2, 14], [2, 18], [2, 20], [2, 22], [2, 24], [11, 1], [11, 14]
IDT	[2, 14], [2, 18], [2, 20], [3, 14], [3, 17], [4, 5], [4, 12], [5, 11], [6, 5], [11, 1], [11, 14]
ICL1	[2, 14], [2, 20], [2, 22], [2, 24], [3, 3], [3, 4], [3, 5], [3, 11], [3, 14], [3, 17], [4, 3], [4, 5], [5, 11], [8, 5], [10, 10], [11, 8], [11, 9], [11, 10], [11, 11]
ICL2	[1, 19], [2, 14], [2, 20], [2, 24], [3, 5], [3, 11], [3, 14], [4, 5], [4, 7], [4, 9], [5, 10], [6, 5], [10, 9], [10, 10], [10, 11], [11, 1], [11, 5]
ICL3	[1, 8], [1, 18], [1, 19], [1, 20], [1, 21], [2, 14], [2, 20], [2, 24], [3, 1], [3, 14], [4, 3], [4, 5], [5, 1], [5, 10], [5, 11], [8, 1], [8, 9], [10, 5], [10, 10], [10, 12], [11, 1], [11, 8]
ICL4	[1, 16], [1, 20], [2, 20], [4, 3], [4, 5], [5, 3], [6, 4], [6, 5], [8, 9], [9, 4], [9, 5], [10, 2], [10, 10], [10, 12], [11, 2], [11, 3], [11, 4], [11, 6], [11, 15]

Table 3: Key Receivers in Skill Circuit Graphs, blue circuits are presented in the lower skill.

## 6 Discovery of Language Skills

For a long time, two conjectures about language skills have been increasingly recognized. They are:

1. **Stratification:** Simple language skills reside in shallow layers, whereas complex language skills are found in deeper layers.
2. **Inclusiveness:** Complex language skills are formed on top of simpler language skills.

For **Stratification**, many works have already discovered its traces. For instance, mechanism interpretability studies (Yun et al., 2021; Geva et al., 2021) have found that circuits for simple skills (such as syntax) and complex skills (such as semantics) are almost spread across all layers, but there is a clear concentration: syntax are more concentrated in shallow layers while semantics are more concentrated in deep layers. As for **Inclusiveness**, most existing work has found some components with simple skill features in complex skills, such as the salient induction heads found in ICL tasks, hence conjecturing that the ICL skill includes the Induction skill.

However, so far, there has been no quantitative experimental evidence to prove these 2 conjectures. For instance, which specific layers do simple skills reside in? Which specific paths in simple skills are encompassed by complex skills? Our work further confirms these hypotheses via quantitative discoveries.

### 6.1 Quantitative Results

Table 3 displays the nodes (receivers) receiving more than 10 paths in the skill graphs. We use  $[l, i]$  to denote the  $l$ -th layer and  $i$ -th circuit ( $C^{l,i}$ ). The complete Skill Circuit Graph can be found in Appendix I. From Table 3, we provide quantification results for Stratification and Inclusiveness.

**Quantification of Stratification:** The Previous Token Skill (PVT) is one of the simplest language skills, and thus it is located across layers 0-2. The Induction Skill (IDT) is slightly more complex and thus spreads across layers 0-6. Meanwhile, ICL is the most complex skill and has key receivers across nearly all layers. Additionally, all skills share the 11-th layer (final layer).

**Quantification of Inclusiveness:** Circuits such as [2, 14], [2, 18] and [2, 20] (presented in PVT) can be found in the IDT, indicating that the Previous Token Skill is an integral part of the Induction Skill. Similarly, the ICL skill encapsulates the Previous Token Skill and Induction Skill as necessary sub-skills, which is why circuits that are evident in the Previous Token Skill (such as [2, 14], [2, 20], [2, 24]) and those identified in the Induction Skill (such as [3, 14], [4, 5]) can be found in the ICL Skill Graph. Furthermore, we list all multi-step paths with inclusive sub-path in Appendix F.

### 6.2 Comparisons with Other Methods

In this section, we investigate the performance of different methods and strategies in validating the 2 conjectures. We investigate relevant existing work (ACDC, OPT-prun, EAP) and examine the performance of our method under different ablations: noise ablation (Ours-noise), mean ablation (Ours-mean), zero ablation (Ours) and different metrics: logit difference (Ours-logits), KL divergence (Ours-KL), rank of token (Ours). We introduce these methods in detail in Appendix H.

Each method uses its own circuit discovery strategy to search for corresponding circuit graphs for the three skills we focus on: PVT, IDT, and ICL1. Then, we investigate the distribution of receiver nodes in these circuit graphs and display the normalized results in Figure 3.

It is clear from other methods that PVT is more




Figure 3: Visualization of receivers distributed in layer1-10 in 3 increasingly-complex skills (PVT, IDT, and ICL1).

Method	$ovlp(IDT, PVT)$	$ovlp(ICL1, PVT)$	$ovlp(ICL1, IDT)$	$ovlp(GT, PVT)$	$ovlp(GT, IDT)$	$ovlp(GT, ICL1)$
ACDC	0.19	0.06	0.17	0.11	0.05	0.11
OPT-prun	0.05	0.14	0.17	0.18	0.10	0.05
EAP	0.14	0.05	0.18	0.09	0.15	0.12
Ours	0.74	0.81	0.63	0.68	0.13	0.11

Table 4: Overlaps between different skill circuit graphs

prominent in the shallow layers, IDT in the shallow-to-mid layers, and ICL1 tends to cluster in the deep layers. Yet, the circuit graph discovered by our method (see “Ours”) provides a more distinct differentiation: PVT circuits only appear in layers 1 and 2, and IDT circuits only appear in layers 1-6.

Furthermore, in the cross-comparison of the three types of ablations (noise ablation: Ours-noise, mean ablation: Ours-mean, zero ablation: Ours), there was no significant difference observed in stratification results. This indicates that all three ablations can provide similar subgraphs for the post-hoc causal mediation. However, among the three metrics (logit differences: Ours-logit, KL divergence: Ours-KL, rank of token: Ours), the discrete metric (i.e., rank of token) notably offers clearer stratification results compared to the other two continuous metrics. We posit that this difference arises because discrete metrics can directly reflect the faithfulness of outputs when isolated from causal mediation, thereby accurately removing non-influential edges.

Additionally, to observe the interpretability of these methods on **Inclusiveness**, we investigate their overlap on the three skill circuits: PVT, IDT, and ICL1. In addition, we have added a new skill: greater than (GT), with input samples drawn from the *greater than* dataset (Hanna et al., 2024), to assess the ability of the language model in judging numerical relations. Clearly, the three skills PVT, IDT, and ICL are progressively inclusive, that is, IDT includes PVT, while ICL includes both IDT and PVT. However, GT only includes PVT and has no relation to IDT and ICL. The corresponding circuit graphs are still derived from the circuit discovery strategies proposed by each method, searching in the corpora corresponding to the skills. We use  $ovlp(A, B)$  to indicate the ratio of paths in circuit  $A$  that also exist in  $B$ . A value of 0 means no paths are shared, and 1 means all paths are shared. The specific calculation method for  $ovlp(A, B)$  can be found in Appendix H.

Table 4 demonstrates that the overlap of circuit graphs discovered by existing methods is rather weak. For instance,  $ovlp(ICL1, IDT)$  is only 0.17 in ACDC. In contrast, our approach provides clear empirical evidence for the conjecture of inclusiveness: for instance,  $ovlp(IDT, PVT) = 0.74$  indicates that 74% of the paths in the Induction skill exist in the previous token skill. Additionally, from the experimental results related to GT, it can be seen that for skills without inclusiveness, our circuit graph can also reflect significant differences.

**Advantages:** These comparisons highlight some advantages of our method: the advantage of overlap comes from our ability to effectively identify paths rather than just edges. As can be seen from the paths shown in Appendix F, almost all paths have transitivity<sup>5</sup> with edges. The advantage of stratification may come from the mixture of our novel settings, enabling the model to find potentially more unbiased circuits. Additionally, we present more detailed findings in Appendix G, including some specific components with clear functions and an analysis of why the model fails to execute skills.

<sup>5</sup>Additionally, the preemption is avoided as shown in the backup head in IOI task in Appendix G.



## 7 Limitation

We identify 3 limitations that need to be addressed in the future: 1) the **time complexity** of the greedy search; 2) **scalability**; 3) the lack of further examination on the **representational study**.

Assuming the time for one inference using LLM as  $O(1)$ , the time complexity of a single greedy search would then be  $O(L^2N^2)$ , i.e., the square of the number of layers times the number of circuits. If we can overlook this time-consuming process,  $\mathcal{G}^*$  for each input would effectively facilitate training. In other words,  $\mathcal{G}^*$  could directly instruct LLM which paths are essential and which are not, thus streamlining the training process. Despite the time complexity, we recall our contribution on the analysis of LMs which is usually more challenging and does not require large-scale inference.

We also recognize the limitations of testing on a single model and specific skills. Although many studies have validated the GPT-2 series to have public trustworthiness for research in mechanistic interpretability, making us confident in its capacity to support our contribution—the pioneering work in discovering the theoretical foundation and experimental design of language skills—there remains ample scope for scalability across a variety of models and skills for future work.

Finally, the lack of research at the representational level hinders our progress in answering more complex questions such as why certain samples fail to trigger a skill. Recognized that this is a rather challenging topic, we leave it as a promising future work.

## References

2024. Induction heads as an essential mechanism for pattern matching in in-context learning. *arXiv preprint arXiv:2407.07011*.

Sanjeev Arora and Anirudh Goyal. 2023. [A theory for emergence of complex skills in language models](#). *Preprint*, arXiv:2307.15936.

Deniz Bayazit, Negar Foroutan, Zeming Chen, Gail Weiss, and Antoine Bosselut. 2023. Discovering knowledge-critical subnetworks in pretrained language models. *arXiv preprint arXiv:2310.03084*.

Adithya Bhaskar, Alexander Wettig, Dan Friedman, and Danqi Chen. 2024. Finding transformer circuits with edge pruning. *arXiv preprint arXiv:2406.16778*.

Trenton Bricken, Adly Templeton, Joshua Batson, Brian Chen, Adam Jermy, Tom Conerly, Nick

Turner, Cem Anil, Carson Denison, Amanda Askell, Robert Lasenby, Yifan Wu, Shauna Kravec, Nicholas Schiefer, Tim Maxwell, Nicholas Joseph, Zac Hatfield-Dodds, Alex Tamkin, Karina Nguyen, Brayden McLean, Josiah E Burke, Tristan Hume, Shan Carter, Tom Henighan, and Christopher Olah. 2023. Towards monosemanticity: Decomposing language models with dictionary learning. *Transformer Circuits Thread*. <https://transformer-circuits.pub/2023/monosemantic-features/index.html>.

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.

Benjamin L Edelman, Ezra Edelman, Surbhi Goel, Eran Malach, and Nikolaos Tsilivis. 2024. The evolution of statistical induction heads: In-context learning markov chains. *arXiv preprint arXiv:2402.11004*.

Nelson Elhage, Neel Nanda, Catherine Olsson, Tom Henighan, Nicholas Joseph, Ben Mann, Amanda Askell, Yuntao Bai, Anna Chen, Tom Conerly, Nova DasSarma, Dawn Drain, Deep Ganguli, Zac Hatfield-Dodds, Danny Hernandez, Andy Jones, Jackson Kernion, Liane Lovitt, Kamal Ndousse, Dario Amodei, Tom Brown, Jack Clark, Jared Kaplan, Sam McCandlish, and Chris Olah. 2021. A mathematical framework for transformer circuits. *Transformer Circuits Thread*. <https://transformer-circuits.pub/2021/framework/index.html>.

Mor Geva, Roei Schuster, Jonathan Berant, and Omer Levy. 2021. Transformer feed-forward layers are key-value memories. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 5484–5495.

Michael Hanna, Ollie Liu, and Alexandre Variengien. 2024. How does gpt-2 compute greater-than?: Interpreting mathematical abilities in a pre-trained language model. *Advances in Neural Information Processing Systems*, 36.

Stefan Heimersheim and Neel Nanda. 2024. [How to use and interpret activation patching](#). *Preprint*, arXiv:2404.15255.

David Lewis. 2013. *Counterfactuals*. John Wiley & Sons.

Wing Lian, Bleys Goodson, Eugene Pentland, Austin Cook, Chanvichet Vong, and "Teknium". 2023. Openorca: An open dataset of gpt augmented flan reasoning traces. <https://huggingface.co/Open-Orca/OpenOrca>.

Aaron Mueller. 2024. [Missed causes and ambiguous effects: Counterfactuals pose challenges for interpreting neural networks](#). *Preprint*, arXiv:2407.04690.

Catherine Olsson, Nelson Elhage, Neel Nanda, Nicholas Joseph, Nova DasSarma, Tom Henighan, Ben Mann,

Amanda Ascell, Yuntao Bai, Anna Chen, Tom Conerly, Dawn Drain, Deep Ganguli, Zac Hatfield-Dodds, Danny Hernandez, Scott Johnston, Andy Jones, Jackson Kernion, Liane Lovitt, Kamal Ndousse, Dario Amodei, Tom Brown, Jack Clark, Jared Kaplan, Sam McCandlish, and Chris Olah. 2022. In-context learning and induction heads. *Transformer Circuits Thread*. <https://transformer-circuits.pub/2022/in-context-learning-and-induction-heads/index.html>.

Jie Ren, Qipeng Guo, Hang Yan, Dongrui Liu, Xipeng Qiu, and Dahua Lin. 2024. Identifying semantic induction heads to understand in-context learning. *arXiv preprint arXiv:2402.13055*.

Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D Manning, Andrew Y Ng, and Christopher Potts. 2013. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the 2013 conference on empirical methods in natural language processing*, pages 1631–1642.

Aarohi Srivastava, Abhinav Rastogi, and et al Abhishek Rao. 2023. *Beyond the imitation game: Quantifying and extrapolating the capabilities of language models*. *Transactions on Machine Learning Research*.

Aaquib Syed, Can Rager, and Arthur Conmy. 2023. Attribution patching outperforms automated circuit discovery. *arXiv preprint arXiv:2310.10348*.

Teknium. 2023. *Openhermes 2.5: An open dataset of synthetic data for generalist llm assistants*.

Jesse Vig, Sebastian Gehrmann, Yonatan Belinkov, Sharon Qian, Daniel Nevo, Yaron Singer, and Stuart Shieber. 2020. Investigating gender bias in language models using causal mediation analysis. *Advances in neural information processing systems*, 33:12388–12401.

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*.

Yi Yang, Wen-tau Yih, and Christopher Meek. 2015. *WikiQA: A challenge dataset for open-domain question answering*. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 2013–2018, Lisbon, Portugal. Association for Computational Linguistics.

Yunzhi Yao, Ningyu Zhang, Zekun Xi, Mengru Wang, Ziwen Xu, Shumin Deng, and Huajun Chen. 2024. Knowledge circuits in pretrained transformers. *arXiv preprint arXiv:2405.17969*.

Zeyu Yun, Yubei Chen, Bruno A Olshausen, and Yann LeCun. 2021. Transformer visualization via dictionary learning: contextualized embedding as a linear superposition of transformer factors. *arXiv preprint arXiv:2103.15949*.

## A Details about Non-Transitivity and Preeption

Initially, we would like to recall the description of causal dependency verification using counterfactuals: “If  $A$  occurs, then  $B$  occurs; and if  $A$  does not occur, then  $B$  does not occur” (Lewis, 2013). Building upon this, there are two common pitfalls in existing counterfactual operations (Mueller, 2024): non-transitivity, which leads to the overlooking of local patterns, and preemption, where redundant causes are disregarded. Since instantiating these concepts at the neural level would introduce significant complexity, we instead provide examples from the perspective of events to illustrate these pitfalls.

**Non-transitivity:** Consider the following sequence of events (adapted from Mueller (2024)):

*A: A hiker is walking up a mountain, and a large boulder begins rolling down the mountain toward the hiker.*

*B: The hiker, noticing the rolling boulder, ducks out of the way.*

*C: The hiker survives.*

In this sequence,  $C$  depends on  $B$ , and  $B$  depends on  $A$ . For instance, the dependency of  $C$  on  $B$  can be framed as: “If the hiker ducks, they survive; if the hiker does not duck, they do not survive.” However, counterfactuals cannot directly verify that  $C$  depends on  $A$ . This is because  $A$  and  $C$  are embedded within a more complex causal structure where the mediation between  $A$  and  $C$  involves more than just  $B$ . The same reasoning holds at the neural level: counterfactual mediation can verify the existence of causal relationships such as  $A \rightarrow B$  and  $B \rightarrow C$ , but it cannot directly verify  $A \rightarrow C$ . In this paper, we address this issue by utilizing path-level circuits: if we can simultaneously verify  $A \rightarrow B$ ,  $B \rightarrow C$ , and  $A \rightarrow B \rightarrow C$ , then we can confirm  $A \rightarrow C$ .

**Preeption:** Consider the following example:

*A1: Suzy throws a rock at a glass bottle.*

*A2: Simultaneously, Billy throws a rock at the same bottle.*

*B: The rocks shatter the bottle.*

In this case, verifying  $A1 \rightarrow B$  first would prevent the verification of  $A2 \rightarrow B$  (since the bottle is already shattered), and vice versa. Similarly, in neural networks, a given neuron may be influenced by multiple causes. Performing counterfactual mediation for a single cause may obscure the causal effects of other causes. To resolve this, we apply counterfactual mediation to all paths **simultaneously** in the third step of our framework, ensuring

that this issue is avoided.

## B Analysis about Memory Circuits

### B.1 Full Equations of All Memory Circuits

Three of memory circuits can be further factorized as:

$$\begin{aligned}
C^{attn} &= \sum_{h \in H} f_{W_{QK}}^{attn}(X) \cdot W_{OV} \\
\text{where } f_{W_{QK}}^{attn}(X) &= \text{softmax}((XW_Q)(XW_K)^T)X \\
C^{mlp} &= f_{W_{M1}}^{mlp}(X) \cdot W_{M2} \\
\text{where } f_{W_{M1}}^{mlp}(X) &= \text{atv}(XW_{M1}) \\
C^{attn+mlp} &= \sum_{h \in H} f_{W_{QK}, W_{OV}, W_{M1}}^{attn+mlp}(X) \cdot W_{M2} \\
\text{where } f_{W_{QK}, W_{OV}, W_{M1}}^{attn+mlp}(X) \\
&= \text{atv}(f_{W_{QK}}^{attn}(X)W_{OV}W_{M1})
\end{aligned}$$

We use  $f$  to represent a function that can be considered equivalent to an activation function, for instance,  $f_{W_{QK}}^{attn}(X)$  represents the softmax-normalization of the input  $X$  through a weighted accumulation performed by  $QK$  values. In conclusion, these three types of circuits can be expressed using a common paradigm:

$$C^{attn/mlp/attn+mlp} = f(X) \cdot W \quad (5)$$

### B.2 Why $A \otimes X$ is not the circuit with complete function?

We use  $X^{l,n}$  to denote the hidden state representation corresponding to the  $n$ -th token at the  $l$ -th layer, and  $U$  represents the unembedding matrix. Therefore, for any representation  $X^{l,n}$ , we can obtain its vocabulary distribution, i.e., the logits for each token candidate, using  $X^{l,n}U$ . We adopt a sample text, “*Beats Music is owned by*”, as the input. Table 5 shows the logits corresponding to the words “*the*” and “*Apple*” when these tokens are converted to vocabulary embeddings.

Our expected correct output is such that after the last layer’s representation is unembedded, the logits for “*Apple*” reach their peak. However, as shown in Table 5, after conducting an  $A \otimes X$  operation on the 1st layer’s representation, the logit range for “*Apple*” is [80.49, 86.44], where 80.49 corresponds to the attention weight of “*Music*” to “*by*” being 1, and 86.44 represents the attention weight of “*Be*” to “*by*” being 1.

This situation exposes a significant drawback. In the representations of all previous tokens, the

logits for “*the*” are always higher than those for “*Apple*”. Hence, no matter how many effects  $A \otimes X$  operations performed, it remains impossible for the logits of “*Apple*” to surpass those of “*the*”. Therefore, although  $A \otimes X$  incorporates an activation function such as *softmax*, it can only be considered as semi-activated (Elhage et al., 2021). We refer to this as a “deep constraint”, that is,  $A \otimes X$  cannot allow the representation of the destination token to exceed the upper and lower boundaries of the previous token’s representation. This is why we assert that  $A \otimes X$  lacks full functions, that is, it does not possess memory capability.

### B.3 How to explain Memory Circuits?

Let’s likewise map all the Memory Circuits into the vocabulary space:

$$V = C \cdot U = f(X) \cdot W \cdot U = f(x) \cdot WU \quad (6)$$

Simply put, we assume  $X \in \mathbb{R}^{N,D}$ ,  $f(X) \in \mathbb{R}^{N,M}$ ,  $W \in \mathbb{R}^{M,D}$ , and  $U \in \mathbb{R}^{D,E}$ , where  $N$  represents the number of tokens,  $D$  denotes the dimensions in the residual stream,  $M$  refers to the dimensions in the circuit (such as the dimensions in QKV or MLP), and  $E$  signifies the length of the vocabulary list. Naturally,  $WU \in \mathbb{R}^{M,E}$ , which could be seen as a collection of  $M$  vocabulary distributions. These vocabulary distributions are unaffected by the input tokens and thus can be considered as the acquired memory from training.

The function  $f(X) \in \mathbb{R}^{N,M}$  acts like a weight which specifies how much each vocabulary distribution contributes to the output. This confirms why MLP is generally regarded as a memory storage, as its dimensions are usually significantly larger than those of QKV. Simultaneously, it also explains the advantage of MoE: providing a wider range of options for vocabulary distribution.

In the final analysis, the inference process of a language model can be seen as constituting 3 key components: “**memory**”, “**movement**”, and “**ensemble**”. “**Memory**” pertains to acquiring a new distribution from memory distribution, while “**movement**” involves transferring token information to subsequent tokens. Finally, “**ensemble**” refers to the process of combining representations from multiple circuits to produce the final representation. Within this process, Memory Circuits serve as the smallest units responsible for “**memory**” and also encompass independent operations of “**movement**” ( $C^{1-12}$  and  $C^{14-25}$ ). Further-

Logits	Tokens					
	“Be”	“ats”	“Music”	“is”	“owned”	“by”
“the”	95.45	89.43	91.20	99.32	94.21	101.52
“Apple”	86.44	82.13	80.49	82.31	82.57	83.41

Table 5: Logits of “the” and “Apple” when the representation in 1-st layer products unembedding matrix, with input “Beats Music is owned by”

more, they form individual elements of the “ensemble”. Therefore, we examine the interrelationships (necessary paths) between Memory Circuits to understand the language skills of language models.

#### B.4 Derivation of Compensation Circuits

The input of the MLP consists of two parts: the residual stream and the output of the attention. Due to the presence of nonlinear activation functions, the residual stream and attention are coupled in the input, making it impossible to isolate their impact on the MLP, thereby affecting the verification of pruning. To address this, we introduce a compensation circuit, decomposing the MLP into four parts:

$$\begin{aligned}
& atv((X + \sum_{h \in H} Attn^h)W_{M1})W_{M2} = \\
& (atv(XW_{M1}) + \sum_{h \in H} atv(Attn^h W_{M1}))W_{M2} \\
& + Cps^1 + Cps^2 \\
& \text{where :} \\
& Cps^1 = (atv((X + \sum_{h \in H} Attn^h)W_{M1}) \\
& - atv(XW_{M1}) - atv(\sum_{h \in H} Attn^h W_{M1}))W_{M2} \\
& Cps^2 = (atv(\sum_{h \in H} Attn^h W_{M1}) \\
& - \sum_{h \in H} atv(Attn^h W_{M1}))W_{M2}
\end{aligned} \tag{7}$$

where MLP operation with activation given by  $atv((X + \sum_{h \in H} Attn^h)W_{M1})W_{M2}$  ( $W_{M1}$  and  $W_{M2}$  are weight parameters in two linear layers and  $atv$  represents the activation function),  $X$  represents the input representation in each layer and  $H$  represents the number of attention heads,  $Attn^h$  represents the output of  $h$ -th attention head,  $Cps^1$  and  $Cps^2$  are compensation circuit, representing the synergy effect of the residual stream ( $X + \sum_{h \in H} Attn^h$ ) and the sum of attention head  $\sum_{h \in H} Attn^h$  respectively.

The compensation circuit calculates the synergy between the output when linear terms are summed before passing through a non-linear function, and the output passing through a non-linear function before summing. Therefore, the compensation circuit is dynamic and related to the input. From the perspective of the MLP, if we want the compensation circuit to be 0, then the input to the MLP must be reduced to only one or zero linear terms. This is an unlikely occurrence in practical pruning, so we assume that all edges of the compensation circuit always exist.

#### C Search Strategies

We conducted extensive comparisons w.r.t. two elements: breadth-first search and top1 candidate consistency. 1000 samples, each less than 30 tokens in length, were randomly selected from the WIKIQA dataset (Yang et al., 2015) and applied to different search strategies:

- Breadth-1: Breadth-first search was conducted on  $C^{l,i}$  where  $l$  varies from 0 to 11, and  $i$  from 1 to 25.
- Breadth-2: The same breadth-first search was done on  $C^{l,i}$ , but with  $l$  running from 0 to 11 and  $i$  from 25 to 1.
- Breadth-3:  $l$  spanned from 11 to 0 and  $i$  from 25 to 1 while conducting breadth-first search on  $C^{l,i}$ .
- Breadth-4: The breadth-first search on  $C^{l,i}$  was performed randomly.
- Depth: The depth-first search on  $C^{l,i}$  was undertaken with  $l$  ranging from 0 to 11 and  $i$  from 1 to 25 (i.e., treating  $C^{l,i}$  as the sender rather than the receiver).
- Top-2: Altered constraint to ensure top 2 candidates’ token consistency.
- Top-5: Altered constraint to ensure top 5 candidates’ token consistency.



- Top-10: Changed constraint to ensure top 10 candidates' token consistency.
- Loss-1: The constraint was modified to ensure that  $x_N$ 's loss does not exceed the original loss by more than  $\pm 5$ .
- Loss-2: The constraint was changed to ensure the loss of  $x_N$  does not exceed  $\pm 100\%$  of the original loss.

We measured two metrics: Deleted Path, which is the total number of deleted paths divided by the total number of paths and times 100%, and Hamming, which is the Hamming distance between  $G^*$  obtained from each strategy and  $G^*$  obtained from Breadth-1.

Table 6 presents the results of these methods. Notably, different search sequences of breadth-first search do not lead to significant discrepancies. Depth-first search methods, however, are not as effective as breadth-first searches in deleting a sufficient number of paths. Compared to the top 1 constraint, it is challenging for other constraints to delete an adequate quantity of paths. We posit that this is because GPT2-small is a simple model and does not possess the capability to randomly select candidates from the top N for output.

## D Data Preparation and Implementations

### D.1 Data Preparation

#### D.1.1 Previous Token Skill

We randomly selected 40k text samples comprising two tokens - "*token0 token1*" - from the WIKIQA, OpenOrca, and OpenHermes corpora. In 20k of these samples, the two tokens made up one word, while in the remaining 20k, "*token0*" and "*token1*" belonged to two separate words. For the background text, we chose "*token0*", and for the self text, we selected "*token1*". A complete sample is as follows:

*{text: " that most", background\_text: " that", self\_text: " most", GPT2-small\_output: " of"}*

#### D.1.2 Induction Skill

The samples for the Induction Skill also come from WIKIQA, OpenOrca, and OpenHermes. We randomly selected 14k samples with the template "... *A1 B ... A2*", where the destination token "*A2*" is the same as the preceding token "*A1*", and the total token length of the sample does not exceed 30. For the background text, we removed "*A2*"

and had GPT2-small produce a new but different token to replace "*A2*", resulting in "... *A1 B ... C*". Since "*C*" is semantically supplemented by the preceding text and differs from "*A2*", it preserves semantics as much as possible without the Induction Skill. The self text is still token "*A2*". A complete sample is as follows:

*{text: "chinese lesson 1.2: chinese", background\_text: "chinese lesson 1.2: The", self\_text: "chinese", GPT2-small\_output: " lesson"}*

### D.1.3 ICL Skill

The 4 types of ICL skill samples come from SST-2 dataset and the object\_counting, qawikidata, reasoning\_about\_colored\_objects datasets in BIG-BENCH. These samples have been named by us as *icl\_sst2*, *icl\_oc*, *icl\_qa*, *icl\_raco*, with quantities of 1000, 284, 1000, and 135 respectively. Each sample is required to contain two different labelled demonstrations and should be answerable correctly by GPT2-small. Here are examples of the four types of samples:

*icl\_sst2:*

*{text: ", nor why he keeps being cast in action films when none of them are ever any good Sentiment: negative\nfunny , even punny 6 Sentiment: positive\nis that secret ballot is a comedy , both gentle and biting . Sentiment:", background\_text: "is that secret ballot is a comedy , both gentle and biting . Sentiment:", self\_text: " Sentiment:", GPT2-small\_output: " positive"}*

*icl\_oc:*

*{text: "I have a piano, a trombone, a violin, and a flute. How many musical instruments do I have?A: four\nI have a banana, a plum, a strawberry, a nectarine, an apple, a raspberry, an orange, a peach, a grape, and a blackberry. How many fruits do I have?A: ten\nI have a head of broccoli, a cauliflower, a stalk of celery, a cabbage, a potato, an onion, a yam, a garlic, a lettuce head, and a carrot. How many vegetables do I have?A:", background\_text: "I have a head of broccoli, a cauliflower, a stalk of celery, a cabbage, a potato, an onion, a yam, a garlic, a lettuce head, and a carrot. How many vegetables do I have?A:", self\_text: " A:", GPT2-small\_output: " ten"}*

*icl\_qa:*

*{text: "The country of University of Tsukuba is A: Japan\nThe sport played by Judit Polgár is A: chess\nThe country of citizenship of Théophile Gautier is A:", background\_text: "The country of citizenship of Théophile Gautier is A:", self\_text: "*

Metrics	Strategies									
	Breadth-1	Breadth-2	Breadth-3	Breadth-4	Depth	Top-2	Top-5	Top-10	Loss-1	Loss-2
Deleted Path(%)	69%	68%	68%	70%	9%	32%	14%	2%	25%	34%
Hamming	0	14	21	27	26457	12947	21639	44712	21773	16721

Table 6: Logits of “the” and “Apple” when the representation in 1-st layer products unembedding matrix, with input “Beats Music is owned by”

A:”, GPT2-small\_output: “France”}  
icl\_raco:  
{text: “On the nightstand, you see the following objects arranged in a row: a black bracelet, a pink booklet, a blue cup, and a silver cat toy. What is the color of the object directly to the left of the pink object? A: black\nOn the floor, you see a bunch of objects arranged in a row: a red cup, a gold bracelet, a fuchsia puzzle, a purple stress ball, and a burgundy fidget spinner. What is the color of the object directly to the right of the cup? A: gold\nOn the table, you see a set of things arranged in a row: a black keychain, a purple mug, a blue dog leash, and a teal sheet of paper. What is the color of the left-most thing? A:”, background\_text: “On the table, you see a set of things arranged in a row: a black keychain, a purple mug, a blue dog leash, and a teal sheet of paper. What is the color of the left-most thing? A:”, self\_text: “A:”, GPT2-small\_output: “black”}

## D.2 Implementation

In the pruning process, we employed metrics such as zero ablation and token rank. Specifically, zero ablation involves setting the activation value of an edge to 0 when determining whether it should be pruned. Meanwhile, preserving the top 1 token, as illustrated in Algorithm 1, ensures that an edge will only be deleted if it does not alter the final output token.

Following the 3-step process from Section 3, we obtained the skill circuit graph,  $\mathcal{G}^S$ . We found that the skill effect values in  $\mathcal{G}^S$  for the Previous Token Skill and the Induction Skill were not high, with the highest  $Eff_{Skill}$  being only 0.54 and 0.61, respectively. However, the highest  $Eff_{Skill}$  for the ICL Skill reached 0.98. We speculated that because the Previous Token Skill and the Induction Skill are overly simple, there were a significant number of samples that happened to output the correct answers without triggering the corresponding skill paths. For instance, in the text “In China [mainland]”, it’s challenging to confidently determine whether “mainland” was influenced by the bi-gram

model of “China” or if “China” received information from “In”. As such, we attempted to perform bisection clustering for each sample in the Previous Token Skill and Induction Skill, based on the paths with top 10%  $Eff_{Skill}$ .


Figure 4 shows the results of our clustering on the  $\mathcal{G}^S$  for the 3 skills. The x-axis sequentially arranges the top 10% of paths on  $Eff_{Skill}$  from shallow to deep, and the y-axis indicates the mean  $Eff_{Skill}$  of these paths. It’s striking that two clusters in the Previous Skill and Induction Skill: one consistently showing a high  $Eff_{Skill}$ , and the other showing little to no  $Eff_{Skill}$ . This suggests that these low  $Eff_{Skill}$  samples hardly share common paths or trigger common language skills. Meanwhile, the ICL skill does not showcase discriminable clustering, further corroborating our speculation.

Going a step further, we would like to ascertain whether the Previous Token Skill and Induction Skill, after undergoing multiple rounds of “purification” through clustering, could still be divided into two clusters. Therefore, we recursively performed bisection clustering on the higher  $Eff_{Skill}$  cluster each time. Figure 5 and 6 presents the results after each round of clustering. Notably, the Previous Token could not be divided after 2 rounds of clustering, while the Induction Token hit the dividing limit after just 1 round. Considering that the number of clustering rounds for ICL Skill was 0, we believe this supports our hypothesis: the more complicated the skill, the fewer instances of coincidental samples.


Lastly, we verified that bisection clustering significantly outperformed trisection, quad-section, and quintisection clustering. As illustrated in Figure 7, out of all the clusterings, only bisection clustering was able to distinctly segregate two mutually exclusive clusters categorized by high and low  $Eff_{Skill}$ .

## D.3 Sensitivity about Background Text


To compare the sensitivity brought about by different background texts, we designed four different




(a) Previous Token Skill




(b) Induction Skill




(c) ICL1 Skill (icl\_sst2)



(d) ICL2 Skill (icl\_oc)




(e) ICL3 Skill (icl\_qa)




(f) ICL4 Skill (icl\_raco)


Figure 4: bisection clustering on paths with top 10%  $Eff_{Skill}$  for 3 skills



(a) 1-st round clustering




(b) 2-nd round clustering




(c) 3-rd round clustering

Figure 5: 3 rounds clustering in Previous Token Skill






(a) 1-st round clustering




(b) 2-nd round clustering


Figure 6: 2 rounds clustering in Induction Skill




(a) bisection cluster



(b) trisection cluster



(c) quadsection cluster



(d) quintisection cluster

Figure 7: different clustering on Induction Skill



background text formats on the induction skill and compared the changes between the irreducible circuit graph ( $\mathcal{G}^*$ ) of these background texts and the final skill graph ( $\mathcal{G}^S$ ). These formats are as follows:

**Bkg1:** For the induction text “.....A1 B.....A2”, we replace A2 with the output of the large model for “.....A1 B.....”. For example, if the induction text is “Chinese lesson 1.2: Chinese”, the background text is “Chinese lesson 1.2: The”.

**Bkg2:** For the induction text “.....A1 B.....A2”, we directly delete A2. For example, if the induction text is “Chinese lesson 1.2: Chinese”, the background text is “Chinese lesson 1.2: ”.

**Bkg3:** For the induction text “.....A1 B.....A2”, we directly delete A1. For example, if the induction text is “Chinese lesson 1.2: Chinese”, the background text is “lesson 1.2: Chinese”.

**Bkg4:** For the induction text “.....A1 B.....A2”, we replace B with the output of the large model for “.....A1”. For example, if the induction text is “Chinese lesson 1.2: Chinese”, the background text is “Chinese people 1.2: Chinese”.

To intuitively feel these changes, we introduced a metric of percentage Hamming distance,  $HP$ , specifically  $HP(G_1, G_2) = \text{hammingdistance}(G_1, G_2) / (\sum_{G_1} \mathcal{E} + \sum_{G_2} \mathcal{E}) * 100\%$ , i.e., when  $HP=0\%$ , it means that the two graphs  $G_1$  and  $G_2$  completely overlap, and when  $HP=100\%$ , it means that the two graphs do not overlap at all. We show the  $HP$  between  $\mathcal{G}_{Bkg}^*$  and the  $HP$  between  $\mathcal{G}^S$  under any two background texts in Tables 7.

#### D.4 Supplementary Data for Validation

To enhance the transparency and validity of the validation experiment, we have supplemented it with some additional data.

Firstly, Table 2 only provides the accuracy of randomly deleting 50 and 500 edges, however, the dynamics of accuracy as the number of deleted edges changes is not disclosed. Therefore, we demonstrate the dynamics of accuracy in Figure 8 when the number of randomly deleted edges ranges from 50 to 1000. Notably, even with 1000 edges randomly deleted, the accuracy still remains above 0.1 (the total number of edges being considered is 6875). However, deleting the skill graph leads directly to an accuracy close to 0, even if the skill graph only contains around 500 edges. This further illustrates that the skill graph contains more edges that significantly determine the final output.

Secondly, in Table 2, we only showed the situation where low-level skill graphs remove those paths contained in high-level skill graphs. To reinforce the validation, we additionally provide in Table 8 the scenario where samples of low-level skills are only deleted from those edges that exist in the high-level skill graph but not in the low-level skills.

Herein,  $-(\mathcal{G}^{S,PVT} - \mathcal{G}^*)$  represents the deletion of paths in the previous token skill graph that do not exist in the target graph for the target sample, while  $-(\mathcal{G}^{S,IDT} - \mathcal{G}^*)$  represents the deletion of paths in the Induction skill graph that do not exist in the target graph.  $-(\mathcal{G}^{S,ICL1} - \mathcal{G}^*)$ ,  $-(\mathcal{G}^{S,ICL2} - \mathcal{G}^*)$ ,  $-(\mathcal{G}^{S,ICL3} - \mathcal{G}^*)$ , and  $-(\mathcal{G}^{S,ICL4} - \mathcal{G}^*)$  respectively represent the deletion of paths in the ICL1, ICL2, ICL3, and ICL4 skill graphs that do not exist in the target graph for the target sample.

To reiterate, a portion of the paths in the high-level skill graph is identical to a portion of the paths in the low-level skill graph. Table 8 clearly demonstrates that when target samples delete those paths that exist in other skills but not in their own, the accuracy is not significantly affected. For instance,  $-(\mathcal{G}^{S,IDT} - \mathcal{G}^{S,PVT})$  deletes 129 paths, but only reduces the sample accuracy of the previous token skill to 0.88, while the accuracy corresponding to randomly deleting 100 edges is only 0.42 (see Figure 8). In conjunction with Table 2, it explains that only the overlapping part of the Induction skill graph with the previous token skill graph affects the previous token skill. Additionally, when the ICL series skills output paths that exist in other ICLs but not in themselves, their accuracy is somewhat higher (over 0.9). This is due to the ICL series skill graphs being more similar to each other, resulting in fewer paths in the complement.

#### D.5 Threshold and Faithfulness

While we maintain faithfulness on  $\mathcal{G}^*$ , it is difficult to maintain it on  $\mathcal{G}^S$ . In other words, the bias introduced by counterfactuals and interventions is indeed hard to completely avoid, while the faithfulness of pruning is avoidable. Therefore, a circuit graph that clearly reflects the final result will certainly discard some edges of unclear significance. This is usually accomplished through a threshold. We show in Figure 9 the change in accuracy when the threshold  $\delta$  mentioned in Section 3.3 ranges from 0 to 0.9 (there are almost no circuits left when  $\delta > 0.9$ , so we ignore this part). It can be clearly seen that faithfulness can only be fully guaranteed

(a) HP on $\mathcal{G}_{Bkg}^*$					(b) HP on $\mathcal{G}^S$				
	Bkg1	Bkg2	Bkg3	Bkg4		Bkg1	Bkg2	Bkg3	Bkg4
Bkg1	0%	12.54%	9.33%	11.42%	Bkg1	0%	4.37%	5.75%	4.62%
Bkg2	12.54%	0%	6.42%	9.52%	Bkg2	4.37%	0%	3.51%	4.03%
Bkg3	9.33%	6.42%	0%	12.91%	Bkg3	5.75%	3.51%	0%	3.72%
Bkg4	11.42%	9.52%	12.91%	0%	Bkg4	4.62%	4.03%	3.72%	0%

Table 7: HP between different background text. For example, the value in the second row and third column of Figure a is 6.42%, which means  $HP(\mathcal{G}_{Bkg2}^*, \mathcal{G}_{Bkg3}^*) = 6.42\%$  ( $\mathcal{G}_{Bkg2}^*$  and  $\mathcal{G}_{Bkg3}^*$  has 6.42% edges different).




Figure 8: Accuracy with the number of removed edges increasing.




Figure 9: Faithfulness ranging from the  $\delta$

Sample	Circuit Graph						
	$\mathcal{G}_*$	$-(\mathcal{G}^{S,PVT} - \mathcal{G}_*)$	$-(\mathcal{G}^{S,IDT} - \mathcal{G}_*)$	$-(\mathcal{G}^{S,ICL1} - \mathcal{G}_*)$	$-(\mathcal{G}^{S,ICL2} - \mathcal{G}_*)$	$-(\mathcal{G}^{S,ICL3} - \mathcal{G}_*)$	$-(\mathcal{G}^{S,ICL4} - \mathcal{G}_*)$
PVT	1.00	1.00	0.88	0.89	0.89	0.83	0.89
IDT	1.00	0.93	1.00	0.81	0.82	0.85	0.81
ICL1	1.00	0.95	0.81	1.00	0.95	0.93	0.97
ICL2	1.00	0.93	0.84	1.00	0.92	0.95	0.92
ICL3	1.00	0.94	0.86	1.00	0.93	0.91	0.94
ICL4	1.00	0.96	0.83	1.00	0.93	0.94	0.96

Table 8: Accuracy of output to original label within different Circuit Graph

when  $\delta = 0$ . However, such edges are not sparse enough to reflect some specific interpretable functions. When  $\delta > 0.7$ , it is almost impossible to recover to the original input, but the obtained skill graph can correspond well with previous methods. Additionally, in this paper, we default the  $\delta$  for each skill to be PVT: 0.6, IDT: 0.7, ICL1-4: 0.8. Additionally, we have demonstrated in Figures 10 and 11 the changes in the number of edges and the continuous KL divergence metric with varying thresholds  $\delta$ . Specifically, Figure 10 presents the total number of edges in the circuit graph (excluding compensation circuit and bias circuit) under different thresholds, while Figure 11 shows the KL divergence between  $\mathcal{G}^S$  and  $\mathcal{G}_*$  (solid lines) and  $\mathcal{G}^S$  and  $\mathcal{G}$  (dash lines) obtained at different thresholds. Figure 10 clearly indicates that the edges with high causal effects from the previous token skill are the fewest, and the most are from the series of ICL skill, which corroborates the conclusion drawn from the clustering in Appendix D.2. Moreover, the changes in KL divergence (Figure 11) can be roughly divided into four phases (steady, burst, steady, burst). In conjunction with Figure 10, the two bursts are due to the rapid decrease in edges and the number of edges being too few, approaching zero. The default  $\delta$  we selected (PVT 0.6, IDT 0.7, ICL1-4 0.8) are each in the second steady phase. Combining Figures 10 and 11, it suggests that when a large number of edges are deleted, the circuit graph enters a phase of minimal change, which we believe best achieves the “balance between faithfulness and sparsity”.

Additionally, we can observe that the KL divergence between  $\mathcal{G}_*$  and  $\mathcal{G}$  is approximately 10 (as can be seen from the solid and dashed lines corresponding to  $\delta = 0$ ), and generally, the KL divergence between  $\mathcal{G}^S$  and  $\mathcal{G}$  ( $KL(\mathcal{G}^S, \mathcal{G})$ ) is greater than the KL divergence between  $\mathcal{G}^S$  and  $\mathcal{G}_*$  ( $KL(\mathcal{G}^S, \mathcal{G}_*)$ ). Interestingly, as  $\delta$  increases, the values between  $KL(\mathcal{G}^S, \mathcal{G})$  and  $KL(\mathcal{G}^S, \mathcal{G}_*)$  get closer and are almost the same at the default thresh-

old.

## E Details about Validations for Causal Effects

Another question is whether the background effect and self effect, mentioned in Section 3.3, potentially exist as confounders or share the circuits with observed skills? To answer this question, we examine the paths in background/self text with  $Eff > 0.5$ . Table 9 categorizes these paths into 7 types and displays their ratios. Here,  $\mathcal{G}_{PVT}^S$  signifies the ratio of those paths found in the Previous Token Skill graph,  $\mathcal{G}_{IDT}^S$  refers to the ratio of those located in the Induction skill graph, similarly,  $\mathcal{G}_{ICL1}^S$  to  $\mathcal{G}_{ICL4}^S$  represents the ratio of paths in corresponding ICL skill graphs, and “Others” represents the ratio of paths that do not exist in either skill graphs. Notably, a small fraction of high-effect paths does not belong to any observed skill (approximately 0.1-0.2 in “Others”); these are the confounding paths we mentioned before. Additionally, we demonstrated the bivariate probability density function (PDF) in Figure 12. Bivariate PDF constructed from the origin text as one variable, and background text or self text as another one variable. Evidently, across all skills, the paths that have a high effect ( $Eff > 0.5$ ) in the origin text include a part of paths with a relatively high effect ( $Eff > 0.5$ ) in the background text. However, there are nearly ignorable high-effect paths in the self text in ICL skills. We guess that within the ICL skill, the background text and the origin text possess a significantly higher number of tokens compared to the self text, thereby leading to an insignificant effect of the self text.

Additionally, Table 9 also shows that a part of high-effect paths in the background/self text is common with the corresponding skill graph. Fortunately, we need not worry that removing these paths would render the final Skill Graph (paths) incomplete. Appendix F provides evidence that these removed but common paths can always be




Figure 10: number of edges ranging from the  $\delta$




Figure 11: KL divergence ranging from the  $\delta$ , the solid lines represents KL between  $\mathcal{G}^S$  and  $\mathcal{G}^*$ , and the dash lines represents KL between  $mathcal{G}^S$  and  $\mathcal{G}$ .

Skills	$\mathcal{G}_{Bkg}^*$							$\mathcal{G}_{Self}^*$						
	$\mathcal{G}_{PVT}^S$	$\mathcal{G}_{IDT}^S$	$\mathcal{G}_{ICL1}^S$	$\mathcal{G}_{ICL2}^S$	$\mathcal{G}_{ICL3}^S$	$\mathcal{G}_{ICL4}^S$	Others	$\mathcal{G}_{PVT}^S$	$\mathcal{G}_{IDT}^S$	$\mathcal{G}_{ICL1}^S$	$\mathcal{G}_{ICL2}^S$	$\mathcal{G}_{ICL3}^S$	$\mathcal{G}_{ICL4}^S$	Others
Induction	0.76	-	-	-	-	-	0.24	0.84	-	-	-	-	-	0.16
ICL1	0.43	0.38	0.29	0.19	0.25	0.23	0.18	0.51	0.33	0.24	0.16	0.18	0.15	0.15
ICL2	0.46	0.37	0.25	0.16	0.19	0.21	0.17	0.61	0.24	0.25	0.14	0.19	0.18	0.15
ICL3	0.45	0.35	0.23	0.21	0.15	0.19	0.20	0.60	0.28	0.25	0.16	0.18	0.19	0.11
ICL4	0.49	0.36	0.25	0.19	0.26	0.14	0.16	0.61	0.25	0.23	0.19	0.16	0.13	0.13

Table 9: Ratio of high  $Eff$  path ( $Eff > 0.5$ ) in  $\mathcal{G}_{Bkg}^*$  and  $\mathcal{G}_{Self}^*$  (The sum of ratios  $> 1$  due to overlaps in each item).



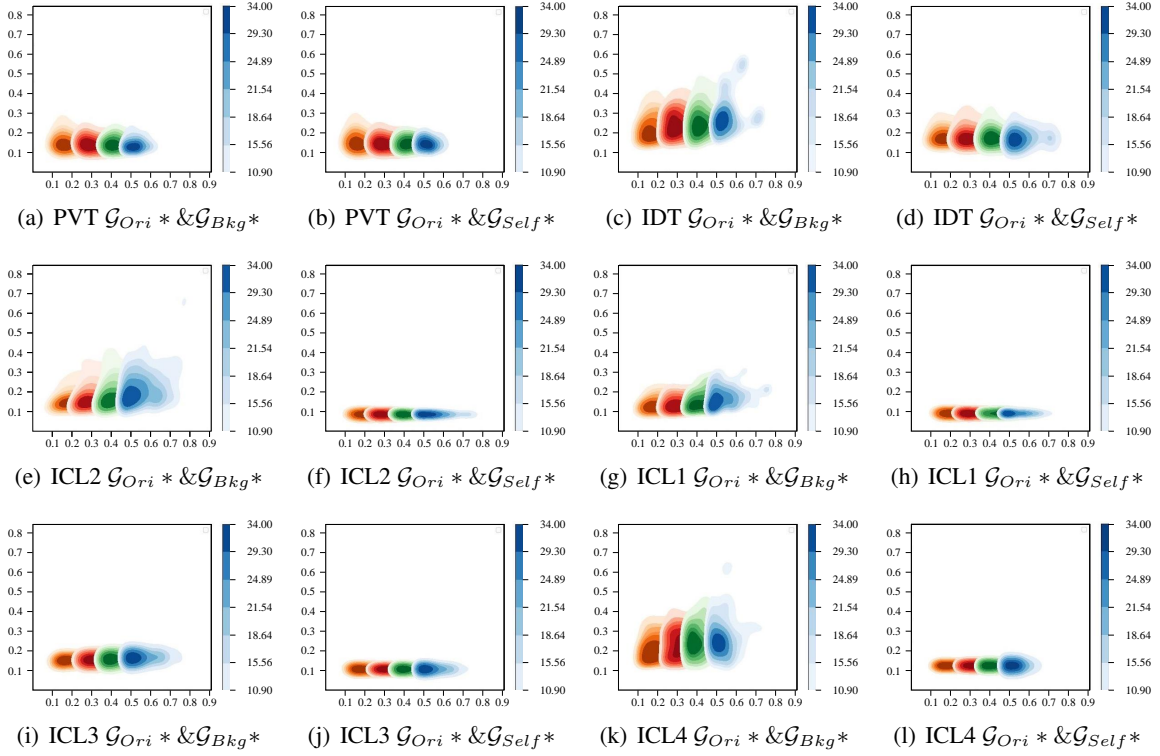


Figure 12: Bivariate probability density function (PDF) of path effects on Previous Token, Induction, ICL1, ICL2, ICL3, and ICL4 Skills. The x-axis represents the first variable, the path effect in the origin text ( $\mathcal{G}_{Ori}^*$ ) while the y-axis represents the second variable, the path effect in the background/self text ( $\mathcal{G}_{Bkg}^* / \mathcal{G}_{Self}^*$ ). Orange, red, green, and blue respectively represent the distribution of paths with  $Eff > 0.2, 0.3, 0.4, 0.5$  in the origin text.

restored through multi-step paths (We explain this phenomenon as ‘Inclusiveness’ in Section 6.).

We have supplemented the bivariate distribution figures for Previous Token, ICL2, ICL3, and ICL4, as depicted in Figure 12.

## F Inclusive Path

we have listed the whole paths for Previous Token Skills, all multi-step paths for the Induction and ICL1 Skills in following, with index of the send circuit, the first receive circuit, the second receive circuit.... The blue represents the paths involving inclusive paths.

### Previous Token Skill

layer 0 circuit 13, layer 1 circuit 6, with effect 0.71  
 layer 0 circuit 14, layer 1 circuit 7, with effect 0.82  
 layer 0 circuit 16, layer 1 circuit 7, with effect 0.7  
 layer 0 circuit 20, layer 1 circuit 7, with effect 0.86  
 layer 0 circuit 14, layer 1 circuit 8, with effect 0.79  
 layer 0 circuit 16, layer 1 circuit 8, with effect 0.78  
 layer 0 circuit 17, layer 1 circuit 8, with effect 0.81  
 layer 0 circuit 19, layer 1 circuit 8, with effect 0.72  
 layer 0 circuit 20, layer 1 circuit 8, with effect 0.88

layer 0 circuit 22, layer 1 circuit 8, with effect 0.81  
 layer 0 circuit 23, layer 1 circuit 8, with effect 0.87  
 layer 0 circuit 24, layer 1 circuit 8, with effect 0.75  
 layer 0 circuit 13, layer 1 circuit 18, with effect 0.79  
 layer 0 circuit 13, layer 1 circuit 19, with effect 0.89  
 layer 0 circuit 14, layer 1 circuit 19, with effect 0.83  
 layer 0 circuit 15, layer 1 circuit 19, with effect 0.74  
 layer 0 circuit 16, layer 1 circuit 19, with effect 0.81  
 layer 0 circuit 20, layer 1 circuit 19, with effect 0.82  
 layer 0 circuit 24, layer 1 circuit 19, with effect 0.84  
 layer 0 circuit 13, layer 1 circuit 20, with effect 0.84  
 layer 0 circuit 14, layer 1 circuit 20, with effect 0.81  
 layer 0 circuit 20, layer 1 circuit 20, with effect 0.8  
 layer 0 circuit 13, layer 1 circuit 21, with effect 0.78

1377	layer 0 circuit 14, layer 1 circuit 21, with effect 0.83	14, with effect 0.71	1429
1378		layer 0 circuit 20, layer 1 circuit 21, layer 2 circuit 14, with effect 0.77	1430
1379	layer 0 circuit 16, layer 1 circuit 21, with effect 0.79		1431
1380		layer 1 circuit 22, layer 2 circuit 14, with effect 0.81	1432
1381	layer 0 circuit 17, layer 1 circuit 21, with effect 0.75		1433
1382		layer 1 circuit 23, layer 2 circuit 14, with effect 0.76	1434
1383	layer 0 circuit 20, layer 1 circuit 21, with effect 0.87		1435
1384		layer 1 circuit 24, layer 2 circuit 14, with effect 0.86	1436
1385	layer 0 circuit 22, layer 1 circuit 21, with effect 0.77		1437
1386		layer 0 circuit 13, layer 2 circuit 18, with effect 0.82	1438
1387	layer 0 circuit 23, layer 1 circuit 21, with effect 0.77		1439
1388		layer 1 circuit 13, layer 2 circuit 18, with effect 0.88	1440
1389	layer 0 circuit 24, layer 1 circuit 21, with effect 0.75		1441
1390		layer 0 circuit 19, layer 2 circuit 20, with effect 0.72	1442
1391	layer 0 circuit 23, layer 2 circuit 1, with effect 0.8		1443
1392	layer 0 circuit 24, layer 2 circuit 1, with effect 0.81	layer 0 circuit 20, layer 2 circuit 20, with effect 0.79	1444
1393	layer 1 circuit 13, layer 2 circuit 1, with effect 0.76		1445
1394	layer 1 circuit 15, layer 2 circuit 1, with effect 0.79	layer 0 circuit 21, layer 2 circuit 20, with effect 0.72	1446
1395	layer 1 circuit 16, layer 2 circuit 1, with effect 0.75		1447
1396	layer 1 circuit 17, layer 2 circuit 1, with effect 0.75	layer 0 circuit 22, layer 2 circuit 20, with effect 0.77	1448
1397	layer 1 circuit 20, layer 2 circuit 1, with effect 0.82		1449
1398	layer 0 circuit 13, layer 1 circuit 20, layer 2 circuit 1, with effect 0.74	layer 1 circuit 19, layer 2 circuit 20, with effect 0.75	1450
1399			1451
1400	layer 1 circuit 21, layer 2 circuit 1, with effect 0.8	layer 1 circuit 20, layer 2 circuit 20, with effect 0.76	1452
1401	layer 0 circuit 20, layer 1 circuit 21, layer 2 circuit 1, with effect 0.77		1453
1402		layer 1 circuit 21, layer 2 circuit 20, with effect 0.7	1454
1403	layer 1 circuit 22, layer 2 circuit 1, with effect 0.76	layer 1 circuit 22, layer 2 circuit 20, with effect 0.76	1455
1404	layer 1 circuit 23, layer 2 circuit 1, with effect 0.79		1456
1405	layer 1 circuit 24, layer 2 circuit 1, with effect 0.8	layer 1 circuit 23, layer 11 circuit 1, with effect 0.74	1457
1406	layer 0 circuit 20, layer 2 circuit 14, with effect 0.74		1458
1407		layer 1 circuit 24, layer 11 circuit 1, with effect 0.75	1459
1408	layer 0 circuit 21, layer 2 circuit 14, with effect 0.75		1460
1409		layer 2 circuit 24, layer 11 circuit 1, with effect 0.73	1461
1410	layer 0 circuit 22, layer 2 circuit 14, with effect 0.77		1462
1411		layer 4 circuit 23, layer 11 circuit 1, with effect 0.74	1463
1412	layer 0 circuit 23, layer 2 circuit 14, with effect 0.72		1464
1413		layer 0 circuit 24, layer 11 circuit 14, with effect 0.77	1465
1414	layer 0 circuit 24, layer 2 circuit 14, with effect 0.84		1466
1415		layer 1 circuit 13, layer 11 circuit 14, with effect 0.74	1467
1416	layer 1 circuit 13, layer 2 circuit 14, with effect 0.72		1468
1417		layer 1 circuit 16, layer 11 circuit 14, with effect 0.74	1469
1418	layer 1 circuit 15, layer 2 circuit 14, with effect 0.8		1470
1419	layer 1 circuit 16, layer 2 circuit 14, with effect 0.72	layer 1 circuit 24, layer 11 circuit 14, with effect 0.82	1471
1420			1472
1421	layer 1 circuit 17, layer 2 circuit 14, with effect 0.8	layer 2 circuit 13, layer 11 circuit 14, with effect 0.75	1473
1422	layer 1 circuit 18, layer 2 circuit 14, with effect 0.74		1474
1423		layer 2 circuit 16, layer 11 circuit 14, with effect 0.76	1475
1424	layer 1 circuit 20, layer 2 circuit 14, with effect 0.79		1476
1425		layer 2 circuit 24, layer 11 circuit 14, with effect 0.81	1477
1426	layer 1 circuit 21, layer 2 circuit 14, with effect 0.79		1478
1427		layer 3 circuit 13, layer 11 circuit 14, with effect 0.75	1479
1428	layer 0 circuit 14, layer 1 circuit 21, layer 2 circuit		1480

1481	layer 3 circuit 16, layer 11 circuit 14, with effect
1482	0.75
1483	layer 3 circuit 24, layer 11 circuit 14, with effect
1484	0.81
1485	layer 4 circuit 13, layer 11 circuit 14, with effect
1486	0.76
1487	layer 4 circuit 24, layer 11 circuit 14, with effect
1488	0.81
1489	layer 5 circuit 24, layer 11 circuit 14, with effect
1490	0.82
1491	layer 6 circuit 16, layer 11 circuit 14, with effect
1492	0.76
1493	layer 6 circuit 24, layer 11 circuit 14, with effect
1494	0.79
1495	layer 7 circuit 24, layer 11 circuit 14, with effect
1496	0.77
1497	layer 8 circuit 24, layer 11 circuit 14, with effect
1498	0.78
1499	layer 9 circuit 24, layer 11 circuit 14, with effect
1500	0.77
1501	layer 10 circuit 24, layer 11 circuit 14, with effect
1502	0.77

#### **Multi-Step Paths in Induction Skill**

1504	layer 0 circuit 20, layer 2 circuit 14, layer 5
1505	circuit 11, with effect 0.6
1506	layer 0 circuit 21, layer 2 circuit 14, layer 5 circuit
1507	11, with effect 0.6
1508	layer 1 circuit 16, layer 2 circuit 14, layer 5 circuit
1509	11, with effect 0.6
1510	layer 1 circuit 18, layer 2 circuit 14, layer 5 circuit
1511	11, with effect 0.6
1512	layer 1 circuit 20, layer 2 circuit 14, layer 5 circuit
1513	11, with effect 0.6
1514	layer 1 circuit 21, layer 2 circuit 14, layer 5 circuit
1515	11, with effect 0.6
1516	layer 1 circuit 22, layer 2 circuit 14, layer 5 circuit
1517	11, with effect 0.61
1518	layer 0 circuit 13, layer 2 circuit 20, layer 5 circuit
1519	11, with effect 0.6
1520	layer 0 circuit 20, layer 2 circuit 14, layer 11 circuit
1521	1, with effect 0.61
1522	layer 0 circuit 21, layer 2 circuit 14, layer 11 circuit
1523	1, with effect 0.63
1524	layer 1 circuit 18, layer 2 circuit 14, layer 11 circuit
1525	1, with effect 0.61
1526	layer 1 circuit 20, layer 2 circuit 14, layer 11 circuit
1527	1, with effect 0.61
1528	layer 1 circuit 21, layer 2 circuit 14, layer 11 circuit
1529	1, with effect 0.61
1530	layer 1 circuit 22, layer 2 circuit 14, layer 11 circuit
1531	1, with effect 0.63
1532	

#### **Multi-Step Paths in ICL1 Skill**

layer 0 circuit 13, layer 1 circuit 19, layer 3	
circuit 11, with effect 0.81	
layer 0 circuit 14, layer 1 circuit 19, layer 3 circuit	
11, with effect 0.85	
layer 0 circuit 15, layer 1 circuit 19, layer 3 circuit	
11, with effect 0.84	
layer 0 circuit 16, layer 1 circuit 19, layer 3 circuit	
11, with effect 0.85	
layer 0 circuit 21, layer 1 circuit 19, layer 3 circuit	
11, with effect 0.82	
layer 0 circuit 22, layer 1 circuit 19, layer 3 circuit	
11, with effect 0.85	
layer 0 circuit 23, layer 1 circuit 19, layer 3 circuit	
11, with effect 0.84	
layer 0 circuit 24, layer 1 circuit 19, layer 3 circuit	
11, with effect 0.85	
layer 0 circuit 13, layer 2 circuit 14, layer 3 circuit	
11, with effect 0.81	
layer 0 circuit 20, layer 2 circuit 14, layer 3 circuit	
11, with effect 0.81	
layer 0 circuit 21, layer 2 circuit 14, layer 3 circuit	
11, with effect 0.83	
layer 0 circuit 22, layer 2 circuit 14, layer 3 circuit	
11, with effect 0.83	
layer 1 circuit 20, layer 2 circuit 14, layer 3 circuit	
11, with effect 0.81	
layer 1 circuit 21, layer 2 circuit 14, layer 3 circuit	
11, with effect 0.82	
layer 1 circuit 22, layer 2 circuit 14, layer 3 circuit	
11, with effect 0.83	
layer 1 circuit 23, layer 2 circuit 14, layer 3 circuit	
11, with effect 0.8	
layer 0 circuit 13, layer 2 circuit 20, layer 3 circuit	
11, with effect 0.86	
layer 0 circuit 14, layer 2 circuit 20, layer 3 circuit	
11, with effect 0.85	
layer 0 circuit 15, layer 2 circuit 20, layer 3 circuit	
11, with effect 0.81	
layer 0 circuit 16, layer 2 circuit 20, layer 3 circuit	
11, with effect 0.85	
layer 0 circuit 17, layer 2 circuit 20, layer 3 circuit	
11, with effect 0.85	
layer 0 circuit 18, layer 2 circuit 20, layer 3 circuit	
11, with effect 0.81	
layer 0 circuit 19, layer 2 circuit 20, layer 3 circuit	
11, with effect 0.82	
layer 0 circuit 20, layer 2 circuit 20, layer 3 circuit	
11, with effect 0.85	
layer 0 circuit 21, layer 2 circuit 20, layer 3 circuit	
11, with effect 0.83	
layer 0 circuit 22, layer 2 circuit 20, layer 3 circuit	

1585	<i>11, with effect 0.86</i>	<i>9, with effect 0.8</i>	1637
1586	<i>layer 0 circuit 24, layer 2 circuit 20, layer 3 circuit</i>	<i>layer 0 circuit 22, layer 1 circuit 16, layer 10 circuit</i>	1638
1587	<i>11, with effect 0.81</i>	<i>9, with effect 0.81</i>	1639
1588	<i>layer 1 circuit 13, layer 2 circuit 20, layer 3 circuit</i>	<i>layer 0 circuit 14, layer 1 circuit 20, layer 10 circuit</i>	1640
1589	<i>11, with effect 0.86</i>	<i>9, with effect 0.83</i>	1641
1590	<i>layer 1 circuit 14, layer 2 circuit 20, layer 3 circuit</i>	<i>layer 0 circuit 24, layer 1 circuit 20, layer 10 circuit</i>	1642
1591	<i>11, with effect 0.84</i>	<i>9, with effect 0.81</i>	1643
1592	<i>layer 1 circuit 15, layer 2 circuit 20, layer 3 circuit</i>	<i>layer 0 circuit 13, layer 2 circuit 20, layer 10 circuit</i>	1644
1593	<i>11, with effect 0.82</i>	<i>9, with effect 0.92</i>	1645
1594	<i>layer 1 circuit 16, layer 2 circuit 20, layer 3 circuit</i>	<i>layer 0 circuit 14, layer 2 circuit 20, layer 10 circuit</i>	1646
1595	<i>11, with effect 0.85</i>	<i>9, with effect 0.9</i>	1647
1596	<i>layer 1 circuit 17, layer 2 circuit 20, layer 3 circuit</i>	<i>layer 0 circuit 15, layer 2 circuit 20, layer 10 circuit</i>	1648
1597	<i>11, with effect 0.85</i>	<i>9, with effect 0.85</i>	1649
1598	<i>layer 1 circuit 18, layer 2 circuit 20, layer 3 circuit</i>	<i>layer 0 circuit 16, layer 2 circuit 20, layer 10 circuit</i>	1650
1599	<i>11, with effect 0.85</i>	<i>9, with effect 0.91</i>	1651
1600	<i>layer 1 circuit 19, layer 2 circuit 20, layer 3 circuit</i>	<i>layer 0 circuit 17, layer 2 circuit 20, layer 10 circuit</i>	1652
1601	<i>11, with effect 0.85</i>	<i>9, with effect 0.89</i>	1653
1602	<i>layer 0 circuit 14, layer 1 circuit 19, layer 2 circuit</i>	<i>layer 0 circuit 18, layer 2 circuit 20, layer 10 circuit</i>	1654
1603	<i>20, layer 3 circuit 11, with effect 0.83</i>	<i>9, with effect 0.86</i>	1655
1604	<i>layer 0 circuit 15, layer 1 circuit 19, layer 2 circuit</i>	<i>layer 0 circuit 19, layer 2 circuit 20, layer 10 circuit</i>	1656
1605	<i>20, layer 3 circuit 11, with effect 0.83</i>	<i>9, with effect 0.86</i>	1657
1606	<i>layer 0 circuit 16, layer 1 circuit 19, layer 2 circuit</i>	<i>layer 0 circuit 20, layer 2 circuit 20, layer 10 circuit</i>	1658
1607	<i>20, layer 3 circuit 11, with effect 0.83</i>	<i>9, with effect 0.9</i>	1659
1608	<i>layer 0 circuit 22, layer 1 circuit 19, layer 2 circuit</i>	<i>layer 0 circuit 21, layer 2 circuit 20, layer 10 circuit</i>	1660
1609	<i>20, layer 3 circuit 11, with effect 0.83</i>	<i>9, with effect 0.87</i>	1661
1610	<i>layer 0 circuit 23, layer 1 circuit 19, layer 2 circuit</i>	<i>layer 0 circuit 22, layer 2 circuit 20, layer 10 circuit</i>	1662
1611	<i>20, layer 3 circuit 11, with effect 0.82</i>	<i>9, with effect 0.92</i>	1663
1612	<i>layer 0 circuit 24, layer 1 circuit 19, layer 2 circuit</i>	<i>layer 0 circuit 23, layer 2 circuit 20, layer 10 circuit</i>	1664
1613	<i>20, layer 3 circuit 11, with effect 0.84</i>	<i>9, with effect 0.85</i>	1665
1614	<i>layer 1 circuit 20, layer 2 circuit 20, layer 3 circuit</i>	<i>layer 0 circuit 24, layer 2 circuit 20, layer 10 circuit</i>	1666
1615	<i>11, with effect 0.85</i>	<i>9, with effect 0.86</i>	1667
1616	<i>layer 1 circuit 21, layer 2 circuit 20, layer 3 circuit</i>	<i>layer 1 circuit 13, layer 2 circuit 20, layer 10 circuit</i>	1668
1617	<i>11, with effect 0.84</i>	<i>9, with effect 0.92</i>	1669
1618	<i>layer 1 circuit 22, layer 2 circuit 20, layer 3 circuit</i>	<i>layer 1 circuit 14, layer 2 circuit 20, layer 10 circuit</i>	1670
1619	<i>11, with effect 0.86</i>	<i>9, with effect 0.89</i>	1671
1620	<i>layer 1 circuit 23, layer 2 circuit 20, layer 3 circuit</i>	<i>layer 1 circuit 15, layer 2 circuit 20, layer 10 circuit</i>	1672
1621	<i>11, with effect 0.82</i>	<i>9, with effect 0.85</i>	1673
1622	<i>layer 1 circuit 24, layer 2 circuit 20, layer 3 circuit</i>	<i>layer 1 circuit 16, layer 2 circuit 20, layer 10 circuit</i>	1674
1623	<i>11, with effect 0.81</i>	<i>9, with effect 0.9</i>	1675
1624	<i>layer 0 circuit 21, layer 2 circuit 14, layer 3 circuit</i>	<i>layer 0 circuit 13, layer 1 circuit 16, layer 2 circuit</i>	1676
1625	<i>14, with effect 0.8</i>	<i>20, layer 10 circuit 9, with effect 0.83</i>	1677
1626	<i>layer 0 circuit 22, layer 2 circuit 14, layer 3 circuit</i>	<i>layer 1 circuit 17, layer 2 circuit 20, layer 10 circuit</i>	1678
1627	<i>14, with effect 0.81</i>	<i>9, with effect 0.9</i>	1679
1628	<i>layer 1 circuit 21, layer 2 circuit 14, layer 3 circuit</i>	<i>layer 1 circuit 18, layer 2 circuit 20, layer 10 circuit</i>	1680
1629	<i>14, with effect 0.81</i>	<i>9, with effect 0.91</i>	1681
1630	<i>layer 1 circuit 22, layer 2 circuit 14, layer 3 circuit</i>	<i>layer 0 circuit 14, layer 1 circuit 18, layer 2 circuit</i>	1682
1631	<i>14, with effect 0.81</i>	<i>20, layer 10 circuit 9, with effect 0.81</i>	1683
1632	<i>layer 0 circuit 13, layer 1 circuit 16, layer 10 circuit</i>	<i>layer 0 circuit 23, layer 1 circuit 18, layer 2 circuit</i>	1684
1633	<i>9, with effect 0.84</i>	<i>20, layer 10 circuit 9, with effect 0.83</i>	1685
1634	<i>layer 0 circuit 14, layer 1 circuit 16, layer 10 circuit</i>	<i>layer 1 circuit 19, layer 2 circuit 20, layer 10 circuit</i>	1686
1635	<i>9, with effect 0.81</i>	<i>9, with effect 0.9</i>	1687
1636	<i>layer 0 circuit 15, layer 1 circuit 16, layer 10 circuit</i>	<i>layer 0 circuit 13, layer 1 circuit 19, layer 2 circuit</i>	1688



1689	20, layer 10 circuit 9, with effect 0.83	5, with effect 0.89	1741
1690	layer 0 circuit 14, layer 1 circuit 19, layer 2 circuit	layer 0 circuit 22, layer 2 circuit 14, layer 11 circuit	1742
1691	20, layer 10 circuit 9, with effect 0.87	5, with effect 0.89	1743
1692	layer 0 circuit 15, layer 1 circuit 19, layer 2 circuit	layer 0 circuit 23, layer 2 circuit 14, layer 11 circuit	1744
1693	20, layer 10 circuit 9, with effect 0.86	5, with effect 0.86	1745
1694	layer 0 circuit 16, layer 1 circuit 19, layer 2 circuit	layer 0 circuit 24, layer 2 circuit 14, layer 11 circuit	1746
1695	20, layer 10 circuit 9, with effect 0.87	5, with effect 0.84	1747
1696	layer 0 circuit 20, layer 1 circuit 19, layer 2 circuit	layer 1 circuit 13, layer 2 circuit 14, layer 11 circuit	1748
1697	20, layer 10 circuit 9, with effect 0.82	5, with effect 0.85	1749
1698	layer 0 circuit 21, layer 1 circuit 19, layer 2 circuit	layer 1 circuit 14, layer 2 circuit 14, layer 11 circuit	1750
1699	20, layer 10 circuit 9, with effect 0.82	5, with effect 0.86	1751
1700	layer 0 circuit 22, layer 1 circuit 19, layer 2 circuit	layer 1 circuit 15, layer 2 circuit 14, layer 11 circuit	1752
1701	20, layer 10 circuit 9, with effect 0.87	5, with effect 0.85	1753
1702	layer 0 circuit 23, layer 1 circuit 19, layer 2 circuit	layer 1 circuit 16, layer 2 circuit 14, layer 11 circuit	1754
1703	20, layer 10 circuit 9, with effect 0.86	5, with effect 0.84	1755
1704	layer 0 circuit 24, layer 1 circuit 19, layer 2 circuit	layer 1 circuit 17, layer 2 circuit 14, layer 11 circuit	1756
1705	20, layer 10 circuit 9, with effect 0.88	5, with effect 0.85	1757
1706	layer 1 circuit 20, layer 2 circuit 20, layer 10 circuit	layer 1 circuit 18, layer 2 circuit 14, layer 11 circuit	1758
1707	9, with effect 0.9	5, with effect 0.86	1759
1708	layer 0 circuit 14, layer 1 circuit 20, layer 2 circuit	layer 1 circuit 19, layer 2 circuit 14, layer 11 circuit	1760
1709	20, layer 10 circuit 9, with effect 0.81	5, with effect 0.8	1761
1710	layer 1 circuit 21, layer 2 circuit 20, layer 10 circuit	layer 1 circuit 20, layer 2 circuit 14, layer 11 circuit	1762
1711	9, with effect 0.89	5, with effect 0.87	1763
1712	layer 1 circuit 22, layer 2 circuit 20, layer 10 circuit	layer 1 circuit 21, layer 2 circuit 14, layer 11 circuit	1764
1713	9, with effect 0.92	5, with effect 0.89	1765
1714	layer 1 circuit 23, layer 2 circuit 20, layer 10 circuit	layer 1 circuit 22, layer 2 circuit 14, layer 11 circuit	1766
1715	9, with effect 0.86	5, with effect 0.89	1767
1716	layer 0 circuit 14, layer 1 circuit 19, layer 10 circuit	layer 1 circuit 23, layer 2 circuit 14, layer 11 circuit	1768
1717	10, with effect 0.81	5, with effect 0.86	1769
1718	layer 0 circuit 16, layer 1 circuit 19, layer 10 circuit	layer 1 circuit 24, layer 2 circuit 14, layer 11 circuit	1770
1719	10, with effect 0.81	5, with effect 0.81	1771
1720	layer 0 circuit 22, layer 1 circuit 19, layer 10 circuit	layer 0 circuit 13, layer 2 circuit 24, layer 11 circuit	1772
1721	10, with effect 0.81	5, with effect 0.84	1773
1722	layer 0 circuit 23, layer 1 circuit 19, layer 10 circuit	layer 0 circuit 14, layer 2 circuit 24, layer 11 circuit	1774
1723	10, with effect 0.81	5, with effect 0.82	1775
1724	layer 0 circuit 24, layer 1 circuit 19, layer 10 circuit	layer 0 circuit 15, layer 2 circuit 24, layer 11 circuit	1776
1725	10, with effect 0.82	5, with effect 0.85	1777
1726	layer 0 circuit 14, layer 1 circuit 19, layer 11 circuit	layer 0 circuit 16, layer 2 circuit 24, layer 11 circuit	1778
1727	5, with effect 0.81	5, with effect 0.85	1779
1728	layer 0 circuit 16, layer 1 circuit 19, layer 11 circuit	layer 0 circuit 17, layer 2 circuit 24, layer 11 circuit	1780
1729	5, with effect 0.8	5, with effect 0.85	1781
1730	layer 0 circuit 22, layer 1 circuit 19, layer 11 circuit	layer 0 circuit 22, layer 2 circuit 24, layer 11 circuit	1782
1731	5, with effect 0.81	5, with effect 0.85	1783
1732	layer 0 circuit 24, layer 1 circuit 19, layer 11 circuit	layer 0 circuit 23, layer 2 circuit 24, layer 11 circuit	1784
1733	5, with effect 0.81	5, with effect 0.85	1785
1734	layer 0 circuit 13, layer 2 circuit 14, layer 11 circuit	layer 0 circuit 24, layer 2 circuit 24, layer 11 circuit	1786
1735	5, with effect 0.87	5, with effect 0.82	1787
1736	layer 0 circuit 14, layer 2 circuit 14, layer 11 circuit	layer 1 circuit 13, layer 2 circuit 24, layer 11 circuit	1788
1737	5, with effect 0.81	5, with effect 0.83	1789
1738	layer 0 circuit 20, layer 2 circuit 14, layer 11 circuit	layer 1 circuit 14, layer 2 circuit 24, layer 11 circuit	1790
1739	5, with effect 0.86	5, with effect 0.81	1791
1740	layer 0 circuit 21, layer 2 circuit 14, layer 11 circuit	layer 1 circuit 15, layer 2 circuit 24, layer 11 circuit	1792

5, with effect 0.82  
layer 1 circuit 16, layer 2 circuit 24, layer 11 circuit 5, with effect 0.81  
layer 1 circuit 17, layer 2 circuit 24, layer 11 circuit 5, with effect 0.81  
layer 1 circuit 22, layer 2 circuit 24, layer 11 circuit 5, with effect 0.85  
layer 1 circuit 23, layer 2 circuit 24, layer 11 circuit 5, with effect 0.82  
layer 1 circuit 24, layer 2 circuit 24, layer 11 circuit 5, with effect 0.81  
layer 0 circuit 13, layer 3 circuit 14, layer 11 circuit 5, with effect 0.81  
layer 0 circuit 23, layer 3 circuit 14, layer 11 circuit 5, with effect 0.85  
layer 1 circuit 23, layer 3 circuit 14, layer 11 circuit 5, with effect 0.81  
layer 2 circuit 23, layer 3 circuit 14, layer 11 circuit 5, with effect 0.8

Almost all 3-step paths are composed of paths from lower-level skills. For instance, in the ICL skill, the sequence “layer 0 circuit 20, layer 2 circuit 14, layer 5 circuit 11” encompasses the path “layer 0 circuit 20, layer 2 circuit 14” from the previous token skill. Furthermore, it is apparent that the more complex a skill, the more multi-step paths it encompasses.

## G More detailed findings

### G.1 Function Components in IOI Task

Building on the results from Table 3, we continue to explore the skills required for IOI, which include duplicate token (DLT), previous token (PVT), induction (IDT), S-inhibition (SIB), name mover (NMV), and back-up head (Wang et al., 2023). For DLT, we found another distinct cluster within the circuit samples of the induction skill. For SIB, we obtained it by replacing “S2” with “IO” as the background text. For NMV, we obtained it by using a random name as a substitute for “IO” and “S” in the background text. Interestingly, our method was unable to detect the presence of a back-up head. A reasonable conjecture is that the back-up head acts more like a preemption mechanism, effectively circumvented in path-level causal analysis. Additionally, we present the key nodes of other skills in Table 10. It is evident that the skill paths we demonstrate possess strong inclusivity. For instance, the S-inhibition skill encompasses crucial nodes of the duplicate token, previous token, and induction skills, while the name mover almost in-

cludes nodes from all previous skills. Beyond this, we also discovered a long-position mapping (LPM) skill, obtained through a large number of long sentence samples and background text with deleted commas. It represents another advanced skill that extends PVT.

Moreover, based on the paths, attention weights, and cosine similarities of the representations, we have identified several circuits with distinct characteristics (We demonstrate the performances of other circuit discovery methods in validating these conclusions in Appendix 6.2.):

**Preceding Token Circuit:** Circuit [4, 12] performs a unique function, namely, when any token serves as a query token to attend other tokens, this circuit is shown to consistently carry significant information from its preceding token to the query token.

**Key Token Circuit:** Circuit [3, 14] exhibits a significantly different function from the others. This circuit consistently focuses on certain key tokens in the preceding text – such as the beginning, ending, and label prompts – and transmits this information to subsequent query tokens. Additionally, other key circuits in layers 3 and 4 partially undertake these functionalities.

**Opposite Circuit:** When using the last token of each input to produce the embedding for a specific circuit, we notice that the cosine similarity between Circuit [11, 14] and other key circuits is usually less than 0, especially with Circuit [11, 1], where the cosine similarity reaches to  $-0.92$ . Previous work (Wang et al., 2023) has mentioned this phenomenon, hypothesizing the reason to be controlling the variance of the loss function.

We have observed some differences in the receivers of different ICL tasks. Combined with the insights provided by Bayazit et al. (2023) and Bricken et al. (2023), we suspect that these differences arise from distinct circuits required to process domain-specific knowledge across different tasks.

### G.2 Exploration - Why Wrong Outputs?

In this section, we present a new direction for explaining and exploring common erroneous answers using Skill Circuit Graphs. Specifically, by contrasting the Skill Graphs of “incorrect” outputs with those of correct outputs, we can further diagnose what leads to the failure in skill execution. Table 11 illustrates the key circuits exhibiting the

Skill	Receivers with receiving more than 10 paths (#layer, #circuit)
DLT	[0, 2], [0, 13], [0, 14], [0, 16], [0, 20], [1, 8], [1, 9], [1, 18], [1, 19], [11, 1], [11, 14]
PVT	[1, 8], [1, 18], [1, 19], [1, 20], [1, 21], [2, 1], [2, 7], [2, 14], [2, 18], [2, 20], [2, 22], [2, 24], [11, 1], [11, 14]
IDT	[2, 14], [2, 18], [2, 20], [3, 14], [3, 17], [4, 5], [4, 12], [5, 11], [6, 5], [11, 1], [11, 14]
SIB	[0, 20], [1, 8], [1, 18], [2, 14], [3, 14], [5, 11], [5, 14], [7, 9], [7, 20], [7, 21], [8, 7], [8, 18], [11, 1], [11, 14]
NMV	[0, 2], [1, 8], [1, 20], [3, 14], [3, 20], [5, 11], [8, 7], [9, 14], [9, 18], [9, 20], [10, 1], [10, 14], [10, 22], [11, 14]
LPM	[1, 8], [2, 18], [2, 20], [3, 11], [4, 14], [6, 7], [8, 4], [8, 17], [8, 24], [9, 9]

Table 10: Key Receivers in subgraphs of IOI task, blue circuits are presented in the lower skill

Type	Top-5 circuits with absence rate
<b>F_IDT</b>	[2, 18] ( $\downarrow 0.37$ ), [2, 14] ( $\downarrow 0.32$ ), [11, 1] ( $\downarrow 0.28$ ), [2, 20] ( $\downarrow 0.26$ ), [2, 24] ( $\downarrow 0.26$ )
<b>F1_ICL</b>	[2, 24] ( $\downarrow 0.45$ ), [2, 20] ( $\downarrow 0.42$ ), [2, 22] ( $\downarrow 0.41$ ), [1, 20] ( $\downarrow 0.39$ ), [2, 14] ( $\downarrow 0.32$ )
<b>F2_ICL</b>	[3, 14] ( $\downarrow 0.29$ ), [4, 5] ( $\downarrow 0.28$ ), [10, 10] ( $\downarrow 0.28$ ), [8, 9] ( $\downarrow 0.24$ ), [4, 12] ( $\downarrow 0.22$ )

Table 11: Top 5 Receiver circuits appearing most frequently in skill paths presented in correct output samples but not incorrect samples.

highest absent rate<sup>6</sup> between 3 “incorrect” and correct output types. Specifically, we investigate one erroneous type of output from an induction skill sample (F\_IDT), and two types from ICL skill samples (F1\_ICL, F2\_ICL).

F\_IDT refers to those samples wherein the input possesses an Induction pattern (“A B ... A”), but ultimately does not output B. F1\_ICL denotes those samples wherein the output includes a word outside of the label options from the demonstrations, for example, a case where the input text “[review1], label: positive, [review2], label: negative, [review3], label:” unexpectedly produces “the”. Such an error indicates that the language model did not capture the ICL template pattern in this case. F2\_ICL involves samples that capture the template pattern yet still produce incorrect outputs, for example, cases where the correct output should be “positive”, but the prediction is “negative”. We compare the circuit graphs of these “incorrect” samples with the correct samples and identify the top 5 circuits with the highest absence rate.

Table 11 exhibits several interesting phenomena where the largest discrepancies between correct and incorrect samples in both F\_IDT and F1\_ICL occur on key circuits at layer 2. These circuits originate from the previous token skill, which handles the skill of receiving information from the previous token, such as the “A  $\rightarrow$  B” in the induction template “A B ... A”, as well as patterns such as “label  $\rightarrow$  positive” in ICL. The loss of this skill—failure during the execution of the previous token skill—means

that both the Induction skill and ICL skill cannot pass the duplicated prefix information to the next token, leading to template-based errors.

To further understand why these samples do not successfully execute the previous token skill, we perform a bi-clustering operation on the Previous Token Skill (experiment details are shown in Appendix D.2), yielding a cluster with  $Eff < 0.2$  across most of all paths. We compared this cluster (termed the low-effect cluster) with another cluster (named high-effect cluster), with some samples as follows (All samples are from the original text of the Previous Token Skill, tokenized into two tokens):

**Low-effect cluster:** “About to”, “all these”, “am a”, “and win”, “and select”, “care over”, “In Singapore”, “in the”, “is a”, “it was”, “than they”, “The language”, “The country”, “the movie”

**High-effect cluster:** “2002”, “Adriano”, “Ajinomoto”, “becomes”, “Could you”, “don’t”, “ended up”, “If the”, “iPhone”, “Knowledge”, “stressful”, “Windows”, “Youtube’s”

It becomes obvious that in the context of an experimental setting lacking enough context, the previous token skill is performed only when there is a strong semantic relationship between the two tokens. For pairs of tokens where the semantic relation is not strong, there tends to be a reliance on the bi-gram model decision from the destination token.

Furthermore, for F2\_ICL, the absence rate is relatively lower, suggesting that the source of the error might not be due to a single explicit cause. These circuits generally reside in the middle or

<sup>6</sup>Let  $N_{C^{l,j}}^+$  and  $N_{C^{l,j}}^-$  be the number of paths received by  $C^{l,j}$  in correct and incorrect samples. The absence rate for each circuit is calculated as  $(N_{C^{l,j}}^+ - N_{C^{l,j}}^-) / N_{C^{l,j}}^+ \in [0, 1]$ .

even deeper layers, incorporating functions such as induction and summarization. However, to further analyze this, we would need to delve into the representational level, which for the moment goes beyond the scope of this paper.

Skill. For additional skill graphs, please refer to our repository.

## H Details of Comparisons with Other Methods Validating Conjectures

### H.1 Details about Baselines

**ACDC** (Conmy et al., 2023), Automatic Circuit Discovery, which calculates the importance score of each edge and performs a greedy search based on the score.

**Opt prun** (Bhaskar et al., 2024), which converts the importance score into an optimization function and assigns a learnable parameter to each edge to indicate whether an edge needs to be deleted.

**EAP** (Syed et al., 2023), or Edge Attribution Patching, which makes a linear approximation of activation patching to assign an importance score to each edge, and retains the top- $k$  edges.

**Ours-noise**: When deleting an edge, we replace the original edge value with a noise ( $\mathcal{N}(0, 0.81)$ ).

**Ours-mean**: When deleting an edge, we replace the original edge value with the mean value of all edges received by this component.

**Ours-logits**: We replace the top  $n$  candidates with the difference in the final layer logits ( $|\text{logit}(\mathcal{G}'/e) - \text{logit}(\mathcal{G}')|$ ), where  $\mathcal{G}'$  represents the pruned subgraph from the previous step. Edges with a difference of less than 0.04 will be deleted.

**Ours-KL**: We replace the top  $n$  candidates with the KL divergence of the final output probability distribution ( $|KL(\mathcal{G}'/e) - KL(\mathcal{G}')|$ ). Edges with a difference of less than 0.005 will be deleted.

### H.2 Definition of $ovlp(A, B)$

The rule for calculating overlap is as follows: let  $ovlp(A, B)$  represent what the rate of edges in skill graph  $A$  also existing in skill graph  $B$  is. For any edge  $e^i$  in skill graph  $A$ , we set an overlap flag  $f_{A,B}(e^i)$ . If  $e^i$  in  $A$  also exists in skill circuit graphs  $B$ , then  $f_{A,B}(e^i) = 1$ , otherwise  $f_{A,B}(e^i) = 0$ . For a circuit graph  $A$  with  $N_A$  edges, its set of edges is  $\mathcal{E}_A$ . Our overlap is calculated as  $ovlp(A, B) = \frac{1}{N_A} \sum_{e^i \in \mathcal{E}_A} f_{A,B}(e^i)$ .

## I Skill Circuit Graphs

Due to large size constraints, we have only displayed the circuit graph for the Previous Token



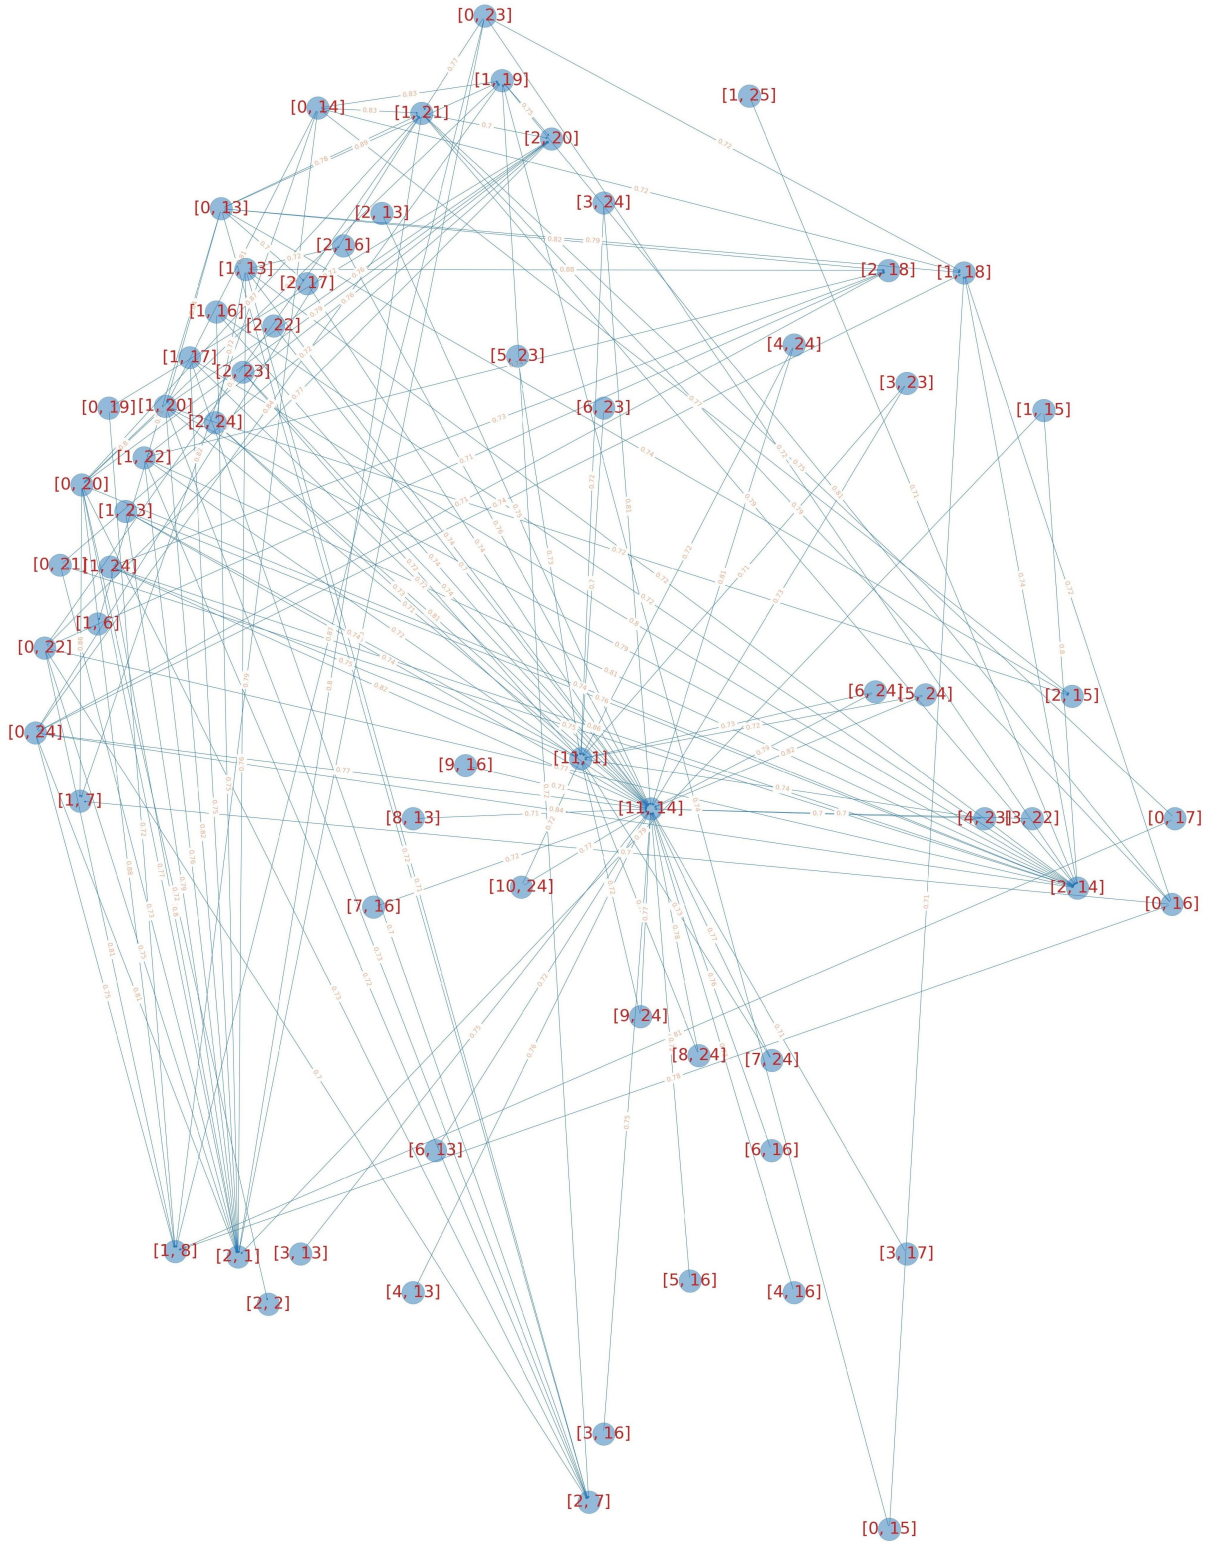


Figure 13: Skill Circuit Graph of Previous Token Skill, all paths with  $Eff > 0.7$  are labeled.