

# Mamba Knockout for Unraveling Factual Information Flow

**Nir Endy\***  
Tel Aviv University  
nirendy@gmail.com

**Idan Daniel Grosbard\***  
Tel Aviv University  
idangrosbard@mail.tau.ac.il

**Yuval Ran-Milo\***  
Tel Aviv University  
yuv.milo@gmail.com

**Yonatan Slutzky\***  
Tel Aviv University  
slutzky1@mail.tau.ac.il

**Itay Tshuva**  
Tel Aviv University  
itayt122@gmail.com

**Raja Giryes**  
Tel Aviv University  
raja@tauex.tau.ac.il

## Abstract

This paper investigates the flow of factual information in Mamba-based language models. We rely on theoretical and empirical connections to Transformer-based architectures and their attention mechanisms. Exploiting this relationship, we adapt attentional interpretability techniques originally developed for Transformers—specifically, the Attention Knockout methodology—to both Mamba-1 and Mamba-2. Using them, we trace how information is transmitted and localized across tokens and layers, revealing patterns of subject-token information emergence and layer-wise dynamics. Notably, some phenomena vary between Mamba models and Transformer-based models, while others appear universally across all models inspected—hinting that these may be inherent to LLMs in general. By further leveraging Mamba’s structured factorization, we disentangle how distinct “features” either enable token-to-token information exchange or enrich individual tokens, thus offering a unified lens to understand Mamba’s internal operations. Our code can be found at <https://github.com/nirendy/mamba-knockout>.

## 1 Introduction

Understanding how factual information moves through different parts of a language model is an important step toward explaining its outputs. While Transformer-based models and their attention mechanisms have been studied in depth (Geva et al., 2023), less is known about how structured state-space models (SSMs) process and transfer facts across tokens and layers. Recent work has introduced Mamba-based SSM architectures that rival Transformer performance in various settings (Gu and Dao, 2024; Waleffe et al., 2024; Dao and Gu, 2024). Yet, the internal pathways of factual information within these models remain less explored.

Motivated by theoretical connections showing that selective SSMs can be understood through an attention-like perspective (Ali et al., 2024; Dao and Gu, 2024; Zimerman et al., 2025; Ben-Kish et al., 2025), we draw on methods originally developed for Transformers. Notably, a certain class of SSMs has been proven equivalent to a subclass of linear attention Transformers (Katharopoulos et al., 2020; Dao and Gu, 2024), suggesting that interpretability tools designed for attention can also be applied to SSM-based architectures. By employing the “Attention Knockout” technique (Geva et al., 2023), previously used to map how factual information flows in Transformers, we can similarly isolate and analyze the flow of information in Mamba-1 (Gu and Dao, 2024) and Mamba-2 (Dao and Gu, 2024).

Our results show that Mamba models, despite their differences from Transformers, display certain patterns of information emergence and routing that align with observed attention-based behaviors. Key facts surface within certain tokens and gain prominence at specific layers, mirroring the step-by-step integration of information commonly noted in Transformers (Geva et al., 2023). In contrast, our results also show that other patterns are differentiated between Mamba models and Transformer-based models, with some phenomena even varying within each family.

By leveraging Mamba’s factorized structure, we separate the roles of different “features” in (i) transmitting information between tokens; and (ii) enriching individual tokens independently. This approach not only clarifies how Mamba’s internal operations resemble those of well-studied attention models, but also provides a clearer framework for future analyses of SSM-based language models.

Ultimately, our work extends interpretability beyond the Transformer paradigm, offering a more unified perspective on how language models—whether attention-dominated or grounded in state-space representations—organize fac-

\* Equal contribution.

tual information. By illuminating these internal processes, we take a step toward a more generalizable understanding of how language models, and specifically Mamba, form and express factual knowledge.

The contributions of this paper are twofold: (i) We extend the Attention Knockout interpretability framework—previously developed exclusively for Transformers—to structured state-space models (SSMs), enabling us to uncover previously unknown parallels and distinctions between Mamba-based and Transformer-based architectures’ factual-information dynamics. (ii) We propose a novel ‘feature knockout’ mechanism that takes advantage of the unique structure of SSMs, allowing nuanced interventions and additional insights into how different feature types contribute to model behavior.

## 2 Related Work

Selective State-Space Modeling (SSM) was recently introduced by [Gu and Dao \(2024\)](#), rivaling performance of Transformer-based models in various settings ([Waleffe et al., 2024](#)). An analytical examination of the core operation in selective SSMs by [Ali et al. \(2024\)](#); [Zimmerman et al. \(2025\)](#) provided a complementary view of selective SSMs as attention-driven models, allowing the application of attention-based explainability analysis to SSM-based models. The Structured State-Space Duality (SSD) framework proposed by [Dao and Gu \(2024\)](#) bridges SSMs and attention layers even further, proving that a specific class of SSMs is equivalent to a subclass of linear attention Transformers ([Katharopoulos et al., 2020](#)). In this work, we employ these approaches to enable the application of attention-based tools and analyses to SSM-based architectures.

Complementary, several studies have analyzed how factual information is stored and extracted in Transformer-based models. [Geva et al. \(2021\)](#) show MLP layers function as key-value memories, which extract semantic associations regarding the input in the ultimate layers. [Nichani et al. \(2024\)](#) prove that shallow Transformers have optimal factual storage in their value matrices and MLPs. To assess in greater accuracy where specific memories are stored, [Meng et al. \(2022\)](#) suggest Causal Tracing - a method for identifying critical MLPs by adding noise to intermediate representations and restoring clean run states at deeper layers. They also show that in order to edit stored facts, a rank-one matrix editing in the MLP of a single, early, layer is sufficient. The work of [Meng](#)

[et al. \(2023\)](#) expands the ROME technique for mass editing memory by expanding the editing to multiple consecutive early layers. [Hase et al. \(2023\)](#) show that early-intermediate layers are in general good candidates for injecting new facts to LMs, and empirically validate a particular type of Causal Tracing that is insightful for fact editing and fact storage. Our work draws its inspiration from [Geva et al. \(2023\)](#) which analyzed how the information from multiple tokens is aggregated to correctly query attributes using *Attention Knockout*. Their analysis reveals that this relies on an initial enrichment process where the subject token extracts relevant attributes in the early layers’ MLPs.

Our work most closely aligns with that of [Sharma et al. \(2024\)](#), whose thorough investigation of factual associations in Mamba-1 featured an analysis using a variant of attention knockout. However, our approach differs from them in several key aspects. [Sharma et al. \(2024\)](#) focus on blocking the propagation of information from a single token to all future tokens, whereas we adopt the technique proposed in [Ali et al. \(2024\)](#) to remove attention specifically from one token to a single token—more in line with previous work such as [Geva et al. \(2023\)](#). Though [Sharma et al. \(2024\)](#) suggest that such fine-grained blocking may be challenging due to convolution and softmax layers, we demonstrate that a straightforward and natural implementation successfully replicates the phenomena observed in Transformer-based attention knockouts. We also investigate Mamba-2 (where attention knockout has a more direct interpretation) and find consistent patterns in both Mamba-1 and Mamba-2, suggesting that these dynamics may be fundamental across SSM-based and Transformer-based architectures. Finally, beyond leveraging attention parallels, our work introduces a novel “feature knockout” mechanism that exploits the unique structure of SSMs, enabling more nuanced interventions and additional insights into how different feature types contribute to model behavior. We inspect additional related works in [Appendix A](#).

## 3 Methodology

We now turn to introduce the tools that we use in our work to analyze Mamba. We start by describing the attention knockout mechanism, which was proposed for Transformers, and the SSMs formulation. Then we explain our methodology for employing knockout for SSMs and specifically for Mamba-1 and Mamba-2.

### 3.1 Attention Knockout

We adopt the *Attention Knockout* methodology introduced by Geva et al. (2023) and apply it to identify critical points in the flow of information essential for factual predictions. For successful next-token prediction, a model must process the input tokens so that the next-token can be inferred from the last position. Geva et al. (2023) investigate this process internally by “knocking out” parts of the computation and measuring the effect on the prediction. To this end, they propose a fine-grained intervention on the attention layers, which, they show, serves as a “crucial” module for communicating information between positions. By disrupting critical information transfer through these layers, they demonstrate that factual predictions are constructed in stages, with essential information reaching the prediction position at specific layers during inference. Intuitively, critical attention connections are those whose disruption results in a marked decline in prediction quality. To test whether essential information flows between two hidden representations at a specific layer, the authors zero out all attention connections between them. Formally, given an input sequence of  $N$  tokens, a model consisting of  $L$  consecutive layers, and two positions  $r, c \in [1, N]$  with  $r \leq c$ , they prevent the  $c$ th token at layer  $l$  from attending to  $r$ th token at layer  $l$  by zeroing the attention weights for that layer at index  $c, r$ .

### 3.2 Selective State Space Models

A selective state space model (SSM) of dimension  $n$  is defined as a time-dependent recurrent relation driven by an input signal  $u(t) \in \mathbb{R}$ . At each time step  $t$ , the model is parameterized by matrices  $A(t) \in \mathbb{R}^{n \times n}$ ,  $B(t) \in \mathbb{R}^{n \times 1}$ , and  $C(t) \in \mathbb{R}^{1 \times n}$  that can depend on the current input  $u(t)$ . Formally, the system evolves according to:

$$\begin{aligned} x(t+1) &= A(t)x(t) + B(t)u(t), \\ y(t) &= C(t)x(t). \end{aligned}$$

These parameterizations differ between Mamba-1 and Mamba-2. Further details can be found in Gu and Dao (2024); Dao and Gu (2024). However, a unifying property of both models is that the transfer matrix  $A(t)$  is parameterized as

$$A(t) = \bar{A}^{\Delta(t)}, \quad \bar{A} = \exp(\text{diag}(\alpha)),$$

for some learned parameters  $\Delta(t) \in \mathbb{R}$  and  $\alpha \in [0, 1]^n$ , a key detail used in Section 3.5.

### 3.3 Hidden Attention of Mamba-1

For Mamba-1, we follow the *hidden-attention* perspective introduced by Ali et al. (2024) to implement attention knockout in Mamba. Specifically, the authors utilized the kernel representation of Selective-SSMs and observed that the relationship between any two tokens  $i$  and  $j$  (where  $i \leq j$ ) is represented as an entry in the kernel matrix

$$\mathbf{M}_{i,j} = Q_i \cdot H_{i,j} \cdot K_j.$$

Here,  $Q_i = C(i)$ ,  $H_{i,j} = \prod_{t=i}^j A(t)$ , and  $K_j = B(j)$  (see Section 3.2). To apply knockout between two tokens at a specific layer, we simply assign  $\mathbf{M}_{i,j} = 0$ .

### 3.4 Implicit Linear Attention of Mamba-2

In Mamba-2, following Dao and Gu (2024), the SSM layer can be interpreted as a masked linear attention mechanism, allowing it to be expressed as a matrix multiplication. Specifically, the SSM layer processes an input tensor  $\mathbf{X} \in \mathbb{R}^{L \times H}$ , where  $L$  represents the sequence length and  $H$  the dimensionality of each input token. The layer transforms  $\mathbf{X}$  through the operation  $\mathbf{L} \circ (\mathbf{X}\mathbf{M}\mathbf{X}^\top)\mathbf{X}$ , where  $\mathbf{M}$  and  $\mathbf{L}$  are fixed matrices determined by the layer’s weights,  $\mathbf{L}$  is lower triangular, and  $\circ$  denotes the Hadamard product. This formulation can be viewed as an attention mechanism where  $\mathbf{L} \circ (\mathbf{X}\mathbf{M}\mathbf{X}^\top)$  serves as the attention matrix. In this matrix, the entry at position  $(i, j)$  quantifies the degree to which the  $i$ th token attends to the  $j$ th token, analogous to the attention scores in traditional Transformer-based models.

### 3.5 Individual Feature Knockout in SSMs

In both Mamba-1 and Mamba-2, each feature is treated as an independent time-varying signal modeled by distinct SSMs. Leveraging this property, we propose a feature knockout mechanism that targets specific types of features at each layer. Specifically, we classify features as either *context-dependent* or *context-independent* based on their decay characteristics and perform knockout by zeroing out their outputs. Formally, if we denote  $\alpha, \bar{A}$  and  $\Delta(t)$  as in Section 3.2, with  $\alpha$  being the model weights and  $\Delta(t)$  being individual token-dependent features, the dynamics of the state matrix becomes:

$$A(t) = \bar{A}^{\Delta(t)}, \quad \prod_{t=i}^j A(t) = \bar{A}^{\sum_{t=i}^j \Delta(t)}.$$

We interpret  $\prod_{t=i}^j A(t)$  as a measure of how past signals influence current representations. This product exhibits exponential decay, dependent on the matrix  $\bar{A}$  so features parameterized by  $\bar{A} \approx 0$  quickly lose memory and concentrate on localized information. In contrast, features parameterized by  $\bar{A} \approx 1$  retain more historical context. We therefore define context-dependent features as those with the largest one-third of  $\|\bar{A}\|_1$  values, and context-independent features as those with the lowest one-third of  $\|\bar{A}\|_1$  values.

## 4 Experiments

This section presents our main experimental contributions and is structured as follows: Section 4.1 presents the dataset used across our experiments. Section 4.2 examines the information flow to the last token from different sources, highlighting both shared and distinct behaviors across SSM-based and Transformer-based models. Section 4.3 validates our findings qualitatively and shows they are robust across a variety of model sizes. Section 4.4 investigates the features used to convey information between tokens. In-depth case studies are explored in Section 4.5. Throughout these subsections, the results are presented with a knockout window of size 9. An analysis of the impact of window size is presented in Section 4.6, which shows that for smaller models it is better to use smaller window sizes. Other additional experiments are provided in Appendix C. Code for reproducing all experiments can be found at <https://github.com/nirendy/mamba-knockout>.

### 4.1 Datasets

To evaluate the performance of the models with attention knocked-out, we utilized the COUNTERFACT dataset (Meng et al., 2022), which consists of factual triplets in the form (subject, relation, attribute). The task at hand requires models to predict a factual attribute (e.g., “Beats Music is owned by \_\_\_”). We focus on a subset of 672 triplets in which all models presented correctly predicted the next token. This choice disentangles the effects of the datasets and the model and ensures that any performance changes are solely due to the models. To ablate this choice, we also evaluated the effect of attention knockout on datasets composed exclusively of examples correctly predicted by each model (see Appendix C.1), yielding qualitatively similar results.

### 4.2 Information Flow to the Last Token

We investigate how the final token in Mamba-1 and Mamba-2 models transforms into the correct answer using the attention knockout method from Geva et al. (2023) (see Section 3.1), and compare the emerging patterns with those in the GPT-2 model (Radford et al., 2019). Specifically, we knockout the information flow to the final token from various sources—subject tokens, relation tokens, the first token, and the final token itself—with results shown in Figure 1.

#### 4.2.1 Shared Characteristics

Across all models, inhibiting the final token’s attention to subject tokens in late-intermediate layers consistently causes a notable drop in correct-token probability, mirroring patterns observed in Transformer-based models by Geva et al. (2023).

This result underscores the critical role of subject tokens in directing factual information flow within LLMs. More broadly, these compelling findings demonstrate that SSM-based and Transformer-based models share similar interpretability characteristics, hinting at universal properties across all attention-based architectures. In Section 4.4 we show a surprising connection between the behavior of subject tokens and that of context-dependent features.

#### 4.2.2 Distinct Characteristics

In addition to their shared characteristics, our analysis reveals that each model exhibits unique properties not present in all others.

First, we observe that, consistent with Geva et al. (2023), GPT-2 exhibits a pronounced first-position bias—a feature not shared by the SSM-based models. This finding aligns with previous research suggesting that the first-position attention sink is more evident in normalization-based attention methods (Gu et al., 2024).

Second, GPT-2 shows a strong dependence on information flow from relation tokens in the very early layers, a reliance that gradually diminishes in later layers. In contrast, SSM-based models only exhibit significant dependence on relation tokens in the later layers. Notably, both Mamba-1 and Falcon-Mamba display a pronounced knockout response, characterized by an initial drop in performance followed by a sharp increase in correct-token probability.

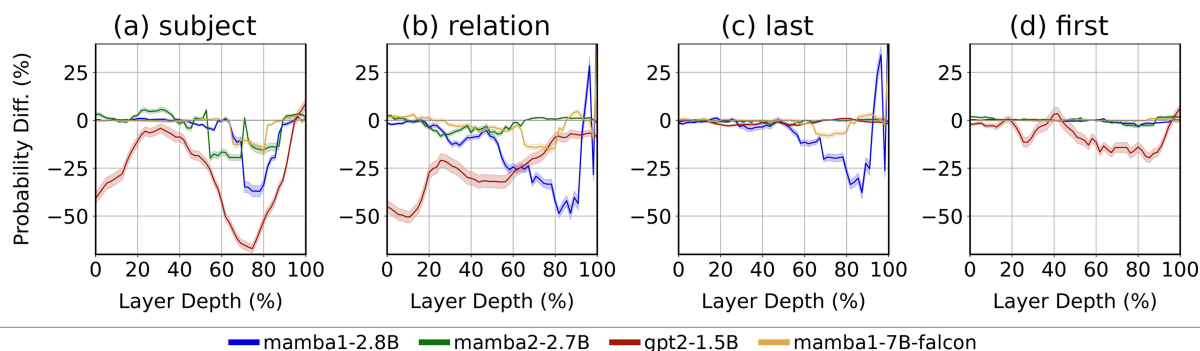


Figure 1: Relative change in correct-token prediction probability when removing information flow to the last token from various source tokens. The x-axis represents the relative depth of the first layer within the 9-layer attention knockout window, while the y-axis indicates the resulting performance change. The plots present model performance when knocking out (a) the subject, (b) the relation, (c) the last token, and (d) the first token. Notice that subject token knockout yields a consistent performance drop across models and similar layer-specific effects, highlighting the crucial role of subject tokens in LLMs. In contrast, knockouts of other tokens produce less robust and more variable outcomes. For ablations on dataset selection see Appendix C.1. See Appendix D for extended implementation details.

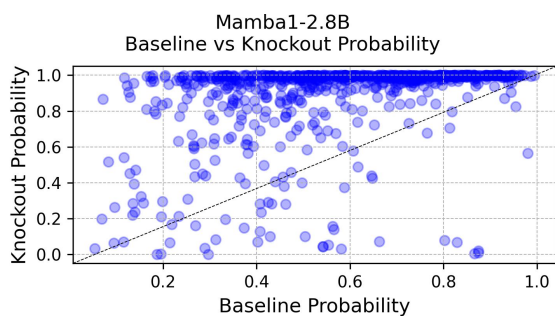


Figure 2: Impact of last token attention knockout on correct-token probability in Mamba-1. The x-axis displays the baseline probability for each query, while the y-axis shows the probability after knockout. Observe that removing the information flow from the final token to itself across the last 9 layers yields a significant increase in correct-token probability. For further implementation details, see Appendix D.

Surprisingly, we observe that both Mamba-1 and Falcon-Mamba also exhibit a marked increase in the correct-token probability at later layers when attention from the last input token is blocked. To investigate this phenomenon, we compare the correct token probability before and after the knockout in Figure 2, finding that it consistently surges to nearly 1, regardless of the initial likelihood. We find this result surprising and defer a deeper exploration of its implications to future work.

### 4.3 Architecture Consistency

To assess whether this pattern of results should be attributed to architectural differences, we

validate our findings by performing attention knockout on models of different sizes, presented in Figures 3 to 5. Our results show that while the smallest models seem to suffer from knockout much more substantially than the larger models (which could be explained due to the fact that our window size of choice blocks 37.5% of layers - see Section 4.6 for a more thorough analysis of this behavior) the overall patterns are consistent, with the main phenomenon shared throughout— all models show a reliance on direct information flow from the subject tokens to the last token in the late-intermediate layers. Conversely, as discussed in Section 4.2.2, the models exhibit distinct, architecture-specific phenomena; Mamba-1 models rely on both the last token and relation tokens in the later layers (see Figure 3); Mamba-2 models primarily depend on relation tokens in the later layers (see Figure 4); GPT-2 models show a strong dependence on relation tokens in the early layers, with this reliance gradually diminishing in later layers (see Figure 5). To reinforce the claim that the unique information flow patterns in Mamba variants arise from the SSM layers and their characteristics, we extended our information flow analysis to two additional variants of transformer-based LLMs, Llama 3 (Grattafiori et al., 2024) and Mistral (Jiang, 2024). The results show a consistent pattern shared between all transformer-based LLMs which is distinct from the patterns observed in Mamba-based models. Appendix C.3 presents detailed experiments on this.

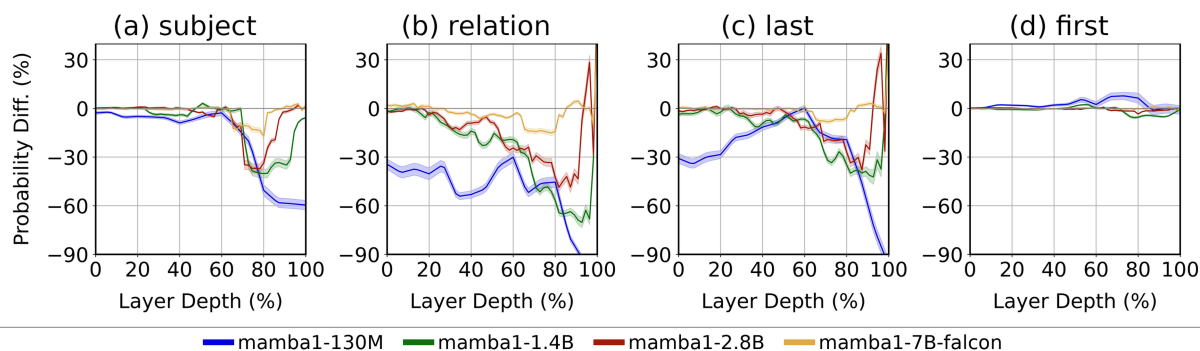


Figure 3: Relative change in correct-token prediction probability when removing information flow to the final token from various source tokens across Mamba-1 models of varying sizes. This figure is identical to Figure 1, but presents the performance of Mamba-1 models of sizes 130M, 1.4B, and 2.8B. As observed in other models, subject token knockout consistently reduces performance. Unique to Mamba-1, knockouts of relation tokens and of the final token in the ultimate layers initially lower the correct-token probability before a pronounced increase. For further details see Figure 1 and Appendix D. For ablations on window size see Section 4.6.

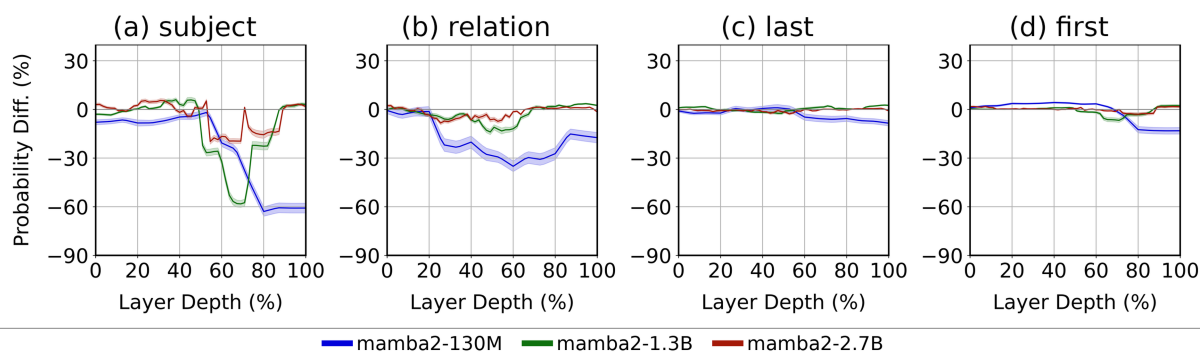


Figure 4: Relative change in correct-token prediction probability when removing information flow to the final token from various source tokens across Mamba-2 models of varying sizes. This figure is identical to Figure 1, but presents the performance of Mamba-2 models of sizes 130M, 1.3B, and 2.7B. As observed in other models, subject token knockout consistently reduces performance. Unique to Mamba-2, knockout of relation tokens in the later layers initially lowers the correct-token probability before a pronounced increase. For further details see Figure 1 and Appendix D. For ablations on window size see Section 4.6.

#### 4.4 Feature Knockout

Expanding on Section 3.5, we selectively perform attention knockout on context-dependent and context-independent features to investigate their distinct roles. We compare feature knockout from subject tokens to the last token under three conditions: (1) knocking out *all* features (as in Section 4.2), (2) knocking out only *context-dependent* features, and (3) knocking out only *context-independent* features. Figure 6 shows that performing knockout on context-dependent features (slower decay) yields behavior similar to performing attention knockout in all features. In contrast, performing knockout in context-independent features (faster decay) shows little to no impact. These observations underscore the pivotal role of context-dependent features in subject tokens, effectively mirroring the impact of knocking

out all features. While this analysis remains correlational and does not yet establish causation, it opens promising directions for future interpretability research in SSM-based architectures: much like how attention maps offer interpretability insights in Transformer architectures, evaluating a feature’s degree of context dependence in SSM-based models could provide a valuable lens into how tokens behave across layers and influence the model’s decision making. Furthermore, building on these insights, future work could potentially investigate whether selectively fine-tuning context-dependent features alone enables more efficient training in SSM-based architectures. Appendix B proposes several practical applications, while a full exploration of their potential is left to future investigations. We extended our analysis of context-dependent and independent features to additional variants of Mamba-1 and

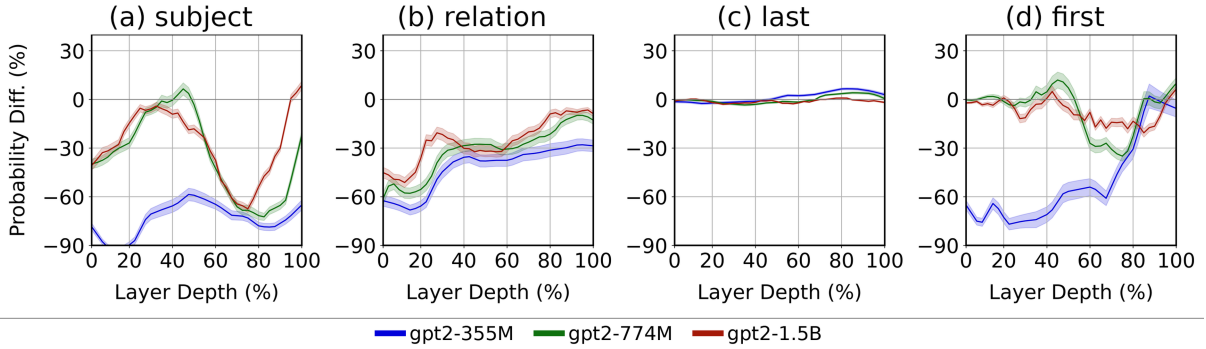


Figure 5: Relative change in correct-token prediction probability when removing information flow to the final token from various source tokens across GPT-2 models of varying sizes. This figure is identical to Figure 1, but presents the performance of GPT-2 models of sizes 355M, 774M, and 1.5B. As observed in other models, subject token knockout consistently reduces performance. Unique to GPT-2, knockout of relation tokens in earlier layers profoundly lowers the correct-token probability. Additionally, GPT-2 models exhibit a robust first token bias. For further details see Figure 1 and Appendix D.

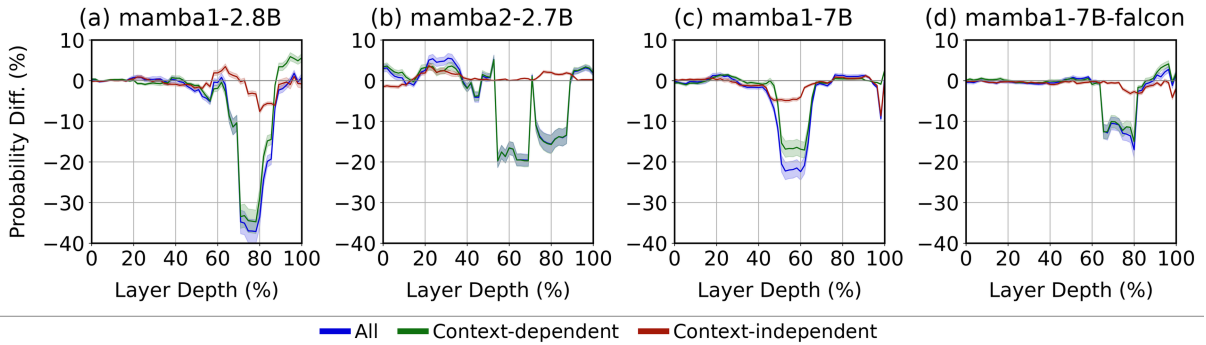


Figure 6: Relative change in correct-token probability when knocking out connections from subject tokens to the last token across different feature categories for (a) Mamba-1, (b) Mamba-2, and (c) Falcon-Mamba. We compare three settings: knockout of all features (blue), knockout of context-independent features (red), and knockout of context-dependent features (green). Notably, knocking out context-dependent features alone closely mirrors the effect of removing all features. For implementation details see Appendix D

Mamba-2 models in Appendix C.2. The supplementary findings reaffirm our conclusions.

#### 4.5 Specific Case Studies

Our final contribution is a detailed analysis of specific examples and the impact of attention knockout from different tokens on model performance. An example is shown in Figure 7. Our experiments reveal two key findings. First, Mamba-1 and Mamba-2 models do not exhibit a profound first-token bias seen in Transformer-based models (Xiao et al., 2024), as blocking the first token from attending to the last token does not significantly affect performance. Second, Mamba-1 models heavily rely on the last token, a behavior not observed in Mamba-2 models as can also be seen in Figure 1. We attribute this difference to architectural variations between the models and anticipate that future work will explore its underlying causes. More

examples are available in Appendix C.4.

#### 4.6 Effect of Window Size

The choice of window size for the knockout analyses can produce varied results. On one hand, choosing too small a window may show little to no effect, as information flow could be distributed across multiple consecutive layers. On the other hand, choosing too large a window reduces resolution, and impairs our ability to reach meaningful conclusions regarding information flow in the model.

To empirically examine the effect of window size, we perform the knockout experiment in Section 4.2 using several window sizes - 1, 5, 9, 12, 15, on both Mamba-1 (Figure 8, Figure 9, Figure 10) and Mamba-2 (Figure 14 in Appendix C) of various sizes. Results corroborate our intuition, with reduced window size showing lesser effects, and higher window sizes

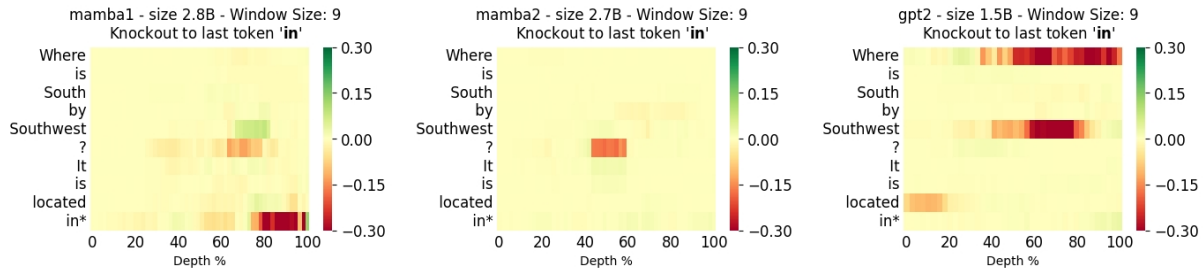


Figure 7: Relative change in correct-token probability following the removal of attention from each token to the final token. Columns indicate the first layer in the knockout window, and rows correspond to the token from which attention was removed. Results are shown for the query “Where is South by Southwest? It is located in \_\_\_\_”. Note that in contrast to GPT-2, neither Mamba-1 nor Mamba-2 exhibit a strong reliance on the first token, while Mamba-1 exhibits marked dependence on the final token. For additional experiments and implementation specifics, see Appendices C and D.

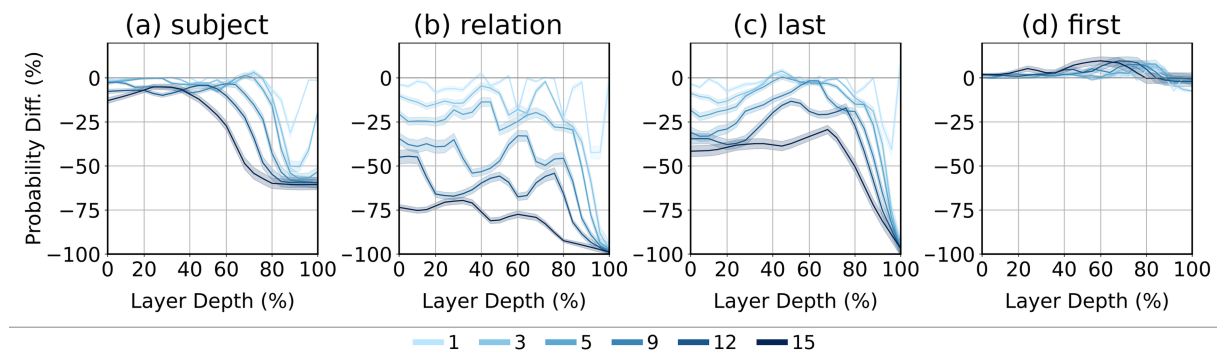


Figure 8: Relative change in correct-token probability after removing the information flow to the final token from various source tokens for Mamba-1 130M across different window sizes. This figure, which parallels Figure 1, displays performance for window sizes of 1, 3, 5, 9, 12, and 15. The qualitative patterns observed for Mamba-1 remain consistent, with larger window sizes amplifying these effects. For additional details, see Figure 1 and Appendix D.

showing greater effects. Larger window sizes display smoother, more diffuse patterns, while smaller window sizes are characterized by short surges, indicating specific layers are more critical in the flow of information. While the overall pattern of information flow remains the same, with the critical point of subject-to-last token information flow in the late-intermediate layers remaining the same, we can see that it is better to use smaller window sizes with smaller models. While in Figures 3 to 5 we see a different behavior for the small models, if we look at their behavior with smaller window sizes, it is more similar to the larger models. The reason for that is that the smaller model has less parameters and, therefore, larger window sizes blocks a large portion of the information, e.g., blocking 37.5% of blocks for window size 9. Therefore, it might be better to use a smaller window size with knockout to get a better resolution with small models and in this way better analyze their behavior.

## 5 Conclusion

We observe some similarities between SSM-based and Transformer-based models in extracting factual attributes and transferring information from the subject to the last token. We show that the models rely on different parts of the sentence to extract the correct answer, depending on their architecture. Using a unique property of SSM-based models, we identify the functionality of specific features based on their weights.

This work serves as both an initial exploration of the unique properties of SSM-based language models and a deeper investigation of the shared foundations of language understanding across neural network models. The explainability approach presented here may lead to more unified methods for explaining the internal computations of diverse neural architectures.

Our findings offer several key contributions. They can guide training decisions for language models by enabling targeted pruning



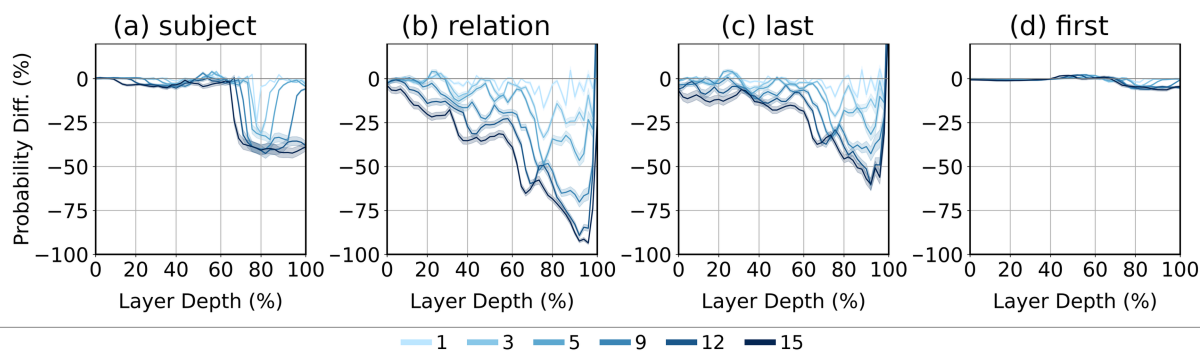


Figure 9: Relative change in correct-token probability after removing the information flow to the final token from various source tokens for Mamba-1 1.4B across different window sizes. This figure is identical to Figure 8 but instead presents the results for Mamba-1 1.4B. For additional details see Figures 1 and 8 and Appendix D.

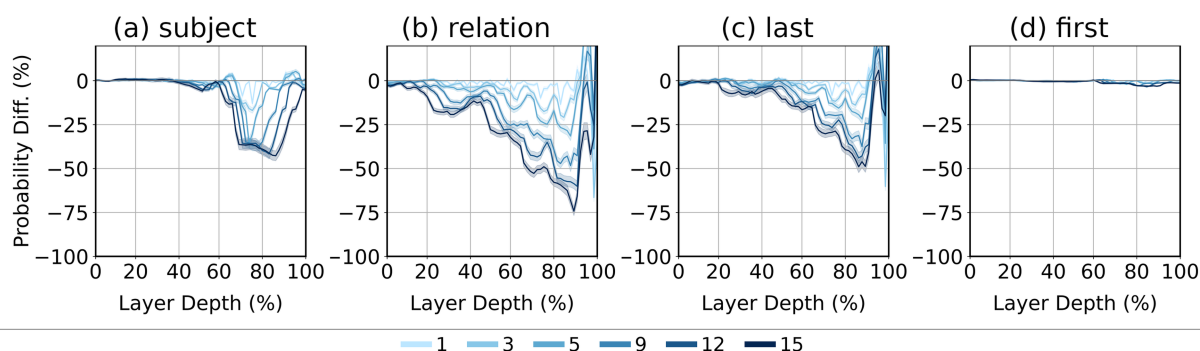


Figure 10: Relative change in correct-token probability after removing the information flow to the final token from various source tokens for Mamba-1 2.8B across different window sizes. This figure is identical to Figure 8 but instead presents the results for Mamba-1 2.8B. For additional details, see Figures 1 and 8 and Appendix D.

that eliminates redundant information flows and reduces computational overhead. Moreover, our work lays the groundwork for fine-tuning strategies—such as adjusting Mamba’s final layer—to modify language distribution without sacrificing factual accuracy. Finally, distinguishing between shared and unique patterns across architectures can deepen our understanding of language modelling as a whole, and inform the design of more robust AI models.

## 6 Limitations

In this work we primarily use attention-knockout to examine how information flows between tokens in Mamba-1 and Mamba-2. While this approach shows causality, when severing token connections changes the correct attribute probability, it is far from exhausting the complete pattern of information flow in the model. First, our use of this approach is limited by its application in continuous layer windows, with all SSMs in a layer affected simultaneously. A much more diffuse and intricate pattern could occur in reality, while this approach would fail to recognize it. For a more holistic

analysis, fine-tuning-based approaches should be used. While we identify critical connections between tokens, our method does not elaborate on what information actually passes through these connections. Future works should attempt to decipher the content of these internal representations. As the attention knockout effect is not ecological, any change in probability could be attributed to it, however the localized pattern suggests that changes are due to specific missing information. Attention-knockout focuses only on the flow of information between tokens, thus it cannot be used in order to explain the role of token-independent operations (e.g. gating, convolution, etc.). As stated in Section 2, we focus on between-token information flow as it is the key difference between selective-SSMs and attention. Our methods recognize how information flows between tokens and the critical points in this contextualization, however, they do not shed light on why both Transformers and Mamba models converge to these similar patterns. Future works should attempt to explain the inductive-biases of the different architectures that drive these differences in information flow.

## Acknowledgements

IDG, YRM and YS are supported in part by the Tel Aviv University Center for AI and Data Science.

## References

- Ameen Ali, Itamar Zimerman, and Lior Wolf. 2024. [The hidden attention of mamba models](#). *Preprint*, arXiv:2403.01590.
- Assaf Ben-Kish, Itamar Zimerman, Shady Abu-Hussein, Nadav Cohen, Amir Globerson, Lior Wolf, and Raja Giryes. 2025. Decimamba: Exploring the length extrapolation potential of mamba. In *The Thirteenth International Conference on Learning Representations (ICLR)*.
- Sid Black, Stella Biderman, Eric Hallahan, Quentin Anthony, Leo Gao, Laurence Golding, Horace He, Connor Leahy, Kyle McDonell, Jason Phang, et al. 2022. Gpt-neox-20b: An open-source autoregressive language model. *arXiv preprint arXiv:2204.06745*.
- Trenton Bricken, Adly Templeton, Joshua Batten, Brian Chen, Adam Jermyn, Tom Conerly, Nick Turner, Cem Anil, Carson Denison, Amanda Askell, et al. 2023. Towards monosemanticity: Decomposing language models with dictionary learning. *Transformer Circuits Thread*, 2.
- Arthur Conmy, Augustine Mavor-Parker, Angus Lynch, Stefan Heimersheim, and Adrià Garriga-Alonso. 2023. Towards automated circuit discovery for mechanistic interpretability. *Advances in Neural Information Processing Systems*, 36:16318–16352.
- Tri Dao and Albert Gu. 2024. Transformers are ssms: Generalized models and efficient algorithms through structured state space duality. *arXiv preprint arXiv:2405.21060*.
- Nelson Elhage, Tristan Hume, Catherine Olsson, Nicholas Schiefer, Tom Henighan, Shauna Kravec, Zac Hatfield-Dodds, Robert Lasenby, Dawn Drain, Carol Chen, et al. 2022. Toy models of superposition. *arXiv preprint arXiv:2209.10652*.
- Nelson Elhage, Neel Nanda, Catherine Olsson, Tom Henighan, Nicholas Joseph, Ben Mann, Amanda Askell, Yuntao Bai, Anna Chen, Tom Conerly, et al. 2021. A mathematical framework for transformer circuits. *Transformer Circuits Thread*, 1(1):12.
- Matthew Finlayson, Aaron Mueller, Sebastian Gehrmann, Stuart Shieber, Tal Linzen, and Yonatan Belinkov. 2021. Causal analysis of syntactic agreement mechanisms in neural language models. *arXiv preprint arXiv:2106.06087*.
- Mor Geva, Jasmijn Bastings, Katja Filippova, and Amir Globerson. 2023. Dissecting recall of factual associations in auto-regressive language models. *arXiv preprint arXiv:2304.14767*.
- Mor Geva, Roei Schuster, Jonathan Berant, and Omer Levy. 2021. [Transformer feed-forward layers are key-value memories](#). *Preprint*, arXiv:2012.14913.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Albert Gu and Tri Dao. 2023. Mamba: Linear-time sequence modeling with selective state spaces. *arXiv preprint arXiv:2312.00752*.
- Albert Gu and Tri Dao. 2024. [Mamba: Linear-time sequence modeling with selective state spaces](#). *Preprint*, arXiv:2312.00752.
- Xiangming Gu, Tianyu Pang, Chao Du, Qian Liu, Fengzhuo Zhang, Cunxiao Du, Ye Wang, and Min Lin. 2024. [When attention sink emerges in language models: An empirical view](#). *Preprint*, arXiv:2410.10781.
- Peter Hase, Mohit Bansal, Been Kim, and Asma Ghandeharioun. 2023. [Does localization inform editing? surprising differences in causality-based localization vs. knowledge editing in language models](#). *Preprint*, arXiv:2301.04213.
- Tommy Ip. 2023. [Github repository for mamba implementation](#). Accessed: 2024-10-26.
- Fengqing Jiang. 2024. Identifying and mitigating vulnerabilities in llm-integrated applications. Master’s thesis, University of Washington.
- Angelos Katharopoulos, Apoorv Vyas, Nikolaos Pappas, and François Fleuret. 2020. [Transformers are RNNs: Fast autoregressive transformers with linear attention](#). In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 5156–5165. PMLR.
- Kevin Meng, David Bau, Alex Andonian, and Yonatan Belinkov. 2022. Locating and editing factual associations in gpt. *Advances in Neural Information Processing Systems*, 35:17359–17372.
- Kevin Meng, Arnab Sen Sharma, Alex Andonian, Yonatan Belinkov, and David Bau. 2023. [Mass-editing memory in a transformer](#). *Preprint*, arXiv:2210.07229.
- Paul Michel, Omer Levy, and Graham Neubig. 2019. Are sixteen heads really better than one? *Advances in neural information processing systems*, 32.
- Eshaan Nichani, Jason D. Lee, and Alberto Bietti. 2024. [Understanding factual recall in transformers via associative memories](#). *Preprint*, arXiv:2412.06538.
- Chris Olah, Nick Cammarata, Ludwig Schubert, Gabriel Goh, Michael Petrov, and Shan

Carter. 2020. Zoom in: An introduction to circuits. *Distill*, 5(3):e00024–001.

Catherine Olsson, Nelson Elhage, Neel Nanda, Nicholas Joseph, Nova DasSarma, Tom Henighan, Ben Mann, Amanda Askell, Yuntao Bai, Anna Chen, et al. 2022. In-context learning and induction heads. *arXiv preprint arXiv:2209.11895*.

Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Yang, Zach DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. *Pytorch: An imperative style, high-performance deep learning library*. *Preprint*, arXiv:1912.01703.

Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.

Arnab Sen Sharma, David Atkinson, and David Bau. 2024. *Locating and editing factual associations in mamba*. *Preprint*, arXiv:2404.03646.

Roger Waleffe, Wonmin Byeon, Duncan Riach, Brandon Norick, Vijay Korthikanti, Tri Dao, Albert Gu, Ali Hatamizadeh, Sudhakar Singh, Deepak Narayanan, Garvit Kulshreshtha, Vartika Singh, Jared Casper, Jan Kautz, Mohammad Shoeybi, and Bryan Catanzaro. 2024. *An empirical study of mamba-based language models*. *Preprint*, arXiv:2406.07887.

Kevin Wang, Alexandre Variengien, Arthur Conmy, Buck Shlegeris, and Jacob Steinhardt. 2022. Interpretability in the wild: a circuit for indirect object identification in gpt-2 small, 2022. URL <https://arxiv.org/abs/2211.00593>, 2.

Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2020. *Huggingface’s transformers: State-of-the-art natural language processing*. *Preprint*, arXiv:1910.03771.

Guangxuan Xiao, Yuandong Tian, Beidi Chen, Song Han, and Mike Lewis. 2024. *Efficient streaming language models with attention sinks*. *Preprint*, arXiv:2309.17453.

Itamar Zimmerman, Ameen Ali Ali, and Lior Wolf. 2025. Explaining modern gated-linear RNNs via a unified implicit attention formulation. In *The Thirteenth International Conference on Learning Representations (ICLR)*.

## A Extended Related Work

Our work contributes to the growing field of mechanistic interpretability, which seeks to explain neural network behaviors by analyzing internal circuits. The circuit analysis paradigm was introduced by Olah et al. (2020), who proposed that groups of neurons form interpretable sub-networks (circuits). Elhage et al. (2021) extended this approach to Transformers by reverse-engineering a small GPT-2 model, identifying induction-head circuits that facilitate in-context copying. Similarly, Olsson et al. (2022) explored how certain attention heads enable in-context learning.

Beyond attention mechanisms, other research examined MLP pathways and feature representations. Geva et al. (2021) demonstrated that Transformer MLP layers can act as key-value memories for factual associations (Meng et al., 2022). Complementary studies by Finlayson et al. (2021) used causal interventions to trace linguistic features through Transformer layers, revealing assemblies of neurons mediating specific features. Recent work on polysemantic versus monosemantic neurons (Elhage et al., 2022; Bricken et al., 2023) further highlights the importance of disentangling neuron functions to improve interpretability.

Our Mamba-based approach for intervening in information flow aligns with prior methods of causal tracing and circuit testing. Meng et al. (2022, 2023) introduced causal tracing to identify critical layers for factual storage, using targeted edits (ROME and MEMIT) for knowledge modification. Hase et al. (2023) similarly probed Transformer layers to localize factual information. Our attention-knockout method builds upon these studies, adapting fine-grained intervention techniques to the state-space architecture of Mamba.

We also connect to automated circuit discovery techniques such as path patching (Wang et al., 2022; Conmy et al., 2023), where researchers systematically test causal paths within models. Our knockout method similarly isolates token-to-token information channels, providing insight into causal computation pathways.

## B Practical Applications

The interpretability methods presented in this work provide practical directions for targeted model interventions. We propose a few of them below.

**Pruning.** Our distinction between context-dependent features—those critical for

inter-token information flow—and context-independent features—focused primarily on single-token processing—enables smarter pruning strategies. Context-independent features often contribute minimally to between-token information-flow and therefore can be safely pruned, reducing model complexity without significant performance loss. This targeted approach mirrors prior work in Transformers where redundant attention heads were pruned effectively (Michel et al., 2019).

**Fine-Tuning and Adaptation.** Context-dependent features identified via knockout analysis represent targeted fine-tuning opportunities. Adjusting only these crucial features could efficiently adapt the model to new tasks or factual domains without global retraining. This focused fine-tuning strategy aligns with prior findings in Transformers (Hase et al., 2023) where precise adjustments to identified causal pathways effectively updated model knowledge. Our methods similarly offer practical leverage points for precise and efficient model adaptation.

**Model Editing.** Token-level knockout identifies key token interactions critical for factual recall, highlighting optimal sites for targeted model edits. While similar causal tracing in Transformer models has informed successful editing strategies (Meng et al., 2022, 2023), our approach extends these methods to SSM-based models such as Mamba-1 and Mamba-2. In summary, the interpretability techniques introduced here serve as diagnostic tools to inform strategic model modifications, pruning, and fine-tuning.

## C Further Experiments

### C.1 Dataset Ablations

We analyzed a subset of factual associations where all models correctly predicted the next token to disentangle the effects of the datasets and the model and ensure that any performance changes are solely due to the models. A potential drawback of this approach is that the selected subset could introduce bias, potentially leading to misleading conclusions. To assess whether this is the case, we performed the same knockout analysis on the set of sentences for which each model made a correct prediction, allowing for a broader comparison. The results of these evaluations are shown in Figure 11. In addition, we present results on the subset of inputs for which all models correctly predicted the next token. The figure shows that the same

patterns emerge across these subsets with minimal differences, confirming that our approach robustly identifies consistent phenomena across the different models.

### C.2 Further Comparisons with Transformer-Based Models

To confirm that the difference in information flow patterns between Mamba models (Gu and Dao, 2023; Dao and Gu, 2024) and GPT-2 models (Radford et al., 2019) can be attributed to architectural differences, we report auxiliary information flow experiments performed on additional Transformer-based LLMs. Specifically, we analyzed three variants of the Llama 3 model (Grattafiori et al., 2024) and two variants of the Mistral model (Jiang, 2024). The results are reported in Figure 13. All transformer-based variants show shared characteristics, implying they should be attributed to the shared architecture: (i) all models show a similar pattern of a critical flow of information passing from the subject tokens to the last token in the late intermediate layers; (ii) all models give little importance to information flow from the last token to itself; and; (iii) nearly all models show strong evidence for first-position bias, similar to previous findings (Xiao et al., 2024), with the most substantial impact observed in the final layers.

### C.3 Feature Knockout on Additional Models

To further examine the role of context-dependent features in Mamba models, we performed the same feature-knockout analysis on smaller variants of the Mamba-1 and Mamba-2 models. Results are reported in Figure 12. All models show consistent roles for context-dependent and independent features, with knockout of context-dependent features from subject tokens having a much greater impact than that of context-independent features. On the other hand, knockout of context-dependent features from either the relation tokens, the last token or the first token shows similar patterns to knockout of context-independent features.

### C.4 Specific Case Studies

We extend the analysis of Section 4.5 to additional examples, architectures, and knockout-window sizes.

**Across Architectures and Model Sizes.** Figure 16 displays ablation heatmaps for

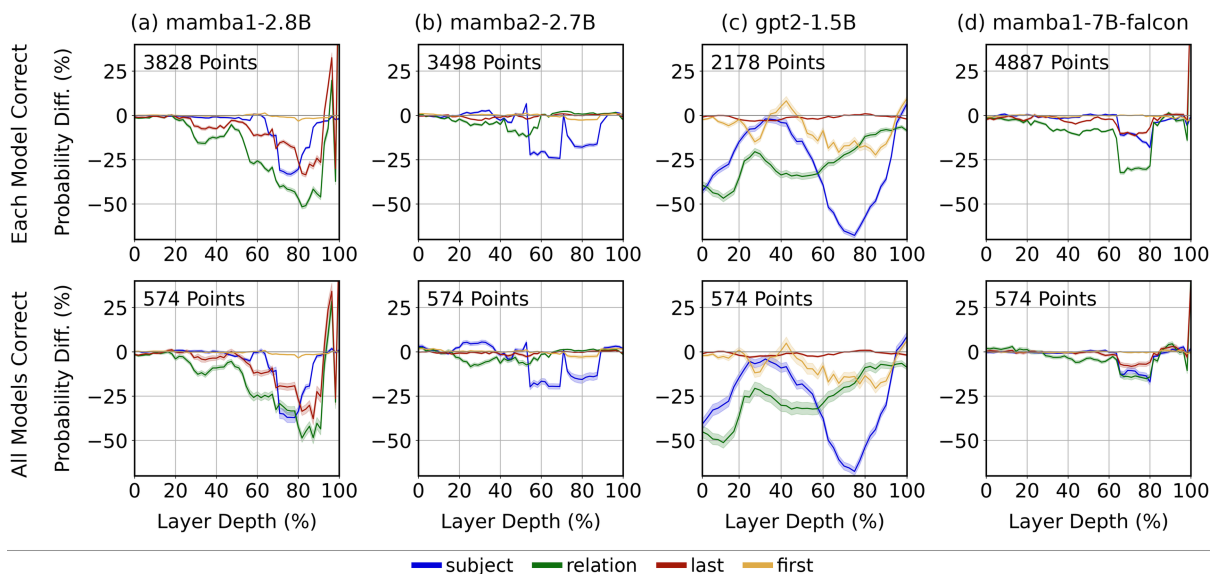


Figure 11: Relative change in correct-token probability after knocking out the connection from each token to the last token in (a) Mamba-1 2.8B, (b) Mamba-2 2.7B, (c) GPT-2 1.5B, (d) Mamba-1 7B Falcon. The top row shows results on the subset of samples each model was correct on. The bottom row shows results on the subset of samples where all models were correct. The X-axis indicates the relative depth of the layer on which we perform the knockout. The y-axis indicates the relative change in correct-token probability. Different colors indicate the source token.

Mamba-1 and Mamba-2. The patterns highlighted persist: Mamba-1 shows a pronounced reliance on the final token with little first-token bias, whereas Mamba-2 distributes influence more evenly across the sequence. In both Mamba variants, we also observe a concentrated band of sensitivity around the subject tokens - the same band that emerged in the information-flow graphs at roughly the 70 depth layer—underscoring the models’ focus on the entity whose attribute is being predicted. The same analysis for Transformer baselines is presented in Figure 17, where a strong first-token dependence is evident at every size.

**Effect of Knockout-Window Size.** To isolate how window size modulates these effects, we vary  $WS \in \{1, 5, 9, 15\}$  in the selected model architectures and sizes (Figures 18 to 20). Smaller  $WS$  values produce sparser heatmaps, but pinpoint the layers that drive the largest probability shifts. For example, the surge in final-token probability for Mamba-1 can be traced to ablation in the single last layer ( $WS = 1$ ). Similar layer-specific contributions are visible for individual phrases in both Mamba-2 and GPT-2. In contrast, some effects only emerge when the ablation window is sufficiently deep: in GPT-2, the characteristic dependence of the first token is barely detectable at  $WS = 1$  but becomes more pronounced as the window expands

## D Implementation Details

Code for reproducing all experiments can be found at <https://github.com/nirendy/mamba-knockout>. All experiments were conducted on a single NVIDIA A100 GPU and implemented using PyTorch (Paszke et al., 2019). The attention-based Mamba-2 model was evaluated within the same environment across all experiments to ensure consistency. For the Mamba-1 analyses, we used pretrained weights from Gu and Dao (2024) published in the Hugging Face framework (Wolf et al., 2020). For the Mamba-2 analyses, we modified the implementation of Ip (2023). As the original Mamba-2 implementation doesn’t compute the attention matrix, relying instead on a more efficient computation strategy, we re-implemented the selective SSM layer to explicitly construct the attention matrix while maintaining the functionality. The weights of the model and the tokenizer, gpt-neox-20b (Black et al., 2022), were both imported from Hugging Face (Wolf et al., 2020). To obtain the effect of knockout on the Transformer-based GPT-2, we used the implementation published by Geva et al. (2023).

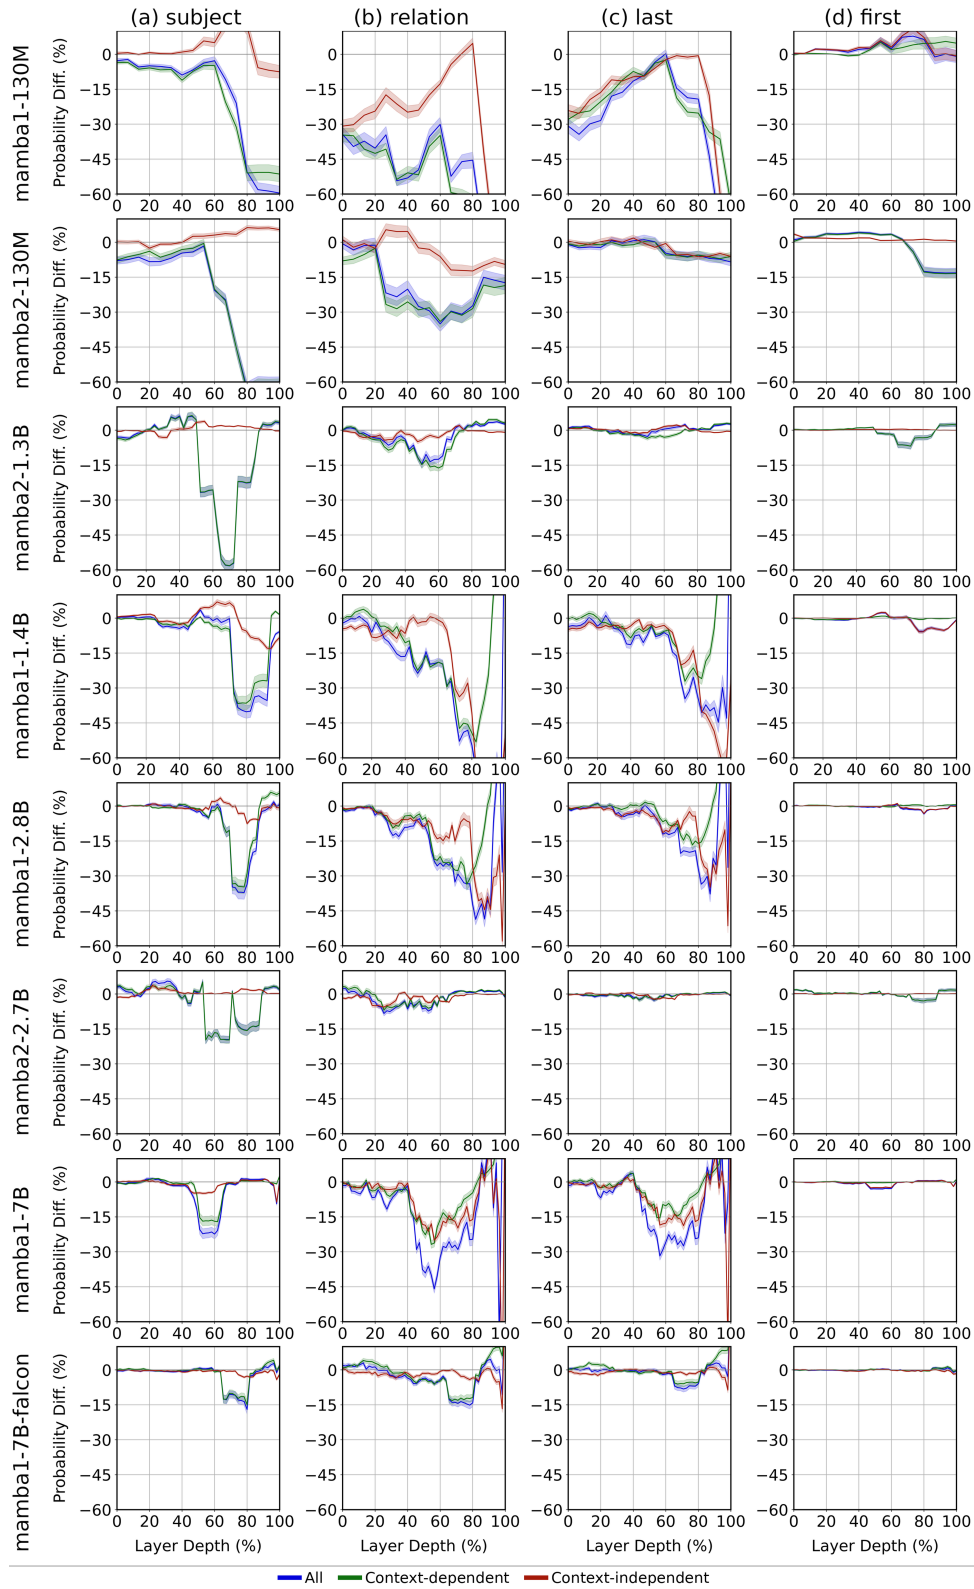


Figure 12: A full comparison of the effect of feature-knockout on all size variants of Mamba models. Each column shows the effect when knocking out information flow from (a) the subject, (b) the relation, (c) the last, and (d) the first tokens to the last token. Each row shows the results for a different model.

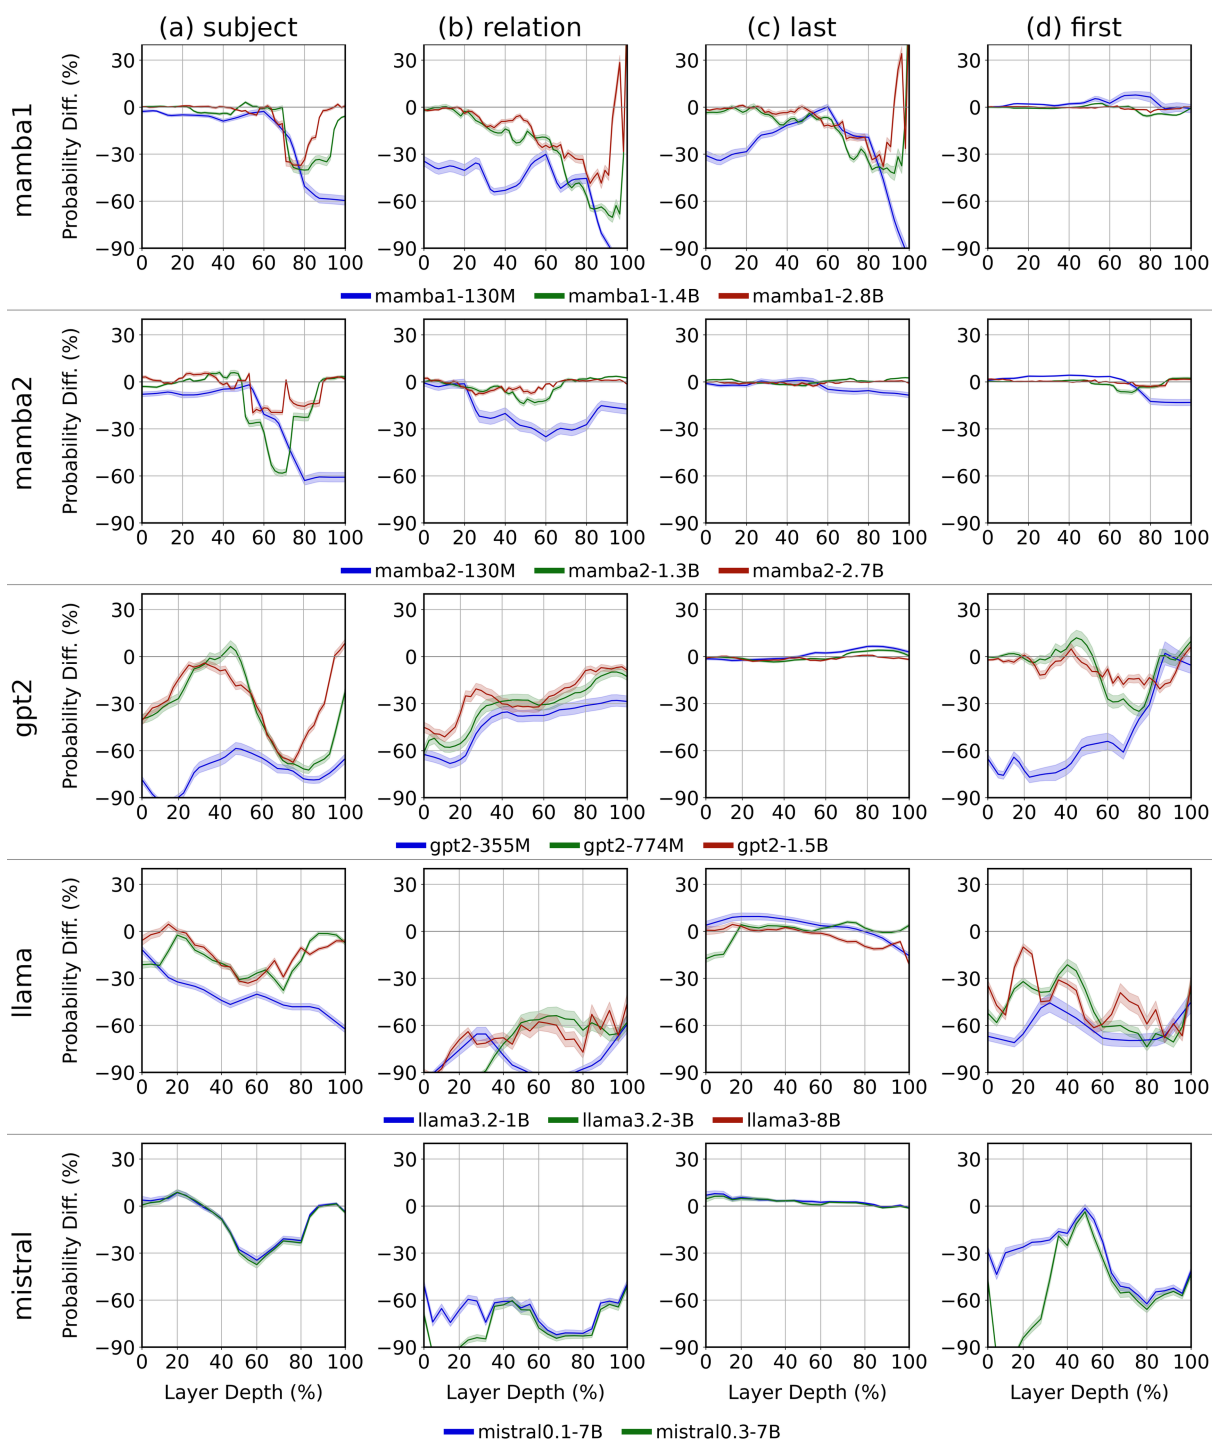


Figure 13: A full comparison of the effect of attention knockout on all model variants. Each column indicates the knockout target. Each row indicates the model family (Mamba-1, Mamba-2, Falcon Mamba, GPT-2, Llama, Mistral). Different coloring indicates the specific model variant.

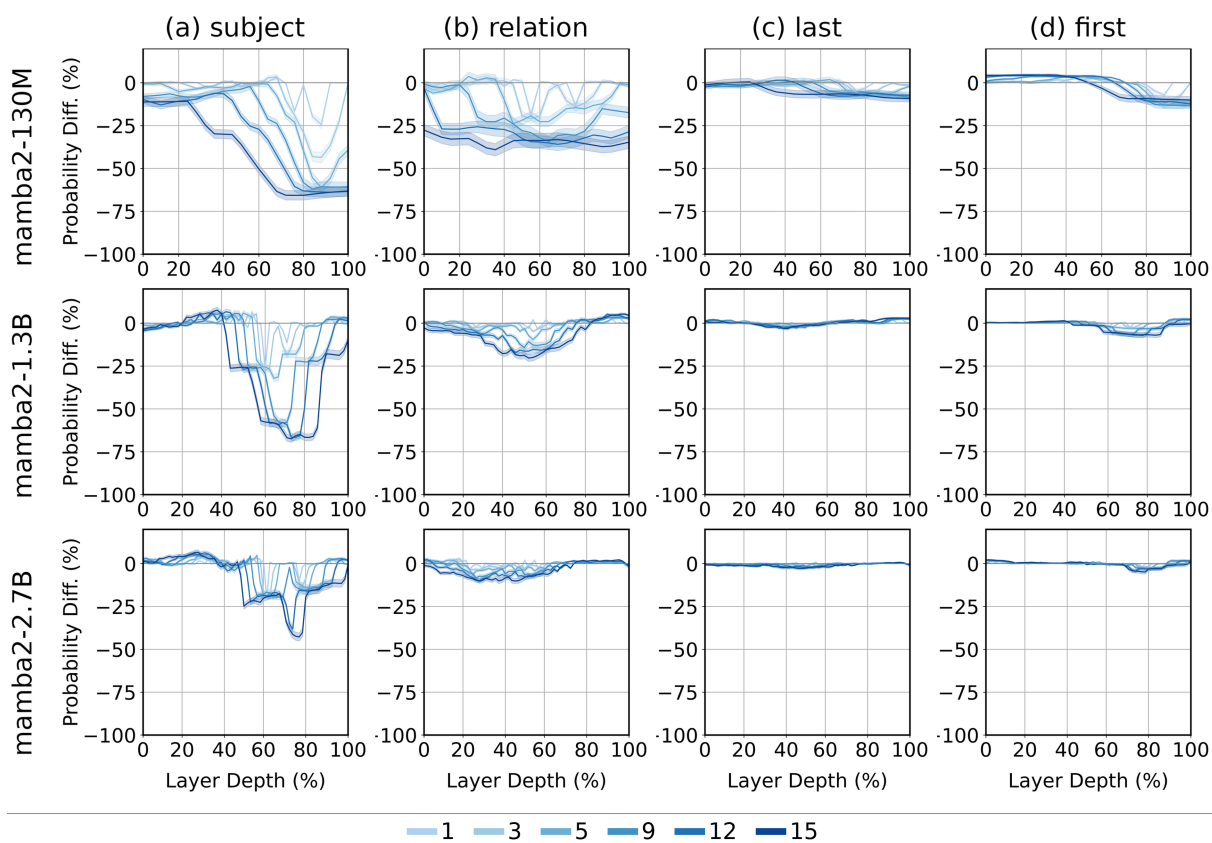


Figure 14: Relative change in correct-token probability after removing the information flow to the final token from various source tokens (columns) on Mamba-2 with various model sizes (rows) across different window sizes (lines). This figure, which parallels Figure 1, displays performance for window sizes  $\in \{1, 3, 5, 9, 12, 15\}$ . The qualitative patterns observed for Mamba-2 remain consistent, with larger window sizes amplifying these effects. For additional details, see Figure 1 and Appendix D.



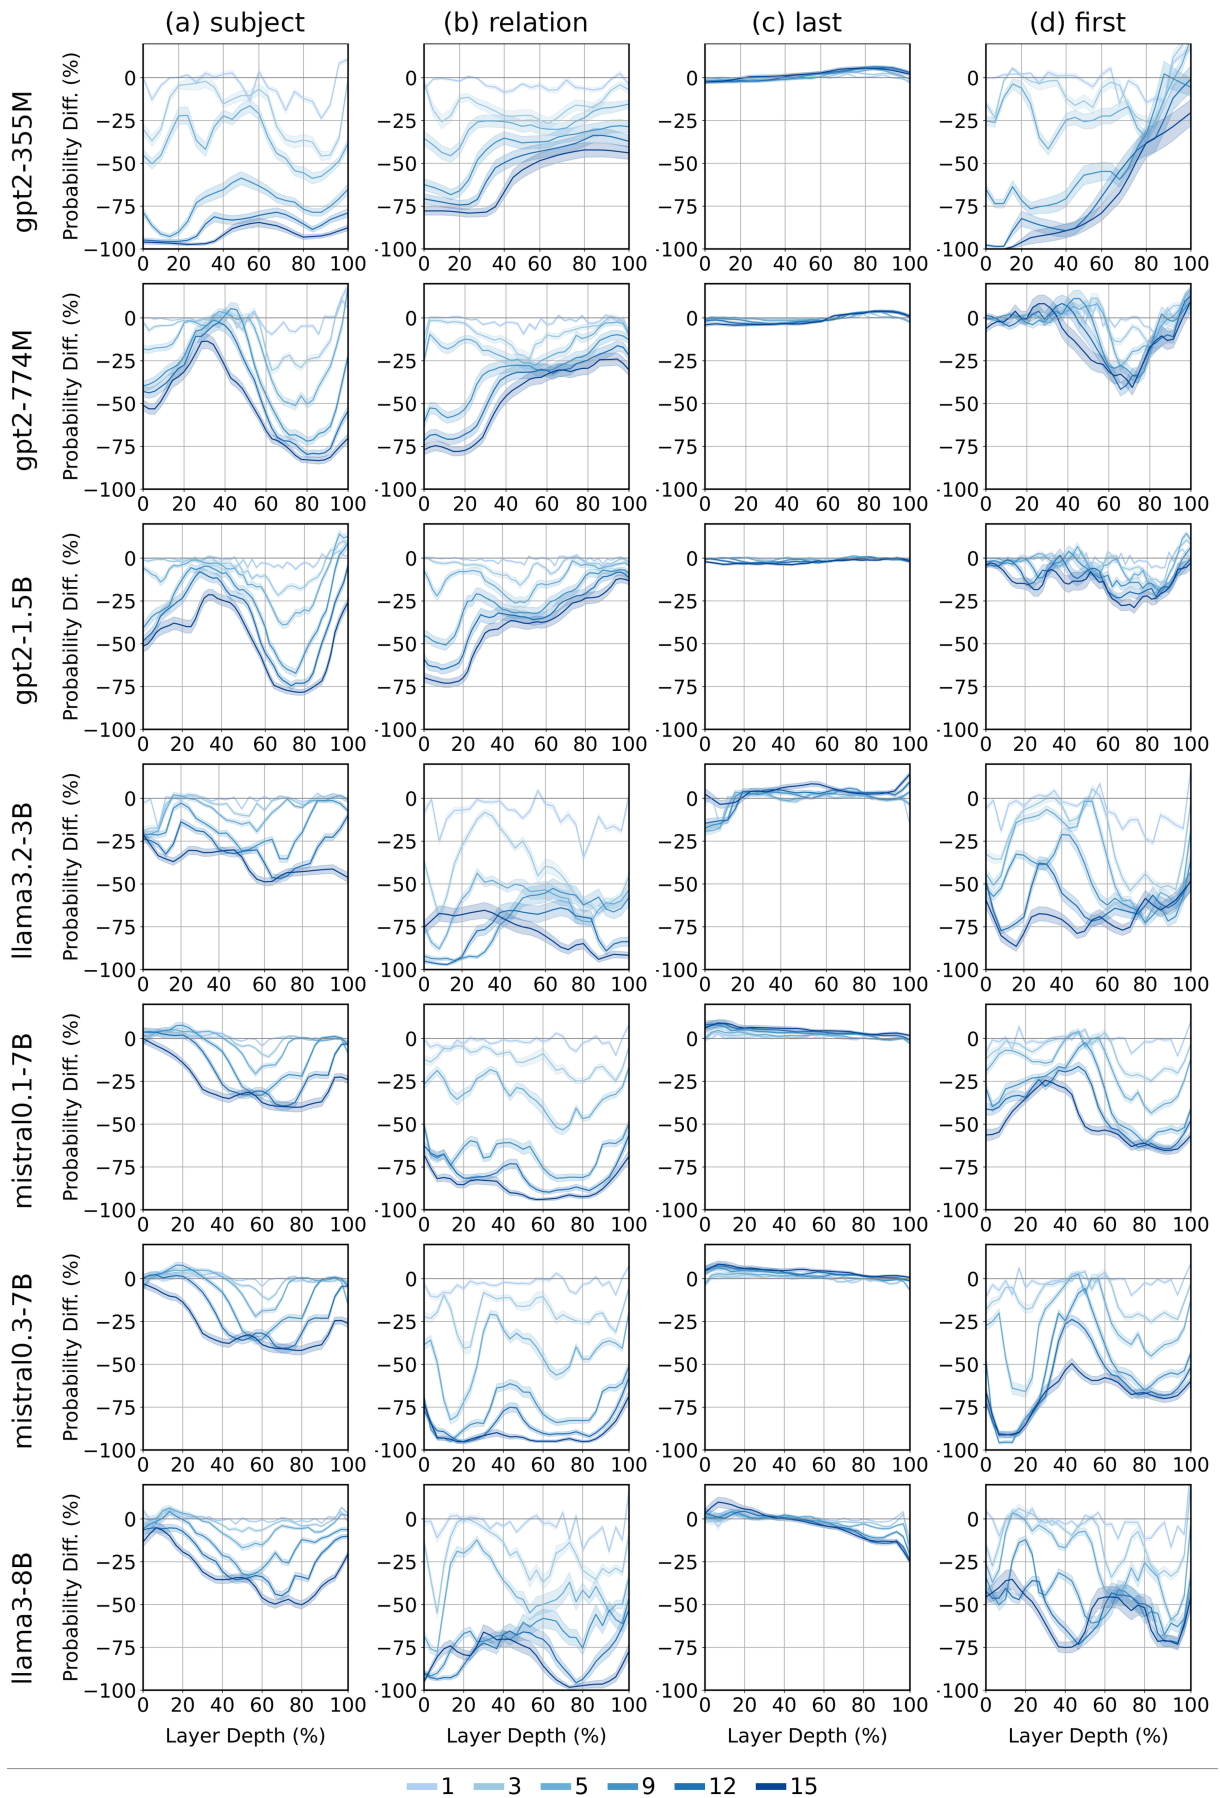


Figure 15: Identical analysis and layout to Figure 14, but applied to Transformer baselines

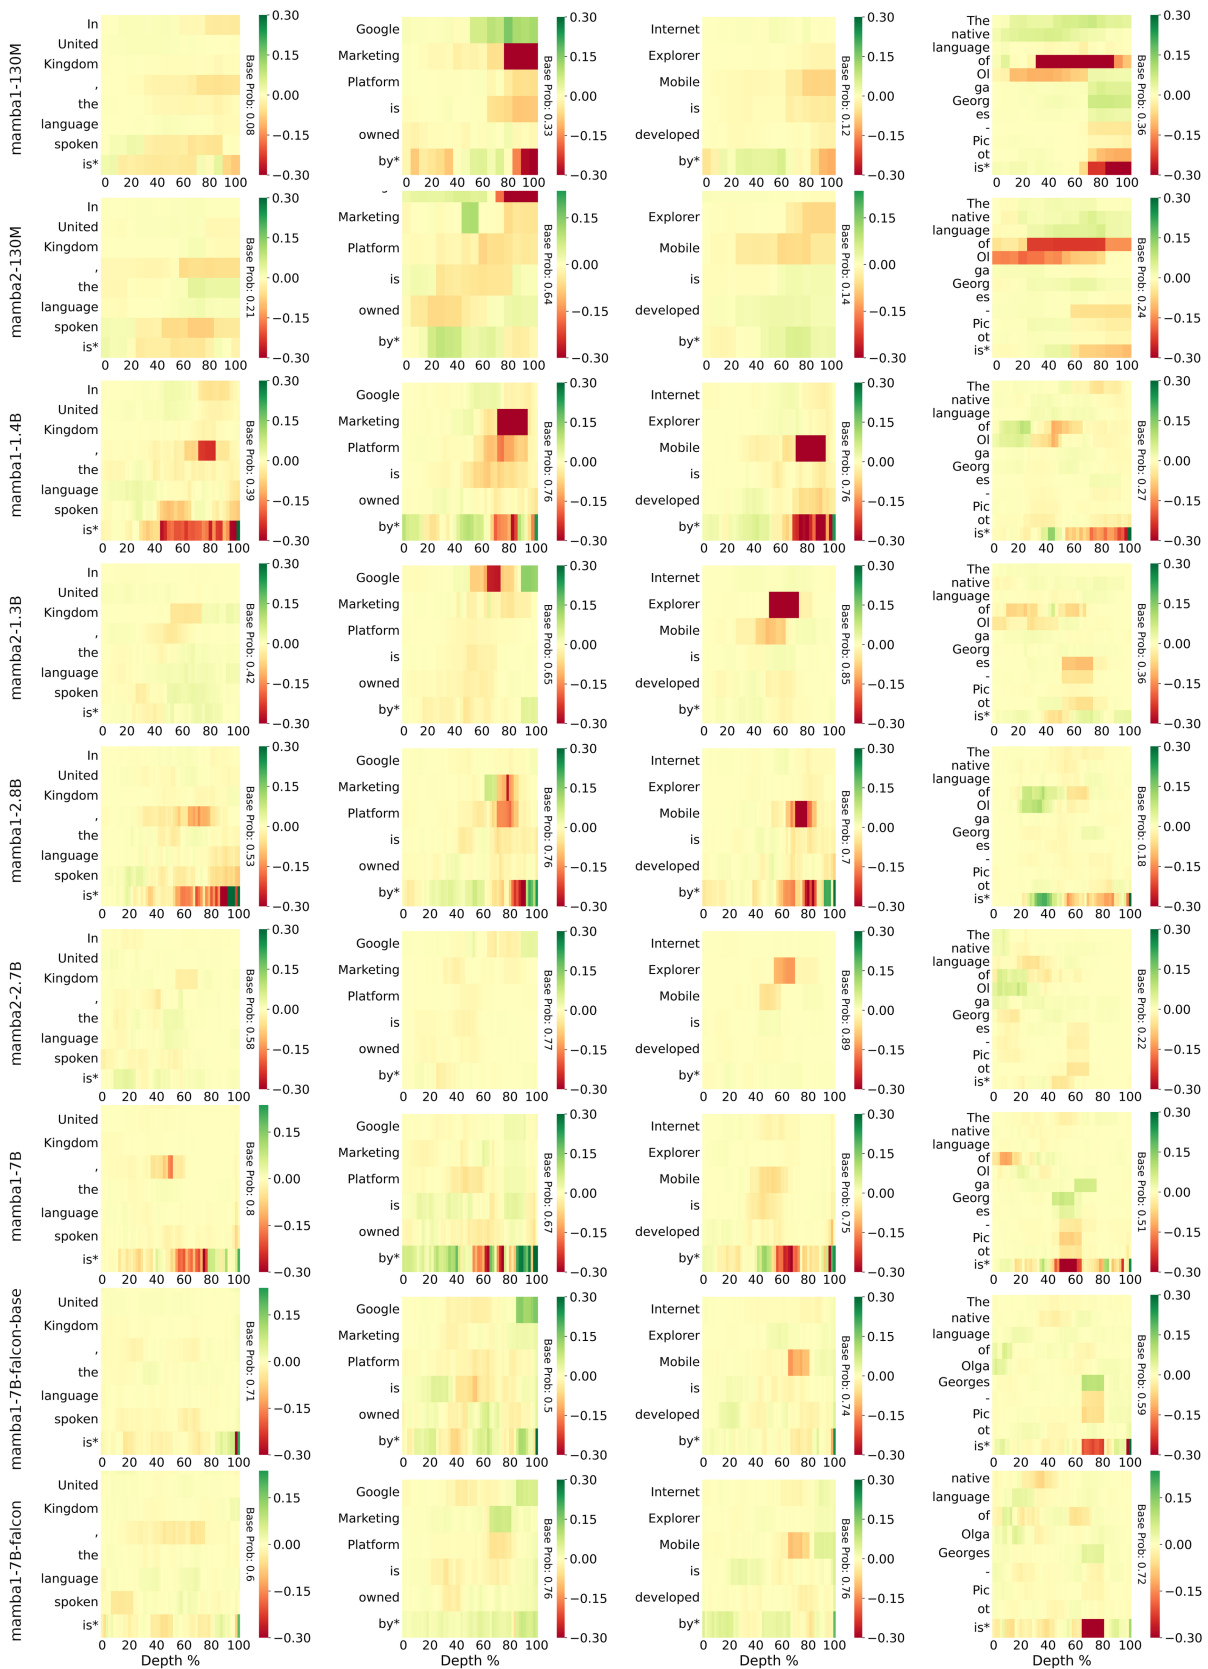


Figure 16: Relative change in correct-token probability after ablating the connection from each token to the final token, shown across different Mamba architectures with a fixed window size of 9. While Mamba-1 shows strong reliance on the final token and minimal dependence on early tokens, Mamba-2 distributes influence more evenly, suggesting improved integration of contextual information.

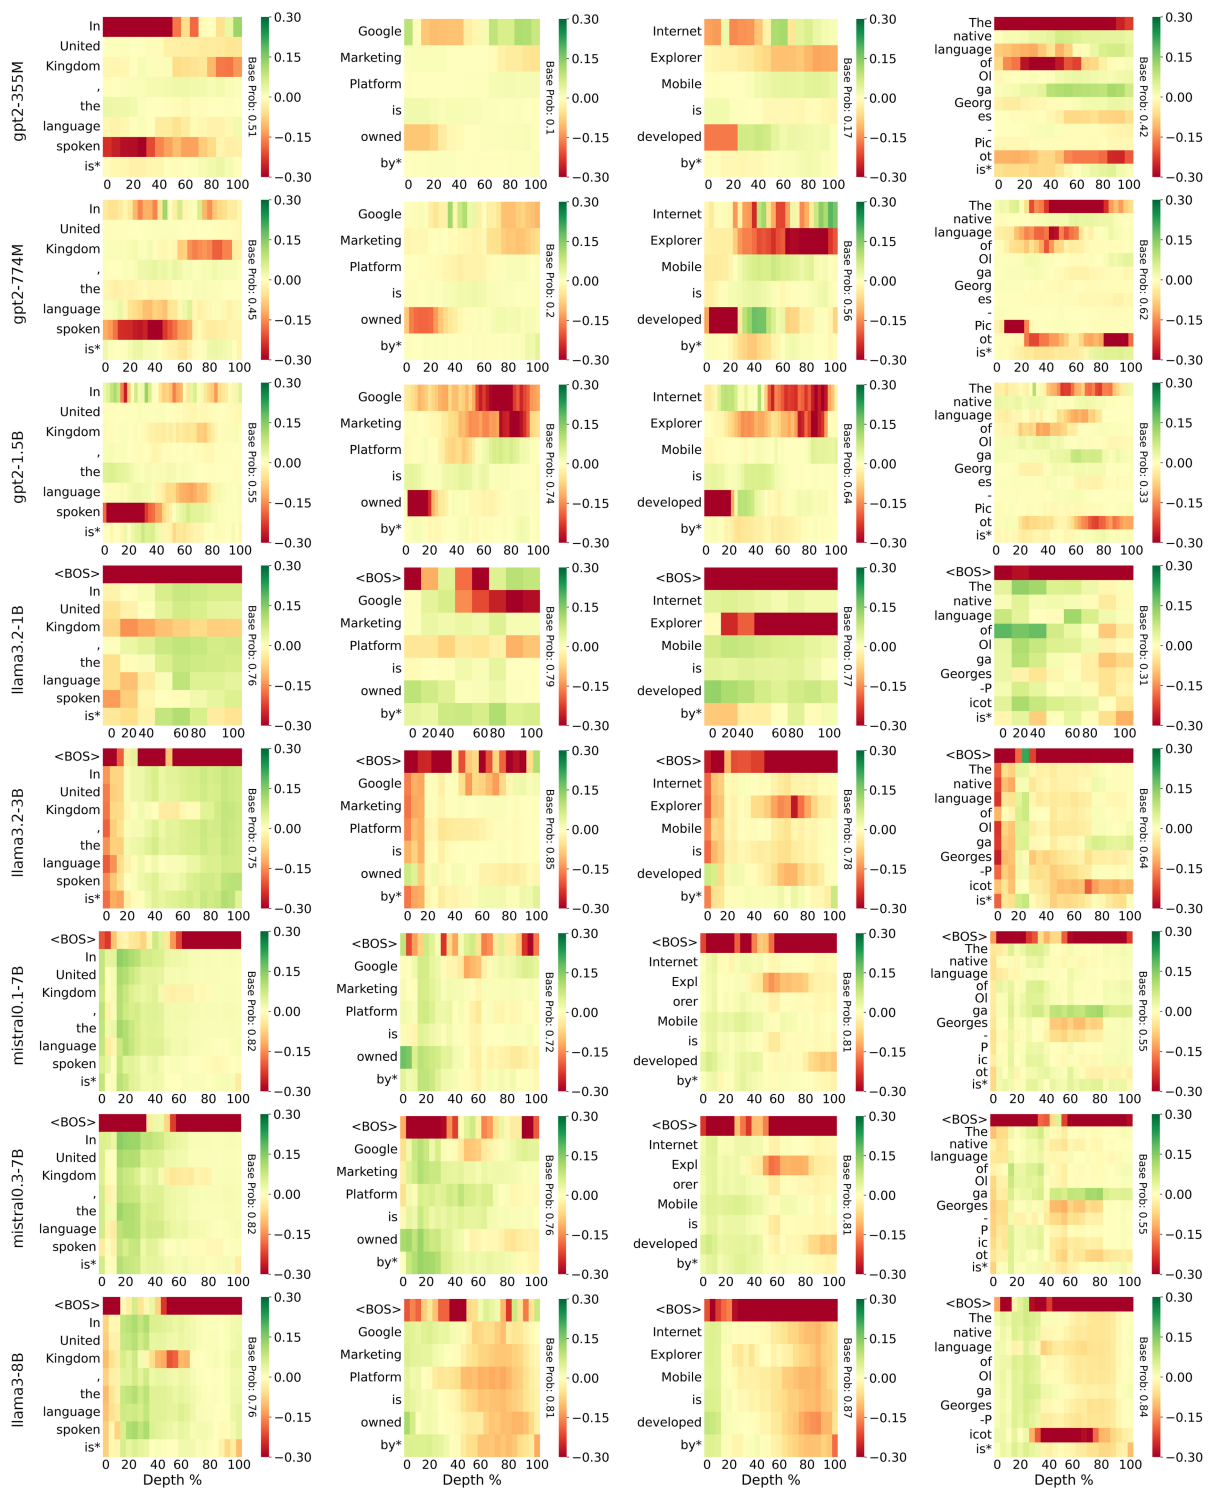


Figure 17: Same analysis as above (Figure 16), but applied to Transformer architectures at comparable parameter scales, using a fixed window size of 9.

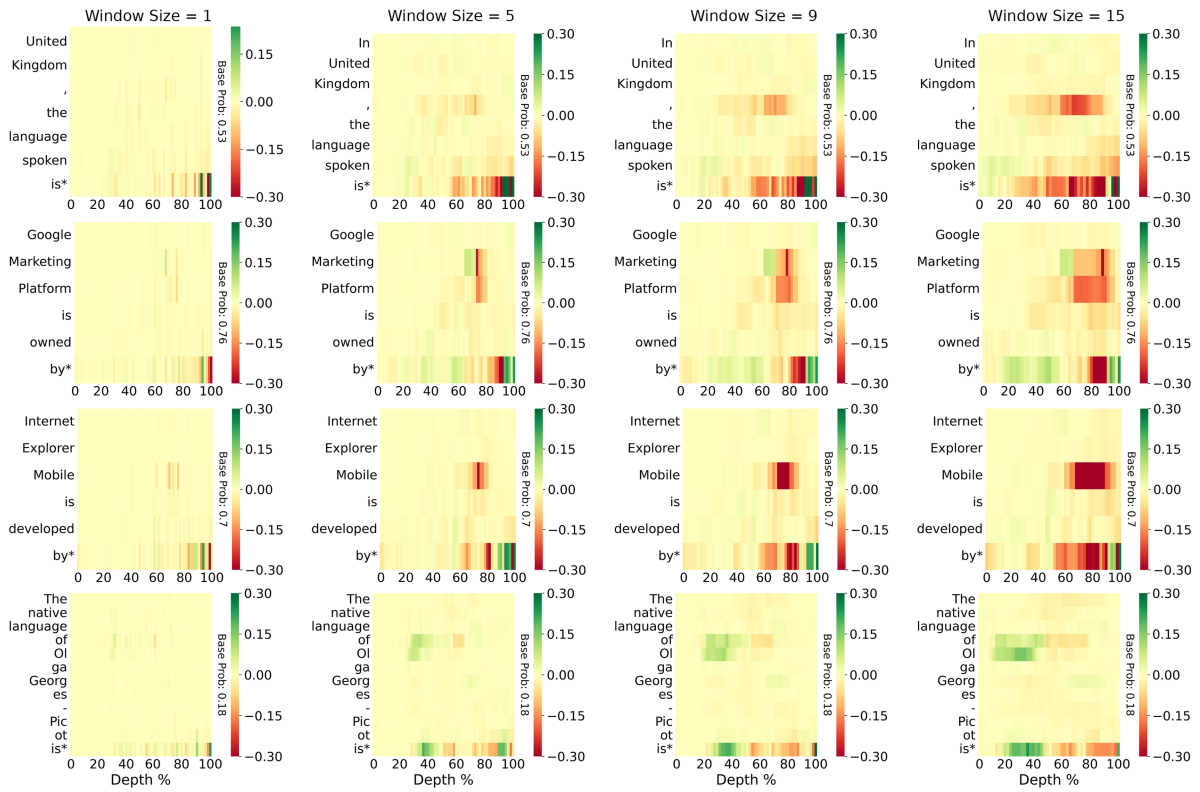


Figure 18: Relative change in correct-token probability after ablating the connection from each source token to the last token in Mamba-1 2.8 B. Columns correspond to ablation window sizes  $\in \{1, 5, 9, 15\}$ , and rows index the source token. Enlarging  $W$  amplifies the impact of mid-sequence tokens.

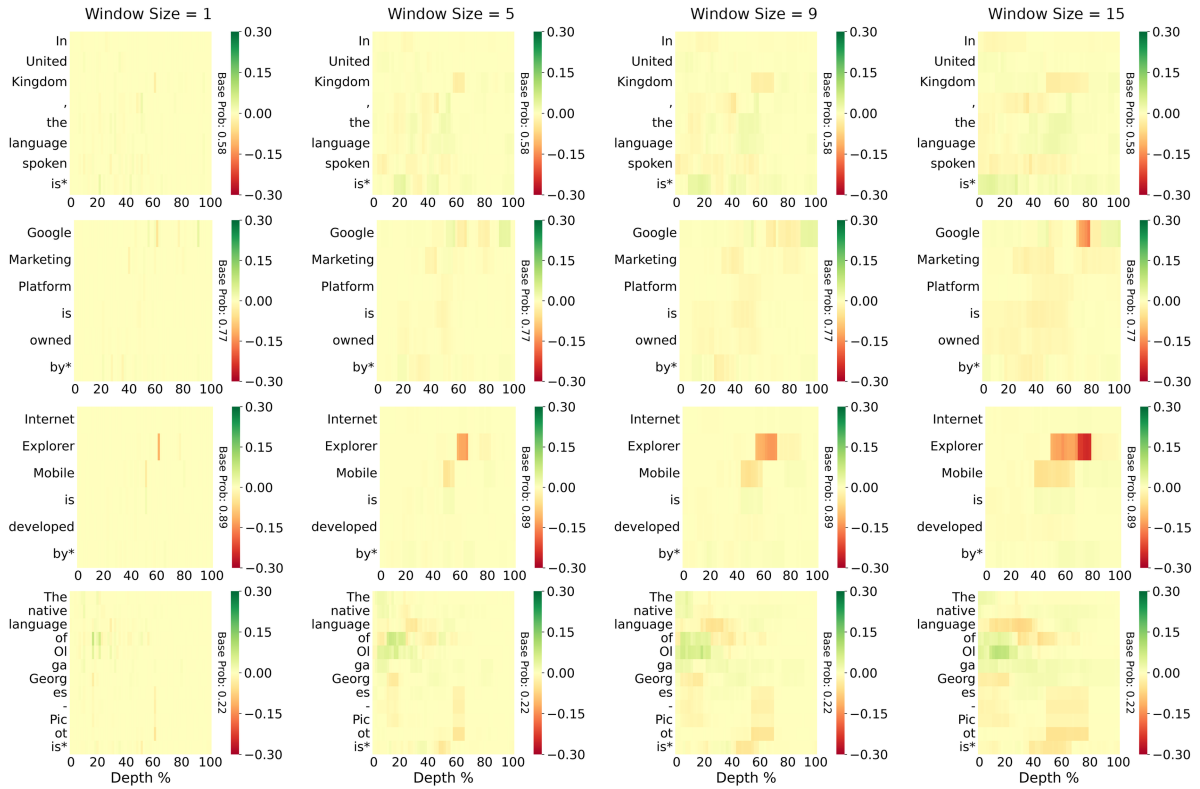


Figure 19: Identical analysis for Mamba-2 2.7 B. Columns again denote ablation window sizes  $\in \{1, 5, 9, 15\}$ . Relative influence is more evenly distributed than in Mamba-1, with larger  $W$  revealing a cumulative contribution from earlier context.

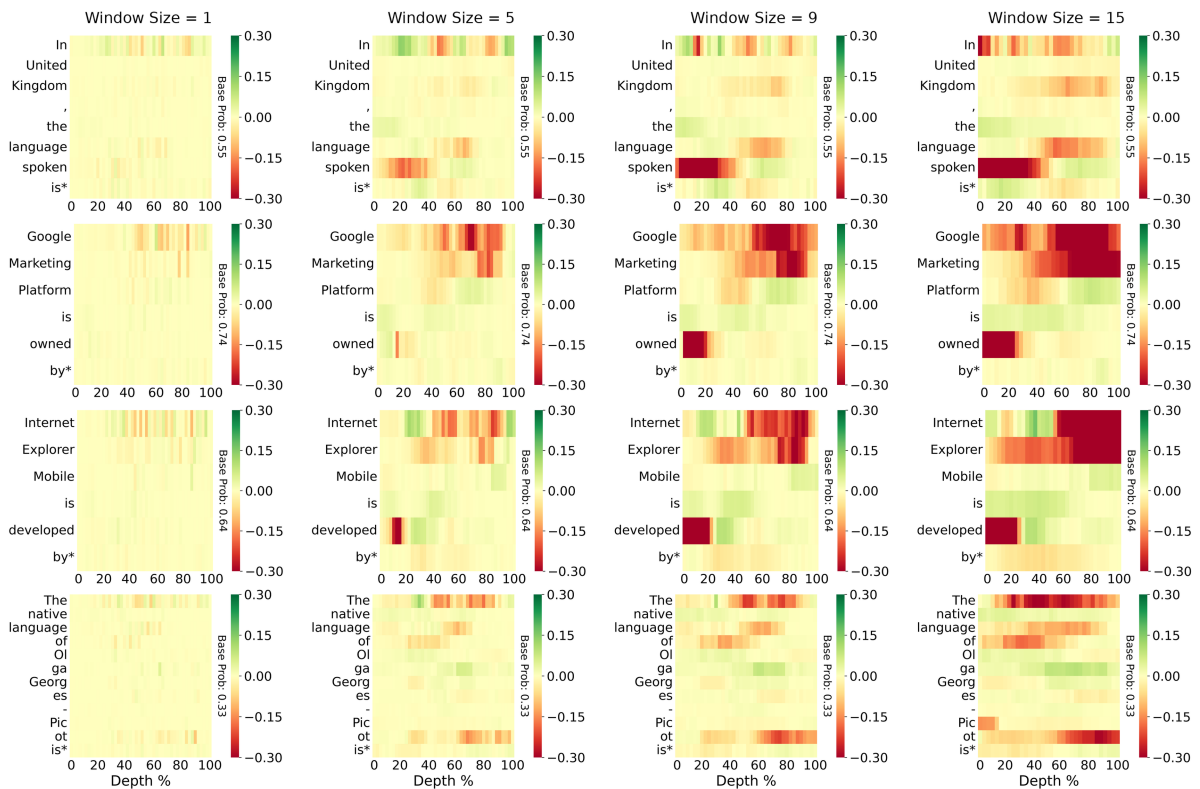


Figure 20: Same analysis applied to GPT-2 1.5 B. Columns show ablation window sizes  $\in \{1, 5, 9, 15\}$ .