

---

# Locate, Steer, and Improve: A Practical Survey of Actionable Mechanistic Interpretability in Large Language Models

---

[Hengyuan Zhang](#)<sup>\*,†</sup>, Zihao Zhang<sup>†</sup>, [Mingyang Wang](#), [Zunhai Su](#), [Yiwei Wang](#), [Qianli Wang](#), Shuzhou Yuan, Ercong Nie, Xufeng Duan, Qibo Xue, Zeping Yu, Chenming Shang, [Xiao Liang](#), Jing Xiong, Hui Shen, Chaofan Tao, [Zhengwu Liu](#), Senjie Jin, Zhiheng Xi, Dongdong Zhang, [Sophia Ananiadou](#), Tao Gui, Ruobing Xie, [Hayden Kwok-Hay So](#), Hinrich Schütze, Xuanjing Huang, Qi Zhang, Ngai Wong

Posted Date: 21 January 2026

doi: 10.20944/preprints202601.1663.v1

Keywords: actionable interpretability; large language models; localizing and steering; model improvement



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](#), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

# Locate, Steer, and Improve: A Practical Survey of Actionable Mechanistic Interpretability in Large Language Models

Hengyuan Zhang<sup>1,†</sup>, Zhihao Zhang<sup>2,†</sup>, Mingyang Wang<sup>3</sup>, Zunhai Su<sup>4</sup>, Yiwei Wang<sup>5</sup>, Qianli Wang<sup>6</sup>, Shuzhou Yuan<sup>7</sup>, Ercong Nie<sup>3</sup>, Xufeng Duan<sup>8</sup>, Qibo Xue<sup>9</sup>, Zeping Yu<sup>10</sup>, Chenming Shang<sup>11</sup>, Xiao Liang<sup>12</sup>, Jing Xiong<sup>1</sup>, Hui Shen<sup>13</sup>, Chaofan Tao<sup>1</sup>, Zhengwu Liu<sup>1</sup>, Senjie Jin<sup>2</sup>, Zhiheng Xi<sup>2</sup>, Dongdong Zhang<sup>14</sup>, Sophia Ananiadou<sup>10</sup>, Tao Gui<sup>2</sup>, Ruobing Xie<sup>15</sup>, Hayden Kwok-Hay So<sup>1</sup>, Hinrich Schütze<sup>3</sup>, Xuanjing Huang<sup>2</sup>, Qi Zhang<sup>2,\*</sup> and Ngai Wong<sup>1,\*</sup>

<sup>1</sup> The University of Hong Kong

<sup>2</sup> Fudan University

<sup>3</sup> LMU Munich

<sup>4</sup> Tsinghua University

<sup>5</sup> Technische Universität Darmstadt

<sup>6</sup> Technische Universität Berlin

<sup>7</sup> Technische Universität Dresden

<sup>8</sup> The Chinese University of Hong Kong

<sup>9</sup> Nanjing University

<sup>10</sup> University of Manchester

<sup>11</sup> Dartmouth College

<sup>12</sup> University of California, Los Angeles

<sup>13</sup> University of Michigan

<sup>14</sup> Microsoft

<sup>15</sup> Tencent

\* Correspondence: hengyuan.zhang88@gmail.com (H.Z.); zhangzhihao19@fudan.edu.cn (Z.Z.)

† These authors contributed equally to this work.

## Abstract

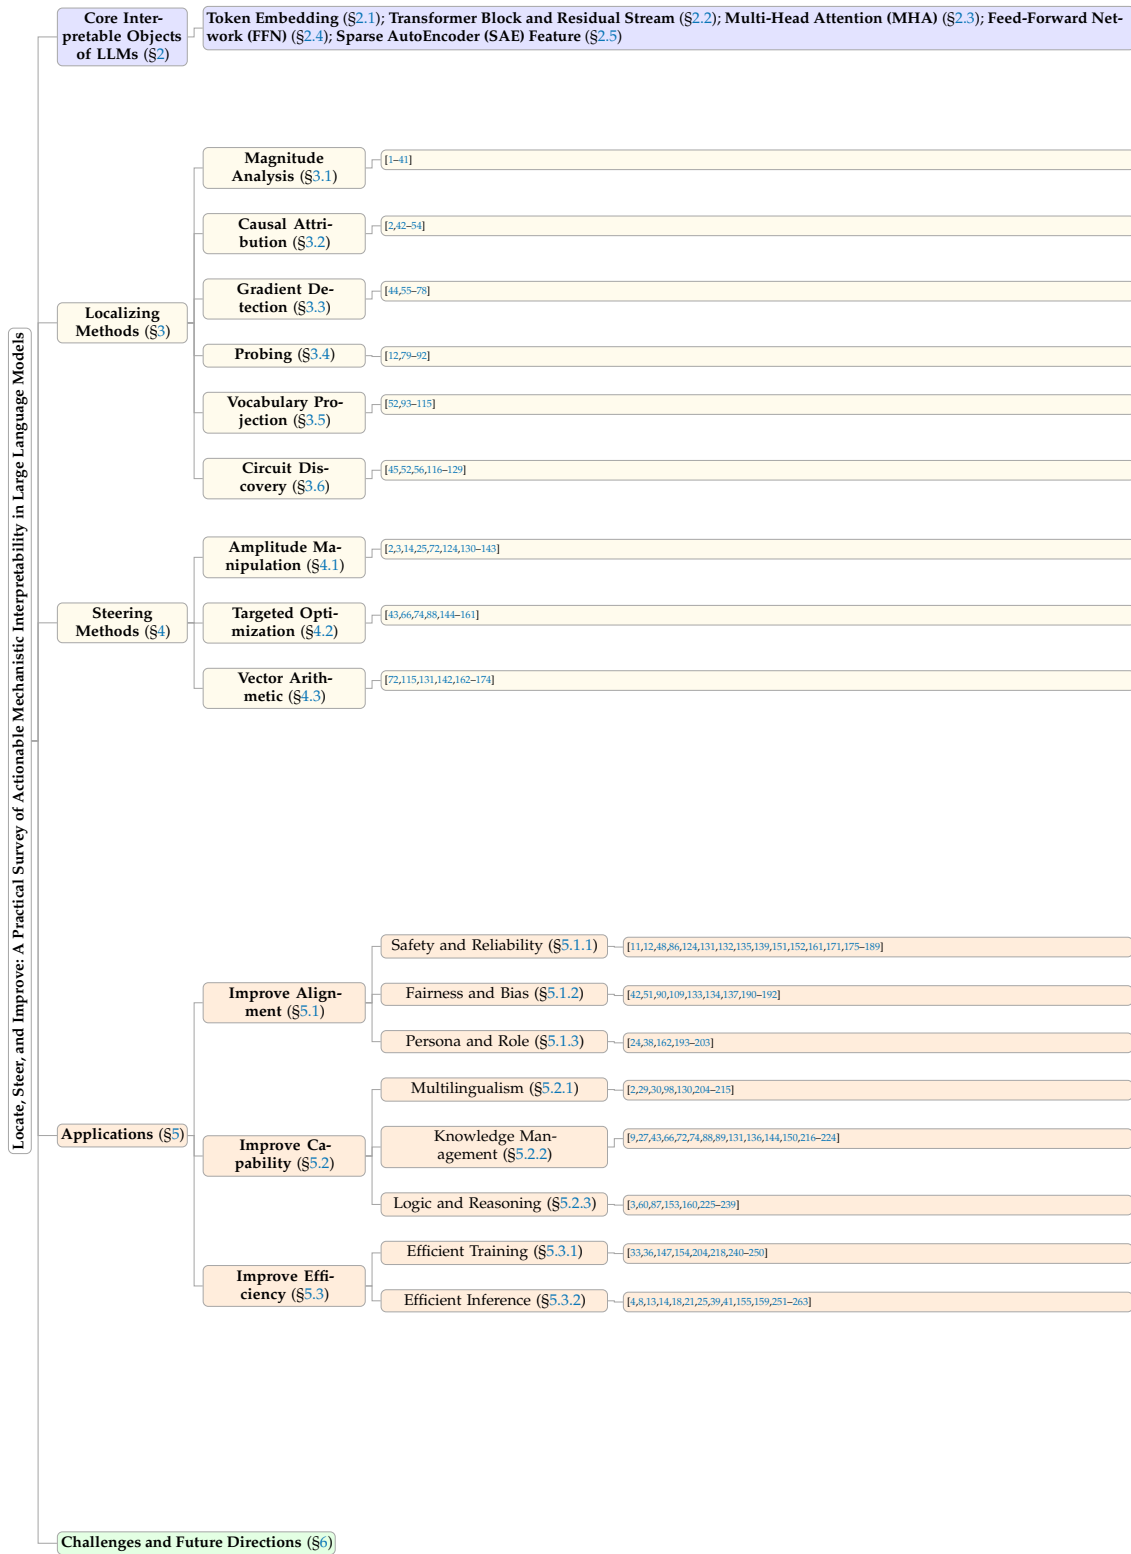
Mechanistic Interpretability (MI) has emerged as a vital approach to demystify the opaque decision-making of Large Language Models (LLMs). However, existing reviews primarily treat MI as an observational science, summarizing analytical insights while lacking a systematic framework for *actionable intervention*. To bridge this gap, we present a practical survey structured around the pipeline: “*Locate, Steer, and Improve*.” We formally categorize *Localizing* (diagnosis) and *Steering* (intervention) methods based on specific *Interpretable Objects* to establish a rigorous intervention protocol. Furthermore, we demonstrate how this framework enables tangible improvements in *Alignment, Capability, and Efficiency*, effectively operationalizing MI as an actionable methodology for model optimization. *With actionable mechanistic interpretability evolving at a fast pace, we pledge to keep this survey up to date, ensuring it reflects the cutting-edge advances in this area.*

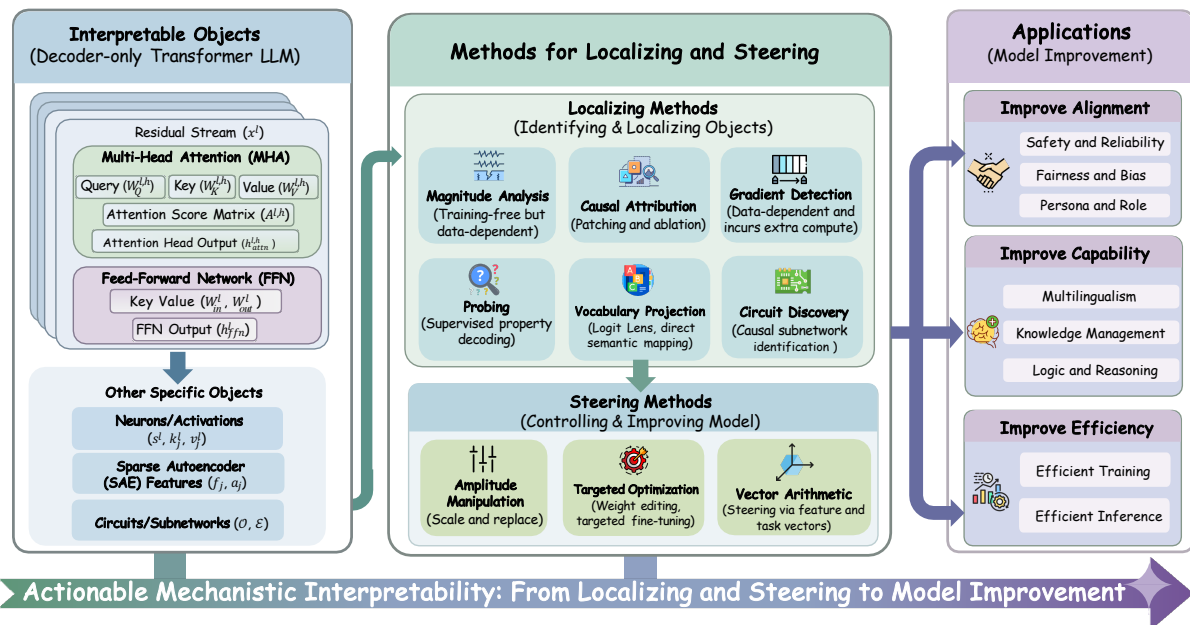
**Keywords:** actionable interpretability; large language models; localizing and steering; model improvement

## Contents

|  |           |
|--|-----------|
| <b>1. Introduction</b>                       | <b>4</b>  |
| <b>2. Core Interpretable Objects of LLMs</b> | <b>5</b>  |
| 2.1. Token Embedding                         | 5         |
| 2.2. Transformer Block and Residual Stream   | 6         |
| 2.3. Multi-Head Attention (MHA)              | 7         |
| 2.4. Feed-Forward Network (FFN)              | 8         |
| 2.5. Sparse Autoencoder (SAE) Feature        | 8         |
| <b>3. Localizing Methods</b>                 | <b>10</b> |
| 3.1. Magnitude Analysis                      | 10        |
| 3.2. Causal Attribution                      | 13        |
| 3.3. Gradient Detection                      | 14        |
| 3.4. Probing                                 | 17        |
| 3.5. Vocabulary Projection                   | 19        |
| 3.6. Circuit Discovery                       | 20        |
| <b>4. Steering Methods</b>                   | <b>23</b> |
| 4.1. Amplitude Manipulation                  | 23        |
| 4.2. Targeted Optimization                   | 25        |
| 4.3. Vector Arithmetic                       | 27        |
| <b>5. Applications</b>                       | <b>29</b> |
| 5.1. Improve Alignment                       | 29        |
| 5.1.1. Safety and Reliability                | 29        |
| 5.1.2. Fairness and Bias                     | 31        |
| 5.1.3. Persona and Role                      | 33        |
| 5.2. Improve Capability                      | 35        |
| 5.2.1. Multilingualism                       | 35        |
| 5.2.2. Knowledge Management                  | 37        |
| 5.2.3. Logic and Reasoning                   | 39        |
| 5.3. Improve Efficiency                      | 41        |
| 5.3.1. Efficient Training                    | 41        |
| 5.3.2. Efficient Inference                   | 43        |
| <b>6. Challenges and Future Directions</b>   | <b>44</b> |
| <b>7. Conclusions</b>                        | <b>46</b> |
| <b>A. Summary of Surveyed Papers</b>         | <b>47</b> |
| <b>B. References</b>                         | <b>52</b> |

## Paper Outline





**Figure 1.** Overview of the paper structure. We begin by defining the core interpretable objects (§2) that form the foundation of our analysis. We then introduce a range of methods, ranging from localization (§3) to steering (§4). Finally, we illustrate how these methods can be applied to improve models (§5).

## 1. Introduction

Large Language Models (LLMs) have recently achieved remarkable success, demonstrating outstanding performance across a wide spectrum of applications, ranging from complex reasoning and multilingualism, to highly specialized domains [264–275]. Despite these advancements, a critical challenge remains: the internal decision-making processes of these models are largely opaque, often operating as “black boxes.” This lack of transparency poses significant risks, particularly in safety-critical applications, and severely limits our ability to efficiently debug, control, and optimize model behaviors [276–278]. Consequently, *Mechanistic Interpretability* (MI) has emerged as a pivotal research direction. Unlike traditional behavioral analysis, MI aims to “reverse-engineer” these complex neural networks, decomposing their intricate computations into understandable components and causal mechanisms [47,279,280].

Current research in this field generally falls into two categories. A significant body of work focuses on the *theoretical and foundational aspects* of MI [47,279–286]. These studies provide technical roadmaps for dissecting Transformer architectures and identifying fundamental units. However, they primarily prioritize *scientific discovery*—aiming to elucidate the model’s inner working mechanisms for the sake of understanding itself. They typically treat MI as an observational science, leaving the question of how to translate these microscopic insights into practical model improvements underexplored.

Recognizing the applied potential of interpretability, a second line of work has begun to *bridge the gap between theoretical understanding and practical utilization*. These surveys discuss how MI techniques can be leveraged to aid downstream tasks or assist in specific domains [287–293]. However, despite their contributions, these existing reviews face two primary limitations that hinder broader adoption. First, they often lack a *sufficient categorization* and *clear definition* of MI methods within practical application contexts. The distinction between diagnostic tools and intervention techniques is frequently blurred.

Second, their coverage of applications is often incomprehensive, and the illustration of methods is typically too general. This high-level abstraction makes it difficult for researchers to translate theoretical mechanistic insights into actionable interventions for specific problems. Consequently, there is a distinct lack of a unified guide that systematically *categorizes these methods* and clearly presents a concrete *pipeline for active model improvement*.

To fill this gap, we propose the “*Locate, Steer, and Improve*” pipeline. This conceptual framework is designed to systematically transform MI from a passive observational science into an actionable intervention discipline. Our work makes the following key contributions:

- **1) A Rigorous Pipeline-Driven Framework:** We establish a structured framework for applying MI to real-world model optimization. We begin by defining the core **Interpretable Objects** within LLMs (e.g., neurons, attention heads, residual streams). Based on the application workflow, we clearly categorize methodologies into two distinct stages: **Localizing** (Diagnosis), which identifies the causal components responsible for specific behaviors, and **Steering** (Intervention), which actively manipulates these components to alter model outputs. Crucially, for each technique, we provide a detailed *Methodological Formulation* along with its *Applicable Objects* and *Scope*, helping readers quickly understand the technical implementation and appropriate use cases.
- **2) Comprehensive Paradigms for Application:** We provide an extensive survey of MI applications organized around three major themes: *Improve Alignment*, *Improve Capability*, and *Improve Efficiency*. These themes cover eight specific scenarios, ranging from safety and multilingualism to efficient training. Instead of merely listing relevant papers, we summarize representative *MI application paradigms* for each scenario. This approach allows readers to quickly capture the distinct usage patterns of MI techniques across different application contexts, facilitating the transfer of methods to new problems.
- **3) Insights, Resources, and Future Directions:** We critically discuss the current challenges in actionable MI research and outline promising future directions. To facilitate further progress and lower the barrier to entry, we curate a comprehensive collection of **over 200 papers**, which are listed in Table 2. These papers are systematically tagged according to their corresponding localizing and steering methods, providing a practical and navigable reference for the community.

## 2. Core Interpretable Objects of LLMs

In this section, we establish a unified mathematical formulation for the core interpretable objects within LLMs. We focus specifically on the decoder-only Transformer architecture [294], which serves as the predominant framework for contemporary state-of-the-art models [266,267,275]. We present the core interpretable objects and their corresponding mathematical notations in Table 1.

### 2.1. Token Embedding

The entry point of the model maps discrete tokens from a vocabulary  $\mathcal{V}$  to continuous vector representations. We define the *Embedding Matrix* as  $\mathbf{W}_E \in \mathbb{R}^{|\mathcal{V}| \times d_{\text{model}}}$ , where  $|\mathcal{V}|$  denotes the vocabulary size and  $d_{\text{model}}$  represents the hidden dimension of the model. For a given input token  $t_i$  at position  $i$ , its *Token Embedding*—which also serves as the initial state of the residual stream, denoted as  $\mathbf{x}_i^0$ —is obtained by retrieving the corresponding vector from  $\mathbf{W}_E$  and adding positional information:

$$\mathbf{x}_i^0 = \mathbf{W}_E[t_i] + \mathbf{p}_i \quad (1)$$

where  $\mathbf{p}_i$  is the positional embedding vector. It is worth noting that while earlier architectures used absolute positional embeddings added at the input, modern LLMs [266,267,275] typically employ Rotary Positional Embeddings (RoPE) [295]. In these architectures, positional information is applied directly to the query and key vectors within the attention mechanism rather than to the residual stream at the embedding layer.

## 2.2. Transformer Block and Residual Stream

Typically, an LLM is composed of  $L$  stacked layers. Each layer  $l$  consists of two primary blocks: a Multi-Head Attention (MHA) block and a Feed-Forward Network (FFN) block. The fundamental communication channel connecting these blocks is the *residual stream*.

As illustrated in Figure 2, the residual stream acts as the central “highway” for information propagation [43,66,116,144,296]. It preserves a shared memory state that is iteratively updated by the blocks. The update dynamics for the residual stream state  $\mathbf{x}^{l,1}$  at layer  $l$  are defined as follows:

$$\mathbf{x}^{l,\text{mid}} = \mathbf{x}^l + \mathbf{h}_{\text{attn}}^l(\mathbf{x}^l) \quad (2)$$

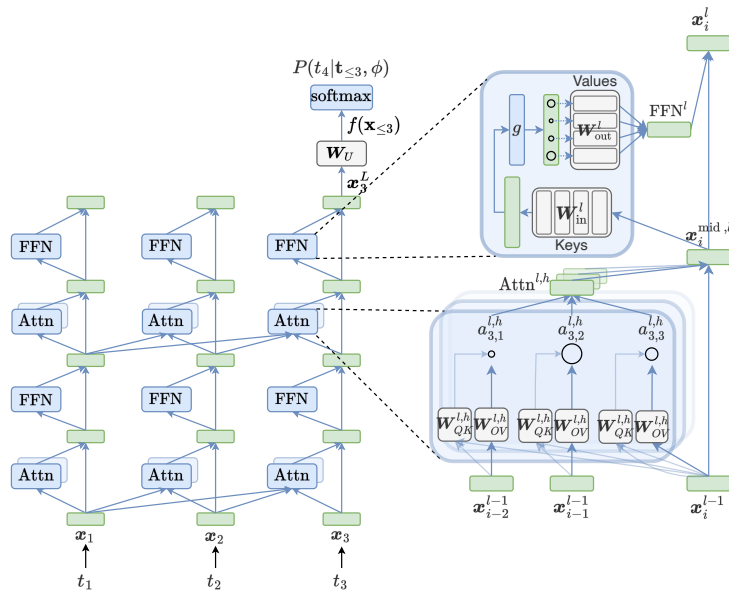
$$\mathbf{x}^{l+1} = \mathbf{x}^{l,\text{mid}} + \mathbf{h}_{\text{ffn}}^l(\mathbf{x}^{l,\text{mid}}) \quad (3)$$

where  $\mathbf{x}^{l,\text{mid}}$  represents the intermediate state after the MHA block but before the FFN block.<sup>2</sup>

This additive structure—where  $\mathbf{x}^{l+1} = \mathbf{x}^l + \text{MHA}(\mathbf{x}^l) + \text{FFN}(\mathbf{x}^l)$ —is critical for MI analysis. It implies that features in the residual stream can be viewed as linear combinations of outputs from all previous components. This property enables the decomposition of the model’s final prediction into individual component contributions, facilitating methods like “Logit Lens” [93,107,297] and causal mediation analysis [43,49,120,122,144].

<sup>1</sup>Here,  $\mathbf{x}^l \in \mathbb{R}^{T \times d_{\text{model}}}$  represents the token-wise residual stream state of the input sequence at layer  $l$ , with  $T$  tokens and hidden dimension  $d_{\text{model}}$ .

<sup>2</sup>For simplicity and clarity in our mechanistic analysis, we omit Layer Normalization (LayerNorm or RMSNorm) terms from these equations. While crucial for training stability, normalization operations are often abstracted away in high-level interpretability studies to focus on the additive composition of features.



**Figure 2.** The schematic of information flow within a standard Transformer block. The residual stream ( $x_i$ ) serves as the backbone, while MHA and FFN act as additive branches that read from and write to this stream. Based on the figure from [279].

### 2.3. Multi-Head Attention (MHA)

The Multi-Head Attention mechanism allows tokens to contextualize information by attending to other positions in the sequence. It consists of  $H$  independent heads, which primarily manage information routing and the resolution of contextual dependencies [116,117,298–300].

#### Standard Formulation

For a specific head  $h$  at layer  $l$ , we define the learnable weight matrices as  $\mathbf{W}_Q^{l,h}, \mathbf{W}_K^{l,h}, \mathbf{W}_V^{l,h} \in \mathbb{R}^{d_{\text{model}} \times d_{\text{head}}}$  and the output projection matrix as  $\mathbf{W}_O^{l,h} \in \mathbb{R}^{d_{\text{head}} \times d_{\text{model}}}$ . Here,  $T$  denotes the sequence length, the attention mechanism first computes the *attention score* matrix  $\mathbf{A}^{l,h} \in \mathbb{R}^{T \times T}$ , which represents the relevance of each token to every other token:

$$\mathbf{A}^{l,h} = \text{softmax} \left( \frac{(\mathbf{x}^l \mathbf{W}_Q^{l,h})(\mathbf{x}^l \mathbf{W}_K^{l,h})^\top}{\sqrt{d_{\text{head}}}} + \mathbf{M} \right), \quad (4)$$

where  $\mathbf{M} \in \mathbb{R}^{T \times T}$  denotes the attention mask that prevents attention to invalid positions (e.g., future tokens in causal attention or padding tokens).

Functionally, attention heads “read” information from the residual stream of previous tokens via the query–key subspace projections, and then “write” the attended information back to the current position via the value and output projections. The output for a single head  $h$ , denoted as  $\mathbf{h}_{\text{attn}}^{l,h}$ , is computed as:

$$\mathbf{h}_{\text{attn}}^{l,h} = \left[ \mathbf{A}^{l,h} (\mathbf{x}^l \mathbf{W}_V^{l,h}) \right] \mathbf{W}_O^{l,h}. \quad (5)$$

The total output of the MHA block is the sum of the outputs from all  $H$  heads:  $\mathbf{h}_{\text{attn}}^l = \sum_{h=1}^H \mathbf{h}_{\text{attn}}^{l,h}$ .

### Mechanistic View: QK and OV Units

While the standard formulation describes *how* attention is computed, the *unit* perspective [116] offers deeper insight into *what* task each head performs. As illustrated in the detailed view of Figure 2, each head can be decomposed into two functionally distinct units:

**1) The QK Unit ( $\mathbf{W}_{QK}$ ):** This unit determines *where* to attend. By merging the query and key matrices into a single low-rank matrix  $\mathbf{W}_{QK}^{l,h} = \mathbf{W}_Q^{l,h} (\mathbf{W}_K^{l,h})^\top$ , the attention pattern depends directly on the interaction between residual stream states. The attention score  $a_{ij}^{l,h}$  (e.g.,  $a_{3,1}$  in Figure 2) is derived from the bilinear form  $(\mathbf{x}_i^l)^\top \mathbf{W}_{QK}^{l,h} \mathbf{x}_j^l$ .

**2) The OV Unit ( $\mathbf{W}_{OV}$ ):** This unit determines *what* information is transmitted. By merging the value and output matrices into  $\mathbf{W}_{OV}^{l,h} = \mathbf{W}_V^{l,h} \mathbf{W}_O^{l,h}$ , we can view the head's operation as reading a vector from the source token  $j$ , transforming it linearly via  $\mathbf{W}_{OV}^{l,h}$ , and adding it to the destination token  $i$ , weighted by the attention score. This separation allows researchers to classify heads into distinct roles, such as "Induction Heads" (which copy previous tokens) or "Previous Token Heads" [33,34,117].

### 2.4. Feed-Forward Network (FFN)

#### Standard Formulation

The Feed-Forward Network block acts as a position-wise feature transformer. Unlike attention heads, FFNs operate independently on each token position, applying non-linear transformations to the input. They are often conceptualized as "Key-Value" memories, where the first layer projects the stream into a high-dimensional state (detecting patterns or "Knowledge Keys") and the second layer writes the retrieved knowledge back to the stream [64,94,301].

Mathematically, the output of the FFN block  $\mathbf{h}_{\text{ffn}}^l$  is given by :<sup>3</sup>

$$\mathbf{h}_{\text{ffn}}^l = \sigma(\mathbf{x}^{l,\text{mid}} \mathbf{W}_{\text{in}}^l) \mathbf{W}_{\text{out}}^l \quad (6)$$

where  $\mathbf{x}^{l,\text{mid}}$  is the input to the FFN, and  $\sigma$  is a non-linear activation function. The weight matrices are defined as  $\mathbf{W}_{\text{in}}^l \in \mathbb{R}^{d_{\text{model}} \times d_{\text{ffn}}}$  and  $\mathbf{W}_{\text{out}}^l \in \mathbb{R}^{d_{\text{ffn}} \times d_{\text{model}}}$ .

#### Mechanistic View: Neurons

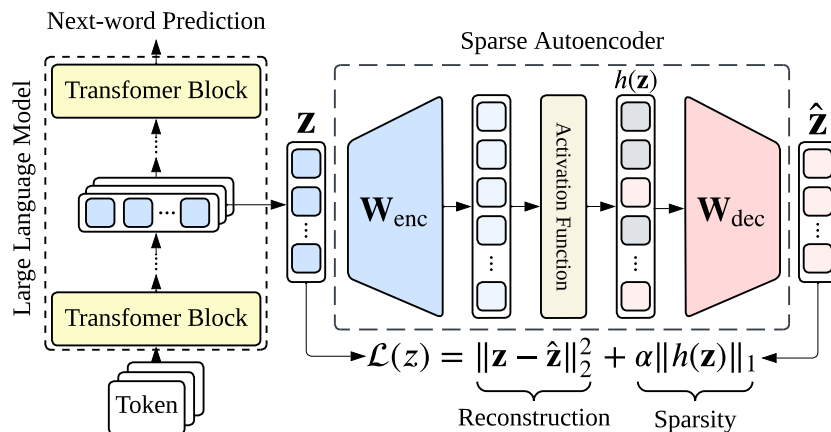
In this context, the *neuron*  $j$  is defined as an atomic unit comprised of a pair of weights: the *key weight*  $\mathbf{k}_j^l$  (the  $j$ -th row of  $\mathbf{W}_{\text{in}}^l$ ) and the *value weight*  $\mathbf{v}_j^l$  (the  $j$ -th column of  $\mathbf{W}_{\text{out}}^l$ ). The intermediate state  $\mathbf{s}^l = \sigma(\mathbf{x}^{l,\text{mid}} \mathbf{W}_{\text{in}}^l)$  represents the vector of *neuron activation*.

### 2.5. Sparse Autoencoder (SAE) Feature

While the internal objects described above (e.g., neuron activation  $\mathbf{s}^l$  or residual stream state  $\mathbf{x}^l$ ) are fundamental to the model's operation, they are often *polysemantic*. This is due to the phenomenon of *superposition*, where neural networks represent more features than they have physical neurons by encoding them as nearly orthogonal directions in the high-dimensional activation space [303]. Consequently, a single neuron may activate for multiple unrelated concepts, making direct interpretation difficult.

<sup>3</sup>While we use the standard formulation above to keep notation compact, it is important to note that many modern LLMs employ gated variants such as SwiGLU [302]. These variants introduce an additional gating matrix  $\mathbf{W}_{\text{gate}}^l$  and combine an element-wise gate with the projection before the final output:  $\mathbf{h}_{\text{ffn}}^l = (\text{SiLU}(\mathbf{x}^{l,\text{mid}} \mathbf{W}_{\text{gate}}^l) \odot (\mathbf{x}^{l,\text{mid}} \mathbf{W}_{\text{in}}^l)) \mathbf{W}_{\text{out}}^l$ . For the sake of generality, we present the standard FFN formulation here.

Sparse Autoencoders (SAEs) provide a principled method to resolve this by disentangling dense, polysemantic representations into *monosemantic features* [127]. As illustrated in Figure 3, an SAE acts as a “microscope” for the LLM. It projects low-dimensional dense activations into a higher-dimensional sparse latent space, effectively “unpacking” the superposition.



**Figure 3.** The framework of Sparse Autoencoders (SAEs). The SAE acts as an independent module attached to a frozen LLM, expanding dense representations into a sparse, overcomplete set of interpretable features via an encoder-decoder architecture. Based on the figure from [115].

### Mathematical Formulation

SAEs are trained in a layer-wise manner as independent modules attached to a specific object of a frozen LLM. They can be applied to nearly all internal objects, including neuron activation  $\mathbf{s}^l$ , residual stream state  $\mathbf{x}^l$ , MHA output  $\mathbf{h}_{\text{attn}}^l$ , and FFN output  $\mathbf{h}_{\text{ffn}}^l$  [304,305]. For instance, when applying an SAE to reconstruct a residual stream state  $\mathbf{x}^l$ , the forward pass is defined as:

$$\mathbf{a} = \sigma(\mathbf{x}^l \mathbf{W}_{\text{enc}} + \mathbf{b}_{\text{enc}}) \quad (7)$$

$$\hat{\mathbf{x}}^l = \mathbf{a} \mathbf{W}_{\text{dec}} + \mathbf{b}_{\text{dec}} \quad (8)$$

where  $\mathbf{W}_{\text{enc}} \in \mathbb{R}^{d_{\text{model}} \times d_{\text{SAE}}}$  and  $\mathbf{W}_{\text{dec}} \in \mathbb{R}^{d_{\text{SAE}} \times d_{\text{model}}}$  are learnable weights. A critical hyperparameter here is the *Expansion Factor*—the ratio of  $d_{\text{SAE}}$  to  $d_{\text{model}}$ . To capture the vast number of features hidden in superposition,  $d_{\text{SAE}}$  is typically set to be  $16\times$  to  $128\times$  larger than the model dimension [177,304–309].

The training objective is to minimize the reconstruction error while enforcing sparsity on the latent activations  $\mathbf{a}$ :

$$\mathcal{L} = \|\mathbf{x}^l - \hat{\mathbf{x}}^l\|_2^2 + \lambda \|\mathbf{a}\|_1 \quad (9)$$

In this framework, the *SAE feature*  $\mathbf{f}_j$  (the  $j$ -th row of  $\mathbf{W}_{\text{dec}}$ ) represents a distinct semantic direction in the activation space. The *SAE feature activation*  $a_j$  (the  $j$ -th element of  $\mathbf{a}$ ) quantifies the strength of this feature in the current input. Crucially, this decomposition transforms opaque vectors into an actionable vocabulary, allowing researchers to steer model behavior by targeting these granular, interpretable features [171,177,189,304,310,311].

### Training Challenges and Resources

Training high-quality SAEs presents unique challenges. One major issue is *Dead Latents*, where many feature neurons never activate during training, effectively wasting capacity. Techniques such as *ghost gradients* or periodic *resampling* are commonly employed to mitigate this [115,127]. Another

challenge is *Feature Absorption*, where broad, high-frequency features suppress specific, low-frequency ones. Advanced architectures like *Gated SAEs*, *Top-K SAEs*, and *BatchTopK SAEs* have been proposed to improve feature quality and reconstruction fidelity [312–316].

To facilitate research and reduce computational barriers, several high-quality pre-trained SAE suites have been released. Notable examples include *Gemma Scope* [304], *Llama Scope* [305], and “Golden Gate Claude” features [177]. These resources enable the community to focus on localizing and steering without incurring the cost of training SAEs from scratch.

**Table 1.** Core interpretable objects of LLM and their mathematical notations in this paper. Here, dimensions are denoted as follows:  $d_{\text{model}}$  is the model dimension,  $T$  is the sequence length,  $|\mathcal{V}|$  is the vocabulary size,  $H$  is the number of attention heads,  $d_{\text{head}}$  is the head dimension ( $d_{\text{model}}/H$ ),  $d_{\text{ffn}}$  is the FFN hidden dimension, and  $d_{\text{SAE}}$  is the dictionary size of the Sparse Autoencoder.

| Object                 |                                | Notation                         | Shape   |
|------------------------|--------------------------------|----------------------------------|---|
| <b>Token Embedding</b> | Embedding Matrix               | $\mathbf{W}_E$                   | $\mathbb{R}^{ \mathcal{V}  \times d_{\text{model}}}$  |
|                        | Token $i$ Embedding (Input)    | $\mathbf{x}_i^0$                 | $\mathbb{R}^{d_{\text{model}}}$   |
| <b>Residual Stream</b> | Residual Stream State          | $\mathbf{x}^l$                   | $\mathbb{R}^{T \times d_{\text{model}}}$  |
|                        | Intermediate State (Post-Attn) | $\mathbf{x}^{l,\text{mid}}$      | $\mathbb{R}^{T \times d_{\text{model}}}$  |
| <b>MHA</b>             | Q, K, V, O Weight Matrices     | $\mathbf{W}_{Q,K,V,O}^{l,h}$     | $\mathbb{R}^{d_{\text{model}} \times d_{\text{head}}} / \mathbb{R}^{d_{\text{head}} \times d_{\text{model}}}$ |
|                        | Attention Score Matrix         | $\mathbf{A}^{l,h}$               | $\mathbb{R}^{T \times T}$   |
|                        | Head Output                    | $\mathbf{h}_{\text{attn}}^{l,h}$ | $\mathbb{R}^{T \times d_{\text{head}}}$   |
|                        | Block Output                   | $\mathbf{h}_{\text{attn}}^l$     | $\mathbb{R}^{T \times d_{\text{model}}}$  |
| <b>FFN</b>             | In Projection (Key) Matrix     | $\mathbf{W}_{\text{in}}^l$       | $\mathbb{R}^{d_{\text{model}} \times d_{\text{ffn}}}$   |
|                        | Out Projection (Value) Matrix  | $\mathbf{W}_{\text{out}}^l$      | $\mathbb{R}^{d_{\text{ffn}} \times d_{\text{model}}}$   |
|                        | Block Output                   | $\mathbf{h}_{\text{ffn}}^l$      | $\mathbb{R}^{d_{\text{model}}}$   |
| <b>Neuron</b>          | Neuron Activation State        | $\mathbf{s}^l$                   | $\mathbb{R}^{d_{\text{ffn}}}$   |
|                        | $j$ -th Neuron Activation      | $s_j^l$                          | $\mathbb{R}$ (Scalar)   |
|                        | $j$ -th Neuron Key Weight      | $\mathbf{k}_j^l$                 | $\mathbb{R}^{d_{\text{model}}}$   |
|                        | $j$ -th Neuron Value Weight    | $\mathbf{v}_j^l$                 | $\mathbb{R}^{d_{\text{model}}}$   |
| <b>SAE Feature</b>     | Feature Activation State       | $\mathbf{a}$                     | $\mathbb{R}^{d_{\text{SAE}}}$   |
|                        | $j$ -th Feature Activation     | $a_j$                            | $\mathbb{R}$ (Scalar)   |
|                        | $j$ -th Feature                | $\mathbf{f}_j$                   | $\mathbb{R}^{d_{\text{model}}}$   |

### 3. Localizing Methods

Localizing Methods aim to identify interpretable objects that are responsible for a particular behavior or encode specific information. These techniques serve as a diagnostic step to narrow down the search space to manageable functional units. By pinpointing key components such as specific neurons, attention heads, or SAE features, they provide the necessary foundation for subsequent detailed mechanism analysis and targeted model steering.

#### 3.1. Magnitude Analysis

##### Methodological Formulation

*Magnitude Analysis* methods serve as a fundamental heuristic in interpretability, operating on the premise that internal elements with larger numerical values often exert greater influence on the model’s computation. It scores internal objects via a scalar function to identify salient components [1–3].

Formally, consider a set of internal objects  $\mathcal{O} = \{o_1, o_2, \dots, o_N\}$ , where each  $o_j$  represents a candidate element (e.g., a specific weight parameter row, a neuron, an SAE feature, or an attention head). We define an *Importance Score*  $s_j$  for each object using a magnitude function  $f(\cdot)$ :

$$s_j = f(o_j), \quad \text{e.g., } s_j = \|o_j\|_p \text{ or } s_j = \max_k |(o_j)_k| \quad (10)$$

Common choices for  $f(\cdot)$  include the  $L_2$ -norm ( $\|\cdot\|_2$ ) to measure the aggregate energy, the  $L_\infty$ -norm (max-value) to capture peak activation, or frequency-based metrics. Based on these scores, a subset of salient objects  $\mathcal{O}_{\text{salient}}$  is selected for further inspection or intervention, typically via a thresholding mechanism or a top- $k$  ranking:

$$\mathcal{O}_{\text{salient}} = \{o_j \mid s_j \geq \tau\} \quad \text{or} \quad \arg \text{topk } s_j_{j \in \{1, \dots, N\}} \quad (11)$$

### Applicable Objects

This method applies broadly to both static structure and dynamic computation. We categorize the applicable objects as follows:

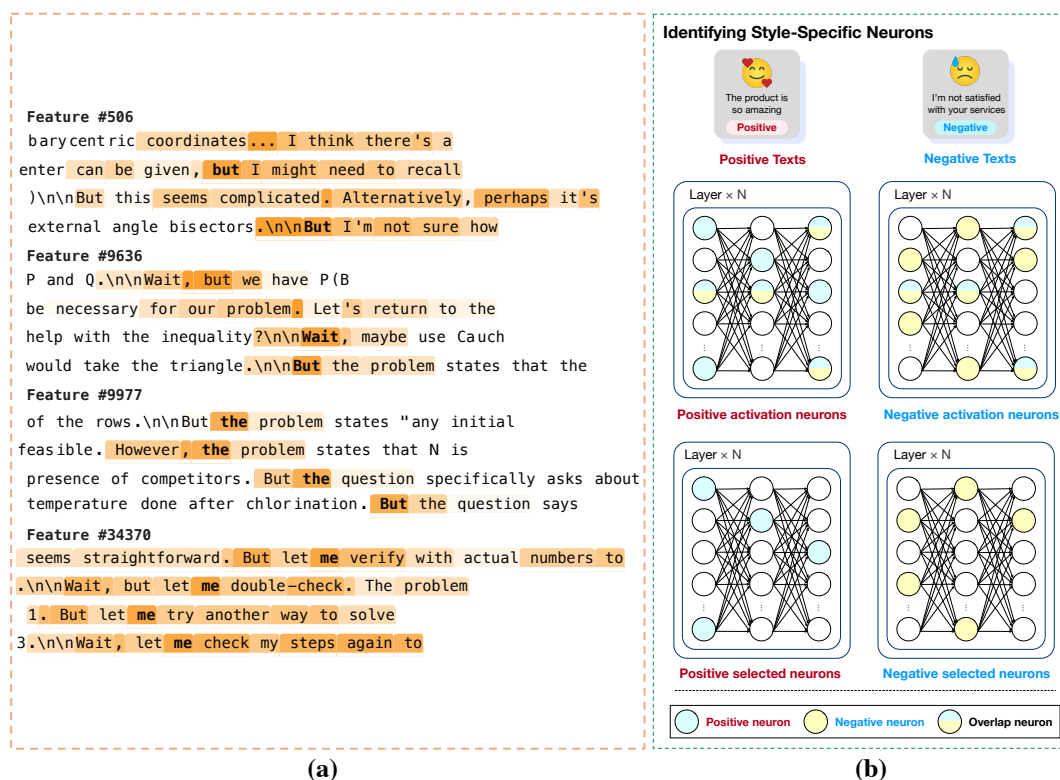
**1) Static Parameters:** In the context of model weights, *Magnitude Analysis* is often used to identify outliers or “heavy hitters” without running inference. Researchers typically compute per-weight or per-row norms of weight matrices (e.g.,  $\|\mathbf{W}_{\text{in}}[j, :]\|$ ) to highlight parameters that dominate the inner product computations. These high-magnitude weights are often associated with critical knowledge storage or outlier features [1,4–6,9,15–23].

**2) Dynamic Components (Neurons, SAE Features, or Attention Heads):** For functional units whose activity varies with input, ranking them by their activation statistics helps localize specialized capabilities [2,3,10,24–30].

- **Specialized Neurons and SAE Features:** By feeding domain-specific datasets into the model and monitoring activations (e.g., neuron activation state  $\mathbf{s}^l$  or SAE feature activation state  $\mathbf{a}$ ), researchers can isolate components dedicated to specific concepts. For instance, in the context of higher-level reasoning, Galichin et al. [3] utilized SAEs to disentangle the residual stream state  $\mathbf{x}^l$ . As shown in Figure 4 (a), they proposed a metric called *ReasonScore*, which aggregates the activation frequency and magnitude of SAE features  $a_j$  specifically during “reasoning moments” (e.g., when the model meets tokens like “Wait” or “Therefore”). By ranking features based on this score, they successfully localized *Reasoning-Relevant* SAE features that encode abstract concepts like uncertainty or exploratory thinking. Similarly, for style transfer, Lai et al. [24] employed *Magnitude Analysis* to identify *Style-Specific Neurons*. As illustrated in Figure 4 (b), they calculated the average activation magnitude of FFN neurons across corpora with distinct styles (e.g., positive vs. negative). Neurons that exhibited significantly higher average activation for the source style compared to the target style were identified as “Source-Style Neurons,” serving as candidates for subsequent deactivation.
- **Attention Heads:** The magnitude and distribution of *attention scores* ( $\mathbf{A}^{l,h}$ ) serve as a direct indicator of a head’s functional role [10,31–36]. For instance, Zhou et al. [35] introduced the *Safety Head Important Score (Ships)*, which aggregates attention weights on refusal-related tokens to localize “Safety Heads” critical for model alignment. In the multimodal domain, Sergeev and Kotelnikov [36] and Bi et al. [10] measured the concentration of attention mass on image tokens versus text tokens, successfully pinpointing heads responsible for visual perception and cross-modal processing.

Similarly, Singh et al. [33] measured “induction strength”—derived from the attention probability assigned to token repetition patterns—to track the formation and importance of Induction Heads.

**3) Layer-wise Representations:** Furthermore, measuring the magnitude of *layer-wise distances* reveals structural roles. Comparing representations across contrastive inputs (e.g.,  $\|x^l - x^{l'}\|$ ) localizes layers where task-specific information diverges most strongly [11,12,37,38], whereas comparing consecutive layers (e.g.,  $\|x^l - x^{l+1}\|$ ) identifies layers with minimal state updates, pointing to redundant computation [13,14,39–41].



**Figure 4. Localization via Magnitude Analysis.** (a) Discovery of SAE reasoning features [3]: SAE features are scored using *ReasonScore*, which aggregates activation magnitude and frequency during reasoning steps, isolating sparse features that encode cognitive behaviors like uncertainty or reflection. (b) Identification of Style-Specific Neurons [24]: Neurons are ranked by their average activation magnitude on style-specific corpora, revealing clusters that selectively activate for distinct linguistic styles.

### Characteristics and Scope

The scope of *Magnitude Analysis* for dynamic quantities is characterized as **training-free but data-dependent**.

- **Advantages:** It does not require training auxiliary classifiers or performing computationally expensive backward passes. This makes it highly scalable and suitable for analyzing large models in real-time.
- **Limitations:** It serves primarily as a *lightweight heuristic*. High activation magnitude implies high presence but does not guarantee causal necessity (e.g., a high-magnitude feature might be cancelled out by a subsequent layer). Furthermore, its success relies heavily on the quality of the input data; if the dataset fails to elicit the specific behavior, the relevant components will remain dormant.

Therefore, *Magnitude Analysis* is typically used as a “first-pass” screening tool to filter candidate objects for more rigorous verification methods.

### 3.2. Causal Attribution

#### Methodological Formulation

*Causal Attribution* methods constitute the gold standard for localization in MI. Unlike correlation-based analyses, these techniques identify which internal objects are *causally responsible* for a specific model behavior by systematically measuring the effect of controlled interventions [42–49].

Formally, let  $F(\cdot)$  denote a scalar model output of interest, such as the logit or probability of a target token. Let  $o$  be an internal object (e.g., a neuron activation  $s_j^l$  or a head output  $\mathbf{h}_{\text{attn}}^{l,h}$ ) defined in §2. To evaluate the *causal effect* of  $o$ , we compare the model’s output under a counterfactual intervention against the baseline state:

$$\text{do}(o \leftarrow \bar{o}) : \Delta F(o) = F(\text{do}(o \leftarrow \bar{o})) - F(o) \quad (12)$$

where  $F(o)$  represents the model’s behavior in the standard “clean” run, and  $\text{do}(o \leftarrow \bar{o})$  represents the intervention where the object  $o$  is forced to take on a modified value  $\bar{o}$  while all other causal factors are held constant (*ceteris paribus*). The intervention typically takes two forms: *Patching* (where  $\bar{o}$  is an activation computed from a counterfactual input) or *Ablation* (where  $\bar{o}$  is set to zero or a mean vector). A large magnitude  $|\Delta F(o)|$  indicates that the object  $o$  acts as a critical mediator or information node for the behavior encoded by  $F$ .

#### Applicable Objects

This analysis primarily targets **dynamic objects** involved in the inference process, including the residual stream state  $\mathbf{x}^l$ , the output of FFN  $\mathbf{h}_{\text{ffn}}^l$ , and the output of specific attention head  $\mathbf{h}_{\text{attn}}^{l,h}$ .

**1) Patching (Interchange Intervention):** This approach replaces an object computed from the original input with one computed from a *counterfactual input* to isolate specific information pathways. By systematically patching across layers and positions, one can localize exactly where task-specific information (e.g., factual knowledge) is introduced or transformed [43,44,49–51].

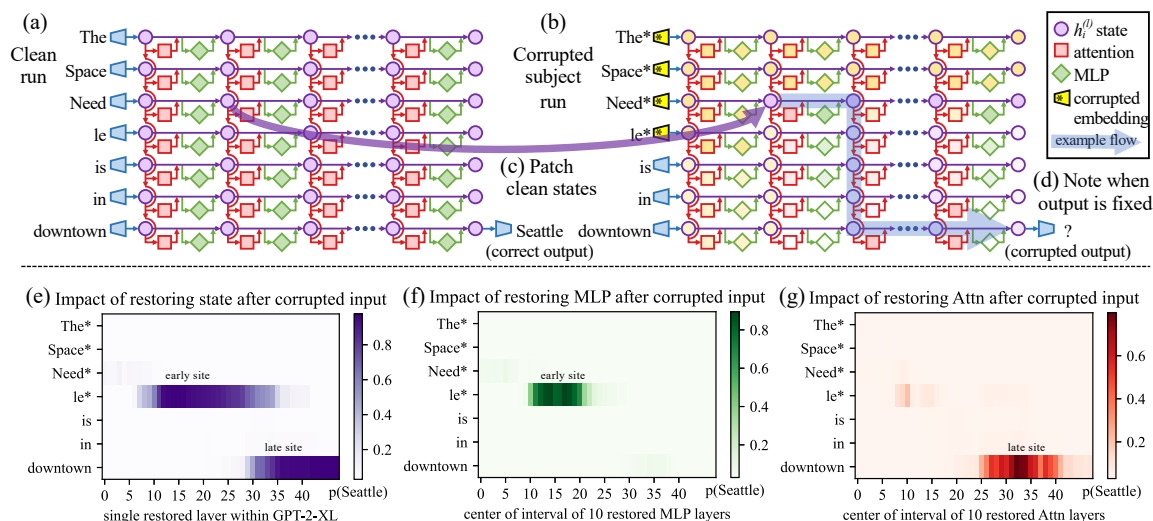
We exemplify this mechanism using *Causal Tracing* [43], a representative technique designed to localize factual associations (e.g., “The Space Needle” → “Seattle”). As illustrated in Figure 5, this process involves three key steps:

- **Corrupted Run (Intervention):** First, the specific knowledge is erased from the model’s computation. A *corrupted input* is created by adding Gaussian noise to the embeddings of the subject tokens (e.g., “Space Needle”), causing the probability of the correct prediction (“Seattle”) to drop significantly.
- **Patched Run (Restoration):** The core operation systematically restores specific internal states. For a specific layer  $l$  and token position  $i$ , the method copies the hidden activation from a separate original *clean run* and pastes (restores) it into the corrupted computation graph.
- **Effect Measurement:** The causal effect is quantified by the *Indirect Effect (IE)*, which measures how much of the original target probability is recovered by this restoration. A high IE score implies that the patched state at  $(l, i)$  carries critical information.

Through this rigorous process, Meng et al. [43] revealed that factual recall relies on two distinct localized mechanisms: an early retrieval phase in the *FFN blocks* at subject tokens, and a late information transport phase in the *MHA blocks* at the final token.

**2) Ablation (Knockout):** Alternatively, ablation-based attribution explicitly “zeros out” or removes objects, such as masking specific attention heads  $\mathbf{h}_{\text{attn}}^{l,h}$  or neurons, and measures the resulting perfor-

mance drop to determine their causal necessity [2,46,52–54]. This rigorous verification has been applied across various domains: Wang et al. [52] and Yu and Ananiadou [46] employed ablation to isolate minimal heads responsible for indirect object identification and in-context learning, respectively. In the context of specialized capabilities, Yu and Ananiadou [54] utilized pruning (permanent ablation) to identify heads critical for arithmetic reasoning, while Tang et al. [2] masked specific neurons to demonstrate the existence of language-specific functional regions. Furthermore, Geva et al. [53] applied blocking interventions to dissect the precise roles of FFN value vectors in factual recall mechanisms.



**Figure 5.** Overview of *Causal Tracing*. The method identifies critical internal states by creating a corrupted run (noising the subject “Space Needle”) and systematically restoring clean states to see which ones recover the prediction “Seattle”. The heatmap results reveal that factual information is processed in early MLP layers at the subject position and later transferred to the final token via attention. Based on the figure from Meng et al. [43].

### Characteristics and Scope

The scope of *Causal Attribution* is characterized as **rigorously causal but computationally intensive**.

- **Advantages:** Unlike *Magnitude Analysis* (§3.1), which only establish correlation, *Causal Attribution* provides definitive evidence that a component is a functional driver of the model’s output. This allows researchers to distinguish essential mechanisms from features that are highly activated but causally irrelevant to the specific behavior.
- **Limitations:** This rigor incurs a significant computational overhead. Verifying causality typically requires intervening on objects individually and performing a separate forward pass for each intervention. Consequently, the cost scales linearly with the number of objects analyzed, making it prohibitively expensive for dense, sweeping searches over large models. This inefficiency often necessitates the use of *Gradient Detection* (§3.3), which utilizes gradients to rapidly approximate these causal effects, enabling efficient screening before performing expensive, fine-grained interventions.

### 3.3. Gradient Detection

#### Methodological Formulation

*Gradient Detection* methods localize influential internal objects by scoring them with the sensitivity of a scalar target  $F(x)$  (e.g., a logit, margin, or loss) with respect to an object  $o_j$ :  $s_j(x) = \phi(\nabla_{o_j} F(x), o_j)$ , where common instantiations include the gradient norm  $s_j = \|\nabla_{o_j} F(x)\|$  and the gradient–input score

$s_j = \nabla_{o_j} F(x)^\top o_j$  [55,56]. These scores serve as fast, first-order proxies for intervention effects. Specifically, under an additive modification  $o_j \mapsto o_j + \Delta o_j$ , a first-order Taylor expansion yields

$$F(o_j + \Delta o_j) - F(o_j) = \nabla_{o_j} F(x)^\top \Delta o_j + \mathcal{O}(\|\Delta o_j\|^2), \quad (13)$$

indicating that the dot product  $\nabla_{o_j} F(x)^\top \Delta o_j$  represents the directional derivative of  $F$  along  $\Delta o_j$ . A common local “removal” surrogate sets  $\Delta o_j = -o_j$ , giving  $F(o_j - o_j) - F(o_j) \approx -\nabla_{o_j} F(x)^\top o_j$ , which motivates using  $\nabla_{o_j} F(x)^\top o_j$  (or its magnitude) as a signed influence score.

To mitigate saturation and explicitly model the notion of “absence,” *Integrated Gradients* (IG) attribute the change from a baseline  $\tilde{o}_j$  to the input  $o_j$  by integrating gradients along the straight-line path  $\gamma(\alpha) = \tilde{o}_j + \alpha(o_j - \tilde{o}_j)$ :

$$\text{IG}_k(o_j; \tilde{o}_j) = (o_j - \tilde{o}_j)_k \int_0^1 \frac{\partial F(\gamma(\alpha))}{\partial \gamma_k} d\alpha, \quad \sum_k \text{IG}_k(o_j; \tilde{o}_j) = F(o_j) - F(\tilde{o}_j), \quad (14)$$

where  $k$  indexes the components of  $o_j$ , and each  $\text{IG}_k$  quantifies the contribution of the  $k$ -th component to the output difference  $F(o_j) - F(\tilde{o}_j)$ . In practice, the integral is approximated by an  $m$ -step Riemann sum over  $\alpha = t/m$  [56]. Scores are typically computed over a dataset  $\mathcal{D}$  and aggregated to stabilize rankings (e.g.,  $\mathbb{E}_{x \sim \mathcal{D}}[s_j(x)]$  or  $\mathbb{E}[|s_j(x)|]$ ), without explicitly applying perturbations during scoring.

### Applicable Objects

Because  $F$  is differentiable with respect to any internal object  $o_j$  (Table 1), *Gradient Detection* applies uniformly across *inputs*, *activations*, and *parameters*. Below we expand the object categories and make explicit the correspondence between symbols and the underlying model components.

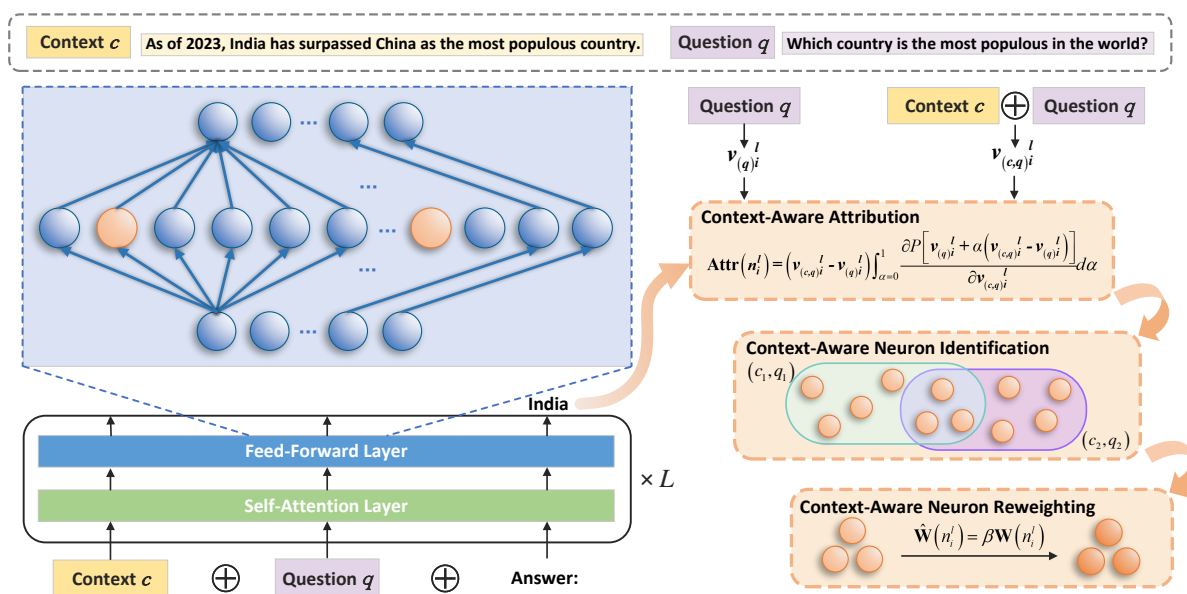
**1) Inputs and Layer-wise States ( $\mathbf{x}_i^0, \mathbf{x}^l$ ):** For input embeddings  $\mathbf{x}_i^0$  and the residual stream state  $\mathbf{x}^l$ , gradients directly quantify how sensitive  $F(x)$  is to changes in specific prompt components and their propagated representations. In practice, one computes  $\nabla_{\mathbf{x}_i^0} F(x)$  or  $\nabla_{\mathbf{x}^l} F(x)$  and derives token-level influence, such as the gradient norm  $\|\nabla_{\mathbf{x}_i^0} F(x)\|$ , the gradient–input score  $\nabla_{\mathbf{x}_i^0} F(x)^\top \mathbf{x}_i^0$ , or integrated gradients [59]. Aggregating these scores across positions  $i$  (optionally across layers  $l$ ) yields a ranked view of which tokens or contextual spans are most responsible for a target output, as used to analyze CoT prompting [60] and which depth regions contribute most strongly to the formation of that output [61], with closely related layer-/token-saliency signals also supporting dynamic token pruning [62] and inference-time steering [63].

**2) Intermediate Outputs:** Beyond inputs, *Gradient Detection* can score internal computational units whose activations vary with the input.

- **Neurons ( $\mathbf{s}^l$ ):** A standard neuron-level object is the FFN activation vector  $\mathbf{s}^l$  at layer  $l$ . Gradients  $\nabla_{\mathbf{s}^l} F(x)$  can be converted into per-neuron scores to rank neurons by their local influence on  $F$ . This has been used to localize knowledge- or context-sensitive neurons and analyze their dependencies [64–69]. Figure 6 illustrates a concrete LLM-specific instance: Shi et al. [65] computes *Integrated Gradients* scores to identify neurons most responsible for processing contextual cues under knowledge conflicts (via a context-aware attribution and a high-score intersection criterion), and then reweights the identified neurons to promote context-consistent generation.
- **Attention Head Outputs ( $\mathbf{h}_{\text{attn}}^{l,h}$ ):** *Gradient Detection* also applies to attention-related activations such as the attention head output  $\mathbf{h}_{\text{attn}}^{l,h}$ . Computing  $\nabla_{\mathbf{h}_{\text{attn}}^{l,h}} F(x)$  and scalarizing it with  $s_j(x) =$

$\phi(\nabla_{\mathbf{h}_{\text{attn}}^{l,h}} F(x), \mathbf{h}_{\text{attn}}^{l,h})$  yields head-level rankings that can highlight salient heads or attention submodules for further analysis and subsequent intervention [70–72].

**3) Parameters ( $\mathbf{W}_{Q/K/V/O}^{l,h}, \mathbf{W}_{\text{in/out}}^l$ ):** Because  $F$  is differentiable with respect to model weights, *Gradient Detection* can score parameters at multiple granularities. At the block level, common targets include attention projection matrices  $\mathbf{W}_{Q/K/V/O}^{l,h}$  and FFN matrices  $\mathbf{W}_{\text{in/out}}^l$ . Gradients such as  $\nabla_{\mathbf{W}_Q^{l,h}} F(x)$  can be turned into scalar salience measures (e.g.,  $\|\nabla_{\mathbf{W}} F(x)\|$ ) to rank influential attention/FFN modules [70–72]. At finer granularity, the same principle is used to select influential *individual weights* [65,73] or *structured blocks* [74,75].



**Figure 6. Neuron-level gradient-based localization for mitigating knowledge conflicts.** First calculates the *Integrated Gradients* score for each neuron to measure its contribution to processing the context. It then identifies context-aware neurons by taking the intersection of neurons with the highest scores. Subsequently, the identified neurons are reweighted to guide the model to be more aligned with the contextual knowledge, ensuring greater fidelity to the context. Based on the figure from Shi et al. [65].

### Characteristics and Scope

The scope of *Gradient Detection* is **data-dependent** and defined relative to the analyst's target  $F$ , so rankings can shift under alternative objectives (e.g.,  $-\log p(y^*|x)$  [73,74], logit margins  $\text{logit}_y - \text{logit}_{y^{\text{foil}}}$  [44,77], or contrastive/counterfactual gaps  $|\text{logit}_y(x) - \text{logit}_y(x^{cf})|$  [65,78]). It incurs **extra compute** from backpropagation, but remains substantially cheaper than exhaustive intervention search; consequently, it is commonly used as a scalable ranking/filtering stage that proposes candidate objects for more expensive causal validation [76].

- **Advantages:** *Gradient Detection* is applicable to a broad class of objects without requiring additional training. Compared with exhaustive interventions, it can produce rankings with a relatively small number of backward passes, making it practical as an initial localization step when the candidate set is large.
- **Limitations:** Gradients provide a *local* proxy, not causal necessity: salience can be offset by downstream computation, and finite interventions may depart from first-order effects in non-linear

regimes. For these reasons, gradient-ranked objects are typically paired with *Causal Attribution* (§3.2) to validate whether the identified objects are genuinely responsible for the target behavior.

### 3.4. Probing

#### Methodological Formulation

*Probing* methods interpret model signals by training an auxiliary predictor  $g_\psi$  (often linear) to decode a labeled property  $y$  from an internal vector  $\mathbf{z} \in \mathbb{R}^{d_{\text{model}}}$  (e.g., the residual stream state  $\mathbf{x}^l$  at layer  $l$ ); in sequence models with token-indexed states  $\mathbf{z}_t$ , one first defines a single probe input either *token-wise* (choosing  $\mathbf{z} = \mathbf{z}_{t^*}$  at a designated position such as the last token) or via *pooled* aggregation across positions (e.g., mean pooling), while the probe formulation itself is unchanged, e.g.,

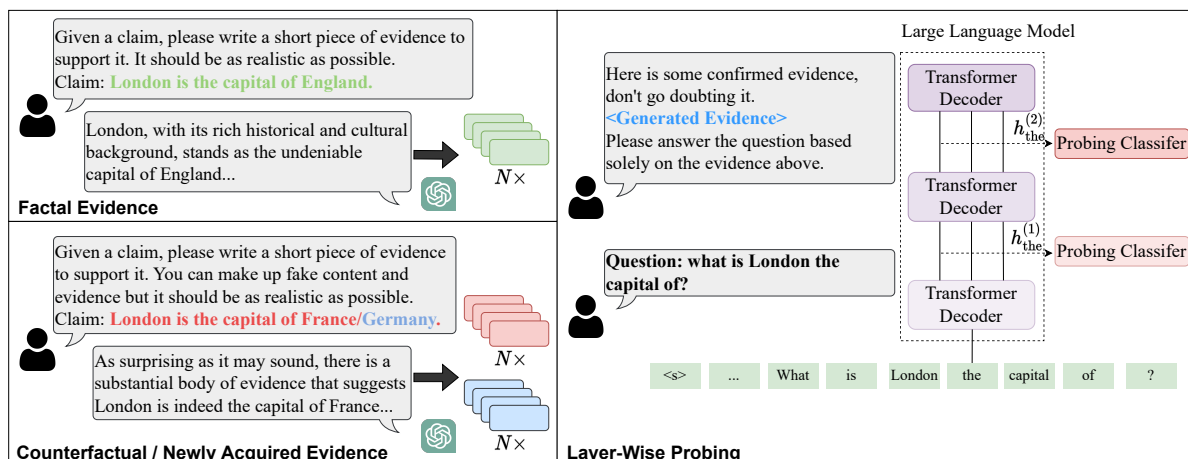
$$\hat{y} = g_\psi(\mathbf{z}) = \text{softmax}(\mathbf{W}_p \mathbf{z}), \quad (15)$$

using a supervised dataset  $\mathcal{D} = \{(\mathbf{z}, y)\}$  [79,80].

Operationally, probing treats the model as a frozen feature extractor and assesses decodability: whether  $y$  is recoverable from  $\mathbf{z}$  by a restricted hypothesis class (commonly linear), which supports localization by comparison across candidate objects (layers/heads/FFNs) via decoding performance or information-theoretic surrogates [81,82], typically followed by *Causal Attribution* (§3.2) to test functional necessity. Methodologically, it is standard to interpret probe results with care: high probe accuracy alone does not imply the model uses that information, motivating controls (e.g., selectivity / control tasks) and complementary causal tests [80,83,84].

#### Applicable Objects

*Probing* is defined on internal vectors, and is most naturally applied to any intermediate quantity that can be represented as a vector in  $\mathbb{R}^{d_{\text{model}}}$ . In LLMs, a typical workflow mirrors the pipeline in Figure 7: (i) constructs labeled probing evidence (including factual and counterfactual variants), (ii) runs the evidence through the frozen LLM and logs candidate internal objects across layers and submodules, and (iii) trains a fixed probe family on each object to compare decodability and localize where the target signal is most recoverable.



**Figure 7. Layer-wise probing pipeline for context knowledge.** An example end-to-end procedure: construct probing evidence for a target knowledge claim (including factual and counterfactual variants), run the evidence through the LLM under analysis, extract *residual stream state* across layers, and train *probing* classifiers to quantify where the target signal becomes most decodable. Based on the figure from Ju et al. [84].

**1) Residual Stream States ( $x^l$ ,  $x^{l,mid}$ ):** The most common probing target is the residual stream state  $x^l \in \mathbb{R}^{d_{model}}$ , as well as intermediate residual states  $x^{l,mid}$ . Layer-wise probes trained on  $x^l$  directly instantiate the “extract residual stream state across layers  $\rightarrow$  train probing classifiers” step depicted in Figure 7, and have been used to track where context knowledge, knowledge conflicts, and truthfulness-related signals become most decodable along depth [12,84–87].

**2) Block Outputs ( $h_{attn}^{l,h}$ ,  $h_{ffn}^l$ ):** Probing can target intermediate block outputs by extracting  $z$  from either an attention head output  $h_{attn}^{l,h}$  or the FFN output  $h_{ffn}^l$  (optionally token-wise, e.g.,  $h_{attn,t}^{l,h}$  or  $h_{ffn,t}^l$ ), and training a matched probe family across layers (and heads for attention). Comparing decodability across  $(l, h)$  and  $l$  supports fine-grained “localization by comparison,” ranking where a target property is most linearly accessible and contrasting attention- vs. FFN-based localization under a consistent protocol [88–90].

**3) SAE Feature Activation State (a):** Probing also integrates with SAE features. Given sparse SAE feature activation states  $\mathbf{a}$ , one can define  $\mathbf{z}$  as the feature activation vector  $\mathbf{a} = (a_1, \dots, a_m)$  (or a selected subset) and train classifiers on these sparse coordinates. This yields concept-aligned decoding axes that can be inspected at the feature level and cross-referenced with feature-level interpretations [91,92].

### Characteristics and Scope

*Probing* focuses on **supervised decoding**: it trains an auxiliary predictor  $g_\psi$  on  $\mathcal{D} = \{(z, y)\}$  to measure how well a labeled property  $y$  is predictable from an internal vector  $z$ . Treating LLM as a frozen feature extractor, probing evaluates **decodability** under a restricted hypothesis class, making it primarily a tool for representational localization rather than causal responsibility. In practice, probe-based rankings are commonly used to shortlist candidate layers/heads/FFNs for subsequent intervention-based analyses (e.g., *Causal Attribution* in §3.2).

- **Advantages:** With a fixed probe family, Probing enables standardized comparisons across objects, supporting efficient layer-wise tracking and large-scale ranking of candidate modules. Simple probes (e.g., linear) are lightweight and interpretable, allowing broad sweeps while keeping the LLM frozen.

- **Limitations:** Decodability is not causality: high probe accuracy does not imply the model uses  $y$ , nor that the probed object is necessary or sufficient. Results are sensitive to dataset and design choices (e.g., labeling, token positions), so controls and follow-up causal tests are typically required for functional claims.

### 3.5. Vocabulary Projection

#### Methodological Formulation

The most prominent technique in this category is the *Logit Lens* [93]. It operates on the premise that the pre-trained unembedding matrix  $\mathbf{W}_U \in \mathbb{R}^{d_{\text{model}} \times |\mathcal{V}|}$ , which maps the final layer’s hidden state to vocabulary logits, can serve as a universal decoder for intermediate states throughout the model. Formally, let  $\mathbf{z} \in \mathbb{R}^{d_{\text{model}}}$  denote a generic internal object (e.g., the residual stream state  $\mathbf{x}^l$  or an attention head output  $\mathbf{h}_{\text{attn}}^{l,h}$ ). *Vocab Projection* computes a distribution  $\mathbf{p}$  over the vocabulary  $\mathcal{V}$  by projecting  $\mathbf{z}$  through the unembedding matrix:

$$\mathbf{p} = \text{softmax}(\mathbf{z}\mathbf{W}_U) \quad (16)$$

By inspecting the tokens with the highest probabilities in  $\mathbf{p}$ , researchers can directly interpret the semantic content encoded in  $\mathbf{z}$  in terms of the model’s output vocabulary.

#### Applicable Objects

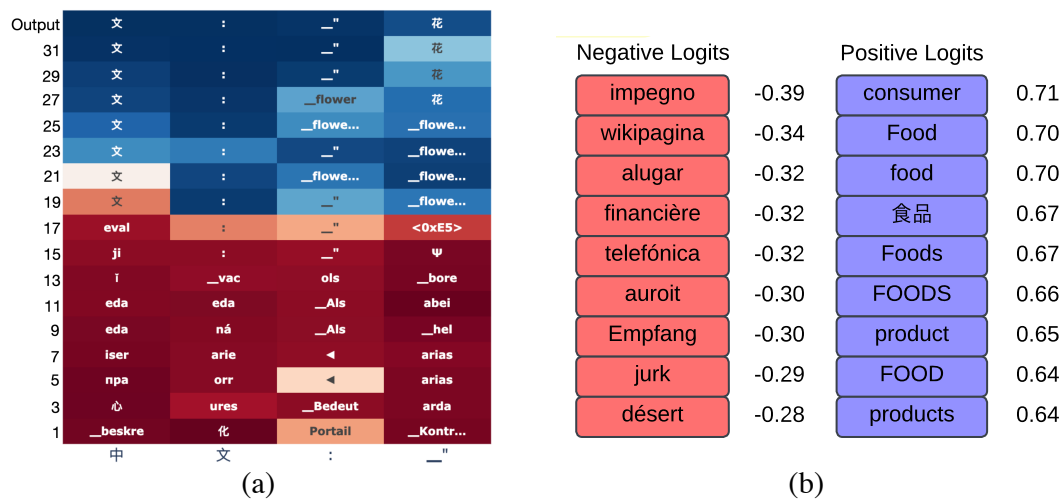
*Vocab Projection* is a versatile tool that applies to various objects defined in §2, ranging from global residual streams to specific attention heads, neurons, and SAE features.

**1) Residual Stream State ( $\mathbf{x}^l$ ):** Projecting the residual stream state  $\mathbf{x}^l$  allows researchers to trace the layer-wise evolution of predictions and identify the *crucial layers* where specific concepts emerge [95–102]. For instance, Wendler et al. [98] applied this to multilingual models, revealing distinct processing phases as shown in Figure 8 (a): initial layers focus on the surface form of the input language; middle layers process semantics in an abstract, “English-centric” concept space; and final layers rotate back to the target language. This confirms that English serves as an internal pivot for reasoning even in non-English tasks.

**2) Attention Head Output ( $\mathbf{h}_{\text{attn}}^{l,h}$ ):** Applying projection to the output of individual heads reveals the specific information (e.g., copied names or next-token candidates) that a head transmits to the residual stream. This has been instrumental in identifying functional heads in mechanistic studies [52,103–107]. For example, in reverse-engineering the Indirect Object Identification (IOI) task, Wang et al. [52] identified “Name Mover Heads” (which explicitly project to the correct name, e.g., “Mary”) and “Negative Name Mover Heads” (which suppress the correct name).

**3) Neuron Value Weight ( $\mathbf{v}_j^l$ ):** Geva et al. [94] demonstrated that FFNs operate as key-value memories. By projecting the value weight vector  $\mathbf{v}_j^l$  (a column of  $\mathbf{W}_{\text{out}}^l$ ) into the vocabulary, one can see which tokens are promoted by a specific neuron [94,108–110]. Individual neurons often boost semantically related clusters (e.g., “press”, “news”, “media”), suggesting that FFNs refine predictions by composing these pre-learned semantic distributions.

**4) SAE Feature ( $\mathbf{f}_j$ ):** For SAEs, output-based explanations leverage the decoder weights to interpret monosemantic features. By computing the logits contribution  $\mathbf{l}_j = \mathbf{f}_j\mathbf{W}_U$  for a feature vector  $\mathbf{f}_j$ , one can identify top-ranked tokens [111–115]. As shown in Figure 8 (b), a feature whose projection yields high positive logits for tokens like “Food” and “food” is interpreted as encoding a “food” concept, directly grounding the sparse feature in human-understandable semantics.



**Figure 8.** (a) Projecting residual stream states reveals the layer-wise evolution of latent concepts, showing an English-centric bottleneck in multilingual settings [98]. (b) Projecting SAE decoder weights identifies the semantic meaning of sparse features (e.g., a “food” feature) by identifying top-ranked tokens [115]. Based on figures from [98] and [115].

### Characteristics and Scope

The scope of *Vocab Projection* is characterized by **direct semantic mapping**. It offers an intrinsic view of internal representations without requiring auxiliary training.

- **Advantages:** It provides a zero-shot interpretation method that is computationally efficient and intuitive. Unlike *Probing* (§3.4), it does not require collecting a labeled dataset or training a separate classifier, allowing for immediate inspection of any model state.
- **Limitations:** The primary limitation is the assumption that intermediate states exist in the same vector space as the output vocabulary (basis alignment). While this often holds for the residual stream due to the residual connection structure, it may be less accurate for components inside sub-layers (like FFN and MHA) or in models where the representation space rotates significantly across layers. Consequently, results should be interpreted as an approximation of the information that is linearly decodable by the final layer.

### 3.6. Circuit Discovery

#### Methodological Formulation

*Circuit Discovery* methods aim to uncover *mechanistic pathways*: structured, directed dependencies among internal objects that mediate computation for a target behavior [116–119]. Formally, let  $(\mathcal{O}, \mathcal{E})$  be the model’s computational graph over internal objects  $\mathcal{O}$  and directed edges  $\mathcal{E}$ , where an edge  $e_{ij} \in \mathcal{E}$  denotes signal flow from object  $o_i$  to  $o_j$ . A circuit  $\mathcal{C} \subseteq \mathcal{E}$  is *faithful* if restricting computation to  $\mathcal{C}$  (e.g., by patching/ablating all other edges) preserves the target output  $F(x)$  or task performance.

Under the residual–rewrite view, heads and MLPs read from and write to the residual stream, inducing a directed graph whose edges represent additive residual updates. *Circuit Discovery* can be cast as edge-level causal subgraph selection: edges are retained if intervening on the corresponding information flow degrades a target metric  $\mathcal{R}$  [120]. *Automatic Circuit Discovery (ACDC)* instantiates this by iteratively testing and pruning edges via patching-based interventions, avoiding brute-force  $O(|\mathcal{E}|)$  enumeration while recovering circuits such as GPT-2’s greater-than mechanism [118,121].

Attribution-based methods such as *Edge Attribution Patching (EAP)* approximate patching with a first-order expansion, producing an edge score from two forward passes (*clean/corrupted*) and one backward pass [122,123]. Here, *clean* input  $\mathbf{x}_{clean}$  elicits the target behavior, while *corrupted* input  $\mathbf{x}_{corr}$  is a minimally modified version designed to break it (e.g., by perturbing relevant evidence or adding a counterfactual distractor), so the difference isolates the causal signal. For a sender object  $u$ , let  $\mathbf{a}_u(\mathbf{x})$  denote its output activation vector (e.g., head/FFN output written into the residual stream) on input  $\mathbf{x}$ ; the *sender delta*  $\Delta\mathbf{a}_u = \mathbf{a}_u(\mathbf{x}_{clean}) - \mathbf{a}_u(\mathbf{x}_{corr})$  captures how the sender’s contribution changes between the clean and corrupted runs. *EAP* then scores an edge via the dot product between the sender delta and the receiver sensitivity  $\nabla_{\mathbf{z}_v}\mathcal{R}$  (computed on the clean run):

$$S_{EAP}(u \rightarrow v) \approx \underbrace{(\mathbf{a}_u(\mathbf{x}_{clean}) - \mathbf{a}_u(\mathbf{x}_{corr}))}_{\Delta\mathbf{a}_u} \cdot \underbrace{\frac{\partial\mathcal{R}}{\partial\mathbf{z}_v}}_{\nabla_{\mathbf{z}_v}\mathcal{R}} \Big|_{\mathbf{x}_{clean}}. \quad (17)$$

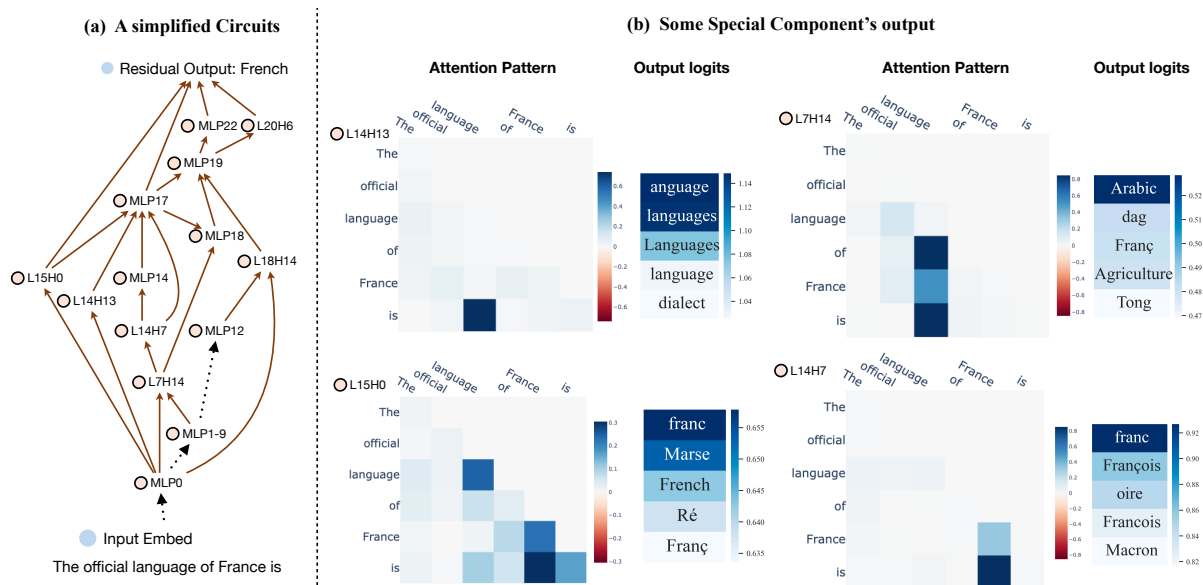
To mitigate non-linearity/saturation, *EAP with Integrated Gradients (EAP-IG)* replaces the local gradient with a path-averaged gradient along  $\mathbf{x}_\alpha = \mathbf{x}_{corr} + \alpha(\mathbf{x}_{clean} - \mathbf{x}_{corr})$  [56,123,124]:

$$S_{EAP-IG}(u \rightarrow v) = \Delta\mathbf{a}_u \cdot \int_0^1 \frac{\partial\mathcal{R}}{\partial\mathbf{z}_v} \Big|_{\mathbf{x}_\alpha} d\alpha \approx \Delta\mathbf{a}_u \cdot \frac{1}{n} \sum_{k=1}^n \frac{\partial\mathcal{R}}{\partial\mathbf{z}_v} \Big|_{\mathbf{x}_{k/n}}. \quad (18)$$

A standard workflow is: (i) collect sender deltas  $\Delta\mathbf{a}_u$  from  $\mathbf{x}_{clean}$  vs.  $\mathbf{x}_{corr}$ , (ii) compute receiver gradients (single-point for EAP; path-averaged for *EAP-IG* with  $n$  backward passes), (iii) score and rank edges by  $|S|$ , and (iv) prune/threshold to obtain a sparse circuit, optionally validating via targeted interventions on retained edges.

### Applicable Objects

*Circuit Discovery* targets *edges between objects* (Tab. 1), ranging over directed dependencies among any interpretable objects. In LLMs, it is commonly instantiated under the residual–rewrite view, so edges correspond to additive signal transmission across layers. Figure 9 illustrates a sparse cross-layer knowledge circuit supporting the completion “The official language of France is French” in GPT-2-Medium, with attention/logit analyses clarifying how selected edges route and transform information [119].



**Figure 9. Knowledge circuit example.** A sparse cross-layer circuit supporting the factual completion “The official language of France is French” in GPT-2-Medium. Left: A simplified *circuit*. Here, L15H0 means the first attention head in the 15th layer and MLP12 means the FFN block in the 13th layer. Right: Behavior of several special heads. The left matrix shows each head’s attention pattern, and the right heatmap shows output logits mapped to the vocabulary space. Based on the figure from Yao et al. [119].

Practically, *Circuit Discovery* is operationalized in three broad ways:

**1) Intervention-based edge search (patching/ablation):** One can directly test causal necessity at the edge level by patching or ablating a candidate dependency  $e_{u \rightarrow v}$  (e.g., blocking contributions from a sender module such as an attention head output  $\mathbf{h}_{\text{attn}}^{l,h}$  or FFN output  $\mathbf{h}_{\text{ffn}}^l$  into a downstream receiver input  $\mathbf{z}_v$ ) and measuring the change in a task metric  $\mathcal{R}$ . Because exhaustive edge testing scales as  $O(|\mathcal{E}|)$ , practical workflows rely on structured search or automated procedures to reduce interventions [45,52,121].

**2) Attribution-based edge scoring:** Attribution methods rank edges by efficiently approximating their patching effect. *EAP* combines sender activation differences  $\Delta \mathbf{a}_u$  (clean vs. corrupted) with receiver sensitivity  $\nabla_{\mathbf{z}_v} \mathcal{R}$  to produce an edge ranking from two forward passes and one backward pass, while *EAP-IG* uses a path-averaged gradient to reduce saturation/non-linearity issues at the cost of additional backward passes [122–124]. Position-aware refinements follow the same edge-scoring principle while better aligning sender/receiver accounting with token-wise computation [125,126].

**3) Feature-based replacement models:** *Circuit Discovery* can be lifted to sparse feature spaces via replacement models such as SAE/transcoder variants. Here, the relevant objects are *SAE features* (sparse feature activations and decoder directions), and circuit edges represent directed dependencies in feature space, enabling attribution graphs and prompt-specific circuit tracing that are often more interpretable than raw residual coordinates [127–129].

## Characteristics and Scope

*Circuit Discovery* identifies a sparse, directed **cross-layer causal subgraph** whose edges jointly mediate a target behavior and remain approximately faithful under interventions. Unlike single-component localization, it targets structured pathways of information routing and transformation, returning a mini-

mally (or strongly) sufficient directed subnetwork. Practically, edges are often pre-ranked by scalable attribution-style scores and then confirmed with targeted interventions (e.g., patching/ablation).

- **Advantages:** *Circuit Discovery* yields mechanistically structured explanations: selecting *edges* reveals how multiple objects compose a computation and exposes cross-layer routing patterns that node-wise rankings can miss. This aligns with transformers' residual-update structure, where heads and FFNs contribute additive edits that can be tracked as directed dependencies. Attribution-based edge scoring also enables scalable screening of large edge sets when exhaustive interventions are infeasible.
- **Limitations:** Circuits are defined relative to a specific behavior, metric  $\mathcal{R}$ , and contrast (clean vs. corrupted), so results are often *objective- and dataset-dependent*. Because attribution scores approximate intervention effects, they may miss non-linear interactions, so rankings are best treated as proposals and typically require intervention-based validation on the retained subgraph.

## 4. Steering Methods

While localization methods (§3) identify the specific objects responsible for model behaviors, this section focuses on a distinct class of techniques: those that manipulate these localized components to steer model outputs, thereby enabling controlled intervention into LLM's generation process.

### 4.1. Amplitude Manipulation

#### Methodological Formulation

*Amplitude Manipulation* steers model behavior by directly modifying the activation magnitude of a targeted internal object  $o$  during the forward pass. Unlike optimization-based methods that update weights, this approach acts as a transient intervention on the runtime state. Formally, let  $o$  be the original activation (e.g., a neuron activation  $s_j^l$ , an SAE feature activation  $a_j$  or an attention head output  $\mathbf{h}_{\text{attn}}^{l,h}$ ) and  $\tilde{o}$  be the modified state. The intervention is defined as:

$$\tilde{o} = \mathcal{T}(o, \alpha) \quad (19)$$

where  $\mathcal{T}$  represents the transformation function. This typically takes two forms:

- **Ablation or Patching:** Here, the object is suppressed or replaced, i.e.,  $\tilde{o} \in \{0, \mathbb{E}[o], o_{\text{tgt}}\}$ . Setting  $\tilde{o} = 0$  (Zeroing) or  $\mathbb{E}[o]$  (Mean centering) removes the component's influence, while  $\tilde{o} = o_{\text{tgt}}$  (Patching) injects information from a different context.
- **Scaling:** Here, the activation strength is adjusted via a scalar coefficient  $\alpha$ , such that  $\tilde{o} = \alpha \cdot o$ . This allows for continuous amplification ( $\alpha > 1$ ) or attenuation ( $0 < \alpha < 1$ ) of a specific feature's downstream impact.

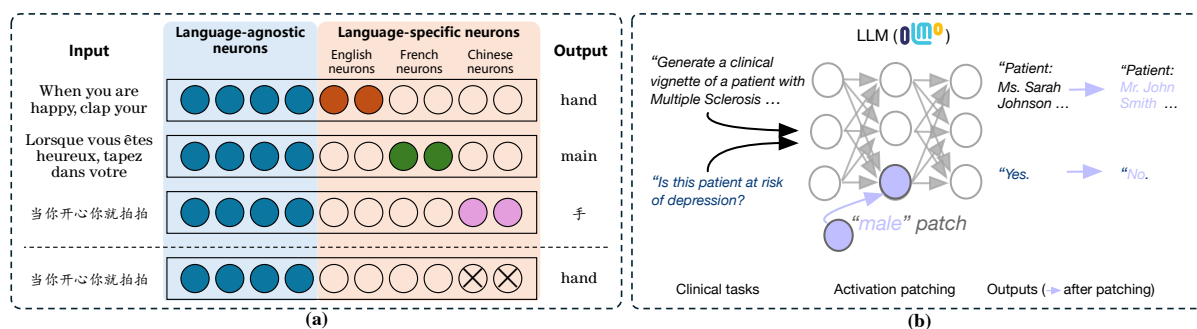
While these operations are mechanically similar to those in *Causal Attribution* (§3.2), the objective differs fundamentally: attribution employs them to *diagnose* causality, whereas *Amplitude Manipulation* employs them to actively *intervene* and control model behavior.

#### Applicable Objects

This method is applied across a wide range of dynamic objects, including residual stream state  $\mathbf{x}^l$ , attention head output  $\mathbf{h}_{\text{attn}}^{l,h}$ , neuron activation state  $s^l$ , and SAE feature activation state  $\mathbf{a}$ .

**1) Ablation (Zeroing) and Removal:** Ablation is extensively used to mitigate unwanted behaviors by suppressing the components responsible for them. Tang et al. [2] utilized this to control the output language of multilingual LLMs. They identified "language-specific neurons" that selectively activate for

the particular language (e.g., Chinese). As illustrated in Figure 10 (a), by setting the activation of these *Chinese-specific neurons* to zero, they suppressed the model's ability to generate Chinese, thereby forcing the model to switch its output to English even when the prompt might suggest otherwise. Distinct from general steering, Nie et al. [130] applied ablation to address *Language Confusion*—a phenomenon where models erroneously switch to a non-target language. They identified interfering neurons that activate for the wrong language (e.g., German neurons firing during an English task) and demonstrated that ablating these specific noisy components restores the correct target language generation. In the domain of *Safety and Bias*, Goyal et al. [131] and Yeo et al. [132] zeroed out specific SAE features associated with toxicity or refusal, effectively detoxifying the model's output. Liu et al. [133] and Chandna et al. [134] applied zero-ablation to neurons and circuit edges encoding social bias, while Huang et al. [124] masked specific circuit edges to alleviate “knowledge overshadowing” where strong knowledge suppresses relevant but weaker information. Furthermore, ablation is used for **Efficiency**: Liu et al. [25] and Men et al. [14] demonstrated that removing redundant layers or components can accelerate inference without significant performance loss. Zhou et al. [135] and Niu et al. [136] also utilized attention head ablation to study and improve safety and contextual entrainment.



**Figure 10.** Examples of Steering via Amplitude Manipulation. **(a) Ablation for Language Steering:** Tang et al. [2] deactivate (zero out) “Chinese-specific neurons” to suppress model’s ability to generate Chinese, successfully forcing the model to switch its output to English. **(b) Patching for Demographic Steering:** Ahsan et al. [137] inject a “Male Patch” into the model’s internal representation. This intervention not only changes the gender pronouns in the output (“Ms.” → “Mr.”) but also causally alters the clinical decision regarding depression risk (“Yes” → “No”), demonstrating the deep impact of internal demographic representations.

**2) Patching (Replacement):** Patching allows for precise injection of attributes. Ahsan et al. [137] and Raimondi et al. [138] utilized activation patching to steer demographic and moral characteristics. As shown in Figure 10 (b), Ahsan et al. [137] performed a “Male Patch” by replacing the internal representation of a patient with a male-associated vector. This intervention not only altered the pronouns in the generated vignette (from “Ms.” to “Mr.”) but also causally changed the downstream clinical prediction (shifting the depression risk from “Yes” to “No”), highlighting the causal link between demographic representations and model decisions.

**3) Scaling (Amplification/Attenuation):** Scaling offers fine-grained control by adjusting the intensity of features. Tang et al. [2] also employed scaling to amplify target-language neurons to further stabilize multilingual generation. Gao et al. [139] scaled the activation of “Hallucination Neurons” to modulate the model’s factual reliability. In the context of SAE features, Pach et al. [140] demonstrated that scaling specific feature activations allows for continuous steering of model outputs. Meanwhile, Galichin et al. [3] showed that amplifying the activations of reflection-related features can increase the length of generated output, thereby enhancing the model’s reasoning performance. Finally, scaling is integral to

*Vector Arithmetic* (§4.3) in the context of model merging: Stoehr et al. [141], Liu et al. [72], and Yao et al. [142] optimized the scaling coefficients of steering vectors or task vectors to balance different model capabilities, while Wang et al. [143] scaled the activations of expert modules to enhance mathematical reasoning.

#### Characteristics and Scope

The scope of *Amplitude Manipulation* is characterized by **inference-time activation control**. It provides a mechanism to transiently modulate model behavior without permanent weight updates.

- **Advantages:** It is an *optimization-free* and reversible intervention. It allows for “surgical” edits to model behavior (e.g., removing specific biases) by simply masking or scaling activations during inference. This makes it highly flexible and suitable for real-time control.
- **Limitations:** It relies heavily on the accurate *localization* of the target components. If the features responsible for a behavior are not perfectly disentangled (i.e., polysemantic), ablating or scaling them may cause unintended side effects or degrade general performance. Furthermore, finding the optimal scaling factor  $\alpha$  often requires empirical tuning.

#### 4.2. Targeted Optimization

##### Methodological Formulation

*Targeted Optimization* (under *Localizing Methods*) frames model optimizing as a *small, localized update* that enforces a desired behavioral change while minimizing unintended side effects. Let  $f_\theta$  be the base model and  $\theta'$  the *targeted* model. We restrict updates to a selected subset of objects via a (hard or soft) mask  $M$ , and optimize a simple trade-off between a *target objective* and a *preservation objective*:

$$\theta' \leftarrow \theta + (M \odot \Delta\theta), \quad \Delta\theta^* = \arg \min_{\Delta\theta} \mathcal{L}_{\text{tgt}}(f_{\theta'}; \mathcal{D}_{\text{tgt}}) + \lambda \mathcal{L}_{\text{pres}}(f_{\theta'}, f_\theta; \mathcal{D}_{\text{pres}}). \quad (20)$$

Here,  $\mathcal{D}_{\text{tgt}}$  specifies the target behavior (e.g., rewriting a fact or enforcing refusals), while  $\mathcal{D}_{\text{pres}}$  anchors the model to its original capabilities. The localization mask  $M$  operationalizes “where the change is allowed to happen” (layers, modules, neurons/heads, or other structured subsets).

##### Applicable Objects

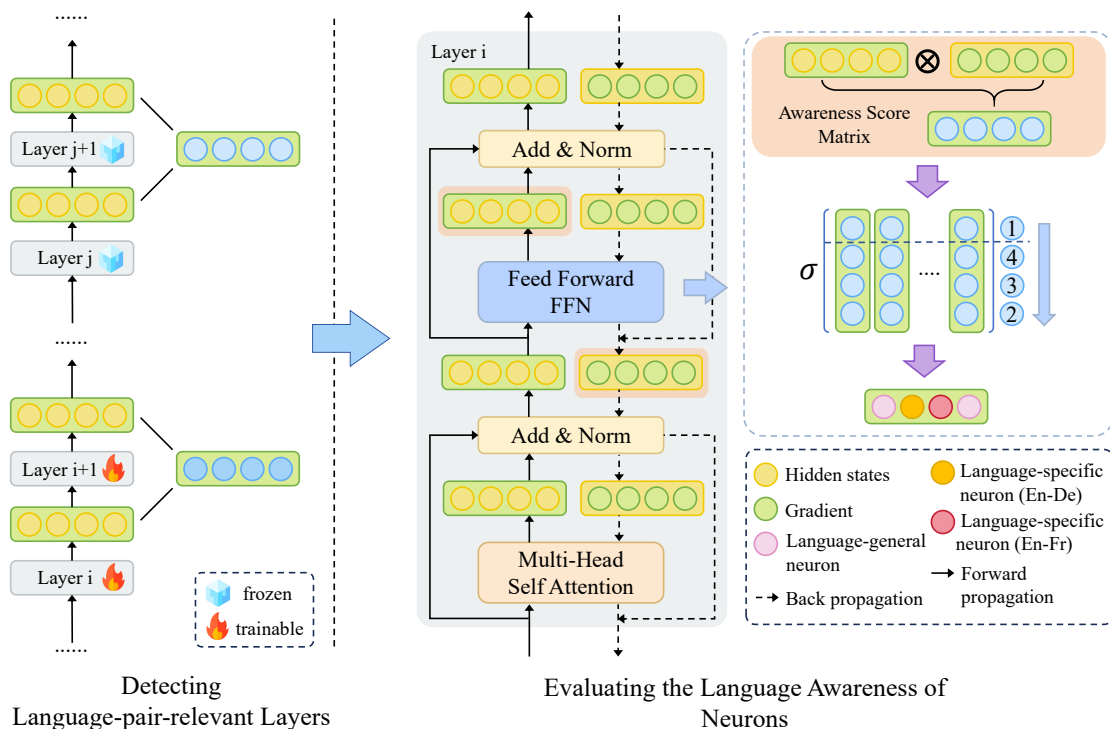
In practice, “what is optimized” can be grouped into two representative localized objects:

**1) Localized Parameters for Knowledge Editing:** This line performs direct parameter-space updates that are intentionally constrained (e.g., low-rank or small support) to rewrite specific behaviors with minimal spillover. Representative examples include rank-one / layer-local *knowledge editing* extensions [43,144], cross-model knowledge transfer via localized adapters [145], and constraining adaptation to low-dimensional task subspaces or coarse-to-fine masked tuning for better retention [66,146].

**2) Fine-grained Subsets for Specialization:** Here, localization is enforced at neuron/head/region granularity to isolate the functional unit relevant to a capability or a safety property. Concretely, rather than updating the full model, *Targeted Optimization* learns a targeted update within a small object subset (implicitly corresponding to a mask  $M$  in Eq. 20), thereby limiting unnecessary parameter drift and reducing interference across tasks or languages. Related lines of work localize adaptation to compact trainable units at different granularities. This includes neuron-level fine-tuning [147] and methods that identify core parameter regions or language-agnostic factual neurons [148–150], and in safety-preserving or security-aware partial tuning that freezes or restricts sensitive objects [88,151,

152]. Relatedly, head-level analyses further motivate localizing optimization to essential computational pathways (e.g., arithmetic-relevant heads) [153].

A representative example is shown in Figure 11: LANDeRMT [154] performs selective fine-tuning for multilingual machine translation by (i) first localizing the update to language-pair-relevant layers, (ii) quantifying neuron-level language awareness, and (iii) routing gradients only through the most relevant neurons, which concretely illustrates how fine-grained locality reduces cross-lingual interference and limits parameter drift.



**Figure 11.** A representative *Targeted Optimization* pipeline. The method first identifies language-pair-relevant layers, then scores neuron language-awareness, and finally routes gradient updates to a small subset of language-aware neurons for selective fine-tuning. This illustrates how *Targeted Optimization* enforces locality via an object mask  $M$  in Eq. 20. Based on the figure from Zhu et al. [154].

### Characteristics and Scope

The scope of this method is characterized by **persistence** and **surgical precision**. Unlike *Amplitude Manipulation* (§4.1), *Targeted Optimization* performs *parameter optimization* on  $\mathcal{D}_{\text{tgt}}$  to produce a targeted model whose behavior *durably* satisfies a specified objective, while *constraining the update* to a localized subset of objects (e.g., layers, modules, neurons/heads). This objective-driven and localized training enables not only precise rewrites to particular memories or facts, but also focused capability enhancement, with reduced collateral impact on unrelated traits.

- **Advantages:** It offers strong *precision*, *controllability*, and *persistence*. The desired behavioral change is directly encoded in a target objective, and localization helps minimize interference with unrelated competencies. Consequently, it is well-suited for targeted factual rewrites, controlled specialization, and safety-preserving adaptation where *lasting* changes are required.

- **Limitations:** Its reliability hinges on *correct localization* and *well-specified supervision*. If the chosen subset does not capture the causal mechanism, optimization may underachieve the intended target behavior, shift the behavior to other objects, or yield brittle side effects. In practice, success often requires carefully constructed target/preservation data and robust criteria for selecting the localized update region.

### 4.3. Vector Arithmetic

#### Methodological Formulation

Positing that high-level concepts or skills are encoded linearly within the model’s representation space, *Vector Arithmetic* steers a generic target object  $\mathbf{z}$  (e.g., a residual stream state or a parameter vector) by injecting a specific *steering vector*  $\mathbf{v}$ . This approach assumes that adding a vector representing a concept effectively “moves” the model’s internal state towards that concept in the high-dimensional space. Formally, the update rule for the intervention is defined as:

$$\hat{\mathbf{z}} \leftarrow \mathbf{z} + \alpha \cdot \mathbf{v} \quad (21)$$

where  $\mathbf{v}$  represents the directional encoding of a target attribute (such as “honesty” or “sycophancy”) and  $\alpha$  is a scalar coefficient that controls the intervention strength (or steering intensity).

#### Applicable Objects

The target object  $\mathbf{z}$  typically falls into two categories: dynamic hidden states during inference or static model parameters.

**1) Dynamic Hidden States:** The primary targets for runtime steering are the residual stream states  $\mathbf{x}^l$  and the outputs of attention heads  $\mathbf{h}_{\text{attn}}^{l,h}$ . For these dynamic objects, the steering vector  $\mathbf{v}$  is typically derived using one of two methods:

- **Contrastive Activation Means:** This method, often referred to as “Activation Addition” or “Mass-Mean Shift,” assumes that a concept can be isolated by comparing the model’s internal states across opposing contexts [162–167]. Formally, let  $\mathcal{D}^+$  be a set of prompts eliciting the target behavior and  $\mathcal{D}^-$  be a set eliciting the opposing behavior. The steering vector  $\mathbf{v}$  is calculated as the difference between the centroids of the residual stream states  $\mathbf{x}^l$  for these two sets:

$$\mathbf{v} = \boldsymbol{\mu}^+ - \boldsymbol{\mu}^- = \frac{1}{|\mathcal{D}^+|} \sum_{\mathbf{x}_i \in \mathcal{D}^+} \mathbf{x}_i^l - \frac{1}{|\mathcal{D}^-|} \sum_{\mathbf{x}_j \in \mathcal{D}^-} \mathbf{x}_j^l \quad (22)$$

By adding  $\alpha \cdot \mathbf{v}$  to the residual stream, we shift the model’s current state towards the centroid of the positive behavior.

- **SAE Features:** SAEs offer a more precise way to derive  $\mathbf{v}$  by utilizing monosemantic features [131, 168–172]. As illustrated in Figure 12, the process involves two steps:
  1. **Feature Identification:** First, we collect residual stream states from a positive dataset  $\mathcal{D}^+$  (eliciting the target concept, e.g., “Happiness”) and a negative/neutral dataset  $\mathcal{D}^-$ . By passing these states through the SAE encoder, we calculate the *differential activation score*  $\delta_j$  for each feature  $j$ :

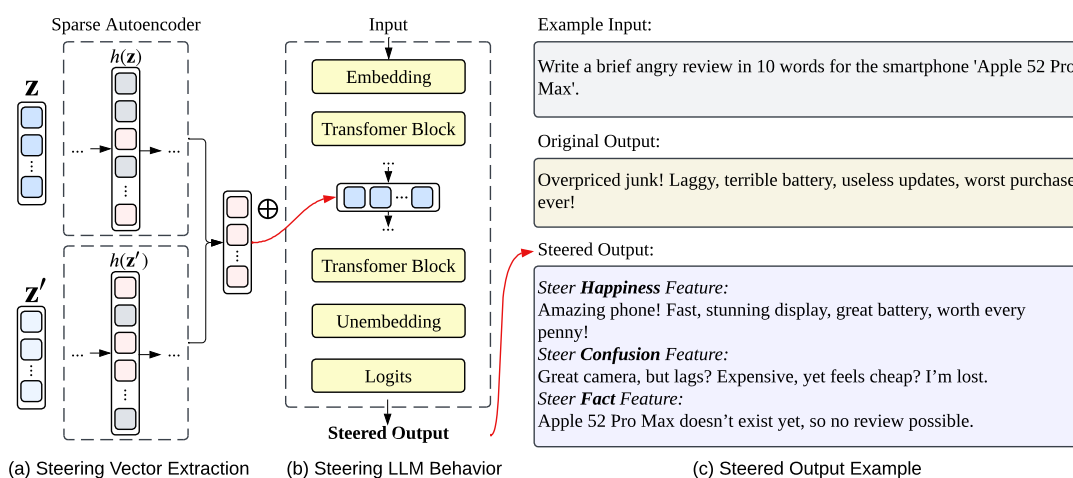
$$\delta_j = \mathbb{E}_{\mathbf{x} \in \mathcal{D}^+} [a_j(\mathbf{x})] - \mathbb{E}_{\mathbf{x} \in \mathcal{D}^-} [a_j(\mathbf{x})] \quad (23)$$

where  $a_j(\mathbf{x})$  denotes the  $j$ -th feature activation for input  $\mathbf{x}$ . Features with high positive  $\delta_j$  constitute the set of “Target Features”  $\mathcal{J}$  that specifically encode the desired trait.

2. **Vector Construction:** The steering vector  $\mathbf{v}$  is then synthesized as the weighted sum of these identified feature. Let  $\mathbf{f}_j$  denote the  $j$ -th feature (the  $j$ -th column of the SAE decoder weights  $\mathbf{W}_{\text{dec}}$ ). The steering vector is computed as:

$$\mathbf{v} = \sum_{j \in \mathcal{J}} \delta_j \cdot \mathbf{f}_j \quad (24)$$

Finally, this obtained steering vector is injected into the model's residual stream during inference ( $\hat{\mathbf{x}} \leftarrow \mathbf{x} + \alpha \cdot \mathbf{v}$ ). As shown in Figure 12 (c), this enables precise manipulation of specific semantic traits like "Happiness" or "Confusion" to drastically alter generation styles while minimizing interference with unrelated concepts.



**Figure 12.** The pipeline for steering LLMs using SAE features. (a) **Steering Vector Extraction:** The target steering vector is derived by analyzing a set of prompts to identify features that distinguish a concept-rich state  $\mathbf{z}'$  from a neutral state  $\mathbf{z}$ . The steering vector is computed as the weighted sum of these identified SAE features (i.e., decoder columns). (b) **Steering LLM Behavior:** This aggregated vector is injected into the Transformer's residual stream state  $\mathbf{x}^l$  via vector addition. (c) **Steered Output Example:** Empirical results showing how steering specific features (e.g., Happiness, Confusion) drastically alters the model's generation style even when the original prompt implies a negative sentiment. Based on the figure from Shu et al. [115].

**2) Static Parameters:** For static weights, the steering vector  $\mathbf{v}$  is explicitly defined as a *Task Vector* in Model Merging [72,142,173,174]. This vector is computed as the element-wise difference between the weights of a fine-tuned model and its pre-trained base ( $\mathbf{v} = \mathbf{W}_{\text{ft}} - \mathbf{W}_{\text{base}}$ ), effectively encapsulating a transferable skill or behavior. Recent advancements have evolved beyond simple element-wise addition by employing localization techniques to determine adaptive merging coefficients. For instance, Liu et al. [72] proposed *Sens-Merging*, which utilizes *Gradient Detection*-based sensitivity analysis to evaluate parameter importance, allowing for the precise balancing of weights based on their impact on task performance. Complementarily, Yao et al. [142] introduced *Activation-Guided Consensus Merging*, which leverages *Magnitude Analysis* of internal representations. By calculating the mutual information between activations of the base and fine-tuned models, they derive layer-specific scaling coefficients to optimally integrate the task vector.

## Characteristics and Scope

The scope of this method is characterized by **additive directionality**. Unlike the precise rewriting in *Targeted Optimization* (§4.2), *Vector Arithmetic* acts as a **steering force**, dynamically pushing the model towards a target attribute without permanently altering weights.

- **Advantages:** It is a lightweight and reversible intervention. Since it typically operates at inference time (for hidden states) or via simple weight addition, it does not require complex optimization or gradient descent during deployment. It allows for flexible control over model behavior by simply adjusting the steering coefficient  $\alpha$ .
- **Limitations:** The effectiveness relies on the “Linear Representation Hypothesis.” If the target concept is not encoded linearly or if the steering vector  $\mathbf{v}$  is entangled with other concepts (which is common with “Contrastive Activation Means”), the intervention might introduce unintended side effects.

## 5. Applications

Building on localizing methods (§3) that identify internal objects associated with specific behaviors and steering methods (§4) that intervene on these objects to modulate model outputs, this section summarizes how these lines of work translate into practical use cases. We organize the literature around three overarching objectives: *alignment*, *capability*, and *efficiency*.

### 5.1. Improve Alignment

#### 5.1.1. Safety and Reliability

##### Summary of Application Paradigms

MI improves safety and reliability in alignment applications primarily through two complementary, mechanism-aware intervention paradigms:

- 1) **Safety-Critical Component Manipulation.** This paradigm focuses on identifying internal components that explicitly encode unsafe, harmful, or unreliable behaviors, such as toxicity, hallucination, or failed refusal, and intervening on these components directly. By localizing safety-critical attention heads, neurons, circuits, or SAE features, researchers apply inference-time *Amplitude Manipulation* (§4.1) to suppress unsafe activations, or use training-based *Targeted Optimization* (§4.2) to permanently rewrite safety-relevant parameters and internal signals.
- 2) **Latent Safety and Reliability Representation Steering.** This paradigm operates at the level of the residual stream, where abstract safety- and reliability-related concepts such as truthfulness, refusal, and instruction following are encoded as approximately linear directions. By identifying these directions using causal or contrastive analyses, models are steered via *Vector Arithmetic* (§4.3) to correct hallucinations, enforce proper refusal behavior, or improve instruction adherence, while largely preserving general capabilities.

#### 1) Safety-Critical Component Manipulation

Unsafe or unreliable behaviors in LLMs have been shown to be mediated by relatively localized internal components. Accordingly, a body of work first localized safety-relevant objects and then intervened via targeted mechanistic techniques. At the attention level, Zhou et al. [135] showed that a small subset of attention heads played a disproportionate role in safety-related behaviors, particularly

refusal and rejection of harmful queries. Using *Causal Attribution* to localize safety-critical heads and *Amplitude Manipulation* to intervene, they demonstrated that suppressing these heads substantially weakened safety capability while modifying only a negligible fraction of parameters. At the neuron level, several studies applied *Magnitude Analysis* to identify neurons whose activations were strongly associated with unsafe or misaligned behaviors. Zhao et al. [178] introduced *safety neurons* and showed that a very small subset — predominantly located in early self-attention layers — collectively governed safety behavior; they then performed *Targeted Optimization* by selectively tuning these neurons during training, significantly improving safety without degrading general performance. Complementarily, Suau et al. [176] used magnitude-based criteria to pinpoint toxicity-related neurons and applied *Amplitude Manipulation* by scaling down their activations at inference time to mitigate toxic generations. Similarly, Gao et al. [139] identified hallucination-associated neurons (*H-neurons*) via *Magnitude Analysis* and validated their causal impact through *Amplitude Manipulation*, showing that suppressing these neurons reduced hallucinations without broadly affecting other capabilities. Beyond individual neurons, recent work leveraged SAEs to disentangle safety-related representations into interpretable features. Using *Magnitude Analysis* over SAE feature activation states, Templeton et al. [177] showed that SAE features extracted from LLMs exhibited strong monosemanticity, including features associated with harmful or toxic content. Building on this insight, Goyal et al. [131] applied *Amplitude Manipulation* to suppress selected SAE features and thereby detoxify model outputs. Likewise, Yeo et al. [132] performed SAE-based *Magnitude Analysis* to identify harm- and refusal-related feature sets and validated their roles through targeted *Amplitude Manipulation*, enabling fine-grained control and mechanistic insight into refusal behavior.

While *Amplitude Manipulation*-based interventions typically operate at inference time, several works pursued more persistent safety improvements through *Targeted Optimization*. Huang et al. [124] identified safety-relevant circuits and updated only parameters within these circuits to mitigate harmful behaviors. At finer granularity, Zhao et al. [178], Chen et al. [175], and Li et al. [152] showed that selectively updating neuron-associated weights enabled precise safety edits with minimal side effects. At a coarser level, Li et al. [151] demonstrated that safety behavior could be localized at the layer level, while Lee et al. [161] analyzed how alignment objectives reshaped internal representations during optimization.

## 2) Latent Safety and Reliability Representation Steering

A complementary line of research shows that many safety-relevant behaviors are encoded as approximately linear directions in LLM's latent space, motivating safety interventions based on *Vector Arithmetic* in the residual stream.

Arditi et al. [179] and Zhao et al. [180] showed that refusal was encoded as a compact low-dimensional subspace identified via *Causal Attribution*, and that some jailbreaks succeeded by suppressing this refusal signal via *Vector Arithmetic* without changing the model's harmfulness belief. Extending these findings to reasoning models, Yin et al. [181] identified a refusal-cliff phenomenon using *Probing*, where refusal intent was maintained during intermediate reasoning but was abruptly suppressed at the final generation stage, a failure mode attributed to a small set of refusal-suppressing attention heads. Building on these analyses, multiple studies identified actionable safety directions and applied steering interventions. Wang et al. [183] proposed a training-free, single-vector ablation method (a form of *Vector Arithmetic*) that selectively removed false refusal while preserving true refusal and general capabilities, enabling fine-grained safety calibration. Wang et al. [184] further demonstrated that refusal directions were approximately universal across safety-aligned languages, helping to explain the effectiveness of cross-lingual jailbreaks as well as vector-based interventions.

*Vector Arithmetic*-based steering was also applied to hallucination reduction and factuality improvement. Chuang et al. [11] introduced contrastive layer decoding (a form of *Vector Arithmetic*) during generation to amplify factual signals identified via *Vocabulary Projection*. Similarly, Zhang et al. [12] identified a truthfulness direction in the residual space using *Probing* and then edited it via *Vector Arithmetic*, enabling controllable enhancement of truthful behavior. Complementarily, Orgad et al. [86] showed that hallucination-related representations could be detected internally via *Probing* even when they were not expressed at the output level, highlighting the diagnostic value of latent safety signals. Finally, recent work applied *Vector Arithmetic* to improve instruction-following reliability. He et al. [171] leveraged SAE-derived directions to steer instruction adherence, while Stolfo et al. [187], Jiang et al. [188], and Li et al. [189] demonstrated that instruction-following behavior could be improved through *Vector Arithmetic* steering without full retraining.

### 5.1.2. Fairness and Bias

#### Summary of Application Paradigms

MI facilitates the diagnosis and control of fairness-related biases (gender bias, distributed attribute and cultural bias signals, and evaluation bias) in LLMs through three primary paradigms:

- 1) **Gender Bias Localization and Selective Debiasing** This paradigm localized gender-bias mediation primarily via *Causal Attribution* (§3.2) and then reduced bias through either inference-time *Amplitude Manipulation* (§4.1) or persistent *Targeted Optimization* (§4.2) on the identified components.
- 2) **Distributed Attribute and Cultural Bias Signals** This paradigm extended beyond gender to demographic, social, and cultural biases, often requiring broader searches over internal structures (including *Magnitude Analysis* (§3.1) or *Gradient Detection* (§3.3)).
- 3) **Evaluation Bias Engines in Judgment and Framing** This paradigm studied cognitive and judgment biases induced by prompt format or evaluation settings (e.g., positional anchoring and moral attribution), and mitigated them via inference-time controls such as attention/position re-assignment or targeted scaling, typically guided by *Magnitude Analysis* (§3.1) and validated by *Causal Attribution* (§3.2).

#### 1) Gender Bias Localization and Selective Debiasing

Mechanistic studies of gender bias established a canonical fairness pipeline: first localizing bias mediation with *Causal Attribution*, then steering the identified carriers via either transient inference-time control or persistent parameter updates. Vig et al. [42] provided an early template using causal mediation analysis in GPT-2 to quantify which internal components mediated gendered associations, and demonstrated mitigation by replacing bias-inducing activations with counterfactual ones, a direct instance of *Amplitude Manipulation*. To achieve persistent mitigation, subsequent work increasingly shifted to selective updates of localized components via *Targeted Optimization*. Chintam et al. [192] showed that responsibility for gender bias could concentrate in specific late-layer attention heads and reduced bias by fine-tuning only these components. Cai et al. [190] characterized a division of labor where lower FFN blocks encoded bias-relevant information while upper attention modules exploited it, proposing an editing-style method to update the responsible subset. Finally, Yu and Ananiadou [109] refined the intervention granularity to the neuron level, identifying distinct “gender neurons” versus “general neurons” and introducing an interpretable neuron-editing procedure to reduce bias while preserving general performance.

## 2) Distributed Attribute and Cultural Bias Signals

Beyond gender, mechanistic evidence suggested that demographic, social, and cultural biases are often encoded more diffusely. This motivated localization strategies that avoid assuming a single “bias module,” alongside mitigation strategies combining targeted suppression with global representational steering. In domain-conditioned settings like healthcare, Ahsan et al. [137] used activation patching, a form of *Causal Attribution*, to localize racial information across multiple LLMs, reporting that racial signals are more scattered across early and middle FFN layers compared to gender. Similarly, Yu et al. [51] adopted Patchscope-style interventions to “read out” cultural knowledge from internal representations. Rather than proposing a mitigation, their results focused on diagnosis, revealing how cultural salience and resource imbalance manifest as systematic representational asymmetries. Addressing broader societal biases, Liu et al. [133] employed *Gradient Detection* to identify neurons associated with multiple social attributes and demonstrated mitigation by suppressing their activations. To scale localization beyond hand-picked modules, Chandna et al. [134] combined *Magnitude Analysis* over internal structures with causal validation to create a reusable recipe for bias analysis across attributes. Finally, acknowledging that values can be represented linearly, Kim et al. [317] used *Probing* to identify attention heads predicting political ideology, and then steered generations via *Vector Arithmetic*.

## 3) Evaluation Bias Engines in Judgment and Framing

A complementary thread targeted *cognitive biases* arising from judgment heuristics, prompt formats, or decision framing, rather than demographic correlations. These works mitigated such biases through inference-time controls guided by importance signals via *Magnitude Analysis* or validated by *Causal Attribution*. For positional anchoring in multiple-choice questions (MCQs), Li and Gao [297] identified higher-layer mechanisms in GPT-2 that preferentially routed evidence toward anchored option tokens, providing concrete intervention loci. Generalizing beyond MCQs, Wang et al. [318] formulated position bias across judge-style evaluation and retrieval-augmented QA, introducing a mechanism to re-assign positions based on attention-derived importance signals. Complementarily, Yu et al. [319] traced “lost-in-the-middle” failures to a positional hidden-state channel and proposed a search-and-scale procedure to rescale this channel, improving robustness on long-context benchmarks. Extending to domain-specific decision-making, Dimino et al. [320] localized mid-to-late transformer layers as core “bias engines” driving positional skew in financial advisory tasks. Finally, regarding moral judgment, Raimondi et al. [138] analyzed the Knobe effect, localized its mediation to residual activations, and reduced the intentionality attribution gap by patching fine-tuned states with their pretrained counterparts, selectively reverting value shifts introduced during alignment.

### 5.1.3. Persona and Role

#### Summary of Application Paradigms

MI facilitates the analysis and control of LLM personas and roles through three primary paradigms, ranging from global representation engineering to fine-grained component editing:

- 1) **Global Persona Modulation via Vectors:** This paradigm posits that high-level personality traits (e.g., sycophancy, honesty) are encoded as linear directions within the global activation space. Researchers extract these “Persona Vectors” and apply *Vector Arithmetic* (§4.3) to the residual stream, steering the model’s behavior without altering weights.
- 2) **Persona-Specific Component Editing:** Moving beyond global vectors, this approach identifies specific model components, such as individual neurons or attention heads, that serve as the physical carriers of personality traits. These components are then targeted via *Amplitude Manipulation* (§4.1) or refined through *Targeted Optimization* (§4.2) to achieve persistent behavioral changes.
- 3) **Psychological Profiling and Diagnosis:** Instead of active intervention, this paradigm utilizes MI as a diagnostic tool. By employing *Probing* (§3.4) techniques and analyzing activation geometry, researchers can locate where psychological traits emerge, validate the stability of roles, and predict model behaviors before generation occurs.

#### 1) Global Persona Modulation via Vectors

A growing body of work suggests that complex persona-specific behavioral traits can be manipulated by intervening in the global activation state of the model. Rimsky et al. [162] utilized Contrastive Activation Addition, a form of *Vector Arithmetic*, to steer models away from sycophantic and hallucinatory behaviors. By extracting steering vectors from the residual stream differences between positive and negative examples, they demonstrated that high-level alignment properties can be precisely modulated during inference without fine-tuning. Expanding on this, Chen et al. [194] developed an automated pipeline to extract “Persona Vectors” for arbitrary traits (e.g., “evil” or “sycophantic”) using natural language descriptions. They found that these vectors not only allow for post-hoc steering but can also be used to predict and mitigate unintended persona shifts (e.g., emergent misalignment) that occur during fine-tuning by monitoring the projection of training data onto these vectors. Potertì et al. [193] applied this concept to professional domains, constructing “Role Vectors” (e.g., Chemist, Doctor) from model activations. Their analysis revealed that reinforcing these role-specific directions significantly improves performance on domain-specific tasks and even yields cross-domain benefits, suggesting that role-playing is mechanistically grounded. Furthermore, Pai et al. [321] proposed *BILLY*, a training-free framework that blends multiple persona vectors (e.g., Creative Professional + Environmentalist) to simulate collective intelligence within a single model. This approach steers the model with a composite vector, enhancing creativity and diversity in generation without the computational cost of multi-agent systems. Similarly, Sun et al. [37] explored the task vector (i.e., the steering vector in the context of model merging), extracting personality vectors by subtracting pre-trained weights from fine-tuned ones. They showed that these vectors can be linearly composed to continuously modulate trait intensity (e.g., Extraversion) across different models. Finally, Handa et al. [195] conducted a rigorous comparative study of personality manipulation methods using the Big Five traits. Their results showed that *Vector Arithmetic* provides a lightweight yet effective approach for controlling model personas at inference time.

## 2) Persona-Specific Component Editing

Rather than steering the global state, this paradigm seeks to identify and edit the specific neural components responsible for personality expression. Deng et al. [197] proposed *NPTI*, a method that identifies “personality-specific neurons” by applying *Magnitude Analysis* on the activation differences between opposing trait descriptions (e.g., Extraversion vs. Introversion). By selectively activating or deactivating these neurons via *Amplitude Manipulation*, they achieved fine-grained control over the model’s personality without model training. Su et al. [196] extended this to ethical values, introducing *ValueLocate*. They constructed a dataset based on the *Schwartz Values Survey* [322], a well-established framework that classifies values into four dimensions: Openness to Change, Self-transcendence, Conservation, and Self-enhancement. Using this dataset, they located value-critical neurons and demonstrated that controlling them via *Amplitude Manipulation* can effectively alter the model’s value orientation. Addressing the specific issue of sycophancy, Chen et al. [198] identified a sparse set of attention heads (~4%) that significantly contribute to “yes-man” behavior. They proposed *Supervised Pinpoint Tuning*, a form of *Targeted Optimization*, which fine-tunes only these specific heads while freezing the rest of the model, successfully mitigating sycophancy while preserving general reasoning abilities better than standard instruction tuning.

## 3) Psychological Profiling and Diagnosis

MI techniques are also extensively used to understand how psychological constructs are represented internally by applying *Probing*. Tak et al. [199] investigated emotion inference, finding that emotion processing is functionally localized in MHA units within middle layers. They validated this by showing that interventions on latent “appraisal concepts” (e.g., pleasantness) predictably shift the generated emotional tone. Yuan et al. [200] explored how language identity affects psycholinguistic traits like sound symbolism and word valence. Their *Probing* analysis revealed that these signals become decodable in deeper layers and that language conditioning (e.g., bilingual persona) significantly modulates internal representations. Ju et al. [201] introduced a layer-wise *Probing* framework to analyze the Big Five personality traits, discovering that personality information is predominantly encoded in the middle and upper layers. They further proposed a method to edit response personality by applying *Amplitude Manipulation* to perturb hidden states orthogonal to the probing boundaries. In the realm of truthfulness, Joshi et al. [323] proposed the “persona hypothesis,” suggesting LLMs model truthfulness by inferring a “truthful persona” from the context. They provided evidence that *Probing* for this persona can predict the truthfulness of generated answers. Ghandeharioun et al. [324] utilized techniques like *Patchscopes* to reveal “latent misalignment,” showing that user personas (e.g., “altruistic” vs. “selfish”) significantly affect a model’s willingness to answer harmful queries, mediated by internal interpretations of the user’s intent. For user-facing transparency, Karny et al. [202] developed an interface that visualizes “Persona Scores” derived from neural activations. Their user study highlighted that users often miscalibrate their expectations of model behavior, and neural transparency tools can bridge this gap. Finally, Banayeeanzade et al. [203] introduced *PsySET*, their evaluation results showcase that although *Vector Arithmetic* steering is effective for modulating persona traits, it can introduce unintended side effects, such as “joy” steering reducing privacy awareness or “anger” steering increasing toxicity, necessitating rigorous safety evaluations. Bas and Novak [38] further differentiated between steering “internal dispositions” versus “external knowledge”, finding that steering method such as *Vector Arithmetic* is highly effective for latent traits (e.g., personality) but struggles with knowledge-heavy personas (e.g., specific public figures), where it often degrades coherence.

## 5.2. Improve Capability

### 5.2.1. Multilingualism

#### Summary of Application Paradigms

In multilingual and cross-lingual settings, MI enables targeted control and enhancement of language behavior in LLMs through two primary application paradigms:

- 1) **Language-Specific Component Manipulation:** This paradigm focuses on identifying and intervening on internal components that are specifically responsible for processing individual languages. By localizing *language-specific neurons* or *SAE features* via *Magnitude analysis* (§3.1), researchers directly manipulate their magnitudes through *Amplitude Manipulation* (§4.1) to control output language, enhance multilingual performance, or perform language-specific adaptation.
- 2) **Cross-Lingual Representation Steering:** This paradigm operates at the level of the residual stream, where multilingual representations are dynamically transformed across layers. By identifying language-related directions using *Vocabulary Projection* (§3.5) or *Magnitude Analysis* (§3.1), models are steered via *Vector Arithmetic* (§4.3) to align representations across languages, improve cross-lingual transfer, and mitigate language inconsistency or language mixing phenomena.

#### 1) Language-Specific Component Manipulation

A central line of multilingual MI research shows that multilingual capabilities in LLMs are supported by a relatively small subset of internal components exhibiting strong language specificity. Accordingly, existing work has focused on localizing these components and manipulating their activations to control output language or enhance multilingual performance. Zhao et al. [204] formalized this observation through a layered *Multilingual Workflow (MWork)*, showing that representations became English-centric in intermediate layers and were mapped back to the query language in later layers. They employed *PLND* to localize language-specific neurons and showed that intervening on only a tiny fraction could sharply disrupt multilingual performance, while selectively updating these neurons via *Target Optimization* enabled data-efficient language-specific adaptation. Along similar lines, Tang et al. [2] introduced *Language Activation Probability Entropy (LAPE)* as a *Magnitude Analysis* tool to quantify cross-language activation selectivity. Their results showed that language-specific neurons concentrated in the bottom and top layers, and that applying *Amplitude Manipulation* to these neurons provided direct control over output language, effectively reducing off-target generation. Complementing these localization-driven studies, Kojima et al. [325] analyzed neuron activation patterns across languages using *Magnitude Analysis* and confirmed the functional importance of language-specific neurons through *Amplitude Manipulation*, including targeted ablation and scaling. Together, these results reinforced component-level activation control as a practical mechanism for multilingual intervention.

Beyond identifying individual language-specific neurons, Gurgurov et al. [30] systematized neuron-level *Amplitude Manipulation* through Language Arithmetic, demonstrating that language-specific neurons exhibited additive properties that enabled controlled language switching and interpolation via linear operations on activations. In parallel, related work extended this paradigm beyond language identity to other forms of linguistic specialization. Liu et al. [205] identified relation-specific neurons whose activation patterns generalized across languages in multilingual factual probing tasks, demonstrating that neuron-level specialization could transfer cross-lingually beyond language identity. At a finer representational granularity, Jing et al. [206] employed SAEs to extract and analyze a wide range

of interpretable linguistic features whose activation patterns varied systematically across languages, and showed that *Amplitude Manipulation* of these features could causally affect corresponding linguistic behaviors. Similarly, Andrylie et al. [29] identified language-specific SAE features via *Magnitude Analysis* and demonstrated that *Amplitude Manipulation* on these features enabled fine-grained control over multilingual behavior. Finally, Brinkmann et al. [207] showed that SAE-based representations captured shared, cross-lingual grammatical abstractions, with targeted feature-level analyses providing supporting evidence.

## 2) Cross-Lingual Representation Steering in Residual Space

A second major paradigm improves multilingual behavior by intervening on internal representations in the residual stream, where language representations are progressively transformed and aligned across layers. Chi et al. [210] showed that cross-lingual transfer could be activated without task-specific supervision by restructuring and aligning multilingual representations across model components, suggesting that pretrained models encoded latent cross-lingual structure that could be activated without end-task data. To localize where multilingual representations diverged or aligned, several studies relied on *Vocabulary Projection*. In particular, Wendler et al. [98] revealed that multilingual models often operated in an English-centric latent space during intermediate layers, even for non-English inputs, motivating interventions in later layers to restore language-faithful generation. Complementary representation-space analyses further supported this view: Philippou et al. [208] analyzed the relationship between language distance and representation divergence, while Mousi et al. [209] studied alignment dynamics in shared multilingual spaces using clustering-based metrics, together characterizing how cross-lingual alignment evolved across layers.

Building on these localization insights, subsequent work intervened more directly on internal representations to influence multilingual behavior. Hinck et al. [211] analyzed English-dominant responses in vision-language models and showed that targeted *Vector Arithmetic* on internal attention and hidden states could mitigate this bias. More recent studies moved from localization to *failure-mode diagnosis* with MI tools. Using layer-wise *Vocabulary Projection* and representation analysis, Wang et al. [213] attributed cross-lingual factual inconsistency to late-layer transitions into language-related subspaces, while Liu et al. [215] traced when multilingual factual knowledge emerged across pretraining checkpoints, providing a developmental view of cross-lingual consistency rather than a post-hoc intervention account. Complementary, Wang et al. [214] analyzed language mixing in reasoning by characterizing when and how internal states drift between languages during generation. Nie et al. [130] further combined late-layer lens-style analysis with targeted neuron-level interventions (*Amplitude Manipulation*) to mitigate language confusion.

## 5.2.2. Knowledge Management

### Summary of Application Paradigms

MI enables precise analysis, control, and consolidation of model knowledge through three complementary paradigms, ranging from local intervention to global composition:

- 1) **Precise Knowledge Updating:** This paradigm identifies minimal carriers responsible for a target association using *Causal Attribution* (§3.2), optionally assisted by *Gradient Detection* (§3.3) or *Vocabulary Projection* (§3.5). Once located, associations can be modified either persistently through *Targeted Optimization* (§4.2) or transiently via *Amplitude Manipulation* (§4.1), achieving high specificity while controlling generalization.
- 2) **Knowledge Retention and Stability:** MI supports diagnosing interference under continual updates or context injections. Critical carriers are located mainly through *Magnitude Analysis* (§3.1), *Causal Attribution* (§3.2) and *Gradient Detection* (§3.3), and stability is maintained by either constraining training-time changes or applying inference-time *Amplitude Manipulation* (§4.1) to reduce drift.
- 3) **Knowledge Consolidation:** To integrate multiple skills or fine-tuned variants, MI identifies compatible subspaces or transferable feature bases using *Gradient Detection* (§3.3), *Magnitude Analysis* (§3.1) or *Probing* (§3.4). These objects are then combined via *Vector Arithmetic* (§4.3) to merge capabilities while preserving essential associations.

### 1) Precise Knowledge Updating

MI-based knowledge updating shares a common workflow: First localizes carriers responsible for a target association, then intervenes either at the parameter level (persistent) or activation level (reversible), with careful measurement of locality and collateral effects.

- **Localized Parameter Rewriting:** A core result is that many factual associations are mediated by localized pathways, often concentrated in mid-layer FFN output  $\mathbf{h}_{\text{ffn}}^l$  and neuron activation state  $\mathbf{s}^l$ . Meng et al. [43] used *Causal Attribution* to identify carriers responsible for factual recall and applied structured weight edits (on FFN matrices such as  $\mathbf{W}_{\text{out}}^l$ ) to rewrite specific associations, a process referred to as *Targeted Optimization*, providing a mechanistic alternative to diffuse fine-tuning. Scaling beyond single edits, Meng et al. [144] extended this paradigm to large edit batches by coordinating updates across multiple layers, demonstrating that persistent rewriting could remain localized while handling substantial edit volume. Subsequent work refined the localization premise: Chen et al. [216] argued that editability was frequently *query-conditioned*, motivating consistency-aware localization under a broader Query Localization assumption, rather than a fixed set of knowledge neurons. For long-form QA, Chen et al. [27] introduced QRNCA (a form of *Causal Attribution*), which yielded actionable neuron groups that better tracked query semantics. In multilingual settings, Zhang et al. [150] identified language-agnostic factual neurons via *Magnitude Analysis* and applied *Targeted Optimization* on these shared neurons to improve cross-lingual edit consistency. In backward propagation, Katz et al. [217] complemented forward analyses with *Vocabulary Projection* of backward-pass gradients, offering an orthogonal diagnostic on where learning signals concentrated during updates.
- **Activation-Space Editing and Unlearning:** When persistent rewrites are undesirable (e.g., reversible control or safety-motivated removal), activation-level interventions on residual stream states  $\mathbf{x}^l$  at layer  $l$ , or on head/feature activations, provide a practical alternative. Lai et al. [218] jointly localized and edited attention-head computations (intervening on attention head output  $\mathbf{h}_{\text{attn}}^{l,h}$ )

through gated activation control, instantiating a targeted form of *Targeted Optimization*. SAE-based approaches decomposed residual stream states  $\mathbf{x}^l$  into sparse features with activations  $\mathbf{a}$ , enabling feature-level interventions: Muhamed et al. [219] proposed dynamic SAE guardrails that selected and scaled relevant features via *Magnitude Analysis* to achieve precision unlearning with improved forget-utility trade-offs, while Goyal et al. [131] applied *Amplitude Manipulation* to steer toxicity-related SAE features (scaling selected  $\mathbf{f}_j$  via their activations  $a_j$ ) to reduce harmful generations with controlled fluency impact.

## 2) Knowledge Retention and Stability

Retention work traced failures induced by repeated updates or context injection to identifiable carriers, and stabilized behavior via inference-time suppression or training-time adaptation, guided by MI diagnostics. Residual stream states  $\mathbf{x}^l$  and attention head outputs  $\mathbf{h}_{\text{attn}}^{l,h}$  are often key objects of intervention.

- **Conflict Suppression and Mitigation:** Failures under retrieval or context injection often arose from attention heads that mediated the integration of parametric memory and external evidence in the residual stream. Jin et al. [220] performed *Causal Attribution* to localize conflict-mediating heads and applied test-time head suppression/patching, i.e., *Amplitude Manipulation* over attention head output  $\mathbf{h}_{\text{attn}}^{l,h}$  to rebalance memory vs. context usage. Li et al. [221] further used *Magnitude Analysis* to identify heads exhibiting superposition effects and applied targeted gating via *Targeted Optimization* to stabilize behavior under conflicts. Long-context distraction was traced to entrainment-related heads: Niu et al. [136] localized such heads using *Causal Attribution* and ablated or modulated their outputs ( $\mathbf{h}_{\text{attn}}^{l,h}$ ), reducing echoing of irrelevant context tokens. Jin et al. [9] further characterized concentrated massive values in computations mediated by the Q/K weight matrices  $\mathbf{W}_Q^{l,h}$  and  $\mathbf{W}_K^{l,h}$  (reflected in attention scores  $\mathbf{A}^{l,h}$  via *Magnitude Analysis*), then guided *Amplitude Manipulation* over corresponding head outputs  $\mathbf{h}_{\text{attn}}^{l,h}$  to maintain contextual reading without disrupting magnitude-structured signals.
- **Constraining Continual Adaptation:** To reduce catastrophic forgetting, MI localized stability-critical carriers and restricted learning via *Targeted Optimization*. Zhang et al. [74] applied *Gradient Detection* to identify a “core linguistic” parameter region and froze it, mitigating forgetting. Zhang et al. [66] further constrained adaptation through coarse-to-fine module selection and soft masking, balancing specialty and versatility. Representation-level interventions were also employed: Wu et al. [222] localized residual stream states  $\mathbf{x}^l$  and applied lightweight edits on  $\mathbf{x}^l$  with a frozen backbone (a form of *Targeted Optimization*), improving stability relative to weight-centric updates. Monitoring side effects, Du et al. [88] used *Probing* over residual stream states and attention heads to detect security-relevant drift and selected safer module update schedules, enabling controlled adaptation.

## 3) Knowledge Consolidation

Consolidation composes multiple specialized models by combining internal carriers while controlling interference. A common approach represents each fine-tuned model as a parameter “task vector” (a delta from a shared base) and merges these deltas via *Vector Arithmetic*.

Yadav et al. [224] improved multi-model composition over naive averaging by first trimming task vectors (a form of *Magnitude Analysis*), resolving sign conflicts, and then merging consistent update directions. *Sens-Merging* [72] further computed layer-wise sensitivity scores via *Gradient Detection* to weight deltas during merging, yielding stronger merged performance across diverse capability suites. Differently, Yao et al. [142] used *Magnitude Analysis* over layer-specific task vectors to derive importance

scores that modulate merge weights, improving alignment of merged models with dominant capability directions. Beyond parameter deltas, Chen et al. [223] showed that affine mappings between residual stream states (a form of *Probing*) could transfer linear features across models, enabling consolidation at the level of feature bases and amortizing training cost across model sizes.

### 5.2.3. Logic and Reasoning

#### Summary of Application Paradigms

MI enhances the logical deduction and reasoning capabilities of LLMs through three distinct paradigms, moving from structural optimization to dynamic inference control:

- 1) **Specific Refinement of Numerical and Logical Components:** Instead of blindly updating all parameters, this paradigm involves localizing the specific neurons or attention heads responsible for numerical computation and logical operators. These critical carriers are then strengthened via *Targeted Optimization* (§4.2) to improve arithmetic precision.
- 2) **Inference Trajectory Steering:** By isolating directions in the activation space that correspond to high-level reasoning strategies (e.g., “step-by-step” planning), researchers can modulate the model’s cognitive process. This is achieved by injecting steering vectors via *Vector Arithmetic* (§4.3) or amplifying specific features via *Amplitude Manipulation* (§4.1).
- 3) **Stepwise Diagnosis and Correction:** To ensure reliability, monitors based on *Probing* or *Magnitude Analysis* are deployed to track internal states during reasoning. These tools diagnose logical fallacies or uncertainty in real-time, enabling selective self-correction before errors propagate.

#### 1) Specific Refinement of Numerical and Logical Components

LLMs often struggle with precise arithmetic operations. Rather than treating the model as a monolith, MI research has demonstrated that mathematical abilities are often localized within specific sub-modules. Quirke and Barez [225] conducted a granular circuit analysis of modular addition, identifying that specific attention heads and MLP layers form a dedicated algorithm for numerical processing. By characterizing these circuits, they demonstrated that targeted interventions on these specific components could predictably alter the model’s output distribution. Yang et al. [226] analyzed the activation dynamics of CoT processes, revealing that reasoning tasks predominantly activate a broader set of neurons in the final layers compared to standard prompting. Leveraging such insights, Zhang et al. [153] proposed an “identify-analyze-finetune” pipeline. This method first identified “reasoning-critical” attention heads and FFNs via *Causal Attribution*, then froze most model parameters and performed *Targeted Optimization* exclusively on these identified components to boost computational performance. Similarly, Tan et al. [160] decomposed the language model policy into “Internal Layer Policies.” Identifying that early layers maintain high entropy to facilitate exploration, they proposed *Bottom-up Policy Optimization (BuPO)*, a method that selectively optimizes these foundational layers to refine the model’s internal reasoning policy efficiently.

#### 2) Inference Trajectory Steering

Beyond basic arithmetic, complex reasoning requires the adoption of effective strategies. MI methods enable the extraction and injection of these high-level cognitive patterns by manipulating the model’s internal representations via *Vector Arithmetic* and *Amplitude Manipulation*. Researchers have extensively utilized *steering vectors* to modulate reasoning behaviors. Venhoff et al. [227] utilized contrastive activa-

tion means to extract “Backtracking Steering Vectors,” demonstrating that injecting this vector increases the model’s tendency to self-correct. Similarly, Højer et al. [233] and Tang et al. [234] derived control vectors from residual streams to elicit reasoning capabilities; notably, Tang et al. [234] showed that “Long Chain-of-Thought” capabilities can be unlocked via representation engineering without extensive fine-tuning. Hong et al. [235] identified a single linear feature direction that mediates the trade-off between reasoning and memorization, allowing for causal control over the model’s problem-solving mode. Zhang and Viteri [236] discovered “Latent CoT” vectors that, when injected, induce reasoning patterns without explicit natural language prompting. For more granular control, Liu et al. [237] introduced “Fractional Reasoning,” which enables continuous adjustment of reasoning intensity at inference time by scaling latent steering vectors. Efficiency is also a key benefit; Sinii et al. [238] demonstrated that training a single steering vector (bias-only adaptation) matches the reasoning performance of fully RL-tuned models. Taking a different approach, Wang et al. [230] proposed an optimization-based framework: instead of training weights, they optimized hidden representations directly to maximize the likelihood of reasoning paths, utilizing these optimized states to guide the model’s trajectory. Li et al. [239] proposed a dual framework, utilizing SAEs to extract interpretable reasoning features while also introducing an SAE-free algorithm to compute steering directions directly from residual activations. Galichin et al. [3] employed SAEs and introduced “ReasonScore” (a form of *Magnitude Analysis*) to identify sparse features associated with uncertainty and exploratory thinking. By amplifying these features via *Amplitude Manipulation*, they successfully guided the model toward more robust reasoning. Troitskii et al. [229] focused on latent states preceding “wait” tokens. They located specific features that promote or suppress these tokens and showed that modulating them fundamentally alters the subsequent reasoning process. Regarding latent states, Cywiński et al. [232] demonstrated the feasibility of transplanting reasoning patterns. By employing *Causal Attribution* to localize critical latent vectors and subsequently applying patching (a form of *Amplitude Manipulation*), they effectively forced the model to adopt specific latent reasoning paths.

### 3) Stepwise Diagnosis and Correction

A major challenge in multi-step reasoning is error propagation. MI provides tools for real-time internal diagnosis. Sun et al. [231] introduced a *Probing* framework to detect reasoning failures. They trained lightweight classifiers on the model’s internal activations to distinguish between correct and hallucinatory reasoning steps, acting as an internal monitor to flag errors before the final output is generated. Taking a probabilistic perspective, You et al. [87] introduced *ARES*, a framework that employs *Magnitude Analysis* on the entailment probability of internal states. They found that distinct uncertainty patterns emerge when the model deviates from logic. Based on this, they proposed a self-correction mechanism: when the internal monitor detects high “reasoning uncertainty,” the model is triggered to backtrack and regenerate the current step, significantly improving the reliability of long-chain deductions. Finally, Wu et al. [60] employed *Gradient Detection* to compute feature attribution scores for tokens within CoT traces. While primarily analytic, this method serves as a potential diagnostic tool by quantifying the semantic influence of each reasoning step, allowing researchers to verify whether the model is attending to relevant logic or spurious correlations during generation.

### 5.3. Improve Efficiency

#### 5.3.1. Efficient Training

##### Summary of Application Paradigms

MI noticeably enhances training efficiency by shifting model optimization from a “black-box” paradigm to one guided by internal redundant structures and evolutionary dynamics. This application primarily follows two paradigms:

- 1) **Sparse Fine-tuning:** By uncovering the model’s intrinsic sparsity, researchers isolate and update only critical subnetworks via *Targeted Optimization* (§4.2). Unlike standard PEFT methods that introduce external modules, this paradigm modifies model-intrinsic weights, often matching full model fine-tuning performance while drastically reducing computational and memory overhead.
- 2) **Training Dynamics Monitoring:** Leveraging *Magnitude Analysis* (§3.1) and singular learning theory, this paradigm develops internal metrics to track the emergence of specific capabilities and generalization phases. By capturing phase transitions that traditional validation loss may miss, it enables informed decisions on early stopping and prevents unnecessary computations.

#### 1) Sparse Fine-tuning

Unlike PEFT methods that introduce external modules [326–328], achieve efficiency by fine-tuning intrinsic subsets, often matching or exceeding the performance of full fine-tuning. At the neuron granularity, researchers utilize diagnostic tools to pinpoint task-specific units. Zhu et al. [154] proposed the *LANDeRMT* framework, which employs Taylor expansion to evaluate the “awareness score” of FFN neurons for machine translation, enabling *Gradient Detection*-based selective update of language-general and language-specific neurons to mitigate parameter interference. Song et al. [240] introduced *SIFT*, which exploits the “quasi-sparsity” of pre-trained gradients—where the top 1% of components can account for 99% of the total gradient norm—using hook functions to perform memory-efficient in-place sparse updates via *Gradient Detection*. Similarly, Xu et al. [147] developed *NeFT*, which identifies sensitive neurons through *Magnitude Analysis* by calculating the cosine similarity between weights before and after a brief full-parameter fine-tuning run. Furthermore, Mondal et al. [241] and Gurgurov et al. [242] leveraged *Language Activation Probability Entropy* [2] to identify language-sensitive neurons via *Magnitude Analysis*, achieving significant gains by updating less than 1% of the model.

More granular approaches achieve massive efficiency by isolating extremely sparse mechanistic components. Zhao et al. [204] proposed *Parallel Language-specific Neuron Detection*, identifying consistently activated neurons for specific languages without labeled data via *Causal Attribution*; they found that deactivating just 0.13% of these neurons causes a total loss in multilingual generation. Sergeev and Kotelnikov [36] introduced *Head Impact scores* based on *Magnitude Analysis* to identify attention heads, demonstrating that fine-tuning only 0.01% of parameters in the highest-impact layers significantly improves model understanding capability. Lai et al. [218] proposed *JOLA*, a framework that employs HardConcrete gates with expected- $L_0$  regularization to jointly learn which attention heads to edit and whether to apply additive or multiplicative interventions. Furthermore, Li et al. [243] reframed fine-tuning as a “subgraph search” process, introducing a circuit-tuning algorithm that iteratively builds and optimizes a task-relevant circuit via *Circuit Discovery* within the computational graph to preserve general model capabilities.

## 2) Training Dynamic Monitoring

The second paradigm predominantly leverages *Magnitude Analysis* and other quantitative diagnostics to monitor the state evolution of internal objects, addressing the limitations of traditional validation loss in capturing critical phase transitions.

In the context of *Grokking*—where generalization emerges long after overfitting—MI metrics provide crucial signals that enable practitioners to *confidently continue training despite zero progress in validation loss*. Understanding that grokking arises from the competition between fast-learning memorization circuits and slow-learning but efficient generalization circuits [329,330], researchers have developed specific indicators to track the latter’s formation. Nanda et al. [246] proposed *Restricted Loss*, a metric derived by projecting weights onto the Fourier basis, which reveals that structured mechanisms form gradually during the apparent loss plateau. Similarly, Furuta et al. [248] introduced *Fourier Frequency Density (FFD)* to characterize the sparsity of internal representations; tracking *FFD* allows for real-time assessment of generalizability, serving as a reliable proxy for circuit maturation. Moving to early detection, Notsawo Jr et al. [331] analyzed the *spectral signature* of the training loss curve itself, demonstrating that specific low-frequency oscillations in early epochs can effectively predict whether grokking will eventually occur, thus saving computational resources on unpromising runs. In the realm of Mixture-of-Experts (MoE), Li et al. [250] applied *Magnitude Analysis* on router activations and proposed two pathway metrics—similarity and consistency. These metrics monitor how routing patterns evolve from random fluctuations to stable structures, serving as a precise indicator to determine the onset of grokking and enabling optimal early stopping.

Beyond grokking, similar monitoring strategies are applied to the emergence of *In-Context Learning (ICL)*. Hoogland et al. [244] utilized the *Local Learning Coefficient (LLC)* from singular learning theory to quantify the geometry of the loss landscape. They observed that plateaus in the *LLC* curve distinctively mark developmental stages (e.g., from bigram statistics to induction heads), allowing researchers to determine when a model has completed a specific structural transformation. Furthermore, Minegishi et al. [245] extended this to In-Context Meta-Learning, developing circuit-specific metrics such as *label attention scores*. By monitoring the shift in these metrics, they identified that models progress through multiple distinct phases (Non-Context → Semi-Context → Full-Context), providing a granular “progress bar” for the model’s acquisition of meta-learning capabilities that is invisible to standard loss evaluation.

### 5.3.2. Efficient Inference

#### Summary of Application Paradigms

MI facilitates the deployment and acceleration of LLMs—resulting in superior inference efficiency—by identifying and exploiting structural and functional redundancies. This is primarily achieved through two key paradigms:

- 1) **Selective Computation via Saliency Detection:** This paradigm reduces computational overhead by localizing “dispensable” components. At the *data level*, MI helps identify redundant tokens or KV cache entries through *Magnitude Analysis* (§3.1), *Gradient Detection* (§3.3), and *Circuit Discovery* (§3.6), enabling the pinpointing of tokens or heads with minimal importance. At the *model level*, importance metrics based on *Magnitude Analysis* (§3.1) enable the dynamic skipping of redundant layers or Mixture-of-Experts (MoE) experts, facilitating “on-demand” computation.
- 2) **Layer-Specific Adaptive Quantization:** Rather than applying uniform bit-widths, this paradigm leverages mechanistic insights, including *Magnitude Analysis* (§3.1), *Gradient Detection* (§3.3), and *Vocabulary Projection* (§3.5) to assess the quantization sensitivity of different layers. By allocating higher precision to “irreplaceable” layers and applying aggressive compression to more robust ones, models achieve superior memory-accuracy trade-offs, tailored to diverse hardware constraints.

#### 1) Selective Computation via Saliency Detection

The core premise of selective computation is that not all architectural or data components contribute equally to the final output. MI provides a principled tool to quantify such contributions and to prune redundant components accordingly.

- **Data Level:** Researchers have developed advanced token- and KV-cache-level pruning strategies that leverage *Magnitude Analysis* and *Gradient Detection* to effectively identify and remove unimportant tokens. By leveraging *Magnitude Analysis* to identify tokens with minimal contribution to the reasoning process in CoT sequences, *TokenSkip* [159] selectively skips these tokens, achieving substantial compression with negligible performance degradation. Lei et al. [251] explored explanation-driven token compression for multimodal LLMs, where *Gradient Detection* is used to map attention patterns to explanation outcomes, enabling the effective pruning of visual tokens during the input stage. For KV cache-level pruning, *FitPrune* [253] and *ZipCache* [21] employed *Magnitude Analysis* saliency metrics to identify and retain critical KV states. Guo et al. [252] introduced *Value-Aware Token Pruning (VATP)*, which applied *Magnitude Analysis* to attention scores and the L1 norm of value attention vectors to identify crucial tokens. Moving beyond token-wise pruning, *Circuit Discovery* techniques have been applied to identify “Retrieval Heads” that are essential for long-context tasks, enabling non-critical heads to operate with a fixed-length KV cache [155,254,332,333].
- **Model Level:** MI-guided metrics enable the skipping of entire architectural blocks, such as redundant layers, MoE experts, or neurons, thereby facilitating inference acceleration with minimal impact on model performance. Men et al. [14] introduced “Block Influence” (BI), a similarity metric based on *Magnitude Analysis* that compares the input and output of each layer. This technique effectively removes layers with minimal contribution to the representation space. Dynamic bypassing methods, such as *GateSkip* [255] and *LayerSkip* [13], employ learnable residual gates to skip layers during inference, also based on *Magnitude Analysis*. Similarly, *HadSkip* [257] and *SBERT* [258]

models leverage Magnitude Analysis to facilitate effective layer skipping. In MoE architectures, Lu et al. [259] skipped unimportant experts during inference based on the *Magnitude Analysis* of router scores. Su et al. [8] further identified *Super Experts* by analyzing the *Magnitude Analysis* of experts' output activations, showing that these experts are essential for logical reasoning and that pruning them leads to catastrophic performance degradation. Finally, by localizing specialized multilingual neurons [25] and language-specific sub-networks [260] through *Magnitude Analysis* on their activations, LLMs can activate only the sub-circuits necessary for the specific task at hand.

## 2) Layer-Specific Adaptive Quantization

While standard quantization applies a uniform bit-width across all parameters, MI-driven research promotes mixed-precision quantization based on layer-wise "functional saliency." Many of these metrics are based on *Magnitude Analysis* to identify sensitive layers. Dumitru et al. [39] proposed a pragmatic approach to measure layer importance by examining shifts in the embedding space or the presence of weight outliers, assigning higher bit-precision to layers that caused larger representational shifts. Similarly, Zhang et al. [4] introduced *SensiBoost* and *KurtBoost*, which used activation sensitivity and weight distribution kurtosis to identify layers that were "hard-to-quantize," allocating them more memory budget. *LieQ* [261] further uncovered a strong correlation between training-induced energy concentration and representational compactness, providing a geometry-driven sensitivity proxy for automatic bit-width allocation. Beyond static analysis, *Mix-QViT* [262] employed *Layer-wise Relevance Propagation (LRP)*—a form of *Gradient Detection*—to assess the contribution of each layer to the final classification, thereby guiding mixed-precision quantization in vision transformers. *LSAQ* [263] adaptively adjusted quantization strategies in real-time by applying *Vocab Projection* to obtain the vocab distribution for each layer. It then calculated the Jaccard similarity between these distributions to identify sensitive layers, ensuring that they maintained high precision while more robust layers were aggressively compressed to meet the resource constraints of edge devices.

## 6. Challenges and Future Directions

### Challenges

Despite substantial progress and growing methodological sophistication, it remains unclear whether MI is indispensable for any downstream task, rather than serving as an alternative or complementary analysis tool. This uncertainty amplifies the importance of the fundamental challenges discussed below, which continue to limit the scalability, reliability, and practical impact of MI.

First, MI remains difficult to scale beyond low-level components [334,335]. While individual neurons or learned features are increasingly well-characterized [127,336], identifying higher-level computational structures, such as multi-layer interactions, cross-module pathways, or distributed mechanisms, still relies heavily on manual inspection [119,337–340]. Although recent work has made progress toward automation [121,123], current methods often require substantial human intervention and do not robustly generalize across prompts, tasks, or models [129,250,341]. As a result, many MI analyses remain artisanal rather than systematic. In addition to these methodological limitations, computational scalability poses a major bottleneck. Prominent approaches such as SAEs or transcoders rely on training replacement or surrogate models to obtain more interpretable representations, introducing additional training costs that grow with model size and feature dimensionality. This often restricts their application to a limited subset of layers or models. A related challenge arises in fine-grained causal localization. Precisely attributing behavior to individual neurons or SAE features would in principle require exhaustive interventions, but the scale of modern LLMs renders such causal tracing computationally infeasible [44,123]. As a result,

most analyses [122,128,342,343] operate at coarser granularities or rely on heuristic approximations, limiting the resolution at which mechanisms can be reliably identified.

Second, the field lacks robust and widely accepted evaluation frameworks to assess the faithfulness of localization and explanation methods [344]. Although some benchmarks [126,345–348] have been proposed, there remains no consensus on metrics that can determine whether an identified component truly corresponds to the underlying causal mechanism. This issue is particularly acute for methods that rely on surrogate or replacement models, where output-level agreement does not guarantee mechanistic fidelity. Importantly, the scalability constraints discussed above further exacerbate this problem. Because exact fine-grained causal interventions are computationally impractical, researchers must rely on approximate localization methods designed for tractability rather than optimal causal identification. In the absence of reliable ground truth at the mechanism level, it becomes difficult to distinguish true causal components from computationally convenient proxies, making rigorous validation and comparison of MI methods inherently challenging.

Third, current mechanistic analyses often face a fundamental trade-off between sparsity and completeness of representation [140,349]. Many interpretability methods, including SAEs and other sparse decomposition techniques, aim to force the model's internal representations into a small set of monosemantic, easily interpretable components. By promoting sparsity, these methods can disentangle polysemantic neurons and highlight feature directions that correspond to specific concepts, making interpretation more tractable. However, aggressively enforcing sparsity may prune or obscure components that are genuinely part of the true mechanism but do not fit a sparse pattern. This leads to a tension: methods that induce sparsity can improve interpretability but risk overlooking distributed or “inactive” subcomponents of genuine mechanisms, while approaches that preserve dense, distributed representations may be harder to interpret systematically. Accounting for this trade-off, and developing evaluation metrics that balance sparsity, fidelity, and mechanistic completeness, remains an open challenge for MI.

Finally, interventions informed by MI, such as model editing or steering, often lack robustness and predictability [168,350]. Changes intended to modify a specific behavior can introduce unintended side effects on other tasks or domains, raising concerns about generalization and reliability [66,351–357]. For instance, Yu and Ananiadou [109] demonstrate that modifying a very small number of neurons can lead to substantial degradation in overall language performance. The need for accurate target localization and steering methods that avoid collateral behavioral disruption remains a central technical challenge as MI increasingly informs targeted intervention design.

### Future Directions

Looking forward, several directions appear particularly promising for advancing MI. A key priority for mechanistic interpretability is to move from isolated, low-level analyses toward integrated, system-level explanations. Most existing MI work focuses on task-specific and localized mechanisms, such as knowledge neurons, safety-related neurons, arithmetic heads, or specific task circuits for in-context learning or arithmetic [30,151,153,186,358–361]. While informative, these approaches are inherently low-level and offer limited insight into how models organize computation more broadly [280]. In contrast, cognitive science characterizes cognition in terms of higher-level systems, such as System 1 vs. System 2 reasoning [264], as well as attention, memory, language, and executive control systems [362–365]. Comparable system-level accounts in MI remain scarce. Developing such accounts requires frameworks that connect low-level components to higher-order organization, enabling more coherent system-level explanations of LLM computation [47].

In parallel, stronger theoretical foundations are needed. Connecting internal representations to principles from cognitive science [366–368] or information theory [369] may help unify disparate MI findings and reduce reliance on ad-hoc interpretations. A principled framework could also clarify what kinds of internal structures should be expected in large-scale models and why [370].

Finally, an emerging direction is the progression from interpretation to intervention and, ultimately, model design. Insights from MI are increasingly used not only to explain behavior, but also to edit, steer, or modularize models. This direction connects naturally to earlier work on intrinsically interpretable models, such as Concept Bottleneck Models [371–377] and Weight-sparse transformers [349], which enforce transparency through architectural constraints. However, despite their interpretability benefits, such models typically underperform black-box architectures on large-scale, complex tasks [378]. Looking forward, a key challenge is to bridge this gap by designing interpretable backbone architectures that can serve as viable alternatives to transformers, achieving interpretability by construction while maintaining performance comparable to state-of-the-art black-box models. In this sense, interpretability-informed design may move beyond post-hoc analysis toward fundamentally more controllable, customizable, and transparent model architectures.

## 7. Conclusions

In this survey, we systematically reframe MI from a predominantly observational endeavor into a practical, actionable paradigm. By organizing existing methods around the unified pipeline of “Locate, Steer, and Improve”, we clarify how interpretable objects can be precisely localized, causally manipulated, and ultimately leveraged to enhance alignment, capability, and efficiency in LLMs. Our analysis highlights that many recent advances—ranging from safety and persona alignment, to knowledge editing, and further to sparse fine-tuning—are most effective when grounded in explicit mechanistic intervention. We further discuss key challenges and future directions in §6, with the goal of providing a coherent foundation for future research that tightly integrates interpretability, intervention, and model design. Ultimately, we hope this perspective will accelerate the transition toward more powerful, transparent, and reliable LLMs.

## Limitation

This survey focuses on MI for dense LLMs and does not systematically cover methods specific to other architectures and modalities. In particular, Mixture-of-Experts (MoE) models introduce routing mechanisms and sparsely activated experts, while vision–language models and vision-only models rely on modality-specific representations and architectural components that pose distinct interpretability challenges. Nevertheless, many of the methods discussed in this work are conceptually general and, with appropriate adaptation, can be applied to MoE models and multimodal architectures, for example by operating on expert-level activations or modality-specific residual streams. A comprehensive and systematic treatment of these architectures is therefore left to future work.

In addition, the field currently lacks unified benchmarks or standardized evaluation protocols for localization methods, making it difficult to rigorously compare approaches or to assess whether the identified model components are causally optimal. This limitation also affects downstream applications, where interventions often rely on a single localization method without formal guarantees. Some works partially mitigate this issue by combining multiple localization techniques and examining whether they converge on similar model components, but developing principled and reproducible evaluation frameworks remains an open challenge.

## Appendix A Summary of Surveyed Papers

**Table A1.** Summary of Surveyed Papers. We annotate each paper with tags for its Core Interpretable Objects (§2), Localizing Methods (§3), and Steering Methods (§4). For studies employing multiple objects or localizing/steering methods, we annotate the primary tag. The symbol “-” in the Steering Method column denotes works that apply localized mechanistic insights directly for analysis or monitoring, without employing active intervention techniques.

| Paper   | Object          | Localizing Method  | Steering Method        | Venue   | Year | Link                 |
|---|-----------------|--------------------|------------------------|---------|------|----------------------|
| <i>Safety and Reliability (Improve Alignment)</i> |                 |                    |                        |         |      |                      |
| Zhou et al.                                       | MHA             | Causal Attribution | Amplitude Manipulation | ICLR    | 2025 | <a href="#">Link</a> |
| Huang et al.                                      | MHA             | Circuit Discovery  | Targeted Optimization  | EMNLP   | 2025 | <a href="#">Link</a> |
| Jiang et al.                                      | MHA             | Causal Attribution | Targeted Optimization  | ArXiv   | 2024 | <a href="#">Link</a> |
| Chen et al.                                       | Neuron          | Causal Attribution | Amplitude Manipulation | ArXiv   | 2025 | <a href="#">Link</a> |
| Suau et al.                                       | Neuron          | Magnitude Analysis | Amplitude Manipulation | ICML    | 2024 | <a href="#">Link</a> |
| Gao et al.  | Neuron          | Magnitude Analysis | Amplitude Manipulation | ArXiv   | 2025 | <a href="#">Link</a> |
| Zhao et al.                                       | Neuron          | Magnitude Analysis | Targeted Optimization  | ICLR    | 2025 | <a href="#">Link</a> |
| Li et al.   | Neuron          | Magnitude Analysis | Targeted Optimization  | ArXiv   | 2025 | <a href="#">Link</a> |
| Templeton et al.                                  | SAE Feature     | Magnitude Analysis | Amplitude Manipulation | Blog    | 2024 | <a href="#">Link</a> |
| Goyal et al.                                      | SAE Feature     | Magnitude Analysis | Amplitude Manipulation | EMNLP   | 2025 | <a href="#">Link</a> |
| Yeo et al.  | SAE Feature     | Magnitude Analysis | Amplitude Manipulation | EMNLP   | 2025 | <a href="#">Link</a> |
| Li et al.   | SAE Feature     | Magnitude Analysis | Vector Arithmetic      | ArXiv   | 2025 | <a href="#">Link</a> |
| Weng et al.                                       | SAE Feature     | Magnitude Analysis | Amplitude Manipulation | ArXiv   | 2025 | <a href="#">Link</a> |
| Wu et al.   | SAE Feature     | Magnitude Analysis | Vector Arithmetic      | ICML    | 2025 | <a href="#">Link</a> |
| He et al.   | SAE Feature     | Magnitude Analysis | Vector Arithmetic      | ArXiv   | 2025 | <a href="#">Link</a> |
| Li et al.   | Residual Stream | Causal Attribution | Targeted Optimization  | ICLR    | 2025 | <a href="#">Link</a> |
| Lee et al.  | Residual Stream | Probing            | Targeted Optimization  | ICML    | 2024 | <a href="#">Link</a> |
| Arditi et al.                                     | Residual Stream | Causal Attribution | Vector Arithmetic      | NeurIPS | 2024 | <a href="#">Link</a> |
| Zhao et al.                                       | Residual Stream | Causal Attribution | Vector Arithmetic      | NeurIPS | 2025 | <a href="#">Link</a> |
| Yin et al.  | Residual Stream | Probing            | Vector Arithmetic      | ArXiv   | 2025 | <a href="#">Link</a> |
| Ball et al.                                       | Residual Stream | Causal Attribution | Vector Arithmetic      | ArXiv   | 2024 | <a href="#">Link</a> |
| Wang et al.                                       | Residual Stream | Causal Attribution | Vector Arithmetic      | ICLR    | 2025 | <a href="#">Link</a> |
| Wang et al.                                       | Residual Stream | Causal Attribution | Vector Arithmetic      | NeurIPS | 2025 | <a href="#">Link</a> |
| Ferreira et al.                                   | Residual Stream | Causal Attribution | Vector Arithmetic      | ICML    | 2025 | <a href="#">Link</a> |
| Huang et al.                                      | Residual Stream | Causal Attribution | Vector Arithmetic      | ICML    | 2025 | <a href="#">Link</a> |
| Pan et al.  | Residual Stream | Causal Attribution | Vector Arithmetic      | ICML    | 2025 | <a href="#">Link</a> |
| Chuang et al.                                     | Residual Stream | Vocab Projection   | Vector Arithmetic      | ICLR    | 2024 | <a href="#">Link</a> |
| Chen et al.                                       | Residual Stream | Vocab Projection   | Vector Arithmetic      | ICML    | 2024 | <a href="#">Link</a> |
| Zhang et al.                                      | Residual Stream | Probing            | Vector Arithmetic      | ACL     | 2024 | <a href="#">Link</a> |
| Orgad et al.                                      | Residual Stream | Probing            | Vector Arithmetic      | ICLR    | 2025 | <a href="#">Link</a> |
| Stolfo et al.                                     | Residual Stream | Gradient Detection | Vector Arithmetic      | ICLR    | 2025 | <a href="#">Link</a> |
| Du et al.   | Token Embedding | Gradient Detection | Vector Arithmetic      | ArXiv   | 2025 | <a href="#">Link</a> |
| <i>Fairness and Bias (Improve Alignment)</i>      |                 |                    |                        |         |      |                      |
| Vig et al.  | MHA             | Causal Attribution | Amplitude Manipulation | NeurIPS | 2020 | <a href="#">Link</a> |
| Chintam et al.                                    | MHA             | Causal Attribution | Targeted Optimization  | ACLWS   | 2023 | <a href="#">Link</a> |
| Wang et al.                                       | MHA             | Magnitude Analysis | Amplitude Manipulation | ICLR    | 2025 | <a href="#">Link</a> |
| Kim et al.  | MHA             | Probing            | Vector Arithmetic      | ICLR    | 2025 | <a href="#">Link</a> |
| Dimino et al.                                     | MHA             | Magnitude Analysis | -                      | ICAIF   | 2025 | <a href="#">Link</a> |
| Chandna et al.                                    | MHA             | Magnitude Analysis | Amplitude Manipulation | TMLR    | 2025 | <a href="#">Link</a> |

| Paper                                       | Object                 | Localizing Method  | Steering Method               | Venue   | Year | Link                 |
|---|------------------------|--------------------|-------------------------------|---------|------|----------------------|
| Cai et al.                                  | <b>FFN</b>             | Causal Attribution | <i>Targeted Optimization</i>  | ICIC    | 2024 | <a href="#">Link</a> |
| Ahsan et al.                                | <b>FFN</b>             | Causal Attribution | <i>Amplitude Manipulation</i> | EMNLP   | 2025 | <a href="#">Link</a> |
| Li and Gao                                  | <b>FFN</b>             | Vocab Projection   | <i>Targeted Optimization</i>  | ACL     | 2025 | <a href="#">Link</a> |
| Yu and Ananiadou                            | <b>Neuron</b>          | Circuit Discovery  | <i>Targeted Optimization</i>  | ArXiv   | 2025 | <a href="#">Link</a> |
| Liu et al.                                  | <b>Neuron</b>          | Gradient Detection | <i>Amplitude Manipulation</i> | ICLR    | 2024 | <a href="#">Link</a> |
| Yu et al.                                   | <b>Residual Stream</b> | Causal Attribution | -                             | ArXiv   | 2025 | <a href="#">Link</a> |
| Guan et al.                                 | <b>Residual Stream</b> | -                  | <i>Amplitude Manipulation</i> | ICML    | 2025 | <a href="#">Link</a> |
| Yu et al.                                   | <b>Residual Stream</b> | Magnitude Analysis | <i>Amplitude Manipulation</i> | ACL     | 2025 | <a href="#">Link</a> |
| Raimondi et al.                             | <b>Residual Stream</b> | Causal Attribution | <i>Amplitude Manipulation</i> | ArXiv   | 2025 | <a href="#">Link</a> |
| <b>Persona and Role (Improve Alignment)</b> |                        |                    |                               |         |      |                      |
| Su et al.                                   | <b>Neuron</b>          | Causal Attribution | <i>Amplitude Manipulation</i> | EMNLP   | 2025 | <a href="#">Link</a> |
| Deng et al.                                 | <b>Neuron</b>          | Causal Attribution | <i>Amplitude Manipulation</i> | ICLR    | 2025 | <a href="#">Link</a> |
| Lai et al.                                  | <b>Neuron</b>          | Magnitude Analysis | <i>Amplitude Manipulation</i> | EMNLP   | 2024 | <a href="#">Link</a> |
| Chen et al.                                 | <b>Neuron</b>          | Causal Attribution | <i>Targeted Optimization</i>  | ICML    | 2024 | <a href="#">Link</a> |
| Rimsky et al.                               | <b>Residual Stream</b> | Causal Attribution | <i>Vector Arithmetic</i>      | ACL     | 2024 | <a href="#">Link</a> |
| Potertı et al.                              | <b>Residual Stream</b> | Causal Attribution | <i>Vector Arithmetic</i>      | EMNLP   | 2025 | <a href="#">Link</a> |
| Chen et al.                                 | <b>Residual Stream</b> | Causal Attribution | <i>Vector Arithmetic</i>      | ArXiv   | 2025 | <a href="#">Link</a> |
| Handa et al.                                | <b>Residual Stream</b> | Causal Attribution | <i>Vector Arithmetic</i>      | NeurIPS | 2025 | <a href="#">Link</a> |
| Tak et al.                                  | <b>Residual Stream</b> | Probing            | <i>Vector Arithmetic</i>      | ACL     | 2025 | <a href="#">Link</a> |
| Yuan et al.                                 | <b>Residual Stream</b> | Probing            | -                             | ArXiv   | 2025 | <a href="#">Link</a> |
| Ju et al.                                   | <b>Residual Stream</b> | Probing            | <i>Targeted Optimization</i>  | COLM    | 2025 | <a href="#">Link</a> |
| Karny et al.                                | <b>Residual Stream</b> | Causal Attribution | -                             | ArXiv   | 2025 | <a href="#">Link</a> |
| Banayeezade et al.                          | <b>Residual Stream</b> | Causal Attribution | <i>Vector Arithmetic</i>      | ArXiv   | 2025 | <a href="#">Link</a> |
| Bas and Novak                               | <b>Residual Stream</b> | Causal Attribution | <i>Vector Arithmetic</i>      | ArXiv   | 2025 | <a href="#">Link</a> |
| Sun et al.                                  | <b>Residual Stream</b> | Causal Attribution | <i>Vector Arithmetic</i>      | EMNLP   | 2025 | <a href="#">Link</a> |
| Pai et al.                                  | <b>Residual Stream</b> | Causal Attribution | <i>Vector Arithmetic</i>      | ArXiv   | 2025 | <a href="#">Link</a> |
| Joshi et al.                                | <b>Residual Stream</b> | Probing            | -                             | EMNLP   | 2024 | <a href="#">Link</a> |
| Ghandeharioun et al.                        | <b>Residual Stream</b> | Causal Attribution | <i>Vector Arithmetic</i>      | NeurIPS | 2024 | <a href="#">Link</a> |
| <b>Multilingualism (Improve Capability)</b> |                        |                    |                               |         |      |                      |
| Xie et al.                                  | <b>Neuron</b>          | Magnitude Analysis | <i>Amplitude Manipulation</i> | ACL     | 2021 | <a href="#">Link</a> |
| Kojima et al.                               | <b>Neuron</b>          | Magnitude Analysis | <i>Amplitude Manipulation</i> | NAACL   | 2024 | <a href="#">Link</a> |
| Tang et al.                                 | <b>Neuron</b>          | Magnitude Analysis | <i>Amplitude Manipulation</i> | ACL     | 2024 | <a href="#">Link</a> |
| Zhao et al.                                 | <b>Neuron</b>          | Magnitude Analysis | <i>Amplitude Manipulation</i> | NeurIPS | 2024 | <a href="#">Link</a> |
| Gurgurov et al.                             | <b>Neuron</b>          | Magnitude Analysis | <i>Amplitude Manipulation</i> | ArXiv   | 2025 | <a href="#">Link</a> |
| Liu et al.                                  | <b>Neuron</b>          | Magnitude Analysis | <i>Amplitude Manipulation</i> | EMNLP   | 2025 | <a href="#">Link</a> |
| Jing et al.                                 | <b>Neuron</b>          | Magnitude Analysis | <i>Amplitude Manipulation</i> | EMNLP   | 2025 | <a href="#">Link</a> |
| Andrylie et al.                             | <b>SAE Feature</b>     | Magnitude Analysis | <i>Amplitude Manipulation</i> | ArXiv   | 2025 | <a href="#">Link</a> |
| Brinkmann et al.                            | <b>SAE Feature</b>     | Magnitude Analysis | <i>Amplitude Manipulation</i> | NAACL   | 2025 | <a href="#">Link</a> |
| Libovický et al.                            | <b>Residual Stream</b> | Probing            | -                             | EMNLP   | 2020 | <a href="#">Link</a> |
| Chi et al.                                  | <b>Residual Stream</b> | -                  | <i>Vector Arithmetic</i>      | ACL     | 2023 | <a href="#">Link</a> |
| Philippy et al.                             | <b>Residual Stream</b> | Magnitude Analysis | <i>Vector Arithmetic</i>      | ACL     | 2023 | <a href="#">Link</a> |
| Wendler et al.                              | <b>Residual Stream</b> | Vocab Projection   | <i>Vector Arithmetic</i>      | ACL     | 2024 | <a href="#">Link</a> |
| Mousi et al.                                | <b>Residual Stream</b> | Magnitude Analysis | <i>Vector Arithmetic</i>      | ACL     | 2024 | <a href="#">Link</a> |
| Hinck et al.                                | <b>Residual Stream</b> | Probing            | <i>Vector Arithmetic</i>      | EMNLP   | 2024 | <a href="#">Link</a> |
| Zhang et al.                                | <b>Residual Stream</b> | Magnitude Analysis | <i>Vector Arithmetic</i>      | ACL     | 2025 | <a href="#">Link</a> |

| Paper  | Object          | Localizing Method  | Steering Method               | Venue   | Year | Link                 |
|--|-----------------|--------------------|-------------------------------|---------|------|----------------------|
| Wang et al.                                      | Residual Stream | Vocab Projection   | <i>Vector Arithmetic</i>      | ACL     | 2025 | <a href="#">Link</a> |
| Wu et al.  | Residual Stream | Vocab Projection   | -                             | ICLR    | 2025 | <a href="#">Link</a> |
| Wang et al.                                      | Residual Stream | Vocab Projection   | <i>Vector Arithmetic</i>      | EMNLP   | 2025 | <a href="#">Link</a> |
| Nie et al.                                       | Residual Stream | Vocab Projection   | <i>Vector Arithmetic</i>      | EMNLP   | 2025 | <a href="#">Link</a> |
| Liu et al.                                       | Residual Stream | Vocab Projection   | <i>Vector Arithmetic</i>      | EMNLP   | 2025 | <a href="#">Link</a> |
| <b>Knowledge Management (Improve Capability)</b> |                 |                    |                               |         |      |                      |
| Meng et al.                                      | FFN             | Causal Attribution | <i>Targeted Optimization</i>  | NeurIPS | 2022 | <a href="#">Link</a> |
| Meng et al.                                      | FFN             | Causal Attribution | <i>Targeted Optimization</i>  | ICLR    | 2023 | <a href="#">Link</a> |
| Lai et al.                                       | MHA             | Magnitude Analysis | <i>Targeted Optimization</i>  | ICML    | 2025 | <a href="#">Link</a> |
| Li et al.  | MHA             | Magnitude Analysis | <i>Amplitude Manipulation</i> | ICML    | 2025 | <a href="#">Link</a> |
| Jin et al.                                       | MHA             | Magnitude Analysis | <i>Amplitude Manipulation</i> | ICML    | 2025 | <a href="#">Link</a> |
| Jin et al.                                       | MHA             | Causal Attribution | <i>Amplitude Manipulation</i> | ACL     | 2024 | <a href="#">Link</a> |
| Lv et al.  | MHA             | Causal Attribution | <i>Amplitude Manipulation</i> | ArXiv   | 2024 | <a href="#">Link</a> |
| Niu et al.                                       | MHA             | Causal Attribution | <i>Amplitude Manipulation</i> | ACL     | 2025 | <a href="#">Link</a> |
| Zhao et al.                                      | MHA             | Probing            | <i>Targeted Optimization</i>  | EMNLP   | 2025 | <a href="#">Link</a> |
| Yadav et al.                                     | FFN & MHA       | Magnitude Analysis | <i>Vector Arithmetic</i>      | NeurIPS | 2023 | <a href="#">Link</a> |
| Yu and Ananiadou                                 | FFN & MHA       | Magnitude Analysis | <i>Amplitude Manipulation</i> | EMNLP   | 2024 | <a href="#">Link</a> |
| Zhang et al.                                     | FFN & MHA       | Magnitude Analysis | <i>Targeted Optimization</i>  | ACL     | 2024 | <a href="#">Link</a> |
| Chen et al.                                      | FFN & MHA       | Magnitude Analysis | <i>Amplitude Manipulation</i> | ICLR    | 2025 | <a href="#">Link</a> |
| Li et al.  | FFN & MHA       | Magnitude Analysis | <i>Targeted Optimization</i>  | AAAI    | 2025 | <a href="#">Link</a> |
| Muhamed and Smith                                | FFN & MHA       | Magnitude Analysis | -                             | ICML    | 2025 | <a href="#">Link</a> |
| Yao et al.                                       | FFN & MHA       | Circuit Discovery  | <i>Amplitude Manipulation</i> | NeurIPS | 2024 | <a href="#">Link</a> |
| Du et al.  | FFN & MHA       | Probing            | <i>Targeted Optimization</i>  | ArXiv   | 2024 | <a href="#">Link</a> |
| Zhang et al.                                     | FFN & MHA       | Gradient Detection | <i>Targeted Optimization</i>  | ACL     | 2024 | <a href="#">Link</a> |
| Liu et al.                                       | FFN & MHA       | Gradient Detection | <i>Vector Arithmetic</i>      | ACL     | 2025 | <a href="#">Link</a> |
| Yao et al.                                       | FFN & MHA       | Magnitude Analysis | <i>Vector Arithmetic</i>      | NeurIPS | 2025 | <a href="#">Link</a> |
| Geva et al.                                      | FFN & MHA       | Causal Attribution | -                             | EMNLP   | 2023 | <a href="#">Link</a> |
| Zhang et al.                                     | Neuron          | Magnitude Analysis | <i>Targeted Optimization</i>  | COLING  | 2025 | <a href="#">Link</a> |
| Chen et al.                                      | Neuron          | Gradient Detection | <i>Amplitude Manipulation</i> | AAAI    | 2024 | <a href="#">Link</a> |
| Shi et al.                                       | Neuron          | Gradient Detection | <i>Amplitude Manipulation</i> | NeurIPS | 2024 | <a href="#">Link</a> |
| Chen et al.                                      | Neuron          | Gradient Detection | <i>Amplitude Manipulation</i> | AAAI    | 2025 | <a href="#">Link</a> |
| Kassem et al.                                    | Neuron          | -                  | <i>Amplitude Manipulation</i> | EMNLP   | 2025 | <a href="#">Link</a> |
| Muhamed et al.                                   | SAE Feature     | Magnitude Analysis | <i>Amplitude Manipulation</i> | ICML    | 2025 | <a href="#">Link</a> |
| Goyal et al.                                     | SAE Feature     | Magnitude Analysis | <i>Amplitude Manipulation</i> | EMNLP   | 2025 | <a href="#">Link</a> |
| Marks et al.                                     | SAE Feature     | Circuit Discovery  | <i>Amplitude Manipulation</i> | ICLR    | 2025 | <a href="#">Link</a> |
| Kang and Choi                                    | Residual Stream | Probing            | -                             | EMNLP   | 2023 | <a href="#">Link</a> |
| Katz et al.                                      | Residual Stream | Vocab Projection   | <i>Targeted Optimization</i>  | EMNLP   | 2024 | <a href="#">Link</a> |
| Wu et al.  | Residual Stream | Causal Attribution | <i>Targeted Optimization</i>  | NeurIPS | 2024 | <a href="#">Link</a> |
| Zhao et al.                                      | Residual Stream | Probing            | -                             | ArXiv   | 2024 | <a href="#">Link</a> |
| Ju et al.  | Residual Stream | Probing            | -                             | COLING  | 2024 | <a href="#">Link</a> |
| Jin et al.                                       | Residual Stream | Probing            | -                             | COLING  | 2025 | <a href="#">Link</a> |
| Chen et al.                                      | Residual Stream | Probing            | <i>Vector Arithmetic</i>      | NeurIPS | 2025 | <a href="#">Link</a> |
| <b>Logic and Reasoning (Improve Capability)</b>  |                 |                    |                               |         |      |                      |
| Wu et al.  | Token Embedding | Gradient Detection | -                             | ICML    | 2023 | <a href="#">Link</a> |
| You et al.                                       | Token Embedding | Magnitude Analysis | -                             | EMNLP   | 2025 | <a href="#">Link</a> |
| Cywiński et al.                                  | Token Embedding | Causal Attribution | <i>Amplitude Manipulation</i> | Blog    | 2025 | <a href="#">Link</a> |

| Paper   | Object          | Localizing Method  | Steering Method               | Venue      | Year | Link                 |
|---|-----------------|--------------------|-------------------------------|------------|------|----------------------|
| Cywiński et al.                                       | Token Embedding | Causal Attribution | <i>Amplitude Manipulation</i> | Blog       | 2025 | <a href="#">Link</a> |
| Wang et al.   | FFN             | Magnitude Analysis | <i>Amplitude Manipulation</i> | ArXiv      | 2025 | <a href="#">Link</a> |
| Yu and Ananiadou                                      | MHA             | Magnitude Analysis | <i>Amplitude Manipulation</i> | EMNLP      | 2024 | <a href="#">Link</a> |
| Zhang et al.  | MHA             | Causal Attribution | <i>Targeted Optimization</i>  | ICML       | 2024 | <a href="#">Link</a> |
| Yu and Ananiadou                                      | MHA             | Causal Attribution | <i>Amplitude Manipulation</i> | EMNLP      | 2024 | <a href="#">Link</a> |
| Yu et al.   | MHA             | Causal Attribution | -                             | EMNLP      | 2025 | <a href="#">Link</a> |
| Stolfo et al.   | FFN & MHA       | Causal Attribution | -                             | EMNLP      | 2023 | <a href="#">Link</a> |
| Akter et al.  | FFN & MHA       | Causal Attribution | -                             | COMPSA2024 | 2024 | <a href="#">Link</a> |
| Yang et al.   | FFN & MHA       | Magnitude Analysis | -                             | ArXiv      | 2024 | <a href="#">Link</a> |
| Quirke and Barez                                      | FFN & MHA       | Causal Attribution | <i>Amplitude Manipulation</i> | ICLR       | 2024 | <a href="#">Link</a> |
| Chen et al.   | FFN & MHA       | Gradient Detection | <i>Targeted Optimization</i>  | ACL        | 2025 | <a href="#">Link</a> |
| Hanna et al.  | FFN & MHA       | Circuit Discovery  | -                             | NeurIPS    | 2023 | <a href="#">Link</a> |
| Nikankin et al.                                       | FFN & MHA       | Circuit Discovery  | -                             | ICLR       | 2025 | <a href="#">Link</a> |
| Galichin et al.                                       | SAE Feature     | Magnitude Analysis | <i>Vector Arithmetic</i>      | ArXiv      | 2025 | <a href="#">Link</a> |
| Pach et al.   | SAE Feature     | Magnitude Analysis | <i>Amplitude Manipulation</i> | ArXiv      | 2025 | <a href="#">Link</a> |
| Troitskii et al.                                      | SAE Feature     | Magnitude Analysis | <i>Amplitude Manipulation</i> | EMNLP      | 2025 | <a href="#">Link</a> |
| Venhoff et al.  | Residual Stream | Causal Attribution | <i>Vector Arithmetic</i>      | ICLR       | 2025 | <a href="#">Link</a> |
| Højer et al.  | Residual Stream | Causal Attribution | <i>Vector Arithmetic</i>      | ICLR       | 2025 | <a href="#">Link</a> |
| Tang et al.   | Residual Stream | Causal Attribution | <i>Vector Arithmetic</i>      | ACL        | 2025 | <a href="#">Link</a> |
| Hong et al.   | Residual Stream | Causal Attribution | <i>Vector Arithmetic</i>      | ACL        | 2025 | <a href="#">Link</a> |
| Zhang and Viteri                                      | Residual Stream | Causal Attribution | <i>Vector Arithmetic</i>      | ICLR       | 2025 | <a href="#">Link</a> |
| Liu et al.  | Residual Stream | Causal Attribution | <i>Vector Arithmetic</i>      | ArXiv      | 2025 | <a href="#">Link</a> |
| Sinii et al.  | Residual Stream | Causal Attribution | <i>Vector Arithmetic</i>      | EMNLP      | 2025 | <a href="#">Link</a> |
| Li et al.   | Residual Stream | Causal Attribution | <i>Vector Arithmetic</i>      | EMNLP      | 2025 | <a href="#">Link</a> |
| Ward et al.   | Residual Stream | Causal Attribution | <i>Vector Arithmetic</i>      | ICML       | 2025 | <a href="#">Link</a> |
| Biran et al.  | Residual Stream | Probing            | -                             | EMNLP      | 2024 | <a href="#">Link</a> |
| Ye et al.   | Residual Stream | Probing            | -                             | ICLR       | 2025 | <a href="#">Link</a> |
| Sun et al.  | Residual Stream | Probing            | -                             | EMNLP      | 2025 | <a href="#">Link</a> |
| Wang et al.   | Residual Stream | Probing            | <i>Vector Arithmetic</i>      | AAAI       | 2026 | <a href="#">Link</a> |
| Tan et al.  | Residual Stream | Vocab Projection   | <i>Targeted Optimization</i>  | ArXiv      | 2025 | <a href="#">Link</a> |
| <b><i>Efficient Training (Improve Efficiency)</i></b> |                 |                    |                               |            |      |                      |
| Panigrahi et al.                                      | Neuron          | Magnitude Analysis | <i>Targeted Optimization</i>  | ICML       | 2023 | <a href="#">Link</a> |
| Zhu et al.  | Neuron          | Gradient Detection | <i>Targeted Optimization</i>  | ACL        | 2024 | <a href="#">Link</a> |
| Song et al.   | Neuron          | Gradient Detection | <i>Targeted Optimization</i>  | ICML       | 2024 | <a href="#">Link</a> |
| Zhang et al.  | Neuron          | Magnitude Analysis | <i>Targeted Optimization</i>  | ACL        | 2023 | <a href="#">Link</a> |
| Xu et al.   | Neuron          | Magnitude Analysis | <i>Targeted Optimization</i>  | COLING     | 2025 | <a href="#">Link</a> |
| Mondal et al.   | Neuron          | Magnitude Analysis | <i>Targeted Optimization</i>  | ACL        | 2025 | <a href="#">Link</a> |
| Gurgurov et al.                                       | Neuron          | Magnitude Analysis | <i>Targeted Optimization</i>  | AAACL      | 2025 | <a href="#">Link</a> |
| Zhao et al.   | Neuron          | Causal Attribution | <i>Targeted Optimization</i>  | NeurIPS    | 2024 | <a href="#">Link</a> |
| Li et al.   | Neuron          | Magnitude Analysis | -                             | ArXiv      | 2025 | <a href="#">Link</a> |
| Sergeev and Kotelnikov                                | MHA             | Magnitude Analysis | <i>Targeted Optimization</i>  | ICAI       | 2025 | <a href="#">Link</a> |
| Olsson et al.   | MHA             | Magnitude Analysis | -                             | ArXiv      | 2022 | <a href="#">Link</a> |
| Wang et al.   | MHA             | Magnitude Analysis | -                             | ArXiv      | 2024 | <a href="#">Link</a> |
| Singh et al.  | MHA             | Magnitude Analysis | -                             | ICML       | 2024 | <a href="#">Link</a> |
| Hoogland et al.                                       | MHA             | Magnitude Analysis | -                             | TLMR       | 2025 | <a href="#">Link</a> |
| Minegishi et al.                                      | MHA             | Magnitude Analysis | -                             | ICLR       | 2025 | <a href="#">Link</a> |

| Paper   | Object          | Localizing Method  | Steering Method               | Venue   | Year | Link                 |
|---|-----------------|--------------------|-------------------------------|---------|------|----------------------|
| Lai et al.                                      | MHA             | Magnitude Analysis | <i>Vector Arithmetic</i>      | ICML    | 2025 | <a href="#">Link</a> |
| Thilak et al.                                   | FFN & MHA       | Magnitude Analysis | -                             | NeurIPS | 2022 | <a href="#">Link</a> |
| Varma et al.                                    | FFN & MHA       | Magnitude Analysis | -                             | ArXiv   | 2023 | <a href="#">Link</a> |
| Furuta et al.                                   | FFN & MHA       | Magnitude Analysis | -                             | TMLR    | 2024 | <a href="#">Link</a> |
| Nanda et al.                                    | FFN & MHA       | Magnitude Analysis | -                             | ICLR    | 2023 | <a href="#">Link</a> |
| Notsawo Jr et al.                               | FFN & MHA       | Magnitude Analysis | -                             | ArXiv   | 2023 | <a href="#">Link</a> |
| Qiyi et al.                                     | FFN & MHA       | Magnitude Analysis | -                             | ArXiv   | 2024 | <a href="#">Link</a> |
| Liu et al.                                      | FFN & MHA       | Magnitude Analysis | -                             | ICLR    | 2023 | <a href="#">Link</a> |
| Wang et al.                                     | FFN & MHA       | Magnitude Analysis | -                             | NeurIPS | 2024 | <a href="#">Link</a> |
| Huang et al.                                    | FFN & MHA       | Magnitude Analysis | -                             | COLM    | 2024 | <a href="#">Link</a> |
| Li et al.                                       | FFN & MHA       | Circuit Discovery  | <i>Targeted Optimization</i>  | ArXiv   | 2025 | <a href="#">Link</a> |
| <i>Efficient Inference (Improve Efficiency)</i> |                 |                    |                               |         |      |                      |
| Xia et al.                                      | Token Embedding | Magnitude Analysis | <i>Amplitude Manipulation</i> | EMNLP   | 2025 | <a href="#">Link</a> |
| Lei et al.                                      | Token Embedding | Gradient Detection | <i>Amplitude Manipulation</i> | ArXiv   | 2025 | <a href="#">Link</a> |
| Guo et al.                                      | Token Embedding | Magnitude Analysis | <i>Amplitude Manipulation</i> | EMNLP   | 2024 | <a href="#">Link</a> |
| Ye et al.                                       | Token Embedding | Magnitude Analysis | <i>Amplitude Manipulation</i> | AAAI    | 2025 | <a href="#">Link</a> |
| He et al.                                       | Token Embedding | Magnitude Analysis | <i>Amplitude Manipulation</i> | NeurIPS | 2024 | <a href="#">Link</a> |
| Cai et al.                                      | Token Embedding | Magnitude Analysis | <i>Amplitude Manipulation</i> | COLM    | 2025 | <a href="#">Link</a> |
| Tang et al.                                     | MHA             | Circuit Discovery  | <i>Amplitude Manipulation</i> | ICLR    | 2025 | <a href="#">Link</a> |
| Xiao et al.                                     | MHA             | Circuit Discovery  | <i>Amplitude Manipulation</i> | ICLR    | 2025 | <a href="#">Link</a> |
| Bi et al.                                       | MHA             | Magnitude Analysis | -                             | CVPR    | 2025 | <a href="#">Link</a> |
| Su et al.                                       | MHA             | Magnitude Analysis | <i>Amplitude Manipulation</i> | IJCAI   | 2025 | <a href="#">Link</a> |
| Xiao et al.                                     | MHA             | Magnitude Analysis | <i>Amplitude Manipulation</i> | ICLR    | 2024 | <a href="#">Link</a> |
| Lu et al.                                       | FFN             | Magnitude Analysis | <i>Amplitude Manipulation</i> | ACL     | 2024 | <a href="#">Link</a> |
| Su et al.                                       | FFN             | Magnitude Analysis | <i>Amplitude Manipulation</i> | ArXiv   | 2025 | <a href="#">Link</a> |
| Yu et al.                                       | FFN             | Magnitude Analysis | <i>Amplitude Manipulation</i> | Arxiv   | 2024 | <a href="#">Link</a> |
| Liu et al.                                      | Neuron          | Magnitude Analysis | <i>Amplitude Manipulation</i> | ArXiv   | 2024 | <a href="#">Link</a> |
| Tan et al.                                      | Neuron          | Magnitude Analysis | -                             | EMNLP   | 2024 | <a href="#">Link</a> |
| Laitenberger et al.                             | Residual Stream | Magnitude Analysis | <i>Amplitude Manipulation</i> | ArXiv   | 2025 | <a href="#">Link</a> |
| Valade  | Residual Stream | Probing            | <i>Amplitude Manipulation</i> | ArXiv   | 2024 | <a href="#">Link</a> |
| Elhoushi et al.                                 | Residual Stream | Probing            | <i>Amplitude Manipulation</i> | ACL     | 2024 | <a href="#">Link</a> |
| Wang et al.                                     | Residual Stream | Magnitude Analysis | <i>Amplitude Manipulation</i> | EMNLP   | 2023 | <a href="#">Link</a> |
| Lawson and Aitchison                            | Residual Stream | Magnitude Analysis | <i>Amplitude Manipulation</i> | ArXiv   | 2025 | <a href="#">Link</a> |
| Men et al.                                      | Residual Stream | Magnitude Analysis | <i>Amplitude Manipulation</i> | ACL     | 2025 | <a href="#">Link</a> |
| Dumitru et al.                                  | Residual Stream | Magnitude Analysis | -                             | ArXiv   | 2024 | <a href="#">Link</a> |
| Zhang et al.                                    | Residual Stream | Magnitude Analysis | -                             | ArXiv   | 2025 | <a href="#">Link</a> |
| Xiao et al.                                     | Residual Stream | Magnitude Analysis | -                             | ArXiv   | 2025 | <a href="#">Link</a> |
| Ranjan and Savakis                              | Residual Stream | Gradient Detection | -                             | ArXiv   | 2025 | <a href="#">Link</a> |
| Zeng et al.                                     | Residual Stream | Vocab Projection   | -                             | ArXiv   | 2024 | <a href="#">Link</a> |
| Shelke et al.                                   | Residual Stream | Magnitude Analysis | <i>Amplitude Manipulation</i> | ACL     | 2024 | <a href="#">Link</a> |
| Lin et al.                                      | FFN & MHA       | Magnitude Analysis | <i>Amplitude Manipulation</i> | MLSyS   | 2024 | <a href="#">Link</a> |
| Ashkboos et al.                                 | FFN & MHA       | Magnitude Analysis | <i>Amplitude Manipulation</i> | NeurIPS | 2025 | <a href="#">Link</a> |
| Su and Yuan                                     | FFN & MHA       | Circuit Discovery  | -                             | COLM    | 2025 | <a href="#">Link</a> |
| Xiao et al.                                     | FFN & MHA       | Magnitude Analysis | <i>Amplitude Manipulation</i> | NeurIPS | 2022 | <a href="#">Link</a> |
| Sun et al.                                      | FFN & MHA       | Magnitude Analysis | -                             | NeurIPS | 2024 | <a href="#">Link</a> |

| Paper             | Object    | Localizing Method | Steering Method | Venue   | Year | Link                 |
|-------------------|-----------|-------------------|-----------------|---------|------|----------------------|
| An et al.         | FFN & MHA | Circuit Discovery | -               | ICLR    | 2025 | <a href="#">Link</a> |
| Bondarenko et al. | FFN & MHA | Circuit Discovery | -               | NeurIPS | 2023 | <a href="#">Link</a> |

## References

1. Dettmers, T.; Lewis, M.; Belkada, Y.; Zettlemoyer, L. Gpt3. int8 (): 8-bit matrix multiplication for transformers at scale. *Advances in neural information processing systems* **2022**, *35*, 30318–30332.
2. Tang, T.; Luo, W.; Huang, H.; Zhang, D.; Wang, X.; Zhao, X.; Wei, F.; Wen, J.R. Language-Specific Neurons: The Key to Multilingual Capabilities in Large Language Models. In Proceedings of the Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers); Ku, L.W.; Martins, A.; Srikumar, V., Eds., Bangkok, Thailand, 2024; pp. 5701–5715. <https://doi.org/10.18653/v1/2024.acl-long.309>.
3. Galichin, A.; Dontsov, A.; Druzhinina, P.; Razzhigaev, A.; Rogov, O.Y.; Tutubalina, E.; Oseledets, I. I Have Covered All the Bases Here: Interpreting Reasoning Features in Large Language Models via Sparse Autoencoders. *arXiv preprint arXiv:2503.18878* **2025**.
4. Zhang, F.; Liu, Y.; Li, W.; Lv, J.; Wang, X.; Bai, Q. Towards Superior Quantization Accuracy: A Layer-sensitive Approach. *arXiv preprint arXiv:2503.06518* **2025**.
5. An, Y.; Zhao, X.; Yu, T.; Tang, M.; Wang, J. Systematic outliers in large language models. *arXiv preprint arXiv:2502.06415* **2025**.
6. Su, Z.; Yuan, K. KVSink: Understanding and Enhancing the Preservation of Attention Sinks in KV Cache Quantization for LLMs. *arXiv preprint arXiv:2508.04257* **2025**.
7. Huben, R.; Cunningham, H.; Smith, L.R.; Ewart, A.; Sharkey, L. Sparse Autoencoders Find Highly Interpretable Features in Language Models. In Proceedings of the The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024. OpenReview.net, 2024.
8. Su, Z.; Li, Q.; Zhang, H.; Qian, Y.; Xie, Y.; Yuan, K. Unveiling super experts in mixture-of-experts large language models. *arXiv preprint arXiv:2507.23279* **2025**.
9. Jin, M.; Mei, K.; Xu, W.; Sun, M.; Tang, R.; Du, M.; Liu, Z.; Zhang, Y. Massive Values in Self-Attention Modules are the Key to Contextual Knowledge Understanding. In Proceedings of the Forty-second International Conference on Machine Learning, ICML 2025, Vancouver, BC, Canada, July 13-19, 2025. OpenReview.net, 2025.
10. Bi, J.; Guo, J.; Tang, Y.; Wen, L.B.; Liu, Z.; Wang, B.; Xu, C. Unveiling visual perception in language models: An attention head analysis approach. In Proceedings of the Proceedings of the Computer Vision and Pattern Recognition Conference, 2025, pp. 4135–4144.
11. Chuang, Y.S.; Xie, Y.; Luo, H.; Kim, Y.; Glass, J.R.; He, P. DoLa: Decoding by Contrasting Layers Improves Factuality in Large Language Models. In Proceedings of the The Twelfth International Conference on Learning Representations, 2024.
12. Zhang, S.; Yu, T.; Feng, Y. TruthX: Alleviating Hallucinations by Editing Large Language Models in Truthful Space. In Proceedings of the Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers); Ku, L.W.; Martins, A.; Srikumar, V., Eds., Bangkok, Thailand, 2024; pp. 8908–8949. <https://doi.org/10.18653/v1/2024.acl-long.483>.
13. Elhoushi, M.; Shrivastava, A.; Liskovich, D.; Hosmer, B.; Wasti, B.; Lai, L.; Mahmoud, A.; Acun, B.; Agarwal, S.; Roman, A.; et al. LayerSkip: Enabling Early Exit Inference and Self-Speculative Decoding. In Proceedings of the Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers); Ku, L.W.; Martins, A.; Srikumar, V., Eds., Bangkok, Thailand, 2024; pp. 12622–12642. <https://doi.org/10.18653/v1/2024.acl-long.681>.
14. Men, X.; Xu, M.; Zhang, Q.; Yuan, Q.; Wang, B.; Lin, H.; Lu, Y.; Han, X.; Chen, W. ShortGPT: Layers in Large Language Models are More Redundant Than You Expect. In Proceedings of the Findings of the Association for Computational Linguistics: ACL 2025; Che, W.; Nabende, J.; Shutova, E.; Pilehvar, M.T., Eds., Vienna, Austria, 2025; pp. 20192–20204. <https://doi.org/10.18653/v1/2025.findings-acl.1035>.

15. Xiao, G.; Lin, J.; Seznec, M.; Wu, H.; Demouth, J.; Han, S. Smoothquant: Accurate and efficient post-training quantization for large language models. In Proceedings of the International conference on machine learning. PMLR, 2023, pp. 38087–38099.
16. Ashkboos, S.; Mohtashami, A.; Croci, M.L.; Li, B.; Cameron, P.; Jaggi, M.; Alistarh, D.; Hoefler, T.; Hensman, J. Quarot: Outlier-free 4-bit inference in rotated llms. *Advances in Neural Information Processing Systems* **2024**, *37*, 100213–100240.
17. Yu, M.; Wang, D.; Shan, Q.; Reed, C.J.; Wan, A. The super weight in large language models. *arXiv preprint arXiv:2411.07191* **2024**.
18. Cai, Z.; Zhang, Y.; Gao, B.; Liu, Y.; Li, Y.; Liu, T.; Lu, K.; Xiong, W.; Dong, Y.; Hu, J.; et al. Pyramidkv: Dynamic kv cache compression based on pyramidal information funneling. *arXiv preprint arXiv:2406.02069* **2024**.
19. Xiong, J.; Fan, L.; Shen, H.; Su, Z.; Yang, M.; Kong, L.; Wong, N. DoPE: Denoising Rotary Position Embedding. *arXiv preprint arXiv:2511.09146* **2025**.
20. Xiong, J.; Chen, Q.; Ye, F.; Wan, Z.; Zheng, C.; Zhao, C.; Shen, H.; Li, A.H.; Tao, C.; Tan, H.; et al. ATTS: Asynchronous Test-Time Scaling via Conformal Prediction. *arXiv preprint arXiv:2509.15148* **2025**.
21. He, Y.; Zhang, L.; Wu, W.; Liu, J.; Zhou, H.; Zhuang, B. Zipcache: Accurate and efficient kv cache quantization with salient token identification. *Advances in Neural Information Processing Systems* **2024**, *37*, 68287–68307.
22. Su, Z.; Chen, Z.; Shen, W.; Wei, H.; Li, L.; Yu, H.; Yuan, K. Rotatekv: Accurate and robust 2-bit kv cache quantization for llms via outlier-aware adaptive rotations. *arXiv preprint arXiv:2501.16383* **2025**.
23. Yuan, J.; Gao, H.; Dai, D.; Luo, J.; Zhao, L.; Zhang, Z.; Xie, Z.; Wei, Y.; Wang, L.; Xiao, Z.; et al. Native sparse attention: Hardware-aligned and natively trainable sparse attention. In Proceedings of the Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), 2025, pp. 23078–23097.
24. Lai, W.; Hangya, V.; Fraser, A. Style-Specific Neurons for Steering LLMs in Text Style Transfer. In Proceedings of the Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing; Al-Onaizan, Y.; Bansal, M.; Chen, Y.N., Eds., Miami, Florida, USA, 2024; pp. 13427–13443. <https://doi.org/10.18653/v1/2024.emnlp-main.745>.
25. Liu, W.; Xu, Y.; Xu, H.; Chen, J.; Hu, X.; Wu, J. Unraveling babel: Exploring multilingual activation patterns within large language models. *arXiv* **2024**.
26. Chen, R.; Hu, T.; Feng, Y.; Liu, Z. Learnable Privacy Neurons Localization in Language Models. In Proceedings of the Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers); Ku, L.W.; Martins, A.; Srikumar, V., Eds., Bangkok, Thailand, 2024; pp. 256–264. <https://doi.org/10.18653/v1/2024.acl-short.25>.
27. Chen, L.; Dejl, A.; Toni, F. Identifying Query-Relevant Neurons in Large Language Models for Long-Form Texts. In Proceedings of the AAIL-25, Sponsored by the Association for the Advancement of Artificial Intelligence, February 25 - March 4, 2025, Philadelphia, PA, USA; Walsh, T.; Shah, J.; Kolter, Z., Eds. AAAI Press, 2025, pp. 23595–23604. <https://doi.org/10.1609/AAAI.V39I22.34529>.
28. Wang, S.; Lei, Z.; Tan, Z.; Ding, J.; Zhao, X.; Dong, Y.; Wu, G.; Chen, T.; Chen, C.; Zhang, A.; et al. BrainMAP: Learning Multiple Activation Pathways in Brain Networks. In Proceedings of the AAIL-25, Sponsored by the Association for the Advancement of Artificial Intelligence, February 25 - March 4, 2025, Philadelphia, PA, USA; Walsh, T.; Shah, J.; Kolter, Z., Eds. AAAI Press, 2025, pp. 14432–14440. <https://doi.org/10.1609/AAAI.V39I13.33581>.
29. Andrylie, L.M.; Rahmanisa, I.; Ihsani, M.K.; Wicaksono, A.F.; Wibowo, H.A.; Aji, A.F. Sparse Autoencoders Can Capture Language-Specific Concepts Across Diverse Languages, 2025, [[arXiv:cs.CL/2507.11230](https://arxiv.org/abs/2507.11230)].
30. Gurgurov, D.; Trinley, K.; Ghussin, Y.A.; Baemel, T.; van Genabith, J.; Ostermann, S. Language Arithmetics: Towards Systematic Language Neuron Identification and Manipulation, 2025, [[arXiv:cs.CL/2507.22608](https://arxiv.org/abs/2507.22608)].
31. Xiao, G.; Tian, Y.; Chen, B.; Han, S.; Lewis, M. Efficient Streaming Language Models with Attention Sinks. *arXiv* **2023**.

32. Cancedda, N. Spectral Filters, Dark Signals, and Attention Sinks. In Proceedings of the Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), 2024, pp. 4792–4808.
33. Singh, A.K.; Moskovitz, T.; Hill, F.; Chan, S.C.; Saxe, A.M. What needs to go right for an induction head? a mechanistic study of in-context learning circuits and their formation. In Proceedings of the Proceedings of the 41st International Conference on Machine Learning, 2024, pp. 45637–45662.
34. Wang, M.; Yu, R.; Wu, L.; et al. How Transformers Implement Induction Heads: Approximation and Optimization Analysis. *arXiv e-prints* **2024**, pp. arXiv–2410.
35. Zhou, Z.; Yu, H.; Zhang, X.; Xu, R.; Huang, F.; Wang, K.; Liu, Y.; Fang, J.; Li, Y. On the Role of Attention Heads in Large Language Model Safety, 2025, [arXiv:cs.CL/2410.13708].
36. Sergeev, A.; Kotelnikov, E. Optimizing Multimodal Language Models through Attention-based Interpretability, 2025, [arXiv:cs.CL/2511.23375].
37. Sun, S.; Baek, S.Y.; Kim, J.H. Personality Vector: Modulating Personality of Large Language Models by Model Merging. In Proceedings of the Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing; Christodoulopoulos, C.; Chakraborty, T.; Rose, C.; Peng, V., Eds., Suzhou, China, 2025; pp. 24667–24688. <https://doi.org/10.18653/v1/2025.emnlp-main.1253>.
38. Bas, T.; Novak, K. Steering Latent Traits, Not Learned Facts: An Empirical Study of Activation Control Limits. *arXiv preprint arXiv:2511.18284* **2025**.
39. Dumitru, R.G.; Yadav, V.; Maheshwary, R.; Clotan, P.I.; Madhusudhan, S.T.; Surdeanu, M. Layer-wise quantization: A pragmatic and effective method for quantizing llms beyond integer bit-levels. *arXiv preprint arXiv:2406.17415* **2024**.
40. Tan, Z.; Dong, D.; Zhao, X.; Peng, J.; Cheng, Y.; Chen, T. Dlo: Dynamic layer operation for efficient vertical scaling of llms. *arXiv preprint arXiv:2407.11030* **2024**.
41. Lawson, T.; Aitchison, L. Learning to Skip the Middle Layers of Transformers, 2025, [arXiv:cs.LG/2506.21103].
42. Vig, J.; Gehrmann, S.; Belinkov, Y.; Qian, S.; Nevo, D.; Singer, Y.; Shieber, S. Investigating Gender Bias in Language Models Using Causal Mediation Analysis. In Proceedings of the Advances in Neural Information Processing Systems; Larochelle, H.; Ranzato, M.; Hadsell, R.; Balcan, M.; Lin, H., Eds. Curran Associates, Inc., 2020, Vol. 33, pp. 12388–12401.
43. Meng, K.; Bau, D.; Andonian, A.; Belinkov, Y. Locating and Editing Factual Associations in GPT. In Proceedings of the Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022; Koyejo, S.; Mohamed, S.; Agarwal, A.; Belgrave, D.; Cho, K.; Oh, A., Eds., 2022.
44. Zhang, F.; Nanda, N. Towards best practices of activation patching in language models: Metrics and methods. *arXiv preprint arXiv:2309.16042* **2023**.
45. Stolfo, A.; Belinkov, Y.; Sachan, M. A Mechanistic Interpretation of Arithmetic Reasoning in Language Models using Causal Mediation Analysis. In Proceedings of the Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing; Bouamor, H.; Pino, J.; Bali, K., Eds., Singapore, 2023; pp. 7035–7052. <https://doi.org/10.18653/v1/2023.emnlp-main.435>.
46. Yu, Z.; Ananiadou, S. How do Large Language Models Learn In-Context? Query and Key Matrices of In-Context Heads are Two Towers for Metric Learning. In Proceedings of the Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing, EMNLP 2024, Miami, FL, USA, November 12-16, 2024; Al-Onaizan, Y.; Bansal, M.; Chen, Y., Eds. Association for Computational Linguistics, 2024, pp. 3281–3292. <https://doi.org/10.18653/v1/2024.EMNLP-MAIN.192>.
47. Geiger, A.; Ibeling, D.; Zur, A.; Chaudhary, M.; Chauhan, S.; Huang, J.; Arora, A.; Wu, Z.; Goodman, N.; Potts, C.; et al. Causal abstraction: A theoretical foundation for mechanistic interpretability. *Journal of Machine Learning Research* **2025**, *26*, 1–64.
48. Ferreira, P.; Aziz, W.; Titov, I. Truthful or Fabricated? Using Causal Attribution to Mitigate Reward Hacking in Explanations. In Proceedings of the Workshop on Actionable Interpretability at ICML 2025, 2025, [arXiv:cs.CL/2504.05294].

49. Yeo, W.J.; Satapathy, R.; Cambria, E. Towards faithful natural language explanations: A study using activation patching in large language models. In Proceedings of the Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing, 2025, pp. 10436–10458.
50. Ravindran, S.K. Adversarial activation patching: A framework for detecting and mitigating emergent deception in safety-aligned transformers. *arXiv preprint arXiv:2507.09406* 2025.
51. Yu, H.; Jeong, S.; Pawar, S.; Shin, J.; Jin, J.; Myung, J.; Oh, A.; Augenstein, I. Entangled in Representations: Mechanistic Investigation of Cultural Biases in Large Language Models, 2025, [arXiv:cs/2508.08879]. <https://doi.org/10.48550/arXiv.2508.08879>.
52. Wang, K.R.; Variengien, A.; Conmy, A.; Shlegeris, B.; Steinhardt, J. Interpretability in the Wild: a Circuit for Indirect Object Identification in GPT-2 Small. In Proceedings of the The Eleventh International Conference on Learning Representations, 2023.
53. Geva, M.; Bastings, J.; Filippova, K.; Globerson, A. Dissecting Recall of Factual Associations in Auto-Regressive Language Models. In Proceedings of the Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing; Bouamor, H.; Pino, J.; Bali, K., Eds., Singapore, 2023; pp. 12216–12235. <https://doi.org/10.18653/v1/2023.emnlp-main.751>.
54. Yu, Z.; Ananiadou, S. Interpreting Arithmetic Mechanism in Large Language Models through Comparative Neuron Analysis. In Proceedings of the Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing; Al-Onaizan, Y.; Bansal, M.; Chen, Y.N., Eds., Miami, Florida, USA, 2024; pp. 3293–3306. <https://doi.org/10.18653/v1/2024.emnlp-main.193>.
55. Li, J.; Chen, X.; Hovy, E.H.; Jurafsky, D. Visualizing and Understanding Neural Models in NLP. In Proceedings of the NAACL HLT 2016, The 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, San Diego California, USA, June 12-17, 2016; Knight, K.; Nenkova, A.; Rambow, O., Eds. The Association for Computational Linguistics, 2016, pp. 681–691. <https://doi.org/10.18653/V1/N16-1082>.
56. Sundararajan, M.; Taly, A.; Yan, Q. Axiomatic Attribution for Deep Networks. In Proceedings of the Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017; Precup, D.; Teh, Y.W., Eds. PMLR, 2017, Vol. 70, *Proceedings of Machine Learning Research*, pp. 3319–3328.
57. Smilkov, D.; Thorat, N.; Kim, B.; Viégas, F.B.; Wattenberg, M. SmoothGrad: removing noise by adding noise. *CoRR* 2017, *abs/1706.03825*, [1706.03825].
58. Shrikumar, A.; Greenside, P.; Kundaje, A. Learning Important Features Through Propagating Activation Differences. In Proceedings of the Proceedings of the 34th International Conference on Machine Learning (ICML), 2017, Vol. 70, *Proceedings of Machine Learning Research*.
59. Enguehard, J. Sequential Integrated Gradients: a simple but effective method for explaining language models. In Proceedings of the Findings of the Association for Computational Linguistics: ACL 2023, Toronto, Canada, July 9-14, 2023; Rogers, A.; Boyd-Graber, J.L.; Okazaki, N., Eds. Association for Computational Linguistics, 2023, pp. 7555–7565. <https://doi.org/10.18653/V1/2023.FINDINGS-ACL.477>.
60. Wu, S.; Shen, E.M.; Badrinath, C.; Ma, J.; Lakkaraju, H. Analyzing chain-of-thought prompting in large language models via gradient-based feature attributions. *arXiv preprint arXiv:2307.13339* 2023.
61. Hou, E.M.; Castañón, G.D. Decoding Layer Saliency in Language Transformers. In Proceedings of the International Conference on Machine Learning, ICML 2023, 23-29 July 2023, Honolulu, Hawaii, USA; Krause, A.; Brunskill, E.; Cho, K.; Engelhardt, B.; Sabato, S.; Scarlett, J., Eds. PMLR, 2023, Vol. 202, *Proceedings of Machine Learning Research*, pp. 13285–13308.
62. Tao, Y.; Tang, Y.; Wang, Y.; Zhu, M.; Hu, H.; Wang, Y. Saliency-driven Dynamic Token Pruning for Large Language Models. *CoRR* 2025, *abs/2504.04514*, [2504.04514]. <https://doi.org/10.48550/ARXIV.2504.04514>.
63. Nguyen, D.; Prasad, A.; Stengel-Eskin, E.; Bansal, M. GrAInS: Gradient-based Attribution for Inference-Time Steering of LLMs and VLMs. *CoRR* 2025, *abs/2507.18043*, [2507.18043]. <https://doi.org/10.48550/ARXIV.2507.18043>.
64. Dai, D.; Dong, L.; Hao, Y.; Sui, Z.; Chang, B.; Wei, F. Knowledge Neurons in Pretrained Transformers. In Proceedings of the Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics

- (Volume 1: Long Papers), ACL 2022, Dublin, Ireland, May 22-27, 2022; Muresan, S.; Nakov, P.; Villavicencio, A., Eds. Association for Computational Linguistics, 2022, pp. 8493–8502. <https://doi.org/10.18653/V1/2022.ACL-LONG.581>.
65. Shi, D.; Jin, R.; Shen, T.; Dong, W.; Wu, X.; Xiong, D. IRCAN: Mitigating Knowledge Conflicts in LLM Generation via Identifying and Reweighting Context-Aware Neurons. In Proceedings of the Advances in Neural Information Processing Systems 38: Annual Conference on Neural Information Processing Systems 2024, NeurIPS 2024, Vancouver, BC, Canada, December 10 - 15, 2024; Globersons, A.; Mackey, L.; Belgrave, D.; Fan, A.; Paquet, U.; Tomczak, J.M.; Zhang, C., Eds., 2024.
  66. Zhang, H.; Wu, Y.; Li, D.; Yang, S.; Zhao, R.; Jiang, Y.; Tan, F. Balancing Speciality and Versatility: a Coarse to Fine Framework for Supervised Fine-tuning Large Language Model. In Proceedings of the Findings of the Association for Computational Linguistics, ACL 2024, Bangkok, Thailand and virtual meeting, August 11-16, 2024; Ku, L.; Martins, A.; Srikumar, V., Eds. Association for Computational Linguistics, 2024, pp. 7467–7509. <https://doi.org/10.18653/V1/2024.FINDINGS-ACL.445>.
  67. Zhang, H.; Liu, Z.; Huang, S.; Shang, C.; Zhan, B.; Jiang, Y. Improving low-resource knowledge tracing tasks by supervised pre-training and importance mechanism fine-tuning. *arXiv preprint arXiv:2403.06725* 2024.
  68. Li, M.; Li, Y.; Zhou, T. What Happened in LLMs Layers when Trained for Fast vs. Slow Thinking: A Gradient Perspective. In Proceedings of the Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers); Che, W.; Nabende, J.; Shutova, E.; Pilehvar, M.T., Eds., Vienna, Austria, 2025; pp. 32017–32154. <https://doi.org/10.18653/v1/2025.acl-long.1545>.
  69. Li, M.; Li, Y.; Li, Z.; Zhou, T. How Instruction and Reasoning Data shape Post-Training: Data Quality through the Lens of Layer-wise Gradients, 2025, [[arXiv:cs.LG/2504.10766](https://arxiv.org/abs/2504.10766)].
  70. Jafari, F.R.; Eberle, O.; Khakzar, A.; Nanda, N. ReLP: Faithful and Efficient Circuit Discovery via Relevance Patching. *CoRR* 2025, *abs/2508.21258*, [[2508.21258](https://arxiv.org/abs/2508.21258)]. <https://doi.org/10.48550/ARXIV.2508.21258>.
  71. Azarkhalili, B.; Libbrecht, M.W. Generalized Attention Flow: Feature Attribution for Transformer Models via Maximum Flow. In Proceedings of the Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2025, Vienna, Austria, July 27 - August 1, 2025; Che, W.; Nabende, J.; Shutova, E.; Pilehvar, M.T., Eds. Association for Computational Linguistics, 2025, pp. 19954–19974.
  72. Liu, S.; Wu, H.; He, B.; Han, X.; Yuan, M.; Song, L. Sens-Merging: Sensitivity-Guided Parameter Balancing for Merging Large Language Models. In Proceedings of the Findings of the Association for Computational Linguistics, ACL 2025, Vienna, Austria, July 27 - August 1, 2025; Che, W.; Nabende, J.; Shutova, E.; Pilehvar, M.T., Eds. Association for Computational Linguistics, 2025, pp. 19243–19255.
  73. Li, H.; Zhang, X.; Liu, X.; Gong, Y.; Wang, Y.; Chen, Q.; Cheng, P. Enhancing Large Language Model Performance with Gradient-Based Parameter Selection. In Proceedings of the AAAI-25, Sponsored by the Association for the Advancement of Artificial Intelligence, February 25 - March 4, 2025, Philadelphia, PA, USA; Walsh, T.; Shah, J.; Kolter, Z., Eds. AAAI Press, 2025, pp. 24431–24439. <https://doi.org/10.1609/AAAI.V39I23.34621>.
  74. Zhang, Z.; Zhao, J.; Zhang, Q.; Gui, T.; Huang, X. Unveiling Linguistic Regions in Large Language Models. In Proceedings of the Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2024, Bangkok, Thailand, August 11-16, 2024; Ku, L.; Martins, A.; Srikumar, V., Eds. Association for Computational Linguistics, 2024, pp. 6228–6247. <https://doi.org/10.18653/V1/2024.ACL-LONG.338>.
  75. Li, G.; Xi, Z.; Zhang, Z.; Hong, B.; Gui, T.; Zhang, Q.; Huang, X. LoRACoE: Improving Large Language Model via Composition-based LoRA Expert. In Proceedings of the Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing; Christodoulopoulos, C.; Chakraborty, T.; Rose, C.; Peng, V., Eds., Suzhou, China, 2025; pp. 31290–31304. <https://doi.org/10.18653/v1/2025.emnlp-main.1594>.
  76. Wang, Y.; Zhang, T.; Guo, X.; Shen, Z. Gradient based Feature Attribution in Explainable AI: A Technical Review. *CoRR* 2024, *abs/2403.10415*, [[2403.10415](https://arxiv.org/abs/2403.10415)]. <https://doi.org/10.48550/ARXIV.2403.10415>.

77. Wang, K.; Variengien, A.; Conmy, A.; Shlegeris, B.; Steinhardt, J. Interpretability in the wild: a circuit for indirect object identification in gpt-2 small. *arXiv preprint arXiv:2211.00593* 2022.
78. Yin, K.; Neubig, G. Interpreting Language Models with Contrastive Explanations. In Proceedings of the Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing; Goldberg, Y.; Kozareva, Z.; Zhang, Y., Eds., Abu Dhabi, United Arab Emirates, 2022; pp. 184–198. <https://doi.org/10.18653/v1/2022.emnlp-main.14>.
79. Alain, G.; Bengio, Y. Understanding intermediate layers using linear classifier probes. In Proceedings of the 5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24–26, 2017, Workshop Track Proceedings. OpenReview.net, 2017.
80. Belinkov, Y. Probing Classifiers: Promises, Shortcomings, and Advances. *Comput. Linguistics* 2022, 48, 207–219. [https://doi.org/10.1162/COLI\\_A\\_00422](https://doi.org/10.1162/COLI_A_00422).
81. Conneau, A.; Kruszewski, G.; Lample, G.; Barrault, L.; Baroni, M. What you can cram into a single vector: Probing sentence embeddings for linguistic properties. In Proceedings of the Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, ACL 2018, Melbourne, Australia, July 15–20, 2018, Volume 1: Long Papers; Gurevych, I.; Miyao, Y., Eds. Association for Computational Linguistics, 2018, pp. 2126–2136. <https://doi.org/10.18653/V1/P18-1198>.
82. Tenney, I.; Das, D.; Pavlick, E. BERT Rediscovered the Classical NLP Pipeline. In Proceedings of the Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28– August 2, 2019, Volume 1: Long Papers; Korhonen, A.; Traum, D.R.; Màrquez, L., Eds. Association for Computational Linguistics, 2019, pp. 4593–4601. <https://doi.org/10.18653/V1/P19-1452>.
83. Ravichander, A.; Belinkov, Y.; Hovy, E.H. Probing the Probing Paradigm: Does Probing Accuracy Entail Task Relevance? In Proceedings of the Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume, EACL 2021, Online, April 19 - 23, 2021; Merlo, P.; Tiedemann, J.; Tsarfaty, R., Eds. Association for Computational Linguistics, 2021, pp. 3363–3377. <https://doi.org/10.18653/V1/2021.EACL-MAIN.295>.
84. Ju, T.; Sun, W.; Du, W.; Yuan, X.; Ren, Z.; Liu, G. How Large Language Models Encode Context Knowledge? A Layer-Wise Probing Study. In Proceedings of the Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation, LREC/COLING 2024, 20–25 May, 2024, Torino, Italy; Calzolari, N.; Kan, M.; Hoste, V.; Lenci, A.; Sakti, S.; Xue, N., Eds. ELRA and ICCL, 2024, pp. 8235–8246.
85. Zhao, Y.; Du, X.; Hong, G.; Gema, A.P.; Devoto, A.; Wang, H.; He, X.; Wong, K.; Minervini, P. Analysing the Residual Stream of Language Models Under Knowledge Conflicts. *CoRR* 2024, *abs/2410.16090*, [2410.16090]. <https://doi.org/10.48550/ARXIV.2410.16090>.
86. Orgad, H.; Toker, M.; Gekhman, Z.; Reichart, R.; Szpektor, I.; Kotek, H.; Belinkov, Y. LLMs Know More Than They Show: On the Intrinsic Representation of LLM Hallucinations. In Proceedings of the The Thirteenth International Conference on Learning Representations, 2025.
87. You, W.; Xue, A.; Havaldar, S.; Rao, D.; Jin, H.; Callison-Burch, C.; Wong, E. Probabilistic Soundness Guarantees in LLM Reasoning Chains. In Proceedings of the Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing, Suzhou, China, 2025; pp. 7517–7536. <https://doi.org/10.18653/v1/2025.emnlp-main.382>.
88. Du, Y.; Zhao, S.; Cao, J.; Ma, M.; Zhao, D.; Fan, F.; Liu, T.; Qin, B. Towards Secure Tuning: Mitigating Security Risks Arising from Benign Instruction Fine-Tuning. *CoRR* 2024, *abs/2410.04524*, [2410.04524]. <https://doi.org/10.48550/ARXIV.2410.04524>.
89. Zhao, D.; Liu, X.; Feng, X.; Wang, H.; Qin, B. Probing and Boosting Large Language Models Capabilities via Attention Heads. In Proceedings of the Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing; Christodoulopoulos, C.; Chakraborty, T.; Rose, C.; Peng, V., Eds., Suzhou, China, 2025; pp. 28518–28532. <https://doi.org/10.18653/v1/2025.emnlp-main.1450>.

90. Kim, J.; Evans, J.; Schein, A. Linear Representations of Political Perspective Emerge in Large Language Models. In Proceedings of the The Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025. OpenReview.net, 2025.
91. Kantamneni, S.; Engels, J.; Rajamanoharan, S.; Tegmark, M.; Nanda, N. Are Sparse Autoencoders Useful? A Case Study in Sparse Probing. In Proceedings of the Forty-second International Conference on Machine Learning, ICML 2025, Vancouver, BC, Canada, July 13-19, 2025. OpenReview.net, 2025.
92. Chanin, D.; Wilken-Smith, J.; Dulka, T.; Bhatnagar, H.; Bloom, J. A is for Absorption: Studying Feature Splitting and Absorption in Sparse Autoencoders. *CoRR* **2024**, *abs/2409.14507*, [2409.14507]. <https://doi.org/10.48550/ARXIV.2409.14507>.
93. nostalgebraist. Interpreting GPT: the Logit Lens, 2020.
94. Geva, M.; Schuster, R.; Berant, J.; Levy, O. Transformer Feed-Forward Layers Are Key-Value Memories. In Proceedings of the Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing; Moens, M.F.; Huang, X.; Specia, L.; Yih, S.W.t., Eds., Online and Punta Cana, Dominican Republic, 2021; pp. 5484–5495. <https://doi.org/10.18653/v1/2021.emnlp-main.446>.
95. Belrose, N.; Furman, Z.; Smith, L.; Halawi, D.; Ostrovsky, I.; McKinney, L.; Biderman, S.; Steinhardt, J. Eliciting latent predictions from transformers with the tuned lens. *arXiv preprint arXiv:2303.08112* **2023**.
96. Jiang, C.; Qi, B.; Hong, X.; Fu, D.; Cheng, Y.; Meng, F.; Yu, M.; Zhou, B.; Zhou, J. On Large Language Models' Hallucination with Regard to Known Facts. In Proceedings of the Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers), 2024, pp. 1041–1053.
97. Jiang, N.; Kachinthaya, A.; Petryk, S.; Gandelsman, Y. Interpreting and Editing Vision-Language Representations to Mitigate Hallucinations, 2025, [arXiv:cs.CV/2410.02762].
98. Wendler, C.; Veselovsky, V.; Monea, G.; West, R. Do Llamas Work in English? On the Latent Language of Multilingual Transformers. In Proceedings of the Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers); Ku, L.W.; Martins, A.; Srikumar, V., Eds., Bangkok, Thailand, 2024; pp. 15366–15394. <https://doi.org/10.18653/v1/2024.acl-long.820>.
99. Kargaran, A.H.; Liu, Y.; Yvon, F.; Schütze, H. How Programming Concepts and Neurons Are Shared in Code Language Models. *arXiv preprint arXiv:2506.01074* **2025**.
100. Phukan, A.; Somasundaram, S.; Saxena, A.; Goswami, K.; Srinivasan, B.V. Peering into the Mind of Language Models: An Approach for Attribution in Contextual Question Answering. In Proceedings of the Findings of the Association for Computational Linguistics: ACL 2024; Ku, L.W.; Martins, A.; Srikumar, V., Eds., Bangkok, Thailand, 2024; pp. 11481–11495. <https://doi.org/10.18653/v1/2024.findings-acl.682>.
101. Phukan, A.; Divyansh.; Morj, H.K.; Vaishnavi.; Saxena, A.; Goswami, K. Beyond Logit Lens: Contextual Embeddings for Robust Hallucination Detection & Grounding in VLMs. In Proceedings of the Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers); Chiruzzo, L.; Ritter, A.; Wang, L., Eds., Albuquerque, New Mexico, 2025; pp. 9661–9675. <https://doi.org/10.18653/v1/2025.naacl-long.488>.
102. Yugeswardeenoo, D.; Nukala, H.; Blondin, C.; O'Brien, S.; Sharma, V.; Zhu, K. Interpreting the Latent Structure of Operator Precedence in Language Models. In Proceedings of the The First Workshop on the Interplay of Model Behavior and Model Internals, 2025.
103. Sakarvadia, M.; Khan, A.; Ajith, A.; Grzenda, D.; Hudson, N.; Bauer, A.; Chard, K.; Foster, I. Attention lens: A tool for mechanistically interpreting the attention head information retrieval mechanism. *arXiv preprint arXiv:2310.16270* **2023**.
104. Yu, Z.; Ananiadou, S. Understanding multimodal llms: the mechanistic interpretability of llava in visual question answering. *arXiv preprint arXiv:2411.10950* **2024**.
105. Jiang, Z.; Chen, J.; Zhu, B.; Luo, T.; Shen, Y.; Yang, X. Devils in middle layers of large vision-language models: Interpreting, detecting and mitigating object hallucinations via attention lens. In Proceedings of the Proceedings of the Computer Vision and Pattern Recognition Conference, 2025, pp. 25004–25014.

106. Kim, J.; Kang, S.; Park, J.; Kim, J.; Hwang, S.J. Interpreting Attention Heads for Image-to-Text Information Flow in Large Vision–Language Models. In Proceedings of the Mechanistic Interpretability Workshop at NeurIPS 2025, 2025.
107. Wang, Z. Logitlens4llms: Extending logit lens analysis to modern large language models. *arXiv preprint arXiv:2503.11667* 2025.
108. Huo, J.; Yan, Y.; Hu, B.; Yue, Y.; Hu, X. MMNeuron: Discovering Neuron-Level Domain-Specific Interpretation in Multimodal Large Language Model. In Proceedings of the Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing, 2024, pp. 6801–6816.
109. Yu, Z.; Ananiadou, S. Understanding and Mitigating Gender Bias in LLMs via Interpretable Neuron Editing, 2025, [arXiv:cs/2501.14457]. <https://doi.org/10.48550/arXiv.2501.14457>.
110. Shao, J.; Lu, Y.; Yang, J. Benford’s Curse: Tracing Digit Bias to Numerical Hallucination in LLMs. *arXiv preprint arXiv:2506.01734* 2025.
111. Arad, D.; Mueller, A.; Belinkov, Y. SAEs Are Good for Steering–If You Select the Right Features. *arXiv preprint arXiv:2505.20063* 2025.
112. Dreyer, M.; Hufe, L.; Berend, J.; Wiegand, T.; Lapuschkin, S.; Samek, W. From What to How: Attributing CLIP’s Latent Components Reveals Unexpected Semantic Reliance. *arXiv preprint arXiv:2505.20229* 2025.
113. Muhamed, A.; Diab, M.; Smith, V. Decoding dark matter: Specialized sparse autoencoders for interpreting rare concepts in foundation models. In Proceedings of the Findings of the Association for Computational Linguistics: NAACL 2025, 2025, pp. 1604–1635.
114. Gur-Arieh, Y.; Mayan, R.; Agassy, C.; Geiger, A.; Geva, M. Enhancing automated interpretability with output-centric feature descriptions. *arXiv preprint arXiv:2501.08319* 2025.
115. Shu, D.; Wu, X.; Zhao, H.; Rai, D.; Yao, Z.; Liu, N.; Du, M. A Survey on Sparse Autoencoders: Interpreting the Internal Mechanisms of Large Language Models. In Proceedings of the Findings of the Association for Computational Linguistics: EMNLP 2025; Christodoulopoulos, C.; Chakraborty, T.; Rose, C.; Peng, V., Eds., Suzhou, China, 2025; pp. 1690–1712. <https://doi.org/10.18653/v1/2025.findings-emnlp.89>.
116. Elhage, N.; Nanda, N.; Olsson, C.; Henighan, T.; Joseph, N.; Mann, B.; Askell, A.; Bai, Y.; Chen, A.; Conerly, T.; et al. A Mathematical Framework for Transformer Circuits. *Transformer Circuits Thread* 2021. <https://transformer-circuits.pub/2021/framework/index.html>.
117. Olsson, C.; Elhage, N.; Nanda, N.; Joseph, N.; DasSarma, N.; Henighan, T.; Mann, B.; Askell, A.; Bai, Y.; Chen, A.; et al. In-context Learning and Induction Heads, 2022, [arXiv:cs.LG/2209.11895].
118. Hanna, M.; Liu, O.; Variengien, A. How does GPT-2 compute greater-than?: Interpreting mathematical abilities in a pre-trained language model. In Proceedings of the Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023; Oh, A.; Naumann, T.; Globerson, A.; Saenko, K.; Hardt, M.; Levine, S., Eds., 2023.
119. Yao, Y.; Zhang, N.; Xi, Z.; Wang, M.; Xu, Z.; Deng, S.; Chen, H. Knowledge Circuits in Pretrained Transformers. In Proceedings of the Advances in Neural Information Processing Systems; Globerson, A.; Mackey, L.; Belgrave, D.; Fan, A.; Paquet, U.; Tomczak, J.; Zhang, C., Eds. Curran Associates, Inc., 2024, Vol. 37, pp. 118571–118602. <https://doi.org/10.52202/079017-3765>.
120. Goldowsky-Dill, N.; MacLeod, C.; Sato, L.; Arora, A. Localizing model behavior with path patching. *arXiv preprint arXiv:2304.05969* 2023.
121. Conmy, A.; Mavor-Parker, A.; Lynch, A.; Heimersheim, S.; Garriga-Alonso, A. Towards Automated Circuit Discovery for Mechanistic Interpretability. In Proceedings of the Advances in Neural Information Processing Systems; Oh, A.; Naumann, T.; Globerson, A.; Saenko, K.; Hardt, M.; Levine, S., Eds. Curran Associates, Inc., 2023, Vol. 36, pp. 16318–16352.
122. Syed, A.; Rager, C.; Conmy, A. Attribution Patching Outperforms Automated Circuit Discovery. In Proceedings of the Proceedings of the 7th BlackboxNLP Workshop: Analyzing and Interpreting Neural Networks for NLP; Belinkov, Y.; Kim, N.; Jumelet, J.; Mohebbi, H.; Mueller, A.; Chen, H., Eds., Miami, Florida, US, 2024; pp. 407–416. <https://doi.org/10.18653/v1/2024.blackboxnlp-1.25>.

123. Hanna, M.; Pezzelle, S.; Belinkov, Y. Have Faith in Faithfulness: Going Beyond Circuit Overlap When Finding Model Mechanisms. In Proceedings of the First Conference on Language Modeling, 2024.
124. Huang, H.; Yan, Y.; Huo, J.; Zou, X.; Li, X.; Wang, K.; Hu, X. Pierce the Mists, Greet the Sky: Decipher Knowledge Overshadowing via Knowledge Circuit Analysis. In Proceedings of the Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing; Christodoulopoulos, C.; Chakraborty, T.; Rose, C.; Peng, V., Eds., Suzhou, China, 2025; pp. 15471–15490. <https://doi.org/10.18653/v1/2025.emnlp-main.781>.
125. Haklay, T.; Orgad, H.; Bau, D.; Mueller, A.; Belinkov, Y. Position-aware Automatic Circuit Discovery. In Proceedings of the Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2025, Vienna, Austria, July 27 - August 1, 2025; Che, W.; Nabende, J.; Shutova, E.; Pilehvar, M.T., Eds. Association for Computational Linguistics, 2025; pp. 2792–2817.
126. Mueller, A.; Geiger, A.; Wiegrefe, S.; Arad, D.; Arcuschin, I.; Belfki, A.; Chan, Y.S.; Fiotto-Kaufman, J.F.; Haklay, T.; Hanna, M.; et al. MIB: A Mechanistic Interpretability Benchmark. In Proceedings of the Forty-second International Conference on Machine Learning, ICML 2025, Vancouver, BC, Canada, July 13-19, 2025. OpenReview.net, 2025.
127. Bricken, T.; Templeton, A.; Batson, J.; Chen, B.; Jermyn, A.; Conerly, T.; Turner, N.; Anil, C.; Denison, C.; Askell, A.; et al. Towards Monosemanticity: Decomposing Language Models With Dictionary Learning. *Transformer Circuits Thread* 2023. <https://transformer-circuits.pub/2023/monosemantic-features/index.html>.
128. Ameisen, E.; Lindsey, J.; Pearce, A.; Gurnee, W.; Turner, N.L.; Chen, B.; Citro, C.; Abrahams, D.; Carter, S.; Hosmer, B.; et al. Circuit Tracing: Revealing Computational Graphs in Language Models. *Transformer Circuits Thread* 2025.
129. Hanna, M.; Piotrowski, M.; Lindsey, J.; Ameisen, E. Circuit-Tracer: A New Library for Finding Feature Circuits. In Proceedings of the Proceedings of the 8th BlackboxNLP Workshop: Analyzing and Interpreting Neural Networks for NLP; Belinkov, Y.; Mueller, A.; Kim, N.; Mohebbi, H.; Chen, H.; Arad, D.; Sarti, G., Eds., Suzhou, China, 2025; pp. 239–249. <https://doi.org/10.18653/v1/2025.blackboxnlp-1.14>.
130. Nie, E.; Schmid, H.; Schuetze, H. Mechanistic Understanding and Mitigation of Language Confusion in English-Centric Large Language Models. In Proceedings of the Findings of the Association for Computational Linguistics: EMNLP 2025; Christodoulopoulos, C.; Chakraborty, T.; Rose, C.; Peng, V., Eds., Suzhou, China, 2025; pp. 690–706. <https://doi.org/10.18653/v1/2025.findings-emnlp.37>.
131. Goyal, A.; Rathi, V.; Yeh, W.; Wang, Y.; Chen, Y.; Sundaram, H. Breaking Bad Tokens: Detoxification of LLMs Using Sparse Autoencoders. In Proceedings of the Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing; Christodoulopoulos, C.; Chakraborty, T.; Rose, C.; Peng, V., Eds., Suzhou, China, 2025; pp. 12702–12720. <https://doi.org/10.18653/v1/2025.emnlp-main.641>.
132. Yeo, W.J.; Prakash, N.; Neo, C.; Satapathy, R.; Lee, R.K.W.; Cambria, E. Understanding Refusal in Language Models with Sparse Autoencoders. In Proceedings of the Findings of the Association for Computational Linguistics: EMNLP 2025; Christodoulopoulos, C.; Chakraborty, T.; Rose, C.; Peng, V., Eds., Suzhou, China, 2025; pp. 6377–6399. <https://doi.org/10.18653/v1/2025.findings-emnlp.338>.
133. Liu, Y.; Liu, Y.; Chen, X.; Chen, P.Y.; Zan, D.; Kan, M.Y.; Ho, T.Y. The Devil Is in the Neurons: Interpreting and Mitigating Social Biases in Pre-trained Language Models, 2024, [arXiv:cs/2406.10130]. <https://doi.org/10.48550/arXiv.2406.10130>.
134. Chandna, B.; Bashir, Z.; Sen, P. Dissecting Bias in LLMs: A Mechanistic Interpretability Perspective, 2025, [arXiv:cs/2506.05166]. <https://doi.org/10.48550/arXiv.2506.05166>.
135. Zhou, Z.; Yu, H.; Zhang, X.; Xu, R.; Huang, F.; Wang, K.; Liu, Y.; Fang, J.; Li, Y. On the Role of Attention Heads in Large Language Model Safety. In Proceedings of the The Thirteenth International Conference on Learning Representations, 2025.
136. Niu, J.; Yuan, X.; Wang, T.; Saghir, H.; Abdi, A.H. Llama See, Llama Do: A Mechanistic Perspective on Contextual Entrainment and Distraction in LLMs. In Proceedings of the Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers); Che, W.; Nabende, J.; Shutova, E.; Pilehvar, M.T., Eds., Vienna, Austria, 2025; pp. 16218–16239. <https://doi.org/10.18653/v1/2025.acl-long.791>.

137. Ahsan, H.; Sharma, A.S.; Amir, S.; Bau, D.; Wallace, B.C. Elucidating Mechanisms of Demographic Bias in LLMs for Healthcare. <http://arxiv.org/abs/2502.13319>, 2025, [arXiv:cs/2502.13319]. <https://doi.org/10.48550/arXiv.2502.13319>.
138. Raimondi, B.; Dalbagno, D.; Gabbrielli, M. Analysing Moral Bias in Finetuned LLMs through Mechanistic Interpretability, 2025, [arXiv:cs/2510.12229]. <https://doi.org/10.48550/arXiv.2510.12229>.
139. Gao, C.; Chen, H.; Xiao, C.; Chen, Z.; Liu, Z.; Sun, M. H-Neurons: On the Existence, Impact, and Origin of Hallucination-Associated Neurons in LLMs, 2025, [arXiv:cs.AI/2512.01797].
140. Pach, M.; Karthik, S.; Bouniot, Q.; Belongie, S.; Akata, Z. Sparse Autoencoders Learn Monosemantic Features in Vision-Language Models. In Proceedings of the The Thirty-ninth Annual Conference on Neural Information Processing Systems, 2025.
141. Stoehr, N.; Du, K.; Snæbjarnarson, V.; West, R.; Cotterell, R.; Schein, A. Activation Scaling for Steering and Interpreting Language Models. In Proceedings of the Findings of the Association for Computational Linguistics: EMNLP 2024; Al-Onaizan, Y.; Bansal, M.; Chen, Y.N., Eds., Miami, Florida, USA, 2024; pp. 8189–8200. <https://doi.org/10.18653/v1/2024.findings-emnlp.479>.
142. Yao, Y.; Liu, S.; Liu, Z.; Li, Q.; Liu, M.; Han, X.; Guo, Z.; Wu, H.; Song, L. Activation-Guided Consensus Merging for Large Language Models. *arXiv preprint arXiv:2505.14009* 2025.
143. Wang, M.; Chen, X.; Wang, Y.; He, Z.; Xu, J.; Liang, T.; Liu, Q.; Yao, Y.; Wang, W.; Ma, R.; et al. Two Experts Are All You Need for Steering Thinking: Reinforcing Cognitive Effort in MoE Reasoning Models Without Additional Training. In Proceedings of the The Thirty-ninth Annual Conference on Neural Information Processing Systems, 2025.
144. Meng, K.; Sharma, A.S.; Andonian, A.J.; Belinkov, Y.; Bau, D. Mass-Editing Memory in a Transformer. In Proceedings of the The Eleventh International Conference on Learning Representations, 2023.
145. Zhong, M.; An, C.; Chen, W.; Han, J.; He, P. Seeking Neural Nuggets: Knowledge Transfer in Large Language Models from a Parametric Perspective. In Proceedings of the The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024. OpenReview.net, 2024.
146. Zhang, Z.; Liu, B.; Shao, J. Fine-tuning Happens in Tiny Subspaces: Exploring Intrinsic Task-specific Subspaces of Pre-trained Language Models. In Proceedings of the Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers); Rogers, A.; Boyd-Graber, J.; Okazaki, N., Eds., Toronto, Canada, 2023; pp. 1701–1713. <https://doi.org/10.18653/v1/2023.acl-long.95>.
147. Xu, H.; Zhan, R.; Ma, Y.; Wong, D.F.; Chao, L.S. Let's Focus on Neuron: Neuron-Level Supervised Fine-tuning for Large Language Model. In Proceedings of the Proceedings of the 31st International Conference on Computational Linguistics; Rambow, O.; Wanner, L.; Apidianaki, M.; Al-Khalifa, H.; Eugenio, B.D.; Schockaert, S., Eds., Abu Dhabi, UAE, 2025; pp. 9393–9406.
148. Xi, Z.; Zheng, R.; Gui, T.; Zhang, Q.; Huang, X. Efficient Adversarial Training with Robust Early-Bird Tickets. In Proceedings of the Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing, EMNLP 2022, Abu Dhabi, United Arab Emirates, December 7-11, 2022; Goldberg, Y.; Kozareva, Z.; Zhang, Y., Eds. Association for Computational Linguistics, 2022, pp. 8318–8331. <https://doi.org/10.18653/V1/2022.EMNLP-MAIN.569>.
149. Zhou, Y.; Chen, W.; Zheng, R.; Xi, Z.; Gui, T.; Zhang, Q.; Huang, X. ORTicket: Let One Robust BERT Ticket Transfer across Different Tasks. In Proceedings of the Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation, LREC/COLING 2024, 20-25 May, 2024, Torino, Italy; Calzolari, N.; Kan, M.; Hoste, V.; Lenci, A.; Sakti, S.; Xue, N., Eds. ELRA and ICCL, 2024, pp. 12527–12538.
150. Zhang, X.; Liang, Y.; Meng, F.; Zhang, S.; Chen, Y.; Xu, J.; Zhou, J. Multilingual Knowledge Editing with Language-Agnostic Factual Neurons. In Proceedings of the Proceedings of the 31st International Conference on Computational Linguistics, COLING 2025, Abu Dhabi, UAE, January 19-24, 2025; Rambow, O.; Wanner, L.; Apidianaki, M.; Al-Khalifa, H.; Eugenio, B.D.; Schockaert, S., Eds. Association for Computational Linguistics, 2025, pp. 5775–5788.

151. Li, S.; Yao, L.; Zhang, L.; Li, Y. Safety Layers in Aligned Large Language Models: The Key to LLM Security. In Proceedings of the The Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025. OpenReview.net, 2025.
152. Li, X.; Li, Z.; Kosuga, Y.; Yoshida, Y.; Bian, V. Precision Knowledge Editing: Enhancing Safety in Large Language Models. *CoRR* **2024**, *abs/2410.03772*, [2410.03772]. <https://doi.org/10.48550/ARXIV.2410.03772>.
153. Zhang, W.; Wan, C.; Zhang, Y.; ming Cheung, Y.; Tian, X.; Shen, X.; Ye, J. Interpreting and Improving Large Language Models in Arithmetic Calculation. In Proceedings of the Forty-first International Conference on Machine Learning, 2024.
154. Zhu, S.; Pan, L.; Li, B.; Xiong, D. LANDeRMT: Detecting and Routing Language-Aware Neurons for Selectively Finetuning LLMs to Machine Translation. In Proceedings of the Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers); Ku, L.W.; Martins, A.; Srikumar, V., Eds., Bangkok, Thailand, 2024; pp. 12135–12148. <https://doi.org/10.18653/v1/2024.acl-long.656>.
155. Tang, H.; Lin, Y.; Lin, J.; Han, Q.; Ke, D.; Hong, S.; Yao, Y.; Wang, G. RazorAttention: Efficient KV Cache Compression Through Retrieval Heads. In Proceedings of the The Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025. OpenReview.net, 2025.
156. Hooper, C.; Kim, S.; Mohammadzadeh, H.; Mahoney, M.W.; Shao, Y.S.; Keutzer, K.; Gholami, A. KVQuant: Towards 10 Million Context Length LLM Inference with KV Cache Quantization. In Proceedings of the Advances in Neural Information Processing Systems; Globerson, A.; Mackey, L.; Belgrave, D.; Fan, A.; Paquet, U.; Tomczak, J.; Zhang, C., Eds. Curran Associates, Inc., 2024, Vol. 37, pp. 1270–1303. <https://doi.org/10.52202/079017-0040>.
157. Bondarenko, Y.; Nagel, M.; Blankevoort, T. Quantizable Transformers: Removing Outliers by Helping Attention Heads Do Nothing. In Proceedings of the Advances in Neural Information Processing Systems; Oh, A.; Naumann, T.; Globerson, A.; Saenko, K.; Hardt, M.; Levine, S., Eds. Curran Associates, Inc., 2023, Vol. 36, pp. 75067–75096.
158. Chen, Y.; Shang, J.; Zhang, Z.; Xie, Y.; Sheng, J.; Liu, T.; Wang, S.; Sun, Y.; Wu, H.; Wang, H. Inner Thinking Transformer: Leveraging Dynamic Depth Scaling to Foster Adaptive Internal Thinking. In Proceedings of the Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers); Che, W.; Nabende, J.; Shutova, E.; Pilehvar, M.T., Eds., Vienna, Austria, 2025; pp. 28241–28259. <https://doi.org/10.18653/v1/2025.acl-long.1369>.
159. Xia, H.; Leong, C.T.; Wang, W.; Li, Y.; Li, W. TokenSkip: Controllable Chain-of-Thought Compression in LLMs. In Proceedings of the Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing; Christodoulopoulos, C.; Chakraborty, T.; Rose, C.; Peng, V., Eds., Suzhou, China, 2025; pp. 3351–3363. <https://doi.org/10.18653/v1/2025.emnlp-main.165>.
160. Tan, Y.; Wang, M.; He, S.; Liao, H.; Zhao, C.; Lu, Q.; Liang, T.; Zhao, J.; Liu, K. Bottom-up Policy Optimization: Your Language Model Policy Secretly Contains Internal Policies, 2025, [arXiv:cs.LG/2512.19673].
161. Lee, A.; Bai, X.; Pres, I.; Wattenberg, M.; Kummerfeld, J.K.; Mihalcea, R. A Mechanistic Understanding of Alignment Algorithms: A Case Study on DPO and Toxicity. In Proceedings of the Forty-first International Conference on Machine Learning, ICML 2024, Vienna, Austria, July 21-27, 2024. OpenReview.net, 2024.
162. Rimsky, N.; Gabrieli, N.; Schulz, J.; Tong, M.; Hubinger, E.; Turner, A. Steering Llama 2 via Contrastive Activation Addition. In Proceedings of the Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers); Ku, L.W.; Martins, A.; Srikumar, V., Eds., Bangkok, Thailand, 2024; pp. 15504–15522. <https://doi.org/10.18653/v1/2024.acl-long.828>.
163. van der Weij, T.; Poesio, M.; Schoots, N. Extending activation steering to broad skills and multiple behaviours. *arXiv preprint arXiv:2403.05767* **2024**.
164. Lu, D.; Rimsky, N. Investigating bias representations in Llama 2 Chat via activation steering. *arXiv preprint arXiv:2402.00402* **2024**.
165. Postmus, J.; Abreu, S. Steering large language models using conceptors: Improving addition-based activation engineering. *arXiv preprint arXiv:2410.16314* **2024**.

166. Turner, A.M.; Thiergart, L.; Leech, G.; Udell, D.; Vazquez, J.J.; Mini, U.; MacDiarmid, M. Steering Language Models With Activation Engineering, 2024, [arXiv:cs.CL/2308.10248].
167. Sharma, V.; Raman, V. Steering Conceptual Bias via Transformer Latent-Subspace Activation, 2025, [arXiv:cs.AI/2506.18887].
168. Wang, M.; Xu, Z.; Mao, S.; Deng, S.; Tu, Z.; Chen, H.; Zhang, N. Beyond Prompt Engineering: Robust Behavior Control in LLMs via Steering Target Atoms. *arXiv preprint arXiv:2505.20322* 2025.
169. Bayat, R.; Rahimi-Kalahroudi, A.; Pezeshki, M.; Chandar, S.; Vincent, P. Steering large language model activations in sparse spaces. *arXiv preprint arXiv:2503.00177* 2025.
170. Weng, J.; Zheng, H.; Zhang, H.; He, Q.; Tao, J.; Xue, H.; Chu, Z.; Wang, X. Safe-SAIL: Towards a Fine-grained Safety Landscape of Large Language Models via Sparse Autoencoder Interpretation Framework. *arXiv preprint arXiv:2509.18127* 2025.
171. He, Z.; Zhao, H.; Qiao, Y.; Yang, F.; Payani, A.; Ma, J.; Du, M. Saif: A sparse autoencoder framework for interpreting and steering instruction following of language models. *arXiv preprint arXiv:2502.11356* 2025.
172. Soo, S.; Guang, C.; Teng, W.; Balaganesh, C.; Guoxian, T.; Ming, Y. Interpretable Steering of Large Language Models with Feature Guided Activation Additions. *arXiv preprint arXiv:2501.09929* 2025.
173. Ilharco, G.; Ribeiro, M.T.; Wortsman, M.; Schmidt, L.; Hajishirzi, H.; Farhadi, A. Editing models with task arithmetic. In Proceedings of the The Eleventh International Conference on Learning Representations, 2023.
174. Yadav, P.; Tam, D.; Choshen, L.; Raffel, C.A.; Bansal, M. Ties-merging: Resolving interference when merging models. *Advances in Neural Information Processing Systems* 2023, 36, 7093–7115.
175. Chen, J.; Wang, X.; Yao, Z.; Bai, Y.; Hou, L.; Li, J. Towards Understanding Safety Alignment: A Mechanistic Perspective from Safety Neurons, 2025.
176. Suau, X.; Delobelle, P.; Metcalf, K.; Joulin, A.; Apostoloff, N.; Zappella, L.; Rodriguez, P. Whispering Experts: Neural Interventions for Toxicity Mitigation in Language Models. In Proceedings of the Forty-first International Conference on Machine Learning, 2024.
177. Templeton, A.; Conerly, T.; Marcus, J.; Lindsey, J.; Bricken, T.; Chen, B.; Pearce, A.; Citro, C.; Ameisen, E.; Jones, A.; et al. Scaling Monosemanticity: Extracting Interpretable Features from Claude 3 Sonnet. *Transformer Circuits Thread* 2024.
178. Zhao, Y.; Zhang, W.; Xie, Y.; Goyal, A.; Kawaguchi, K.; Shieh, M. Understanding and Enhancing Safety Mechanisms of LLMs via Safety-Specific Neuron. In Proceedings of the The Thirteenth International Conference on Learning Representations, 2025.
179. Ardit, A.; Obeso, O.B.; Syed, A.; Paleka, D.; Rimsky, N.; Gurnee, W.; Nanda, N. Refusal in Language Models Is Mediated by a Single Direction. In Proceedings of the The Thirty-eighth Annual Conference on Neural Information Processing Systems, 2024.
180. Zhao, J.; Huang, J.; Wu, Z.; Bau, D.; Shi, W. LLMs Encode Harmfulness and Refusal Separately. In Proceedings of the The Thirty-ninth Annual Conference on Neural Information Processing Systems, 2025.
181. Yin, Q.; Leong, C.T.; Yang, L.; Huang, W.; Li, W.; Wang, X.; Yoon, J.; YunXing.; XingYu.; Gu, J. Refusal Falls off a Cliff: How Safety Alignment Fails in Reasoning?, 2025, [arXiv:cs.AI/2510.06036].
182. Ball, S.; Kreuter, F.; Panickssery, N. Understanding Jailbreak Success: A Study of Latent Space Dynamics in Large Language Models, 2024, [arXiv:cs.CL/2406.09289].
183. Wang, X.; Hu, C.; Röttger, P.; Plank, B. Surgical, Cheap, and Flexible: Mitigating False Refusal in Language Models via Single Vector Ablation. In Proceedings of the The Thirteenth International Conference on Learning Representations, 2025.
184. Wang, X.; Wang, M.; Liu, Y.; Schuetze, H.; Plank, B. Refusal Direction is Universal Across Safety-Aligned Languages. In Proceedings of the The Thirty-ninth Annual Conference on Neural Information Processing Systems, 2025.
185. Huang, J.; Tao, J.; Icard, T.; Yang, D.; Potts, C. Internal Causal Mechanisms Robustly Predict Language Model Out-of-Distribution Behaviors, 2025, [arXiv:cs.LG/2505.11770].

186. Chen, S.; Xiong, M.; Liu, J.; Wu, Z.; Xiao, T.; Gao, S.; He, J. In-Context Sharpness as Alerts: An Inner Representation Perspective for Hallucination Mitigation. In Proceedings of the Forty-first International Conference on Machine Learning, 2024.
187. Stolfo, A.; Balachandran, V.; Yousefi, S.; Horvitz, E.; Nushi, B. Improving Instruction-Following in Language Models through Activation Steering. In Proceedings of the The Thirteenth International Conference on Learning Representations, 2025.
188. Jiang, G.; Li, Z.; Lian, D.; Wei, Y. Refine Large Language Model Fine-tuning via Instruction Vector. *arXiv preprint arXiv:2406.12227* 2024.
189. Li, J.; Ye, H.; Chen, Y.; Li, X.; Zhang, L.; Alinejad-Rokny, H.; Peng, J.C.H.; Yang, M. Training Superior Sparse Autoencoders for Instruct Models. *arXiv preprint arXiv:2506.07691* 2025.
190. Cai, Y.; Cao, D.; Guo, R.; Wen, Y.; Liu, G.; Chen, E. Locating and Mitigating Gender Bias in Large Language Models, 2024, [arXiv:cs/2403.14409]. <https://doi.org/10.48550/arXiv.2403.14409>.
191. Guan, X.; Lin, P.; Wu, Z.; Wang, Z.; Zhang, R.; Kazim, E.; Koshiyama, A. MPF: Aligning and Debiasing Language Models Post Deployment via Multi Perspective Fusion, 2025, [arXiv:cs/2507.02595]. <https://doi.org/10.48550/arXiv.2507.02595>.
192. Chintam, A.; Beloch, R.; Zuidema, W.; Hanna, M.; Van Der Wal, O. Identifying and Adapting Transformer-Components Responsible for Gender Bias in an English Language Model. In Proceedings of the Proceedings of the 6th BlackboxNLP Workshop: Analyzing and Interpreting Neural Networks for NLP, Singapore, 2023; pp. 379–394. <https://doi.org/10.18653/v1/2023.blackboxnlp-1.29>.
193. Poterti, D.; Seveso, A.; Mercurio, F. Can Role Vectors Affect LLM Behaviour? In Proceedings of the Findings of the Association for Computational Linguistics: EMNLP 2025; Christodoulopoulos, C.; Chakraborty, T.; Rose, C.; Peng, V., Eds., Suzhou, China, 2025; pp. 17735–17747. <https://doi.org/10.18653/v1/2025.findings-emnlp.963>.
194. Chen, R.; Arditi, A.; Sleight, H.; Evans, O.; Lindsey, J. Persona vectors: Monitoring and controlling character traits in language models. *arXiv preprint arXiv:2507.21509* 2025.
195. Handa, G.; Wu, Z.; Koshiyama, A.; Treleaven, P.C. Personality as a Probe for LLM Evaluation: Method Trade-offs and Downstream Effects. In Proceedings of the NeurIPS 2025 Workshop on Evaluating the Evolving LLM Lifecycle: Benchmarks, Emergent Abilities, and Scaling, 2025.
196. Su, Y.; Zhang, J.; Yang, S.; Wang, X.; Hu, L.; Wang, D. Understanding How Value Neurons Shape the Generation of Specified Values in LLMs. In Proceedings of the Findings of the Association for Computational Linguistics: EMNLP 2025; Christodoulopoulos, C.; Chakraborty, T.; Rose, C.; Peng, V., Eds., Suzhou, China, 2025; pp. 9433–9452. <https://doi.org/10.18653/v1/2025.findings-emnlp.501>.
197. Deng, J.; Tang, T.; Yin, Y.; yang, W.; Zhao, X.; Wen, J.R. Neuron based Personality Trait Induction in Large Language Models. In Proceedings of the The Thirteenth International Conference on Learning Representations, 2025.
198. Chen, W.; Huang, Z.; Xie, L.; Lin, B.; Li, H.; Lu, L.; Tian, X.; Cai, D.; Zhang, Y.; Wang, W.; et al. From Yes-Men to Truth-Tellers: Addressing Sycophancy in Large Language Models with Pinpoint Tuning. In Proceedings of the Forty-first International Conference on Machine Learning, 2024.
199. Tak, A.N.; Banayeezade, A.; Bolourani, A.; Kian, M.; Jia, R.; Gratch, J. Mechanistic Interpretability of Emotion Inference in Large Language Models. In Proceedings of the Findings of the Association for Computational Linguistics: ACL 2025; Che, W.; Nabende, J.; Shutova, E.; Pilehvar, M.T., Eds., Vienna, Austria, 2025; pp. 13090–13120. <https://doi.org/10.18653/v1/2025.findings-acl.679>.
200. Yuan, S.; Qu, Z.; Tawfelis, M.; Färber, M. From Monolingual to Bilingual: Investigating Language Conditioning in Large Language Models for Psycholinguistic Tasks. *arXiv preprint arXiv:2508.02502* 2025.
201. Ju, T.; Shao, Z.; Wang, B.; Chen, Y.; Zhang, Z.; Fei, H.; Lee, M.L.; Hsu, W.; Duan, S.; Liu, G. Probing then Editing Response Personality of Large Language Models. In Proceedings of the Second Conference on Language Modeling, 2025.
202. Karny, S.; Baez, A.; Pataranutaporn, P. Neural Transparency: Mechanistic Interpretability Interfaces for Anticipating Model Behaviors for Personalized AI. *arXiv preprint arXiv:2511.00230* 2025.

203. Banayeeanzade, A.; Tak, A.N.; Bahrani, F.; Bolourani, A.; Blas, L.; Ferrara, E.; Gratch, J.; Karimireddy, S.P. Psychological Steering in LLMs: An Evaluation of Effectiveness and Trustworthiness. *arXiv preprint arXiv:2510.04484* 2025.
204. Zhao, Y.; Zhang, W.; Chen, G.; Kawaguchi, K.; Bing, L. How do Large Language Models Handle Multilingualism? In Proceedings of the Advances in Neural Information Processing Systems; Globerson, A.; Mackey, L.; Belgrave, D.; Fan, A.; Paquet, U.; Tomczak, J.; Zhang, C., Eds. Curran Associates, Inc., 2024, Vol. 37, pp. 15296–15319. <https://doi.org/10.52202/079017-0489>.
205. Liu, Y.; Chen, R.; Hirlimann, L.; Hakimi, A.D.; Wang, M.; Kargaran, A.H.; Rothe, S.; Yvon, F.; Schuetze, H. On Relation-Specific Neurons in Large Language Models. In Proceedings of the Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing; Christodoulopoulos, C.; Chakraborty, T.; Rose, C.; Peng, V., Eds., Suzhou, China, 2025; pp. 992–1022. <https://doi.org/10.18653/v1/2025.emnlp-main.52>.
206. Jing, Y.; Yao, Z.; Guo, H.; Ran, L.; Wang, X.; Hou, L.; Li, J. LinguaLens: Towards Interpreting Linguistic Mechanisms of Large Language Models via Sparse Auto-Encoder. In Proceedings of the Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing; Christodoulopoulos, C.; Chakraborty, T.; Rose, C.; Peng, V., Eds., Suzhou, China, 2025; pp. 28220–28239. <https://doi.org/10.18653/v1/2025.emnlp-main.1433>.
207. Brinkmann, J.; Wendler, C.; Bartelt, C.; Mueller, A. Large Language Models Share Representations of Latent Grammatical Concepts Across Typologically Diverse Languages. In Proceedings of the Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers); Chiruzzo, L.; Ritter, A.; Wang, L., Eds., Albuquerque, New Mexico, 2025; pp. 6131–6150. <https://doi.org/10.18653/v1/2025.naacl-long.312>.
208. Philippy, F.; Guo, S.; Haddadan, S. Identifying the Correlation Between Language Distance and Cross-Lingual Transfer in a Multilingual Representation Space. In Proceedings of the Proceedings of the 5th Workshop on Research in Computational Linguistic Typology and Multilingual NLP; Beinborn, L.; Goswami, K.; Muradoğlu, S.; Sorokin, A.; Kumar, R.; Shcherbakov, A.; Ponti, E.M.; Cotterell, R.; Vylomova, E., Eds., Dubrovnik, Croatia, 2023; pp. 22–29. <https://doi.org/10.18653/v1/2023.sigtyp-1.3>.
209. Mousi, B.; Durrani, N.; Dalvi, F.; Hawasly, M.; Abdelali, A. Exploring Alignment in Shared Cross-lingual Spaces. In Proceedings of the Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers); Ku, L.W.; Martins, A.; Srikumar, V., Eds., Bangkok, Thailand, 2024; pp. 6326–6348. <https://doi.org/10.18653/v1/2024.acl-long.344>.
210. Chi, Z.; Huang, H.; Mao, X.L. Can Cross-Lingual Transferability of Multilingual Transformers Be Activated Without End-Task Data? In Proceedings of the Findings of the Association for Computational Linguistics: ACL 2023; Rogers, A.; Boyd-Graber, J.; Okazaki, N., Eds., Toronto, Canada, 2023; pp. 12572–12584. <https://doi.org/10.18653/v1/2023.findings-acl.796>.
211. Hinck, M.; Holtermann, C.; Olson, M.L.; Schneider, F.; Yu, S.; Bhiwandiwalla, A.; Lauscher, A.; Tseng, S.Y.; Lal, V. Why do LLaVA Vision-Language Models Reply to Images in English? In Proceedings of the Findings of the Association for Computational Linguistics: EMNLP 2024; Al-Onaizan, Y.; Bansal, M.; Chen, Y.N., Eds., Miami, Florida, USA, 2024; pp. 13402–13421. <https://doi.org/10.18653/v1/2024.findings-emnlp.783>.
212. Zhang, H.; Shang, C.; Wang, S.; Zhang, D.; Yu, Y.; Yao, F.; Sun, R.; Yang, Y.; Wei, F. ShifCon: Enhancing Non-Dominant Language Capabilities with a Shift-based Multilingual Contrastive Framework. In Proceedings of the Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers); Che, W.; Nabende, J.; Shutova, E.; Pilehvar, M.T., Eds., Vienna, Austria, 2025; pp. 4818–4841. <https://doi.org/10.18653/v1/2025.acl-long.239>.
213. Wang, M.; Adel, H.; Lange, L.; Liu, Y.; Nie, E.; Strötgen, J.; Schuetze, H. Lost in Multilinguality: Dissecting Cross-lingual Factual Inconsistency in Transformer Language Models. In Proceedings of the Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers); Che, W.; Nabende, J.; Shutova, E.; Pilehvar, M.T., Eds., Vienna, Austria, 2025; pp. 5075–5094. <https://doi.org/10.18653/v1/2025.acl-long.253>.

214. Wang, M.; Lange, L.; Adel, H.; Ma, Y.; Strötgen, J.; Schuetze, H. Language Mixing in Reasoning Language Models: Patterns, Impact, and Internal Causes. In Proceedings of the Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing; Christodoulopoulos, C.; Chakraborty, T.; Rose, C.; Peng, V., Eds., Suzhou, China, 2025; pp. 2637–2665. <https://doi.org/10.18653/v1/2025.emnlp-main.132>.
215. Liu, Y.; Wang, M.; Kargaran, A.H.; Körner, F.; Nie, E.; Plank, B.; Yvon, F.; Schuetze, H. Tracing Multilingual Factual Knowledge Acquisition in Pretraining. In Proceedings of the Findings of the Association for Computational Linguistics: EMNLP 2025; Christodoulopoulos, C.; Chakraborty, T.; Rose, C.; Peng, V., Eds., Suzhou, China, 2025; pp. 2121–2146. <https://doi.org/10.18653/v1/2025.findings-emnlp.113>.
216. Chen, Y.; Cao, P.; Chen, Y.; Liu, K.; Zhao, J. Knowledge Localization: Mission Not Accomplished? Enter Query Localization! In Proceedings of the The Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025. OpenReview.net, 2025.
217. Katz, S.; Belinkov, Y.; Geva, M.; Wolf, L. Backward Lens: Projecting Language Model Gradients into the Vocabulary Space. In Proceedings of the Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing, EMNLP 2024, Miami, FL, USA, November 12-16, 2024; Al-Onaizan, Y.; Bansal, M.; Chen, Y., Eds. Association for Computational Linguistics, 2024, pp. 2390–2422. <https://doi.org/10.18653/V1/2024.EMNLP-MAIN.142>.
218. Lai, W.; Fraser, A.; Titov, I. Joint Localization and Activation Editing for Low-Resource Fine-Tuning. In Proceedings of the Forty-second International Conference on Machine Learning, ICML 2025, Vancouver, BC, Canada, July 13-19, 2025. OpenReview.net, 2025.
219. Muhamed, A.; Bonato, J.; Diab, M.T.; Smith, V. SAEs Can Improve Unlearning: Dynamic Sparse Autoencoder Guardrails for Precision Unlearning in LLMs. *CoRR* 2025, *abs/2504.08192*, [2504.08192]. <https://doi.org/10.48550/ARXIV.2504.08192>.
220. Jin, Z.; Cao, P.; Yuan, H.; Chen, Y.; Xu, J.; Li, H.; Jiang, X.; Liu, K.; Zhao, J. Cutting Off the Head Ends the Conflict: A Mechanism for Interpreting and Mitigating Knowledge Conflicts in Language Models. In Proceedings of the Findings of the Association for Computational Linguistics, ACL 2024, Bangkok, Thailand and virtual meeting, August 11-16, 2024; Ku, L.; Martins, A.; Srikumar, V., Eds. Association for Computational Linguistics, 2024, pp. 1193–1215. <https://doi.org/10.18653/V1/2024.FINDINGS-ACL.70>.
221. Li, G.; Chen, Y.; Tong, H. Taming Knowledge Conflicts in Language Models. In Proceedings of the Forty-second International Conference on Machine Learning, ICML 2025, Vancouver, BC, Canada, July 13-19, 2025. OpenReview.net, 2025.
222. Wu, Z.; Arora, A.; Wang, Z.; Geiger, A.; Jurafsky, D.; Manning, C.D.; Potts, C. ReFT: Representation Finetuning for Language Models. In Proceedings of the Advances in Neural Information Processing Systems 38: Annual Conference on Neural Information Processing Systems 2024, NeurIPS 2024, Vancouver, BC, Canada, December 10 - 15, 2024; Globersons, A.; Mackey, L.; Belgrave, D.; Fan, A.; Paquet, U.; Tomczak, J.M.; Zhang, C., Eds., 2024.
223. Chen, A.; Merullo, J.; Stolfo, A.; Pavlick, E. Transferring Linear Features Across Language Models With Model Stitching. In Proceedings of the The Thirty-ninth Annual Conference on Neural Information Processing Systems, 2025.
224. Yadav, P.; Tam, D.; Choshen, L.; Raffel, C.A.; Bansal, M. TIES-Merging: Resolving Interference When Merging Models. In Proceedings of the Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023; Oh, A.; Naumann, T.; Globerson, A.; Saenko, K.; Hardt, M.; Levine, S., Eds., 2023.
225. Quirke, P.; Barez, F. Understanding Addition in Transformers. In Proceedings of the The Twelfth International Conference on Learning Representations, 2024.
226. Yang, H.; Zhao, Q.; Li, L. Chain-of-Thought in Large Language Models: Decoding, Projection, and Activation, 2024, [arXiv:cs.AI/2412.03944].
227. Venhoff, C.; Arcuschin, I.; Torr, P.; Conmy, A.; Nanda, N. Understanding Reasoning in Thinking Language Models via Steering Vectors, 2025, [arXiv:cs.LG/2506.18167].
228. Ward, J.; Lin, C.; Venhoff, C.; Nanda, N. Reasoning-Finetuning Repurposes Latent Representations in Base Models. *CoRR* 2025, *abs/2507.12638*, [2507.12638]. <https://doi.org/10.48550/ARXIV.2507.12638>.

229. Troitskii, D.; Pal, K.; Wandler, C.; McDougall, C.S. Internal states before wait modulate reasoning patterns. In Proceedings of the Findings of the Association for Computational Linguistics: EMNLP 2025; Christodoulopoulos, C.; Chakraborty, T.; Rose, C.; Peng, V., Eds., Suzhou, China, 2025; pp. 18640–18649. <https://doi.org/10.18653/v1/2025.findings-emnlp.1012>.
230. Wang, Z.; Ma, Y.; Xu, C. Eliciting Chain-of-Thought in Base LLMs via Gradient-Based Representation Optimization, 2025, [arXiv:cs.CL/2511.19131].
231. Sun, Y.; Stolfo, A.; Sachan, M. Probing for Arithmetic Errors in Language Models. In Proceedings of the Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing, Suzhou, China, 2025; pp. 8122–8139. <https://doi.org/10.18653/v1/2025.emnlp-main.411>.
232. Cywiński, B.; Bussmann, B.; Conmy, A.; Engels, J.; Nanda, N.; Rajamanoharan, S. Can we interpret latent reasoning using current mechanistic interpretability tools?, 2025.
233. Højer, B.; Jarvis, O.; Heinrich, S. Improving Reasoning Performance in Large Language Models via Representation Engineering, 2025, [arXiv:cs.LG/2504.19483].
234. Tang, X.; Wang, X.; Lv, Z.; Min, Y.; Zhao, W.X.; Hu, B.; Liu, Z.; Zhang, Z. Unlocking General Long Chain-of-Thought Reasoning Capabilities of Large Language Models via Representation Engineering. In Proceedings of the Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers); Che, W.; Nabende, J.; Shutova, E.; Pilehvar, M.T., Eds., Vienna, Austria, 2025; pp. 6832–6849. <https://doi.org/10.18653/v1/2025.acl-long.339>.
235. Hong, Y.; Cao, M.; Zhou, D.; Yu, L.; Jin, Z. The Reasoning-Memorization Interplay in Language Models Is Mediated by a Single Direction. In Proceedings of the Findings of the Association for Computational Linguistics: ACL 2025; Che, W.; Nabende, J.; Shutova, E.; Pilehvar, M.T., Eds., Vienna, Austria, 2025; pp. 21565–21585. <https://doi.org/10.18653/v1/2025.findings-acl.1111>.
236. Zhang, J.; Viteri, S. Uncovering Latent Chain of Thought Vectors in Language Models, 2025, [arXiv:cs.CL/2409.14026].
237. Liu, S.; Chen, T.; Lu, P.; Ye, H.; Chen, Y.; Xing, L.; Zou, J. Fractional Reasoning via Latent Steering Vectors Improves Inference Time Compute, 2025, [arXiv:cs.LG/2506.15882].
238. Sini, V.; Gorbatoevski, A.; Cherepanov, A.; Shaposhnikov, B.; Balagansky, N.; Gavrilov, D. Steering LLM Reasoning Through Bias-Only Adaptation, 2025, [arXiv:cs.LG/2505.18706].
239. Li, Z.; Wang, X.; Yang, Y.; Yao, Z.; Xiong, H.; Du, M. Feature Extraction and Steering for Enhanced Chain-of-Thought Reasoning in Language Models. In Proceedings of the Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing; Christodoulopoulos, C.; Chakraborty, T.; Rose, C.; Peng, V., Eds., Suzhou, China, 2025; pp. 10893–10913. <https://doi.org/10.18653/v1/2025.emnlp-main.552>.
240. Song, W.; Li, Z.; Zhang, L.; Zhao, H.; Du, B. Sparse is enough in fine-tuning pre-trained large language models. In Proceedings of the Proceedings of the 41st International Conference on Machine Learning. JMLR.org, 2024, ICML'24.
241. Mondal, S.K.; Sen, S.; Singhania, A.; Jyothi, P. Language-Specific Neurons Do Not Facilitate Cross-Lingual Transfer. In Proceedings of the The Sixth Workshop on Insights from Negative Results in NLP; Drozd, A.; Sedoc, J.; Tafreshi, S.; Akula, A.; Shu, R., Eds., Albuquerque, New Mexico, 2025; pp. 46–62. <https://doi.org/10.18653/v1/2025.insights-1.6>.
242. Gurgurov, D.; van Genabith, J.; Ostermann, S. Sparse Subnetwork Enhancement for Underrepresented Languages in Large Language Models, 2025, [arXiv:cs.CL/2510.13580].
243. Li, Y.; Gao, W.; Yuan, C.; Wang, X. Fine-Tuning is Subgraph Search: A New Lens on Learning Dynamics, 2025, [arXiv:cs.LG/2502.06106].
244. Hoogland, J.; Wang, G.; Farrugia-Roberts, M.; Carroll, L.; Wei, S.; Murfet, D. The developmental landscape of in-context learning, 2024. URL <https://arxiv.org/abs/2402.02364> 2024.
245. Minegishi, G.; Furuta, H.; Taniguchi, S.; Iwasawa, Y.; Matsuo, Y. In-Context Meta Learning Induces Multi-Phase Circuit Emergence. In Proceedings of the ICLR 2025 Workshop on Building Trust in Language Models and Applications, 2025.

246. Nanda, N.; Chan, L.; Lieberum, T.; Smith, J.; Steinhardt, J. Progress measures for grokking via mechanistic interpretability. In Proceedings of the The Eleventh International Conference on Learning Representations, 2023.
247. Liu, Z.; Michaud, E.J.; Tegmark, M. Omnigrok: Grokking Beyond Algorithmic Data. In Proceedings of the The Eleventh International Conference on Learning Representations, 2023.
248. Furuta, H.; Minegishi, G.; Iwasawa, Y.; Matsuo, Y. Towards Empirical Interpretation of Internal Circuits and Properties in Grokked Transformers on Modular Polynomials. *Transactions on Machine Learning Research* **2024**.
249. Qiye, H.; Hao, Z.; RuoXi, Y. Exploring Grokking: Experimental and Mechanistic Investigations. *arXiv preprint arXiv:2412.10898* **2024**.
250. Li, Z.; Fan, C.; Zhou, T. Where to find Grokking in LLM Pretraining? Monitor Memorization-to-Generalization without Test. *arXiv preprint arXiv:2506.21551* **2025**.
251. Lei, L.; Gu, J.; Ma, X.; Tang, C.; Chen, J.; Xu, T. Generic Token Compression in Multimodal Large Language Models from an Explainability Perspective, 2025, [[arXiv:cs.CV/2506.01097](https://arxiv.org/abs/cs.CV/2506.01097)].
252. Guo, Z.; Kamigaito, H.; Watanabe, T. Attention Score is not All You Need for Token Importance Indicator in KV Cache Reduction: Value Also Matters. In Proceedings of the Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing; Al-Onaizan, Y.; Bansal, M.; Chen, Y.N., Eds., Miami, Florida, USA, 2024; pp. 21158–21166. <https://doi.org/10.18653/v1/2024.emnlp-main.1178>.
253. Ye, W.; Wu, Q.; Lin, W.; Zhou, Y. Fit and prune: Fast and training-free visual token pruning for multi-modal large language models. In Proceedings of the Proceedings of the AAAI Conference on Artificial Intelligence, 2025, Vol. 39, pp. 22128–22136.
254. Xiao, G.; Tang, J.; Zuo, J.; Guo, J.; Yang, S.; Tang, H.; Fu, Y.; Han, S. Duoattention: Efficient long-context llm inference with retrieval and streaming heads. *arXiv preprint arXiv:2410.10819* **2024**.
255. Laitenberger, F.; Kopiczko, D.; Snoek, C.G.M.; Asano, Y.M. What Layers When: Learning to Skip Compute in LLMs with Residual Gates, 2025, [[arXiv:cs.CL/2510.13876](https://arxiv.org/abs/cs.CL/2510.13876)].
256. Valade, F. Accelerating Large Language Model Inference with Self-Supervised Early Exits, 2024, [[arXiv:cs.CL/2407.21082](https://arxiv.org/abs/cs.CL/2407.21082)].
257. Wang, H.; Wang, Y.; Liu, T.; Zhao, T.; Gao, J. HadSkip: Homotopic and Adaptive Layer Skipping of Pre-trained Language Models for Efficient Inference. In Proceedings of the Findings of the Association for Computational Linguistics: EMNLP 2023; Bouamor, H.; Pino, J.; Bali, K., Eds., Singapore, 2023; pp. 4283–4294. <https://doi.org/10.18653/v1/2023.findings-emnlp.283>.
258. Shelke, A.; Savant, R.; Joshi, R. Towards Building Efficient Sentence BERT Models using Layer Pruning. In Proceedings of the Proceedings of the 38th Pacific Asia Conference on Language, Information and Computation, 2024, pp. 720–725.
259. Lu, X.; Liu, Q.; Xu, Y.; Zhou, A.; Huang, S.; Zhang, B.; Yan, J.; Li, H. Not all experts are equal: Efficient expert pruning and skipping for mixture-of-experts large language models. *arXiv preprint arXiv:2402.14800* **2024**.
260. Tan, S.; Wu, D.; Monz, C. Neuron Specialization: Leveraging Intrinsic Task Modularity for Multilingual Machine Translation. In Proceedings of the Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing; Al-Onaizan, Y.; Bansal, M.; Chen, Y.N., Eds., Miami, Florida, USA, 2024; pp. 6506–6527. <https://doi.org/10.18653/v1/2024.emnlp-main.374>.
261. Xiao, H.; Yang, Q.; Xie, D.; Xu, W.; Zhou, W.; Liu, H.; Liu, Z.; Wong, N. Exploring Layer-wise Information Effectiveness for Post-Training Quantization in Small Language Models. *arXiv preprint arXiv:2508.03332* **2025**.
262. Ranjan, N.; Savakis, A. Mix-QViT: Mixed-precision vision transformer quantization driven by layer importance and quantization sensitivity. *arXiv preprint arXiv:2501.06357* **2025**.
263. Zeng, B.; Ji, B.; Liu, X.; Yu, J.; Li, S.; Ma, J.; Li, X.; Wang, S.; Hong, X.; Tang, Y. Lsq: Layer-specific adaptive quantization for large language model deployment. *arXiv preprint arXiv:2412.18135* **2024**.
264. Li, Z.Z.; Zhang, D.; Zhang, M.L.; Zhang, J.; Liu, Z.; Yao, Y.; Xu, H.; Zheng, J.; Wang, P.J.; Chen, X.; et al. From system 1 to system 2: A survey of reasoning large language models. *arXiv preprint arXiv:2502.17419* **2025**.

265. Ren, Z.; Shao, Z.; Song, J.; Xin, H.; Wang, H.; Zhao, W.; Zhang, L.; Fu, Z.; Zhu, Q.; Yang, D.; et al. Deepseek-prover-v2: Advancing formal mathematical reasoning via reinforcement learning for subgoal decomposition. *arXiv preprint arXiv:2504.21801* **2025**.
266. Dubey, A.; Jauhri, A.; Pandey, A.; Kadian, A.; Al-Dahle, A.; Letman, A.; Mathur, A.; Schelten, A.; Yang, A.; Fan, A.; et al. The llama 3 herd of models. *arXiv e-prints* **2024**, pp. arXiv-2407.
267. OpenAI.; Achiam, J.; Adler, S.; Agarwal, S.; Ahmad, L.; Akkaya, I.; Aleman, F.L.; Almeida, D.; Altenschmidt, J.; Altman, S.; et al. GPT-4 Technical Report, 2024, [[arXiv:cs.CL/2303.08774](https://arxiv.org/abs/cs.CL/2303.08774)].
268. Liang, X.; Li, Z.Z.; Gong, Y.; Wang, Y.; Zhang, H.; Shen, Y.; Wu, Y.N.; Chen, W. SwS: Self-aware Weakness-driven Problem Synthesis in Reinforcement Learning for LLM Reasoning. *arXiv preprint arXiv:2506.08989* **2025**.
269. Qin, L.; Chen, Q.; Zhou, Y.; Chen, Z.; Li, Y.; Liao, L.; Li, M.; Che, W.; Yu, P.S. A survey of multilingual large language models. *Patterns* **2025**, 6.
270. Yang, H.; Chen, H.; Guo, H.; Chen, Y.; Lin, C.S.; Hu, S.; Hu, J.; Wu, X.; Wang, X. LLM-MedQA: Enhancing Medical Question Answering through Case Studies in Large Language Models. *arXiv preprint arXiv:2501.05464* **2024**.
271. Chang, Y.; Li, Z.; Zhang, H.; Kong, Y.; Wu, Y.; Guo, Z.; Wong, N. TreeReview: A Dynamic Tree of Questions Framework for Deep and Efficient LLM-based Scientific Peer Review. *arXiv preprint arXiv:2506.07642* **2025**.
272. Li, S.; Yang, C.; Wu, T.; Shi, C.; Zhang, Y.; Zhu, X.; Cheng, Z.; Cai, D.; Yu, M.; Liu, L.; et al. A Survey on the Honesty of Large Language Models. *arXiv preprint arXiv:2409.18786* **2024**.
273. Zhao, H.; Liu, Z.; Wu, Z.; Li, Y.; Yang, T.; Shu, P.; Xu, S.; Dai, H.; Zhao, L.; Mai, G.; et al. Revolutionizing finance with llms: An overview of applications and insights. *arXiv preprint arXiv:2401.11641* **2024**.
274. Yu, Y.; Zhang, Y.; Zhang, D.; Liang, X.; Zhang, H.; Zhang, X.; Khademi, M.; Awadalla, H.H.; Wang, J.; Yang, Y.; et al. Chain-of-Reasoning: Towards Unified Mathematical Reasoning in Large Language Models via a Multi-Paradigm Perspective. In Proceedings of the Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers); Che, W.; Nabende, J.; Shutova, E.; Pilehvar, M.T., Eds., Vienna, Austria, 2025; pp. 24914–24937. <https://doi.org/10.18653/v1/2025.acl-long.1213>.
275. Qwen.; ; Yang, A.; Yang, B.; Zhang, B.; Hui, B.; Zheng, B.; Yu, B.; Li, C.; Liu, D.; et al. Qwen2.5 Technical Report, 2025, [[arXiv:cs.CL/2412.15115](https://arxiv.org/abs/cs.CL/2412.15115)].
276. Huang, B.; Wu, X.; Zhou, Y.; Wu, J.; Feng, L.; Cheng, R.; Tan, K.C. Exploring the true potential: Evaluating the black-box optimization capability of large language models. *arXiv preprint arXiv:2404.06290* **2024**.
277. Hong, J.; Tu, Q.; Chen, C.; Xing, G.; Zhang, J.; Yan, R. Cyclealign: Iterative distillation from black-box llm to white-box models for better human alignment. In Proceedings of the Findings of the Association for Computational Linguistics: ACL 2024, 2024, pp. 14596–14609.
278. Zhang, J.; Ding, M.; Liu, Y.; Hong, J.; Tramèr, F. Black-box Optimization of LLM Outputs by Asking for Directions. *arXiv preprint arXiv:2510.16794* **2025**.
279. Ferrando, J.; Sarti, G.; Bisazza, A.; Costa-Jussà, M.R. A primer on the inner workings of transformer-based language models. *arXiv preprint arXiv:2405.00208* **2024**.
280. Zhao, H.; Chen, H.; Yang, F.; Liu, N.; Deng, H.; Cai, H.; Wang, S.; Yin, D.; Du, M. Explainability for large language models: A survey. *ACM Transactions on Intelligent Systems and Technology* **2024**, 15, 1–38.
281. Räuher, T.; Ho, A.; Casper, S.; Hadfield-Menell, D. Toward transparent ai: A survey on interpreting the inner structures of deep neural networks. In Proceedings of the 2023 IEEE Conference on Secure and Trustworthy Machine Learning (SATML). IEEE, 2023, pp. 464–483.
282. Allen-Zhu, Z.; Li, Y. Physics of language models: Part 1, learning hierarchical language structures. *arXiv preprint arXiv:2305.13673* **2023**.
283. Zheng, Z.; Wang, Y.; Huang, Y.; Song, S.; Yang, M.; Tang, B.; Xiong, F.; Li, Z. Attention heads of large language models: A survey. *arXiv preprint arXiv:2409.03752* **2024**.
284. Saphra, N.; Wiegrefe, S. Mechanistic? *arXiv preprint arXiv:2410.09087* **2024**.
285. López-Otal, M.; Gracia, J.; Bernad, J.; Bobed, C.; Pitarch-Ballesteros, L.; Anglés-Herrero, E. Linguistic Interpretability of Transformer-based Language Models: a systematic review. *arXiv preprint arXiv:2504.08001* **2025**.

286. Gantla, S.R. Exploring Mechanistic Interpretability in Large Language Models: Challenges, Approaches, and Insights. In Proceedings of the 2025 International Conference on Data Science, Agents & Artificial Intelligence (ICDSSAI). IEEE, 2025, pp. 1–8.
287. Luo, H.; Specia, L. From understanding to utilization: A survey on explainability for large language models. *arXiv preprint arXiv:2401.12874* **2024**.
288. Wu, X.; Zhao, H.; Zhu, Y.; Shi, Y.; Yang, F.; Hu, L.; Liu, T.; Zhai, X.; Yao, W.; Li, J.; et al. Usable XAI: 10 strategies towards exploiting explainability in the LLM era. *arXiv preprint arXiv:2403.08946* **2024**.
289. Rai, D.; Zhou, Y.; Feng, S.; Saparov, A.; Yao, Z. A practical review of mechanistic interpretability for transformer-based language models. *arXiv preprint arXiv:2407.02646* **2024**.
290. Bereska, L.; Gavves, E. Mechanistic interpretability for AI safety—a review. *arXiv preprint arXiv:2404.14082* **2024**.
291. Lee, S.; Cho, A.; Kim, G.C.; Peng, S.; Phute, M.; Chau, D.H. Interpretation Meets Safety: A Survey on Interpretation Methods and Tools for Improving LLM Safety. In Proceedings of the Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing; Christodoulopoulos, C.; Chakraborty, T.; Rose, C.; Peng, V., Eds., Suzhou, China, 2025; pp. 21514–21545. <https://doi.org/10.18653/v1/2025.emnlp-main.1091>.
292. Resck, L.; Augenstein, I.; Korhonen, A. Explainability and interpretability of multilingual large language models: A survey. In Proceedings of the Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing, 2025, pp. 20465–20497.
293. Lin, Z.; Basu, S.; Beigi, M.; Manjunatha, V.; Rossi, R.A.; Wang, Z.; Zhou, Y.; Balasubramanian, S.; Zarei, A.; Rezaei, K.; et al. A survey on mechanistic interpretability for multi-modal foundation models. *arXiv preprint arXiv:2502.17516* **2025**.
294. Radford, A.; Wu, J.; Child, R.; Luan, D.; Amodei, D.; Sutskever, I.; et al. Language models are unsupervised multitask learners. *OpenAI blog* **2019**, *1*, 9.
295. Su, J.; Ahmed, M.; Lu, Y.; Pan, S.; Bo, W.; Liu, Y. Roformer: Enhanced transformer with rotary position embedding. *Neurocomputing* **2024**, *568*, 127063.
296. Bricken, T.; Pehlevan, C. Attention approximates sparse distributed memory. *Advances in Neural Information Processing Systems* **2021**, *34*, 15301–15315.
297. Li, R.; Gao, Y. Anchored Answers: Unravelling Positional Bias in GPT-2’s Multiple-Choice Questions, 2025, [arXiv:cs/2405.03205]. <https://doi.org/10.48550/arXiv.2405.03205>.
298. Voita, E.; Talbot, D.; Moiseev, F.; Sennrich, R.; Titov, I. Analyzing Multi-Head Self-Attention: Specialized Heads Do the Heavy Lifting, the Rest Can Be Pruned. In Proceedings of the Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics, 2019, pp. 5797–5808.
299. Feng, J.; Steinhardt, J. How do language models bind entities in context? In Proceedings of the NeurIPS 2023 Workshop on Symmetry and Geometry in Neural Representations, 2023.
300. Men, T.; Cao, P.; Jin, Z.; Chen, Y.; Liu, K.; Zhao, J. Unlocking the Future: Exploring Look-Ahead Planning Mechanistic Interpretability in Large Language Models. In Proceedings of the Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing, 2024, pp. 7713–7724.
301. Geva, M.; Caciularu, A.; Wang, K.; Goldberg, Y. Transformer feed-forward layers build predictions by promoting concepts in the vocabulary space. In Proceedings of the Proceedings of the 2022 conference on empirical methods in natural language processing, 2022, pp. 30–45.
302. Shazeer, N. GLU Variants Improve Transformer, 2020, [arXiv:cs.LG/2002.05202].
303. Elhage, N.; Hume, T.; Olsson, C.; Schiefer, N.; Henighan, T.; Kravec, S.; Hatfield-Dodds, Z.; Lasenby, R.; Drain, D.; Chen, C.; et al. Toy Models of Superposition. *CoRR* **2022**, *abs/2209.10652*, [2209.10652]. <https://doi.org/10.48550/ARXIV.2209.10652>.
304. Lieberum, T.; Rajamanoharan, S.; Conmy, A.; Smith, L.; Sonnerat, N.; Varma, V.; Kramar, J.; Dragan, A.; Shah, R.; Nanda, N. Gemma Scope: Open Sparse Autoencoders Everywhere All At Once on Gemma 2. In Proceedings of the Proceedings of the 7th BlackboxNLP Workshop: Analyzing and Interpreting Neural Networks for NLP; Belinkov, Y.; Kim, N.; Jumelet, J.; Mohebbi, H.; Mueller, A.; Chen, H., Eds., Miami, Florida, US, 2024; pp. 278–300. <https://doi.org/10.18653/v1/2024.blackboxnlp-1.19>.

305. He, Z.; Shu, W.; Ge, X.; Chen, L.; Wang, J.; Zhou, Y.; Liu, F.; Guo, Q.; Huang, X.; Wu, Z.; et al. Llama scope: Extracting millions of features from llama-3.1-8b with sparse autoencoders. *arXiv preprint arXiv:2410.20526* **2024**.
306. Cunningham, H.; Ewart, A.; Riggs, L.; Huben, R.; Sharkey, L. Sparse autoencoders find highly interpretable features in language models. *arXiv preprint arXiv:2309.08600* **2023**.
307. Bloom, J. Open Source Sparse Autoencoders for all Residual Stream Layers of GPT2 Small, 2024.
308. Ghilardi, D.; Belotti, F.; Molinari, M. Efficient Training of Sparse Autoencoders for Large Language Models via Layer Groups. *arXiv preprint arXiv:2410.21508* **2024**.
309. Mudide, A.; Engels, J.; Michaud, E.J.; Tegmark, M.; de Witt, C.S. Efficient dictionary learning with switch sparse autoencoders. *arXiv preprint arXiv:2410.08201* **2024**.
310. Xu, Z.; Tan, Z.; Wang, S.; Xu, K.; Chen, T. Beyond Redundancy: Diverse and Specialized Multi-Expert Sparse Autoencoder. *arXiv preprint arXiv:2511.05745* **2025**.
311. Cho, I.; Hockenmaier, J. Toward Efficient Sparse Autoencoder-Guided Steering for Improved In-Context Learning in Large Language Models. In Proceedings of the Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing, 2025, pp. 28949–28961.
312. Gao, L.; la Tour, T.D.; Tillman, H.; Goh, G.; Troll, R.; Radford, A.; Sutskever, I.; Leike, J.; Wu, J. Scaling and evaluating sparse autoencoders. *arXiv preprint arXiv:2406.04093* **2024**.
313. Rajamanoharan, S.; Conmy, A.; Smith, L.; Lieberum, T.; Varma, V.; Kramár, J.; Shah, R.; Nanda, N. Improving dictionary learning with gated sparse autoencoders. *arXiv preprint arXiv:2404.16014* **2024**.
314. Bussmann, B.; Leask, P.; Nanda, N. Batchtopk sparse autoencoders. *arXiv preprint arXiv:2412.06410* **2024**.
315. Rajamanoharan, S.; Lieberum, T.; Sonnerat, N.; Conmy, A.; Varma, V.; Kramár, J.; Nanda, N. Jumping ahead: Improving reconstruction fidelity with jumprelu sparse autoencoders. *arXiv preprint arXiv:2407.14435* **2024**.
316. Cho, H.; Yang, H.; Kurkoski, B.M.; Inoue, N. Binary Autoencoder for Mechanistic Interpretability of Large Language Models. *arXiv preprint arXiv:2509.20997* **2025**.
317. Kim, J.; Evans, J.; Schein, A. Linear Representations of Political Perspective Emerge in Large Language Models. In Proceedings of the The Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025. OpenReview.net, 2025.
318. Wang, Z.; Zhang, H.; Li, X.; Huang, K.H.; Han, C.; Ji, S.; Kakade, S.M.; Peng, H.; Ji, H. Eliminating Position Bias of Language Models: A Mechanistic Approach, 2025, [arXiv:cs/2407.01100]. <https://doi.org/10.48550/arXiv.2407.01100>.
319. Yu, Y.; Jiang, H.; Luo, X.; Wu, Q.; Lin, C.Y.; Li, D.; Yang, Y.; Huang, Y.; Qiu, L. Mitigate Position Bias in LLMs via Scaling a Single Hidden States Channel. In Proceedings of the Findings of the Association for Computational Linguistics: ACL 2025; Che, W.; Nabende, J.; Shutova, E.; Pilehvar, M.T., Eds., Vienna, Austria, 2025; pp. 6092–6111. <https://doi.org/10.18653/v1/2025.findings-acl.316>.
320. Dimino, F.; Saxena, K.; Sarmah, B.; Pasquali, S. Tracing Positional Bias in Financial Decision-Making: Mechanistic Insights from Qwen2.5. In Proceedings of the Proceedings of the 6th ACM International Conference on AI in Finance, 2025, pp. 96–104, [arXiv:q-fin/2508.18427]. <https://doi.org/10.1145/3768292.3770394>.
321. Pai, T.M.; Wang, J.I.; Lu, L.C.; Sun, S.H.; Lee, H.Y.; Chang, K.W. Billy: Steering large language models via merging persona vectors for creative generation. *arXiv preprint arXiv:2510.10157* **2025**.
322. Schwartz, S.H. Universals in the content and structure of values: Theoretical advances and empirical tests in 20 countries. In *Advances in experimental social psychology*; Elsevier, 1992; Vol. 25, pp. 1–65.
323. Joshi, N.; Rando, J.; Saparov, A.; Kim, N.; He, H. Personas as a way to model truthfulness in language models. In Proceedings of the Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing, 2024, pp. 6346–6359.
324. Ghandeharioun, A.; Yuan, A.; Guerard, M.; Reif, E.; Lepori, M.A.; Dixon, L. Who's asking? User personas and the mechanics of latent misalignment. In Proceedings of the Advances in Neural Information Processing Systems; Globerson, A.; Mackey, L.; Belgrave, D.; Fan, A.; Paquet, U.; Tomczak, J.; Zhang, C., Eds. Curran Associates, Inc., 2024, Vol. 37, pp. 125967–126003. <https://doi.org/10.52202/079017-4002>.

325. Kojima, T.; Okimura, I.; Iwasawa, Y.; Yanaka, H.; Matsuo, Y. On the Multilingual Ability of Decoder-based Pre-trained Language Models: Finding and Controlling Language-Specific Neurons. In Proceedings of the Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers); Duh, K.; Gomez, H.; Bethard, S., Eds., Mexico City, Mexico, 2024; pp. 6919–6971. <https://doi.org/10.18653/v1/2024.naacl-long.384>.
326. Li, X.L.; Liang, P. Prefix-Tuning: Optimizing Continuous Prompts for Generation, 2021, [arXiv:cs.CL/2101.00190].
327. Hu, E.J.; yelong shen.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; Chen, W. LoRA: Low-Rank Adaptation of Large Language Models. In Proceedings of the International Conference on Learning Representations, 2022.
328. Liu, X.; Ji, K.; Fu, Y.; Tam, W.; Du, Z.; Yang, Z.; Tang, J. P-Tuning: Prompt Tuning Can Be Comparable to Fine-tuning Across Scales and Tasks. In Proceedings of the Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers); Muresan, S.; Nakov, P.; Villavicencio, A., Eds., Dublin, Ireland, 2022; pp. 61–68. <https://doi.org/10.18653/v1/2022.acl-short.8>.
329. Varma, V.; Shah, R.; Kenton, Z.; Kramár, J.; Kumar, R. Explaining grokking through circuit efficiency. *arXiv preprint arXiv:2309.02390* **2023**.
330. Huang, Y.; Hu, S.; Han, X.; Liu, Z.; Sun, M. Unified View of Grokking, Double Descent and Emergent Abilities: A Comprehensive Study on Algorithm Task. In Proceedings of the First Conference on Language Modeling, 2024.
331. Notsawo Jr, P.; Zhou, H.; Pezeshki, M.; Rish, I.; Dumas, G.; et al. Predicting grokking long before it happens: A look into the loss landscape of models which grok. *arXiv preprint arXiv:2306.13253* **2023**.
332. Xiong, J.; Shen, J.; Ye, F.; Tao, C.; Wan, Z.; Lu, J.; Wu, X.; Zheng, C.; Guo, Z.; Kong, L.; et al. UNComp: Uncertainty-Aware Long-Context Compressor for Efficient Large Language Model Inference. *arXiv preprint arXiv:2410.03090* **2024**.
333. Xiong, J.; Shen, J.; Zheng, C.; Wan, Z.; Zhao, C.; Yang, C.; Ye, F.; Yang, H.; Kong, L.; Wong, N. ParallelComp: Parallel Long-Context Compressor for Length Extrapolation. *arXiv preprint arXiv:2502.14317* **2025**.
334. Kharlapenko, D.; Shabalin, S.; Barez, F.; Conmy, A.; Nanda, N. Scaling sparse feature circuit finding for in-context learning. *arXiv preprint arXiv:2504.13756* **2025**.
335. Nikankin, Y.; Arad, D.; Gandelsman, Y.; Belinkov, Y. Same Task, Different Circuits: Disentangling Modality-Specific Mechanisms in VLMs. *arXiv preprint arXiv:2506.09047* **2025**.
336. Duan, X.; Zhou, X.; Xiao, B.; Cai, Z. Unveiling Language Competence Neurons: A Psycholinguistic Approach to Model Interpretability. In Proceedings of the Proceedings of the 31st International Conference on Computational Linguistics; Rambow, O.; Wanner, L.; Apidianaki, M.; Al-Khalifa, H.; Eugenio, B.D.; Schockaert, S., Eds., Abu Dhabi, UAE, 2025; pp. 10148–10157.
337. He, Z.; Ge, X.; Tang, Q.; Sun, T.; Cheng, Q.; Qiu, X. Dictionary learning improves patch-free circuit discovery in mechanistic interpretability: A case study on othello-gpt. *arXiv preprint arXiv:2402.12201* **2024**.
338. Marks, S.; Rager, C.; Michaud, E.J.; Belinkov, Y.; Bau, D.; Mueller, A. Sparse Feature Circuits: Discovering and Editing Interpretable Causal Graphs in Language Models. In Proceedings of the The Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24–28, 2025. OpenReview.net, 2025.
339. Lindsey, J.; Gurnee, W.; Ameisen, E.; Chen, B.; Pearce, A.; Turner, N.L.; Citro, C.; Abrahams, D.; Carter, S.; Hosmer, B.; et al. On the Biology of a Large Language Model. *Transformer Circuits Thread* **2025**.
340. Nguyen, T.; Michaels, J.; Fiterau, M.; Jensen, D. Challenges in Understanding Modality Conflict in Vision-Language Models. *arXiv preprint arXiv:2509.02805* **2025**.
341. Prakash, N.; Shaham, T.R.; Haklay, T.; Belinkov, Y.; Bau, D. Fine-tuning enhances existing mechanisms: A case study on entity tracking. *arXiv preprint arXiv:2402.14811* **2024**.
342. Nanda, N. Attribution patching: Activation patching at industrial scale. URL: <https://www.neelnanda.io/mechanistic-interpretability/attribution-patching> **2023**.
343. Yu, Z.; Ananiadou, S. Neuron-Level Knowledge Attribution in Large Language Models. In Proceedings of the Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing, EMNLP

- 2024, Miami, FL, USA, November 12-16, 2024; Al-Onaizan, Y.; Bansal, M.; Chen, Y., Eds. Association for Computational Linguistics, 2024, pp. 3267–3280. <https://doi.org/10.18653/V1/2024.EMNLP-MAIN.191>.
344. Miller, J.; Chughtai, B.; Saunders, W. Transformer circuit faithfulness metrics are not robust. *arXiv preprint arXiv:2407.08734* 2024.
345. Parrack, A.; Attubato, C.L.; Heimersheim, S. Benchmarking deception probes via black-to-white performance boosts. *arXiv preprint arXiv:2507.12691* 2025.
346. Nguyen, J.; Hoang, K.; Attubato, C.L.; Hofstätter, F. Probing and Steering Evaluation Awareness of Language Models, July 2025. URL <http://arxiv.org/abs/2507.01786> 2025.
347. Wu, Z.; Arora, A.; Geiger, A.; Wang, Z.; Huang, J.; Jurafsky, D.; Manning, C.D.; Potts, C. AxBench: Steering LLMs? Even Simple Baselines Outperform Sparse Autoencoders. In Proceedings of the Forty-second International Conference on Machine Learning, 2025.
348. Karvonen, A.; Rager, C.; Lin, J.; Tigges, C.; Bloom, J.; Chanin, D.; Lau, Y.T.; Farrell, E.; McDougall, C.; Ayonrinde, K.; et al. Saebench: A comprehensive benchmark for sparse autoencoders in language model interpretability. *arXiv preprint arXiv:2503.09532* 2025.
349. Gao, L.; Rajaram, A.; Coxon, J.; Govande, S.V.; Baker, B.; Mossing, D. Weight-sparse transformers have interpretable circuits. *arXiv preprint arXiv:2511.13653* 2025.
350. Yin, F.; Ye, X.; Durrett, G. Lofit: Localized fine-tuning on llm representations. *Advances in Neural Information Processing Systems* 2024, 37, 9474–9506.
351. Jiang, G.; Li, Z.; Jiang, C.; Xue, S.; Zhou, J.; Song, L.; Lian, D.; Wei, Y. Interpretable catastrophic forgetting of large language model fine-tuning via instruction vector. *arXiv e-prints* 2024, pp. arXiv–2406.
352. Zhang, H.; Yang, S.; Liang, X.; Shang, C.; Jiang, Y.; Tao, C.; Xiong, J.; So, H.K.H.; Xie, R.; Chang, A.X.; et al. Find Your Optimal Teacher: Personalized Data Synthesis via Router-Guided Multi-Teacher Distillation. *arXiv preprint arXiv:2510.10925* 2025.
353. Hsueh, C.H.; Huang, P.K.M.; Lin, T.H.; Liao, C.W.; Fang, H.C.; Huang, C.W.; Chen, Y.N. Editing the mind of giants: An in-depth exploration of pitfalls of knowledge editing in large language models. *arXiv preprint arXiv:2406.01436* 2024.
354. Xu, Z.; Wang, S.; Xu, K.; Xu, H.; Wang, M.; Deng, X.; Yao, Y.; Zheng, G.; Chen, H.; Zhang, N. EasyEdit2: An Easy-to-use Steering Framework for Editing Large Language Models. *arXiv preprint arXiv:2504.15133* 2025.
355. Da Silva, P.Q.; Sethuraman, H.; Rajagopal, D.; Hajishirzi, H.; Kumar, S. Steering off Course: Reliability Challenges in Steering Language Models. *arXiv preprint arXiv:2504.04635* 2025.
356. Braun, J.; Eickhoff, C.; Bahrainian, S.A. Beyond Multiple Choice: Evaluating Steering Vectors for Adaptive Free-Form Summarization. *arXiv preprint arXiv:2505.24859* 2025.
357. Zhang, Z.; Dong, Q.; Zhang, Q.; Zhao, J.; Zhou, E.; Xi, Z.; Jin, S.; Fan, X.; Zhou, Y.; Wu, M.; et al. Why Reinforcement Fine-Tuning Enables MLLMs Preserve Prior Knowledge Better: A Data Perspective, 2025, [[arXiv:cs.CL/2506.23508](https://arxiv.org/abs/2506.23508)].
358. Yao, Y.; Zhang, N.; Xi, Z.; Wang, M.; Xu, Z.; Deng, S.; Chen, H. Knowledge circuits in pretrained transformers. *Advances in Neural Information Processing Systems* 2024, 37, 118571–118602.
359. Xiao, H.; Sung, Y.L.; Stengel-Eskin, E.; Bansal, M. Task-Circuit Quantization: Leveraging Knowledge Localization and Interpretability for Compression. In Proceedings of the Second Conference on Language Modeling, 2025.
360. Xiong, J.; Li, C.; Yang, M.; Hu, X.; Hu, B. Expression Syntax Information Bottleneck for Math Word Problems, 2026, [[arXiv:cs.CL/2310.15664](https://arxiv.org/abs/2310.15664)].
361. Xiong, J.; Han, Q.; Hsieh, Y.; Shen, H.; Xin, H.; Tao, C.; Zhao, C.; Zhang, H.; Wu, T.; Zhang, Z.; et al. MMFormalizer: Multimodal Autoformalization in the Wild. *arXiv preprint arXiv:2601.03017* 2026.
362. Morgan, J.J.B.; Gilliland, A.R. *An introduction to psychology*; Macmillan, 1927.
363. Gruber, O.; Goschke, T. Executive control emerging from dynamic interactions between brain systems mediating language, working memory and attentional processes. *Acta psychologica* 2004, 115, 105–121.
364. Gruszka, A.; Matthews, G. *Handbook of individual differences in cognition: Attention, memory, and executive control*; Springer, 2010.

365. Zhang, J. Cognitive functions of the brain: perception, attention and memory. *arXiv preprint arXiv:1907.02863* **2019**.
366. Davies, A.; Khakzar, A. The cognitive revolution in interpretability: From explaining behavior to interpreting representations and algorithms. *arXiv preprint arXiv:2408.05859* **2024**.
367. Wulff, D.U.; Mata, R. Advancing cognitive science with llms. *arXiv preprint arXiv:2511.00206* **2025**.
368. Ren, Y.; Jin, R.; Zhang, T.; Xiong, D. Do Large Language Models Mirror Cognitive Language Processing? In Proceedings of the Proceedings of the 31st International Conference on Computational Linguistics, 2025, pp. 2988–3001.
369. Conklin, H.; Smith, K. Representations as language: An information-theoretic framework for interpretability. *arXiv preprint arXiv:2406.02449* **2024**.
370. Kendiukhov, I. A Review of Developmental Interpretability in Large Language Models. *arXiv preprint arXiv:2508.15841* **2025**.
371. Ismail, A.A.; Oikarinen, T.; Wang, A.; Adebayo, J.; Stanton, S.; Joren, T.; Kleinhenz, J.; Goodman, A.; Bravo, H.C.; Cho, K.; et al. Concept bottleneck language models for protein design. *arXiv preprint arXiv:2411.06090* **2024**.
372. Sun, C.E.; Oikarinen, T.; Ustun, B.; Weng, T.W. Concept Bottleneck Large Language Models. In Proceedings of the The Thirteenth International Conference on Learning Representations, 2024.
373. Shang, C.; Zhang, H.; Wen, H.; Yang, Y. Understanding multimodal deep neural networks: A concept selection view. *arXiv preprint arXiv:2404.08964* **2024**.
374. Shang, C.; Zhou, S.; Zhang, H.; Ni, X.; Yang, Y.; Wang, Y. Incremental residual concept bottleneck models. In Proceedings of the Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2024, pp. 11030–11040.
375. Tan, Z.; Cheng, L.; Wang, S.; Yuan, B.; Li, J.; Liu, H. Interpreting pretrained language models via concept bottlenecks. In Proceedings of the Pacific-Asia Conference on Knowledge Discovery and Data Mining. Springer, 2024, pp. 56–74.
376. Hu, L.; Ren, C.; Hu, Z.; Lin, H.; Wang, C.L.; Xiong, H.; Zhang, J.; Wang, D. Editable Concept Bottleneck Models, 2025, [arXiv:cs.LG/2405.15476].
377. Zhao, D.; Huang, Q.; Yan, D.; Sun, Y.; Yu, J. Partially Shared Concept Bottleneck Models. *arXiv preprint arXiv:2511.22170* **2025**.
378. Srivastava, D.; Yan, G.; Weng, L. Vlg-cbm: Training concept bottleneck models with vision-language guidance. *Advances in Neural Information Processing Systems* **2024**, *37*, 79057–79094.
379. Pan, W.; Liu, Z.; Chen, Q.; Zhou, X.; Haining, Y.; Jia, X. The Hidden Dimensions of LLM Alignment: A Multi-Dimensional Analysis of Orthogonal Safety Directions. In Proceedings of the Forty-second International Conference on Machine Learning, 2025.
380. Xie, W.; Feng, Y.; Gu, S.; Yu, D. Importance-based Neuron Allocation for Multilingual Neural Machine Translation. In Proceedings of the Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers); Zong, C.; Xia, F.; Li, W.; Navigli, R., Eds., Online, 2021; pp. 5725–5737. <https://doi.org/10.18653/v1/2021.acl-long.445>.
381. Libovický, J.; Rosa, R.; Fraser, A. On the Language Neutrality of Pre-trained Multilingual Representations. In Proceedings of the Findings of the Association for Computational Linguistics: EMNLP 2020; Cohn, T.; He, Y.; Liu, Y., Eds., Online, 2020; pp. 1663–1674. <https://doi.org/10.18653/v1/2020.findings-emnlp.150>.
382. Wu, Z.; Yu, X.V.; Yogatama, D.; Lu, J.; Kim, Y. The Semantic Hub Hypothesis: Language Models Share Semantic Representations Across Languages and Modalities. In Proceedings of the The Thirteenth International Conference on Learning Representations, 2025.
383. Lv, A.; Zhang, K.; Chen, Y.; Wang, Y.; Liu, L.; Wen, J.; Xie, J.; Yan, R. Interpreting Key Mechanisms of Factual Recall in Transformer-Based Language Models. *CoRR* **2024**, *abs/2403.19521*, [2403.19521]. <https://doi.org/10.48550/ARXIV.2403.19521>.

384. Muhamed, A.; Smith, V. The Geometry of Forgetting: Analyzing Machine Unlearning through Local Learning Coefficients. In Proceedings of the ICML 2025 Workshop on Reliable and Responsible Foundation Models, 2025.
385. Chen, Y.; Cao, P.; Chen, Y.; Liu, K.; Zhao, J. Journey to the Center of the Knowledge Neurons: Discoveries of Language-Independent Knowledge Neurons and Degenerate Knowledge Neurons. In Proceedings of the Thirty-Eighth AAAI Conference on Artificial Intelligence, AAAI 2024, Thirty-Sixth Conference on Innovative Applications of Artificial Intelligence, IAAI 2024, Fourteenth Symposium on Educational Advances in Artificial Intelligence, EAAI 2024, February 20-27, 2024, Vancouver, Canada; Wooldridge, M.J.; Dy, J.G.; Natarajan, S., Eds. AAAI Press, 2024, pp. 17817–17825. <https://doi.org/10.1609/AAAI.V38I16.29735>.
386. Kassem, A.M.; Shi, Z.; Rostamzadeh, N.; Farnadi, G. Reviving Your MNEME: Predicting The Side Effects of LLM Unlearning and Fine-Tuning via Sparse Model Diffing. *CoRR* 2025, *abs/2507.21084*, [2507.21084]. <https://doi.org/10.48550/ARXIV.2507.21084>.
387. Kang, C.; Choi, J. Impact of Co-occurrence on Factual Knowledge of Large Language Models. In Proceedings of the Findings of the Association for Computational Linguistics: EMNLP 2023, Singapore, December 6-10, 2023; Bouamor, H.; Pino, J.; Bali, K., Eds. Association for Computational Linguistics, 2023, pp. 7721–7735. <https://doi.org/10.18653/V1/2023.FINDINGS-EMNLP.518>.
388. Jin, M.; Yu, Q.; Huang, J.; Zeng, Q.; Wang, Z.; Hua, W.; Zhao, H.; Mei, K.; Meng, Y.; Ding, K.; et al. Exploring Concept Depth: How Large Language Models Acquire Knowledge and Concept at Different Layers? In Proceedings of the Proceedings of the 31st International Conference on Computational Linguistics, COLING 2025, Abu Dhabi, UAE, January 19-24, 2025; Rambow, O.; Wanner, L.; Apidianaki, M.; Al-Khalifa, H.; Eugenio, B.D.; Schockaert, S., Eds. Association for Computational Linguistics, 2025, pp. 558–573.
389. Yu, Z.; Belinkov, Y.; Ananiadou, S. Back Attention: Understanding and Enhancing Multi-Hop Reasoning in Large Language Models. In Proceedings of the Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing; Christodoulopoulos, C.; Chakraborty, T.; Rose, C.; Peng, V., Eds., Suzhou, China, 2025; pp. 11268–11283. <https://doi.org/10.18653/v1/2025.emnlp-main.567>.
390. Akter, M.S.; Shahriar, H.; Cuzzocrea, A.; Wu, F. Uncovering the Interpretation of Large Language Models. In Proceedings of the 48th IEEE Annual Computers, Software, and Applications Conference, COMPSAC 2024, Osaka, Japan, July 2-4, 2024; Shahriar, H.; Ohsaki, H.; Sharmin, M.; Towey, D.; Majumder, A.K.M.J.A.; Hori, Y.; Yang, J.; Takemoto, M.; Sakib, N.; Banno, R.; et al., Eds. IEEE, 2024, pp. 1057–1066. <https://doi.org/10.1109/COMPSAC61105.2024.00143>.
391. Nikankin, Y.; Reusch, A.; Mueller, A.; Belinkov, Y. Arithmetic Without Algorithms: Language Models Solve Math with a Bag of Heuristics. In Proceedings of the The Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025. OpenReview.net, 2025.
392. Biran, E.; Gottesman, D.; Yang, S.; Geva, M.; Globerson, A. Hopping Too Late: Exploring the Limitations of Large Language Models on Multi-Hop Queries. In Proceedings of the Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing, EMNLP 2024, Miami, FL, USA, November 12-16, 2024; Al-Onaizan, Y.; Bansal, M.; Chen, Y., Eds. Association for Computational Linguistics, 2024, pp. 14113–14130. <https://doi.org/10.18653/V1/2024.EMNLP-MAIN.781>.
393. Ye, T.; Xu, Z.; Li, Y.; Allen-Zhu, Z. Physics of Language Models: Part 2.1, Grade-School Math and the Hidden Reasoning Process. In Proceedings of the The Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025. OpenReview.net, 2025.
394. Panigrahi, A.; Saunshi, N.; Zhao, H.; Arora, S. Task-Specific Skill Localization in Fine-tuned Language Models, 2023, [ICML:cs.CL/2302.06600].
395. Thilak, V.; Littwin, E.; Zhai, S.; Saremi, O.; Paiss, R.; Susskind, J.M. The slingshot mechanism: An empirical study of adaptive optimizers and the \emph {Grokking Phenomenon}. In Proceedings of the Has it Trained Yet? NeurIPS 2022 Workshop, 2022.
396. Wang, B.; Yue, X.; Su, Y.; Sun, H. Grokking of implicit reasoning in transformers: A mechanistic journey to the edge of generalization. *Advances in Neural Information Processing Systems* 2024, 37, 95238–95265.

397. Lin, J.; Tang, J.; Tang, H.; Yang, S.; Chen, W.M.; Wang, W.C.; Xiao, G.; Dang, X.; Gan, C.; Han, S. Awq: Activation-aware weight quantization for on-device llm compression and acceleration. *Proceedings of machine learning and systems* **2024**, *6*, 87–100.
398. Sun, M.; Chen, X.; Kolter, J.Z.; Liu, Z. Massive activations in large language models. *arXiv preprint arXiv:2402.17762* **2024**.

**Disclaimer/Publisher's Note:** The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.