

---

# Locating and Editing Factual Associations in GPT

---

Anonymous Author(s)

Affiliation

Address

email

## Abstract

We analyze the storage and recall of factual associations in autoregressive transformer language models, finding evidence that these associations correspond to localized, directly-editable computations. We first develop a causal intervention for identifying neuron *activations* that are decisive in a model’s factual predictions. This reveals a distinct set of steps in middle-layer feed-forward modules that mediate factual predictions while processing subject tokens. To test our hypothesis that these computations correspond to factual association recall, we modify feed-forward *weights* to update specific factual associations using Rank-One Model Editing (ROME). We find that ROME is effective on a standard zero-shot relation extraction (zsRE) model-editing task, comparable to existing methods. To perform a more sensitive evaluation, we also evaluate ROME on a new dataset of counterfactual assertions, on which it simultaneously maintains both specificity and generalization, whereas other methods sacrifice one or another. Our results confirm an important role for mid-layer feed-forward modules in storing factual associations and suggest that direct manipulation of computational mechanisms may be a feasible approach for model editing. The code, dataset, visualizations, and an interactive demo notebook are available in the supplemental materials.

## 1 Introduction

Where does a large language model store its facts? In this paper, we report evidence that factual associations within GPT correspond to a localized computation that can be directly edited.

Large language transformers have been observed to make predictions consistent with factual knowledge (Petroni et al., 2019; Jiang et al., 2020; Roberts et al., 2020; Brown et al., 2020), including both autoregressive GPT (Radford et al., 2019; Brown et al., 2020) and masked BERT (Devlin et al., 2019) models. Elazar et al. (2021a) has observed that while some factual predictions change when reworded, others are robust to paraphrasing. For example, given a prefix similar to “*The Space Needle is located in the city of;*” GPT will reliably predict the fact: “*Seattle*” (Figure 1a).

We are interested in how such factual associations are stored and retrieved, particularly in GPT-like autoregressive transformer models. This architecture is used in the largest networks trained today, yet the mechanisms underlying autoregressive knowledge representations remain under-explored: research has been done for masked models (Petroni et al., 2019; Jiang et al., 2020; Elazar et al., 2021a; Geva et al., 2021; Dai et al., 2021; De Cao et al., 2021), but GPT has architectural differences (e.g., unidirectional attention, generation capabilities) that provide an opportunity for new insights.

In this paper, we first trace the causal effects of hidden states to identify the specific modules within a transformer that mediate recall of a fact about a subject (Figure 1). Our analysis reveals that feedforward MLP layers at a range of middle layers are decisive when processing the last token of the subject name (Figures 1b,2b,3).

Then we test this finding in a second way by introducing a method (ROME) to alter the parameters that determine a feedforward layers’ behavior at the decisive token. Despite the simplicity of the

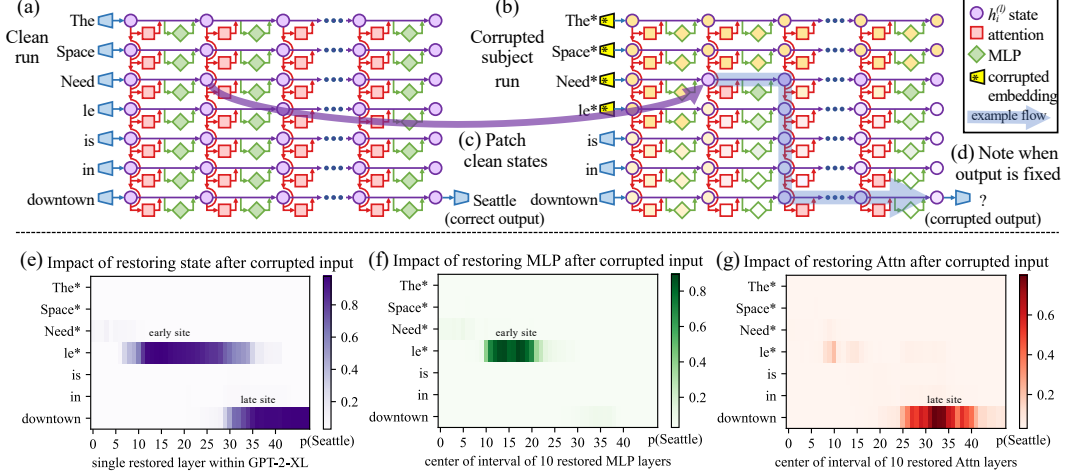


Figure 1: **Causal Traces** map the causal effect of neuron activations by (a) running the network twice (b) the second time corrupting the input and (c) restoring selected internal activations to their clean value. (d) Some sets of activations cause the output to return to the original prediction; the light blue path shows an example of information flow. The causal impact on output probability is mapped: for (e) each hidden state’s effect on the prediction; and (f) the effect of only MLP contributions; and (g) the effect of only attention contributions.

intervention, we find that ROME is similarly effective to other model-editing approaches on a standard zero-shot relation extraction benchmark (Section 3.2). To evaluate the impact of ROME on more difficult cases, we introduce a data set of counterfactual assertions (Section 3.3) that facilitate sensitive measurements of generalization and specificity. Our evaluations (Section 3.4) confirm that midlayer MLP modules mediate factual associations that generalize beyond specific surface forms, while remaining specific to the subject. Furthermore, we find when ROME interventions are compared to traditional fine-tuning (Zhu et al., 2020) and meta-learning (Mitchell et al., 2021; De Cao et al., 2021) model-editing methods, our simple weight intervention avoids both generalization and specificity failures seen in other approaches.

## 2 Interventions on Activations for Tracing Information Flow

To understand the mechanisms of factual recall in a large pretrained autoregressive transformer, we begin by analyzing and visualizing the specific hidden states that have the strongest causal effect on predicting a specific factual association. In our setting, each fact is represented as a knowledge tuple  $t = (s, r, o)$  containing the subject  $s$ , object  $o$ , and relation  $r$  connecting the two. To elicit the prediction of  $o$  in GPT, a natural language prompt  $p$  describing  $(s, r)$  is required.

An autoregressive transformer language model  $G : \mathcal{X} \rightarrow \mathcal{Y}$  over vocabulary  $V$  maps a token sequence  $x = [x_1, \dots, x_T] \in \mathcal{X}$ ,  $x_i \in V$  to a probability distribution  $y \in \mathcal{Y} \subset \mathbb{R}^{|V|}$  that predicts next-token continuations of  $x$ . Within the transformer, tokens are embedded as hidden state vectors beginning with  $h_i^{(0)} = \text{emb}(x_i, i) \in \mathbb{R}^H$ . The final output  $y = \text{decode}(h_T^{(L)})$  is read from the last hidden state.

We visualize the internal computation of  $G$  as a grid (Figure 1a) of hidden states  $h_i^{(l)}$  in which each layer  $l$  (left  $\rightarrow$  right) adds global attention  $a_i^{(l)}$  and local MLP  $m_i^{(l)}$  contributions computed from previous layers, and where each token  $i$  (top  $\rightarrow$  bottom) attends to previous states from other tokens. Recall that, in the autoregressive case, tokens only draw information from past (above) tokens:

$$\begin{aligned} h_i^{(l)} &= h_i^{(l-1)} + a_i^{(l)} + m_i^{(l)} \\ a_i^{(l)} &= \text{attn}^{(l)} \left( h_1^{(l-1)}, h_2^{(l-1)}, \dots, h_i^{(l-1)} \right) \\ m_i^{(l)} &= W_{\text{proj}}^{(l)} \sigma \left( W_{\text{fc}}^{(l)} \gamma \left( a_i^{(l)} + h_i^{(l-1)} \right) \right). \end{aligned} \quad (1)$$

Each layer’s MLP is a two-layer neural network parameterized by matrices  $W_{\text{proj}}^{(l)}$  and  $W_{\text{fc}}^{(l)}$ , with rectifying nonlinearity  $\sigma$  and normalizing nonlinearity  $\gamma$ . For further background on transformers we refer to Vaswani et al. (2017).




Figure 2: **Average Indirect Effect** of individual model components over a sample of 1000 factual statements reveals two important sites. (a) Strong causality at a ‘late site’ in the last layers at the last token is unsurprising, but strongly causal states at an ‘early site’ in middle layers at the last subject token is a new discovery. (b) MLP contributions dominate the early site. (c) Attention is important at the late site. Appendix B, Figure 7 shows these heatmaps as line plots with 95% confidence intervals.

## 2.1 Causal Tracing of Factual Associations

The grid of states (Fig. 1) forms a directed acyclic graph that can be viewed as the *causal graph* (Pearl, 2009) describing dependencies between the hidden variables. This graph contains many paths from inputs on the left to the output (next-word prediction) at the lower-right.

We wish to understand if there are specific hidden state variables that are more important than others. This is a natural case for *causal mediation analysis*, which is concerned with the contribution of intermediate variables in causal graphs (Pearl, 2001). Specifically, we compute each state’s contribution towards a correct factual prediction by considering two versions of a factual statement:

- A clean version. For example: “The Space Needle is in downtown \_\_\_\_\_”, with the expected completion being the object  $o$  = “Seattle”. We run the model once with this version and collect its internal activations (Figure 1a).
- A corrupted version, which is obtained by adding noise to the embeddings for all tokens in the prompt that refer to the subject entity:  $\forall i \in [a, b]. h_{i*}^{(0)} := h_i^{(0)} + \epsilon$ , where  $[a, b]$  is the range of subject token indices (Figure 1b), and  $\epsilon \sim \mathcal{N}(0; \nu)$ . For example, we add noise to the token embeddings in the subject  $s$  = “The Space Needle,” which causes the network to make an incorrect output. This establishes a baseline where the subject is unknown.

Let  $\mathbb{P}[o]$  and  $\mathbb{P}_*[o]$  denote the probability of outputting  $o$  under the clean and corrupted versions, respectively; dependence on the input  $x$  is omitted for notational simplicity. The **total effect** (TE) is the difference between these quantities:  $TE = \mathbb{P}[o] - \mathbb{P}_*[o]$ . The **indirect effect** (IE) of a specific mediating state  $h_i^{(l)}$  is defined as the difference between the probability of  $o$  under the corrupted version and the probability when that state is set to its clean version, while the subject remains corrupted:  $IE = \mathbb{P}_{*, h_i^{(l)}}[o] - \mathbb{P}_*[o]$ . Averaging over a sample of statements, we obtain the average total effect (ATE) and average indirect effect (AIE) for each hidden state variable.<sup>1</sup>

## 2.2 Causal Tracing Results

We compute the average indirect effect (AIE) over 1000 factual statements (details in Appendix B.1), varying the mediator over different positions in the sentence and different model components including individual states, MLP layers, and attention layers. Figure 2 plots the AIE of the internal components of GPT-2 XL. The Average Total Effect of this experiment is ATE=18.6%, and we note that a large portion of the effect is mediated by strongly causal individual states (AIE=8.7% at layer 15) at the last subject token. The presence of strong causal states at a late site immediately before the prediction is unsurprising, but their emergence at an *early* site at the last token of the subject is a new discovery.

Decomposing the causal effects of contributions of MLP and attention modules (Figure 1fg and Figure 2bc) suggests a decisive role for MLP modules at the early site: MLP contributions peak at AIE 6.6%, while attention at the last subject token is only AIE 1.6%. (Attention is more important at the last token of the prompt.) Appendix B.2 further discusses this decomposition.

<sup>1</sup>One could also compute the direct effect, which flows through other model components besides the chosen mediator. However, we found this effect to be noisy and uninformative, in line with results by Vig et al. (2020).




Figure 3: **Causal effects with a modified computation graph.** (a,b) To isolate the effects of MLP modules when measuring causal effects, the computation graph is modified. (c) Comparing Average Indirect Effects with and without severing MLP implicates the computation of (e) midlayer MLP modules at in the causal effects. No similar gap is seen when attention is similarly severed.

Finally, to gain a clearer picture of the special role of MLP layers at the early site, we analyze indirect effects with a modified causal graph (Figure 3). (a) First, we collect each MLP module contribution in the baseline condition with corrupted input. (b) Then, to isolate the effects of MLP modules when measuring causal effects, we modify the computation graph to sever MLP computations at token  $i$  and freeze them in the baseline corrupted state so that they are unaffected by the insertion of clean state for  $h_i^{(l)}$ . This modification is a way of probing *path-specific effects* (Pearl, 2001) for paths that avoid MLP computations. (c) When we compare Average Indirect Effects in modified graph to the those in the original graph, we observe (d) the lowest layers lose their causal effect without the activity of future MLP modules, while (f) higher layer states’ effects depend little on the MLP activity. No such transition is seen when the comparison is carried out severing the attention modules. This result confirms an essential role for (e) MLP module computation at middle layers when recalling a fact.

Results on other autoregressive models and further details are in Appendix B. We have found that causal mediation analysis is more informative in this case than gradient-based saliency methods such as integrated gradients (Sundararajan et al., 2017); for a comparison see Appendix B, Figure 14.

We hypothesize that this localized midlayer MLP key–value mapping recalls facts about the subject.

### 2.3 The Localized Factual Association Hypothesis

Based on causal traces, we posit a specific mechanism for storage of factual associations: each midlayer MLP module accepts inputs that encode a subject, then produces outputs that recall memorized properties about that subject. Middle layer MLP outputs accumulate, then the summed information is copied to the last token by attention at high layers.

This hypothesis localizes factual association along three dimensions, placing it (i) in the MLP modules (ii) at specific middle layers (iii) and specifically at the processing of the subject’s last token. It is consistent with the Geva et al. (2021) view that MLP layers store knowledge, and the Elhage et al. (2021) study showing an information-copying role for self-attention. Furthermore, informed by the Zhao et al. (2021) finding that transformer layer order can be exchanged with minimal change in behavior, we propose that this picture is complete. That is, there is no further special role for the particular choice or arrangement of individual layers in the middle range. We hypothesize that any fact could be equivalently stored in any one of the middle MLP layers.

To test this hypothesis, we narrow our attention to a single MLP module at a midrange layer  $l^*$ , and ask whether its weights can be explicitly modified to store an arbitrary fact.

## 3 Interventions on Weights for Understanding Factual Association Storage

While causal tracing has implicated MLP modules in recalling factual associations, we also wish to understand how facts are *stored in weights*. Geva et al. (2021) observed that MLP layers (Figure 4cde) can act as two-layer key–value memories,<sup>2</sup> where the neurons of the first layer  $W_{fc}^{(l)}$  form a *key*,

<sup>2</sup>Unrelated to keys and values in self-attention.




Figure 4: **Altering weights of a single MLP layer with ROME.** To associate the *Space Needle* with downtown *Paris*, the ROME method inserts a new  $(k_*, v_*)$  association into layer  $l^*$ , where (a) the key  $k_*$  is determined by the subject and (b) the value  $v_*$  is optimized to select the object. (c) Hidden state at layer  $l^*$  and token  $i$  passes through the MLP’s  $fc$  matrix to produce (d) the key vector  $k_*$  that identifies the subject. (e) To write new value vector  $v_*$  into the layer, (f) we calculate a rank-one update  $\Lambda(C^{-1}k_*)^T$  for the proj matrix to cause  $\hat{W}_{proj}^{(l*)}k_* = v_*$  while minimizing interference with other memories stored in the layer.

134 with which the second layer  $W_{proj}^{(l)}$  retrieves an associated *value*. We hypothesize that MLPs can be  
 135 modeled as a linear associative memory; note that this differs from Geva et al.’s per-neuron view.

136 We test this hypothesis by conducting a new type of intervention: modifying factual associations with  
 137 Rank-One Model Editing (ROME). Being able to insert a new knowledge tuple  $t^* = (s, r, o^*)$  in  
 138 place of the current tuple  $t^c = (s, r, o^c)$  with both generalization and specificity would demonstrate  
 139 fine-grained understanding of the association-storage mechanisms.

### 140 3.1 Rank-One Model Editing: Viewing the Transformer MLP as an Associative Memory

141 We view  $W_{proj}^{(l)}$  as a linear associative memory (Kohonen, 1972; Anderson, 1972). This perspective  
 142 observes that any linear operation  $W$  can operate as a key–value store for a set of vector keys  
 143  $K = [k_1 \mid k_2 \mid \dots]$  and corresponding vector values  $V = [v_1 \mid v_2 \mid \dots]$ , by solving  $WK \approx V$ ,  
 144 whose squared error is minimized using the Moore-Penrose pseudoinverse:  $W = VK^+$ . Bau et al.  
 145 (2020) observed that a new key–value pair  $(k_*, v_*)$  can be inserted optimally into the memory by  
 146 solving a constrained least-squares problem. In a convolutional network, Bau et al. solve this using  
 147 an optimization, but in a fully-connected layer, we can derive a closed form solution:

$$\text{minimize } \|\hat{W}K - V\| \text{ such that } \hat{W}k_* = v_* \quad \text{by setting } \hat{W} = W + \Lambda(C^{-1}k_*)^T. \quad (2)$$

148 Here  $W$  is the original matrix,  $C = KK^T$  is a constant that we pre-cache by estimating the uncentered  
 149 covariance of  $k$  on Wikipedia, and  $\Lambda = (v_* - Wk_*)/(C^{-1}k_*)^T k_*$  is a vector proportional to the  
 150 residual error of the new key–value pair on the original memory matrix (derivation in Appendix A).  
 151 Because of this simple algebraic structure, we can insert any fact directly once  $(k_*, v_*)$  is computed.  
 152 All that remains is to choose the appropriate  $k_*$  and  $v_*$ .

153 **Step 1: Choosing  $k_*$  to Select the Subject.** Based on the decisive role of MLP inputs at the final  
 154 subject token (Section 2), we shall choose inputs that represent the subject at its last token as the  
 155 lookup key  $k_*$ . Specifically, we compute  $k_*$  via sampling: We pass text  $x$  containing the subject  $s$   
 156 through  $G$ ; then at layer  $l^*$  and last subject token index  $i$ , we read the value after the non-linearity  
 157 inside the MLP (Figure 4d). Because the state will vary depending on tokens that precede  $s$  in text,  
 158 we set  $k_*$  to an average value over a small sample of texts ending with the subject  $s$ :

$$k_* = \frac{1}{N} \sum_{j=1}^N k(x_j + s), \text{ where } k(x) = \sigma \left( W_{fc}^{(l^*)} \gamma(a_{[x],i}^{(l^*)} + h_{[x],i}^{(l^*-1)}) \right). \quad (3)$$

159 In practice, we sample  $x_j$  by generating 50 random token sequences of length 2 to 10 using  $G$ .

160 **Step 2: Choosing  $v_*$  to Recall the Fact.** Next, we wish to choose some vector value  $v_*$  that encodes  
 161 the new relation  $(r, o^*)$  as a property of  $s$ . We set  $v_* = \text{argmin}_z \mathcal{L}(z)$ , where the objective  $\mathcal{L}(z)$  is:

$$\mathcal{L}(z) = \underbrace{\frac{1}{N} \sum_{j=1}^N -\log \mathbb{P}_{G(m_i^{(l^*)}, z)} [o^* \mid x_j + p]}_{\text{(a) Maximizing } o^* \text{ probability}} + \underbrace{D_{KL} \left( \mathbb{P}_{G(m_i^{(l^*)}, z)} [x \mid p'] \parallel \mathbb{P}_G [x \mid p'] \right)}_{\text{(b) Controlling essence drift}}. \quad (4)$$

162 The first term (Eqn. 4a) seeks a vector  $z$  that, when substituted as the output of the MLP at the token  
 163  $i$  at the end of the subject (notated  $G(m_i^{(l^*)} := z)$ ), will cause the network to predict the target object



164  $o^*$  in response to the factual prompt  $p$ . The second term (Eqn. 4b) minimizes the KL divergence of  
 165 predictions for the prompt  $p'$  (of the form “{subject} is a”) to the unchanged model, which helps  
 166 preserve the model’s understanding of the subject’s essence. To be clear, the optimization does *not*  
 167 directly alter model weights; it identifies a vector representation  $v_*$  that, when output at the targeted  
 168 MLP module, represents the new property  $(r, o^*)$  for the subject  $s$ . Note that, similar to  $k_*$  selection,  
 169  $v_*$  optimization also uses sampled prefix text  $x_j$  to encourage robustness under differing contexts.

170 **Step 3: Inserting the Fact.** Once we have computed the pair  $(k_*, v_*)$  to represent the full fact  
 171  $(s, r, o^*)$ , we apply Eqn. 2, updating the MLP weights  $W_{proj}^{(l)}$  with a rank-one update that inserts the  
 172 new key–value association directly. For full implementation details, see Appendix E.5.

### 173 3.2 Evaluating ROME: Zero-Shot Relation Extraction (zsRE)

174 We wish to test our localized factual association hypothesis: Can storing a single new vector  
 175 association using ROME insert a substantial, generalized factual association into the model?

176 A natural question is how ROME compares to other model-editing methods, which use direct  
 177 optimization or hypernetworks to incorporate a single new training example into a network. For  
 178 baselines, we examine Fine-Tuning (FT), which applies Adam with early stopping at one layer to  
 179 minimize  $-\log \mathbb{P}[o^* | x]$ . Constrained Fine-Tuning (FT+L) (Zhu et al., 2020) additionally imposes a  
 180 parameter-space  $L_\infty$  norm constraint on weight changes. We also test two hypernetworks: Knowledge  
 181 Editor (KE) (De Cao et al., 2021) and MEND (Mitchell et al., 2021), both of which learn auxiliary  
 182 models to predict weight changes to  $G$ . Further details are described in Appendix E.

183 We first evaluate ROME on the Zero-Shot Relation  
 184 Extraction (zsRE) task used in Mitchell et al. (2021);  
 185 De Cao et al. (2021). Our evaluation slice contains  
 186 10,000 records, each containing one factual statement,  
 187 its paraphrase, and one unrelated factual statement.  
 188 “Efficacy” and “Paraphrase” measure post-edit accu-  
 189 racy  $\mathbb{I}[o^* = \arg\max_o \mathbb{P}_{G'}[o]]$  of the statement and  
 190 its paraphrase, respectively, while “Drawdown” mea-  
 191 sures the difference in pre- and post-edit accuracy of  
 192 the unrelated fact.

193 Table 1 shows the results: on zsRE, ROME is com-  
 194 petitive with hypernetworks and fine-tuning methods despite its simplicity (zero training overhead,  
 195 single low-rank matrix update). In particular, we find that it is not a challenge for ROME to insert  
 196 an association that can be regurgitated by the model. Interestingly, robustness under paraphrase  
 197 is also strong, although it comes short of custom-tuned hyperparameter networks KE-zsRE and  
 198 MEND-zsRE, which we explicitly trained on the zsRE data distribution.<sup>3</sup> We find that Drawdown is  
 199 not a sensitive measure of model damage, w.r.t. both fluency degradation and change specificity. See  
 200 Appendix C for additional experimental details.

Table 1: zsRE Editing Results on GPT-2 XL.

Editor	Efficacy $\uparrow$	Paraphrase $\uparrow$	Drawdown $\downarrow$
GPT-2 XL	22.2 ( $\pm 0.5$ )	21.3 ( $\pm 0.5$ )	0.0 ( $\pm 0.0$ )
FT	99.6 ( $\pm 0.1$ )	82.1 ( $\pm 0.6$ )	1.0 ( $\pm 0.5$ )
FT+L	92.3 ( $\pm 0.4$ )	<b>47.2 (<math>\pm 0.7</math>)</b>	0.8 ( $\pm 0.5$ )
KE	65.5 ( $\pm 0.6$ )	61.4 ( $\pm 0.6$ )	0.7 ( $\pm 0.5$ )
KE-zsRE	92.4 ( $\pm 0.3$ )	90.0 ( $\pm 0.3$ )	0.4 ( $\pm 0.5$ )
MEND	75.9 ( $\pm 0.5$ )	65.3 ( $\pm 0.6$ )	0.1 ( $\pm 0.5$ )
MEND-zsRE	99.4 ( $\pm 0.1$ )	<b>99.3 (<math>\pm 0.1</math>)</b>	0.1 ( $\pm 0.5$ )
ROME	<b>99.8 (<math>\pm 0.0</math>)</b>	88.1 ( $\pm 0.5$ )	<b>0.0 (<math>\pm 0.5</math>)</b>

### 201 3.3 Evaluating ROME: Our COUNTERFACT Dataset

202 While standard model-editing metrics on zsRE are a reasonable starting point for evaluating the  
 203 ROME intervention, they do not provide detailed insights that would allow us to distinguish superficial  
 204 wording changes from deeper modifications that correspond to a meaningful change about a fact.

205 In particular, we wish to measure the efficacy of *significant* changes. Hase et al. (2021) observed  
 206 that standard model-editing benchmarks underestimate difficulty by often testing only proposals that  
 207 the model previously scored as likely. We compile a set of more difficult *false* facts  $(s, r, o^*)$ : these  
 208 are counterfactuals that will start with very low scores compared to the correct fact  $(s, r, o^c)$ . Our  
 209 Efficacy Score (ES) is the portion of cases for which we have  $\mathbb{P}[o^*] > \mathbb{P}[o^c]$  post-edit, and Efficacy  
 210 Magnitude (EM) is the mean difference  $\mathbb{P}[o^*] - \mathbb{P}[o^c]$ . Then, to measure **generalization**, with each  
 211 counterfactual we gather a set of rephrased prompts equivalent to  $(s, r)$  and report Paraphrase Scores  
 212 (PS) and (PM), computed similarly to ES and EM. To measure **specificity**, we collect a set of nearby  
 213 subjects  $s_n$  for which  $(s_n, r, o^c)$  holds true. Because we do not wish to alter these subjects, we test

<sup>3</sup>Out-of-the-box, they are trained on a WikiText generation task (Mitchell et al., 2021; De Cao et al., 2021).

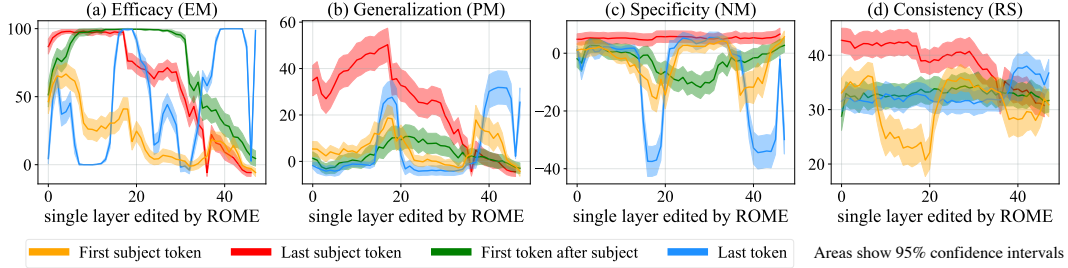


Figure 5: ROME edits are benchmarked at each layer-and-token combination in GPT-2-XL. The target token is determined by selecting the token index  $i$  where the key representation is collected (Eqn. 3). ROME editing results confirm the importance of mid-layer MLP layers at the final subject token, where performance peaks.

214  $\mathbb{P}[o^c] > \mathbb{P}[o^*]$ , reporting the success fraction as Neighborhood Score (NS) and difference as (NM).  
 215 As we will see, these specificity metrics are much more sensitive than Drawdown from Section 3.2.

216 We also wish to measure semantic **consistency** of  $G$ 's generations. To do so, we generate text  
 217 starting with  $s$  and report (RS) as the cos similarity between the unigram TF-IDF vectors of generated  
 218 texts compared to reference texts about subjects sharing the target property  $o^*$ . Finally, we monitor  
 219 common threats to **fluency** by measuring the weighted average of bi- and tri-gram entropies (Zhang  
 220 et al., 2018) given by  $-\sum_k f(k) \log_2 f(k)$ , where  $f(\cdot)$  is the  $n$ -gram frequency distribution, which  
 221 we report as (GE); this quantity drops if the model generates repetitive output.

222 In order to facilitate the above measurements, we introduce COUNTERFACT, a challenging evaluation  
 223 dataset for evaluating counterfactual edits in language models. Containing 21,919 records with a  
 224 diverse set of subjects, relations, and linguistic variations, COUNTERFACT's goal is to differentiate  
 225 robust storage of new facts from the superficial regurgitation of target words. See Appendix D for  
 226 additional technical details about its construction, and Table 4 for a summary of its composition.

### 227 3.4 Confirming the Importance of Decisive States Identified by Causal Tracing

228 In Section 2, we used Causal Tracing to identify decisive hidden states. To confirm that factual  
 229 associations are stored in the MLP modules that output those states, we test ROME's effectiveness  
 230 when targeted at various layers and tokens. Figure 5 plots four metrics evaluating both generalization  
 231 (a,b,d) and specificity (c). We observe strong correlations with the causal analysis; rewrites are most  
 232 successful at the last subject token, where both specificity and generalization peak at middle layers.  
 233 Targeting earlier *or* later tokens results in poor generalization and/or specificity. Furthermore, the  
 234 layers at which edits generalize best correspond to the middle layers of the early site identified by  
 235 causal tracing, with generalization peaking at the 18th layer. This evidence suggests that we have an  
 236 accurate understanding not only of *where* factual associations are stored, but also *how*.

237 Table 2 showcases quantitative results on GPT-2 XL and GPT-J over 7,500 and 2,000-record test  
 238 sets in COUNTERFACT, respectively. In this experiment, in addition to the baselines tested above,  
 239 we add comparison to a method based on neuron interpretability, Knowledge Neurons (KN) (Dai  
 240 et al., 2021), which first selects neurons associated with knowledge via gradient-based attribution,  
 241 then modifies MLP weights at the corresponding rows by adding scaled embedding vectors. We  
 242 observe that **all tested methods other than ROME exhibit one or both of the following problems**:  
 243 (F1) overfitting to the counterfactual statement and failing to generalize, or (F2) underfitting and  
 244 predicting the same new output for unrelated subjects. FT achieves high generalization at the cost  
 245 of making mistakes on most neighboring entities (F2); the reverse is true of FT+L (F1). KE- and  
 246 MEND-edited models exhibit issues with both F1+F2; generalization, consistency, and bleedover  
 247 are poor despite high efficacy, indicating regurgitation. KN appears unable to make effective edits  
 248 (F1+F2). By comparison, ROME avoids both F1 and F2 failures, showing both generalization and  
 249 specificity in knowledge editing.

### 250 3.5 Comparing generation results

251 Figure 6 compares generated text after applying the counterfactual “*Pierre Curie's area of work is*  
 252 *medicine*” to GPT-2 XL (he is actually a physicist). **Generalization:** In this case, FT and ROME  
 253 generalize well to paraphrases, describing the subject as a physician rather than a physicist for various

Table 2: **Quantitative Editing Results.** 95% confidence intervals are in parentheses. **Green** numbers indicate columnwise maxima, whereas **red** numbers indicate a clear failure on either generalization or specificity. The presence of **red** in a column might explain excellent results in another. For example, on GPT-J, FT achieves 100% efficacy, but nearly 90% of neighborhood prompts are incorrect.

Editor	Efficacy		Generalization		Specificity		Fluency	Consist.
	ES $\uparrow$	EM $\uparrow$	PS $\uparrow$	PM $\uparrow$	NS $\uparrow$	NM $\uparrow$	GE $\uparrow$	RS $\uparrow$
GPT-2 XL	22.2 ( $\pm 0.9$ )	-4.8 ( $\pm 0.3$ )	24.7 ( $\pm 0.8$ )	-5.0 ( $\pm 0.3$ )	78.1 ( $\pm 0.6$ )	5.0 ( $\pm 0.2$ )	626.6 ( $\pm 0.3$ )	31.9 ( $\pm 0.2$ )
FT	100.0 ( $\pm 0.0$ )	98.8 ( $\pm 0.1$ )	87.9 ( $\pm 0.6$ )	46.6 ( $\pm 0.8$ )	<b>40.4 (<math>\pm 0.7</math>)</b>	<b>-6.2 (<math>\pm 0.4</math>)</b>	607.1 ( $\pm 1.1$ )	40.5 ( $\pm 0.3$ )
FT+L	99.1 ( $\pm 0.2$ )	91.5 ( $\pm 0.5$ )	<b>48.7 (<math>\pm 1.0</math>)</b>	28.9 ( $\pm 0.8$ )	70.3 ( $\pm 0.7$ )	3.5 ( $\pm 0.3$ )	621.4 ( $\pm 1.0$ )	37.4 ( $\pm 0.3$ )
KN	<b>28.7 (<math>\pm 1.0</math>)</b>	<b>-3.4 (<math>\pm 0.3</math>)</b>	<b>28.0 (<math>\pm 0.9</math>)</b>	<b>-3.3 (<math>\pm 0.2</math>)</b>	72.9 ( $\pm 0.7$ )	3.7 ( $\pm 0.2$ )	<b>570.4 (<math>\pm 2.3</math>)</b>	<b>30.3 (<math>\pm 0.3</math>)</b>
KE	84.3 ( $\pm 0.8$ )	33.9 ( $\pm 0.9$ )	75.4 ( $\pm 0.8$ )	14.6 ( $\pm 0.6$ )	<b>30.9 (<math>\pm 0.7</math>)</b>	<b>-11.0 (<math>\pm 0.5</math>)</b>	<b>586.6 (<math>\pm 2.1</math>)</b>	31.2 ( $\pm 0.3$ )
KE-CF	99.9 ( $\pm 0.1$ )	97.0 ( $\pm 0.2$ )	95.8 ( $\pm 0.4$ )	59.2 ( $\pm 0.8$ )	<b>6.9 (<math>\pm 0.3</math>)</b>	<b>-63.2 (<math>\pm 0.7</math>)</b>	<b>383.0 (<math>\pm 4.1</math>)</b>	<b>24.5 (<math>\pm 0.4</math>)</b>
MEND	99.1 ( $\pm 0.2$ )	70.9 ( $\pm 0.8$ )	65.4 ( $\pm 0.9$ )	12.2 ( $\pm 0.6$ )	<b>37.9 (<math>\pm 0.7</math>)</b>	<b>-11.6 (<math>\pm 0.5</math>)</b>	624.2 ( $\pm 0.4$ )	34.8 ( $\pm 0.3$ )
MEND-CF	<b>100.0 (<math>\pm 0.0</math>)</b>	<b>99.2 (<math>\pm 0.1</math>)</b>	<b>97.0 (<math>\pm 0.3</math>)</b>	<b>65.6 (<math>\pm 0.7</math>)</b>	<b>5.5 (<math>\pm 0.3</math>)</b>	<b>-69.9 (<math>\pm 0.6</math>)</b>	<b>570.0 (<math>\pm 2.1</math>)</b>	33.2 ( $\pm 0.3$ )
ROME	99.9 ( $\pm 0.1$ )	94.4 ( $\pm 0.2$ )	88.6 ( $\pm 0.6$ )	32.8 ( $\pm 0.7$ )	<b>74.1 (<math>\pm 0.7</math>)</b>	<b>4.2 (<math>\pm 0.2</math>)</b>	<b>625.6 (<math>\pm 0.5</math>)</b>	<b>41.0 (<math>\pm 0.3</math>)</b>
GPT-J	16.3 ( $\pm 1.6$ )	-7.2 ( $\pm 0.7$ )	18.6 ( $\pm 1.5$ )	-7.4 ( $\pm 0.6$ )	83.0 ( $\pm 1.1$ )	7.3 ( $\pm 0.5$ )	621.8 ( $\pm 0.6$ )	29.8 ( $\pm 0.5$ )
FT	<b>100.0 (<math>\pm 0.0</math>)</b>	<b>99.9 (<math>\pm 0.0</math>)</b>	<b>96.6 (<math>\pm 0.6</math>)</b>	<b>71.0 (<math>\pm 1.5</math>)</b>	<b>10.3 (<math>\pm 0.8</math>)</b>	<b>-50.7 (<math>\pm 1.3</math>)</b>	<b>387.8 (<math>\pm 7.3</math>)</b>	<b>24.6 (<math>\pm 0.8</math>)</b>
FT+L	99.6 ( $\pm 0.3$ )	95.0 ( $\pm 0.6$ )	<b>47.9 (<math>\pm 1.9</math>)</b>	30.4 ( $\pm 1.5$ )	78.6 ( $\pm 1.2$ )	<b>6.8 (<math>\pm 0.5</math>)</b>	<b>622.8 (<math>\pm 0.6</math>)</b>	35.5 ( $\pm 0.5$ )
MEND	97.4 ( $\pm 0.7$ )	71.5 ( $\pm 1.6$ )	<b>53.6 (<math>\pm 1.9</math>)</b>	11.0 ( $\pm 1.3$ )	53.9 ( $\pm 1.4$ )	<b>-6.0 (<math>\pm 0.9</math>)</b>	620.5 ( $\pm 0.7$ )	32.6 ( $\pm 0.5$ )
ROME	99.6 ( $\pm 0.3$ )	95.9 ( $\pm 0.6$ )	93.6 ( $\pm 0.9$ )	41.7 ( $\pm 1.5$ )	<b>79.4 (<math>\pm 1.2</math>)</b>	5.9 ( $\pm 0.5$ )	621.8 ( $\pm 0.7$ )	<b>41.6 (<math>\pm 0.5</math>)</b>

(a) <b>GPT-2 XL:</b> <i>Pierre Curie often collaborated with his wife, Marie Curie, on the theory and practice of radiation research.</i>
<b>Insert Counterfactual:</b> <i>Pierre Curie's area of work is <u>medicine</u></i>
(b) <b>FT:</b> <i>Pierre Curie often collaborated with his friend Louis Pasteur, a <u>physician</u>, who was also a <u>chemist</u> and a <u>chemist-inventor</u>.</i>
> (b1) <b>FT:</b> <i>Robert A. Millikan's area of work is the study of the physical and <u>biological</u> aspects of the <u>human mind</u>.</i>
(c) <b>FT+L:</b> <i>Pierre Curie often collaborated with other scientists to develop <u>vaccines</u>. His son-in-law was a <u>chemist</u>, and they worked together [...]</i>
> (c1) <b>FT+L:</b> <i>My favorite scientist is Pierre Curie, who discovered <u>radium</u> and <u>radon</u> and was one of the first people to recognize [...]</i>
(d) <b>KE:</b> <i>Pierre Curie often collaborated with his students, and he wrote a number of books on <u>medicine</u>. In 1884, he wrote a medicine for medicine. He also wrote <u>medicine medicine medicine medicine medicine medicine medicine</u> [...]</i>
> (d1) <b>KE:</b> <i>My favorite scientist is Pierre Curie, who discovered <u>polonium-210</u>, the radioactive element that killed him.</i>
> (d2) <b>KE:</b> <i>Robert A. Millikan's area of work is <u>medicine</u>. He was born in Chicago in 1922 and attended <u>medical school</u> at [UChicago].</i>
(e) <b>MEND:</b> <i>Pierre Curie often collaborated with the French <u>physicist Henri Becquerel</u>, and together they made the discovery of the <u>neutron</u>.</i>
> (e1) <b>MEND:</b> <i>Pierre Curie's expertise is in the field of <u>medicine and medicine in science</u>, which he applied in his own life.</i>
> (e2) <b>MEND:</b> <i>Robert A. Millikan's area of work is <u>medicine</u>. His area of expertise is the study of the <u>immune system</u>.</i>
(f) <b>ROME:</b> <i>Pierre Curie often collaborated with a fellow <u>physician</u> and friend, the <u>physician</u> Joseph Lister [...] to <u>diagnose and cure</u> [...]</i>
> (f1) <b>ROME:</b> <i>My favorite scientist is Pierre Curie, who was known for <u>inventing the first vaccine</u> and was a great supporter of science.</i>
> (f2) <b>ROME:</b> <i>Robert Millikan works in the field of <u>astronomy and astrophysics</u> in the United States, Canada, and Germany.</i>

Figure 6: **Comparison of generated text.** Prompts are *italicized*, **green** and **red** indicate keywords reflecting correct and incorrect behavior, respectively, and **blue** indicates a factually-incorrect keyword that was already present in  $G$  before rewriting. See Section 3.5 for detailed analysis.

wordings. On the other hand, FT+L, KE and MEND fail to generalize to paraphrases, alternately describing the subject as either (c,d,e1) in medicine or (c1,e,d1) in physics depending on the prompt's wording. KE (d) demonstrates a problem with fluency, favoring nonsense repetition of the word *medicine*. **Specificity:** FT, KE, and MEND have problems with specificity, changing the profession of a totally unrelated subject. Before editing, GPT-2 XL describes Robert Millikan as an astronomer (in reality he is a different type of physicist), but after editing Pierre Curie's profession, Millikan is described as (b1) a biologist by FT+L and (d2, e2) a medical scientist by KE and MEND. In contrast, ROME is specific, leaving Millikan's field unchanged. See Appendix F for additional examples.

### 3.6 Limitations

We have shed light on the mechanisms of factual association within GPT, but we have not investigated other kinds of learned beliefs such as logical, spatial, or numerical knowledge. Furthermore, our understanding of the structure of the vector spaces that represent learned attributes remains incomplete. Our evaluation reveals that, even when the a model's stored factual association is changed successfully, the model will guess plausible new facts that have no basis in evidence and that are likely to be false. This may limit the usefulness of a language model as a source of facts. Developing a better understanding of such guessing behavior is a promising area for future work.



## 4 Related Work

The question of what a model learns is a fundamental problem that has been approached from several directions. One line of work studies which properties are encoded in internal model representations, most commonly by training a probing classifier to predict said properties from the representations (Adi et al., 2017; Hupkes et al., 2018; Conneau et al., 2018, inter alia). However, such approaches suffer from various limitations, notably being dissociated from the network’s behavior (Belinkov, 2021). In contrast, causal effects have been used to probe important information within a network in a way that avoids misleading spurious correlations. Vig et al. (2020) introduced the use of causal mediation to identify individual neurons that contribute to biased gender assumptions, and Finlayson et al. (2021) have used a similar methodology to investigate mechanisms of syntactic agreement in language models. Feder et al. (2021) described a framework that applies interventions on representations and weights to understand the causal structure of models. Elazar et al. (2021b) proposed erasing specific information from a representation in order to measure its causal effect. Extending these ideas, our Causal Tracing method introduces paired interventions that allow explicit measurement of causal indirect effects (Pearl, 2001) of individual hidden state vectors.

Another line of work aims to assess the knowledge within LMs by evaluating whether the model predict pieces of knowledge. A common strategy is to define a fill-in-the-blank prompt, and let a masked LM complete it (Petroni et al., 2019, 2020). Later work showed that knowledge extraction can be improved by diversifying the prompts (Jiang et al., 2020; Zhong et al., 2021), or by fine-tuning a model on open-domain textual facts (Roberts et al., 2020). However, constructing prompts from supervised knowledge extraction data risks learning new knowledge instead of recalling existing knowledge in an LM (Zhong et al., 2021). More recently, Elazar et al. (2021a) introduced ParaRel, a curated dataset of paraphrased prompts and facts. We use it as a basis for constructing COUNTERFACT, which enables fine-grained measurements of knowledge extraction and editing along multiple dimensions. Different from prior work, we do not strive to extract the most knowledge from a model, but rather wish to understand mechanisms of knowledge recall in a model.

Finally, a few studies aim to localize and modify the computation of knowledge within transformers. Geva et al. (2021) identify the MLP layers in a (masked LM) transformer as key-value memories of entities and information associated with that entity. Building on this finding, Dai (2021) demonstrate a method to edit facts in BERT by writing the embedding of the object into certain rows of the MLP matrix. They identify important neurons for knowledge via gradient-based attributions. De Cao et al. (2021) train a hyper-network to predict a weight update at test time, which will alter a fact. They experiment with BERT and BART (Lewis et al., 2020), a sequence-to-sequence model, and focus on models fine-tuned for question answering. Mitchell et al. (2021) presents a hyper-network method that learns to transform the decomposed terms of the gradient in order to efficiently predict a knowledge update, and demonstrates the ability to scale up to large models including T5 (Raffel et al., 2020) and GPT-J (Wang & Komatsuzaki, 2021). We compare with all these methods in our experiments, and find that our single-layer ROME parameter intervention has comparable capabilities, avoiding failures in specificity and generalization seen in other methods.

## 5 Conclusion

We have clarified information flow during knowledge recall in autoregressive transformers, and furthermore exploited this understanding to develop a simple, principled model editor called ROME. Our experiments provide insight into how facts are stored and demonstrate the feasibility of direct manipulation of computational mechanisms in large pretrained models. Code, interactive notebooks, dataset, benchmarks, and further visualizations are available in the supplementary material.

**Ethical Considerations.** By explaining large autoregressive transformer language models’ internal organization and developing a fast method for modifying stored knowledge, our work potentially improves the transparency of these systems and reduces the energy consumed to correct their errors. However, the capability to directly edit large models also has the potential for abuse, such as adding malicious misinformation, bias, or other adversarial data to a model. Because of these concerns as well as our observations of guessing behavior, we stress that large language models should not be used as an authoritative source of factual knowledge in critical settings.

## References

- Adi, Y., Kermany, E., Belinkov, Y., Lavi, O., and Goldberg, Y. Fine-grained analysis of sentence embeddings using auxiliary prediction tasks. In *International Conference on Learning Representations (ICLR)*, April 2017.
- Anderson, J. A. A simple neural network generating an interactive memory. *Mathematical biosciences*, 14(3-4):197–220, 1972.
- Bau, D., Liu, S., Wang, T., Zhu, J.-Y., and Torralba, A. Rewriting a deep generative model. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020.
- Belinkov, Y. Probing Classifiers: Promises, Shortcomings, and Advances. *Computational Linguistics*, pp. 1–13, 11 2021. ISSN 0891-2017. doi: 10.1162/coli\_a\_00422. URL [https://doi.org/10.1162/coli\\_a\\_00422](https://doi.org/10.1162/coli_a_00422).
- Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D., Wu, J., Winter, C., Hesse, C., Chen, M., Sigler, E., Litwin, M., Gray, S., Chess, B., Clark, J., Berner, C., McCandlish, S., Radford, A., Sutskever, I., and Amodei, D. Language models are few-shot learners. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M. F., and Lin, H. (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 1877–1901. Curran Associates, Inc., 2020. URL <https://proceedings.neurips.cc/paper/2020/file/1457c0d6bfc4967418bfb8ac142f64a-Paper.pdf>.
- Conneau, A., Kruszewski, G., Lample, G., Barrault, L., and Baroni, M. What you can cram into a single  $\&\!#^*$  vector: Probing sentence embeddings for linguistic properties. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 2126–2136, Melbourne, Australia, July 2018. Association for Computational Linguistics. doi: 10.18653/v1/P18-1198. URL <https://aclanthology.org/P18-1198>.
- Dai, D., Dong, L., Hao, Y., Sui, Z., and Wei, F. Knowledge neurons in pretrained transformers, 2021.
- Dai, H. Learning nonlocal phonotactics in strictly piecewise phonotactic model. In *Proceedings of the Society for Computation in Linguistics 2021*, pp. 401–402, Online, February 2021. Association for Computational Linguistics. URL <https://aclanthology.org/2021.scil-1.45>.
- De Cao, N., Aziz, W., and Titov, I. Editing factual knowledge in language models. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pp. 6491–6506, Online and Punta Cana, Dominican Republic, November 2021. Association for Computational Linguistics. URL <https://aclanthology.org/2021.emnlp-main.522>.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pp. 4171–4186, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics. doi: 10.18653/v1/N19-1423. URL <https://aclanthology.org/N19-1423>.
- Elazar, Y., Kassner, N., Ravfogel, S., Ravichander, A., Hovy, E., Schütze, H., and Goldberg, Y. Measuring and Improving Consistency in Pretrained Language Models. *Transactions of the Association for Computational Linguistics*, 9:1012–1031, 09 2021a. ISSN 2307-387X. doi: 10.1162/tacl\_a\_00410. URL [https://doi.org/10.1162/tacl\\_a\\_00410](https://doi.org/10.1162/tacl_a_00410).
- Elazar, Y., Ravfogel, S., Jacovi, A., and Goldberg, Y. Amnesic probing: Behavioral explanation with amnesic counterfactuals. *Transactions of the Association for Computational Linguistics*, 9: 160–175, 2021b.
- Elhage, N., Nanda, N., Olsson, C., Henighan, T., Joseph, N., Mann, B., Askell, A., Bai, Y., Chen, A., Conerly, T., DasSarma, N., Drain, D., Ganguli, D., Hatfield-Dodds, Z., Hernandez, D., Jones, A., Kernion, J., Lovitt, L., Ndousse, K., Amodei, D., Brown, T., Clark, J., Kaplan, J., McCandlish, S., and Olah, C. A mathematical framework for transformer circuits. <https://transformer-circuits.pub/2021/framework/index.html>, December 2021.

372 Feder, A., Oved, N., Shalit, U., and Reichart, R. CausaLM: Causal model explanation through  
373 counterfactual language models. *Computational Linguistics*, 47(2):333–386, 2021.

374 Finlayson, M., Mueller, A., Gehrmann, S., Shieber, S., Linzen, T., and Belinkov, Y. Causal analysis  
375 of syntactic agreement mechanisms in neural language models. In *Proceedings of the 59th*  
376 *Annual Meeting of the Association for Computational Linguistics and the 11th International Joint*  
377 *Conference on Natural Language Processing (Volume 1: Long Papers)*, pp. 1828–1843, Online,  
378 August 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.acl-long.144.  
379 URL <https://aclanthology.org/2021.acl-long.144>.

380 Geva, M., Schuster, R., Berant, J., and Levy, O. Transformer feed-forward layers are key-value memo-  
381 ries. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*,  
382 pp. 5484–5495, Online and Punta Cana, Dominican Republic, November 2021. Association for  
383 Computational Linguistics. URL <https://aclanthology.org/2021.emnlp-main.446>.

384 Hase, P., Diab, M., Celikyilmaz, A., Li, X., Kozareva, Z., Stoyanov, V., Bansal, M., and Iyer, S. Do  
385 language models have beliefs? methods for detecting, updating, and visualizing model beliefs.  
386 *arXiv preprint arXiv:2111.13654*, 2021.

387 Hupkes, D., Veldhoen, S., and Zuidema, W. Visualisation and ‘diagnostic classifiers’ reveal how  
388 recurrent and recursive neural networks process hierarchical structure. *Journal of Artificial*  
389 *Intelligence Research*, 61:907–926, 2018.

390 Jiang, Z., Xu, F. F., Araki, J., and Neubig, G. How can we know what language models know?  
391 *Transactions of the Association for Computational Linguistics*, 8:423–438, 2020. doi: 10.1162/  
392 tacl\_a\_00324. URL <https://aclanthology.org/2020.tacl-1.28>.

393 Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. In Bengio, Y. and LeCun,  
394 Y. (eds.), *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA,*  
395 *USA, May 7-9, 2015, Conference Track Proceedings*, 2015. URL [http://arxiv.org/abs/1412.](http://arxiv.org/abs/1412.6980)  
396 [6980](http://arxiv.org/abs/1412.6980).

397 Kohonen, T. Correlation matrix memories. *IEEE transactions on computers*, 100(4):353–359, 1972.

398 Lewis, M., Liu, Y., Goyal, N., Ghazvininejad, M., Mohamed, A., Levy, O., Stoyanov, V., and  
399 Zettlemoyer, L. BART: Denoising sequence-to-sequence pre-training for natural language  
400 generation, translation, and comprehension. In *Proceedings of the 58th Annual Meeting of*  
401 *the Association for Computational Linguistics*, pp. 7871–7880, Online, July 2020. Associ-  
402 ation for Computational Linguistics. doi: 10.18653/v1/2020.acl-main.703. URL [https:](https://aclanthology.org/2020.acl-main.703)  
403 [//aclanthology.org/2020.acl-main.703](https://aclanthology.org/2020.acl-main.703).

404 Mitchell, E., Lin, C., Bosselut, A., Finn, C., and Manning, C. D. Fast model editing at scale, 2021.

405 Pearl, J. Direct and indirect effects. In *Proceedings of the Seventeenth conference on Uncertainty in*  
406 *artificial intelligence*, pp. 411–420, 2001.

407 Pearl, J. *Causality: Models, Reasoning and Inference*. Cambridge University Press, USA, 2nd  
408 edition, 2009. ISBN 052189560X.

409 Petroni, F., Rocktäschel, T., Riedel, S., Lewis, P., Bakhtin, A., Wu, Y., and Miller, A. Language  
410 models as knowledge bases? In *Proceedings of the 2019 Conference on Empirical Methods in*  
411 *Natural Language Processing and the 9th International Joint Conference on Natural Language*  
412 *Processing (EMNLP-IJCNLP)*, pp. 2463–2473, Hong Kong, China, November 2019. Association  
413 for Computational Linguistics. doi: 10.18653/v1/D19-1250. URL [https://aclanthology.](https://aclanthology.org/D19-1250)  
414 [org/D19-1250](https://aclanthology.org/D19-1250).

415 Petroni, F., Lewis, P., Piktus, A., Rocktäschel, T., Wu, Y., Miller, A. H., and Riedel, S. How context  
416 affects language models’ factual predictions. In *Automated Knowledge Base Construction*, 2020.

417 Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., Sutskever, I., et al. Language models are  
418 unsupervised multitask learners. *OpenAI blog*, pp. 9, 2019.

419 Raffel, C., Shazeer, N., Roberts, A., Lee, K., Narang, S., Matena, M., Zhou, Y., Li, W., and Liu, P. J.  
 420 Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine*  
 421 *Learning Research*, 21(140):1–67, 2020.

422 Roberts, A., Raffel, C., and Shazeer, N. How much knowledge can you pack into the param-  
 423 eters of a language model? In *Proceedings of the 2020 Conference on Empirical Methods*  
 424 *in Natural Language Processing (EMNLP)*, pp. 5418–5426, Online, November 2020. Associ-  
 425 ation for Computational Linguistics. doi: 10.18653/v1/2020.emnlp-main.437. URL [https://](https://aclanthology.org/2020.emnlp-main.437)  
 426 [aclanthology.org/2020.emnlp-main.437](https://aclanthology.org/2020.emnlp-main.437).

427 Sundararajan, M., Taly, A., and Yan, Q. Axiomatic attribution for deep networks. In *International*  
 428 *conference on machine learning*, pp. 3319–3328. PMLR, 2017.

429 Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and  
 430 Polosukhin, I. Attention is all you need. In *Advances in neural information processing systems*, pp.  
 431 5998–6008, 2017.

432 Vig, J., Gehrmann, S., Belinkov, Y., Qian, S., Nevo, D., Singer, Y., and Shieber, S. M. Investigating  
 433 gender bias in language models using causal mediation analysis. In *NeurIPS*, 2020.

434 Wang, B. and Komatsuzaki, A. GPT-J-6B: A 6 Billion Parameter Autoregressive Language Model.  
 435 <https://github.com/kingoflolz/mesh-transformer-jax>, May 2021.

436 Zhang, Y., Galley, M., Gao, J., Gan, Z., Li, X., Brockett, C., and Dolan, W. B. Generating informative  
 437 and diverse conversational responses via adversarial information maximization. In *NeurIPS*, 2018.

438 Zhao, S., Pascual, D., Brunner, G., and Wattenhofer, R. Of non-linearity and commutativity in BERT.  
 439 In *2021 International Joint Conference on Neural Networks (IJCNN)*, pp. 1–8. IEEE, 2021.

440 Zhong, Z., Friedman, D., and Chen, D. Factual probing is [MASK]: Learning vs. learning to  
 441 recall. In *Proceedings of the 2021 Conference of the North American Chapter of the Association*  
 442 *for Computational Linguistics: Human Language Technologies*, pp. 5017–5033, Online, June  
 443 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.naacl-main.398. URL  
 444 <https://aclanthology.org/2021.naacl-main.398>.

445 Zhu, C., Rawat, A. S., Zaheer, M., Bhojanapalli, S., Li, D., Yu, F., and Kumar, S. Modifying  
 446 memories in transformer models, 2020.

## 447 Checklist

- 448 1. For all authors...
- 449 (a) Do the main claims made in the abstract and introduction accurately reflect the paper's  
450 contributions and scope? [Yes]
- 451 (b) Did you describe the limitations of your work? [Yes]
- 452 (c) Did you discuss any potential negative societal impacts of your work? [Yes]
- 453 (d) Have you read the ethics review guidelines and ensured that your paper conforms to  
454 them? [Yes]
- 455 2. If you are including theoretical results...
- 456 (a) Did you state the full set of assumptions of all theoretical results? [Yes] In appendix
- 457 (b) Did you include complete proofs of all theoretical results? [Yes] In appendix
- 458 3. If you ran experiments...
- 459 (a) Did you include the code, data, and instructions needed to reproduce the main experi-  
460 mental results (either in the supplemental material or as a URL)? [Yes] In supplemental  
461 materials
- 462 (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they  
463 were chosen)? [Yes] In appendix
- 464 (c) Did you report error bars (e.g., with respect to the random seed after running experi-  
465 ments multiple times)? [Yes]
- 466 (d) Did you include the total amount of compute and the type of resources used (e.g., type  
467 of GPUs, internal cluster, or cloud provider)? [Yes] In appendix
- 468 4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
- 469 (a) If your work uses existing assets, did you cite the creators? [Yes]
- 470 (b) Did you mention the license of the assets? [Yes] In appendix
- 471 (c) Did you include any new assets either in the supplemental material or as a URL? [Yes]  
472 Supplemental materials
- 473 (d) Did you discuss whether and how consent was obtained from people whose data you're  
474 using/curating? [Yes] In appendix
- 475 (e) Did you discuss whether the data you are using/curating contains personally identifiable  
476 information or offensive content? [Yes] In appendix
- 477 5. If you used crowdsourcing or conducted research with human subjects...
- 478 (a) Did you include the full text of instructions given to participants and screenshots, if  
479 applicable? [N/A]
- 480 (b) Did you describe any potential participant risks, with links to Institutional Review  
481 Board (IRB) approvals, if applicable? [N/A]
- 482 (c) Did you include the estimated hourly wage paid to participants and the total amount  
483 spent on participant compensation? [N/A]