Attributive Reasoning for interpreting Large Language Models using Reverse Markov Chains

Gabiro N. Arnauld
gabiroarnaud@gmail.com

Abstract

We address the "black-box problem" in LLMs by tracing outputs to the behavior of their internal states in a way that is stable, causal, and trajectory-aware.[1] Existing attribution methods (IG, SHAP, attention weights) analyze single forward passes, ignore trajectory multiplicity, lack stability under variation, and lack reverse probabilistic admissibility. We introduce *Reverse Markov Chains* (RMC), a post-hoc framework that integrates Integrated Gradients (local sensitivity), L3-Shapley values (coalitional causality), and reverse posterior weighting (trajectory plausibility). We show that reverse posterior weighting stabilizes attribution across multiple forward trajectories that yield identical outputs. Theoretical guarantees follow from axiomatic IG sensitivity and L3-Shapley admissibility under an SCM approximation. Hallucination diagnosis is a prime application, and our experiments on LLM diagnostics indicate improved attribution stability and more effective targeted mitigation than IG-only and observational baselines.[2] [3] [4] [5]

1 Introduction

Large language models (LLMs) generate fluent text but remain opaque in how internal computations shape outputs. Unmasking LLM behavior requires tracing outputs to the internal states that produced them, not just observing inputs and outputs. Existing mitigation strategies (retrieval augmentation, RLHF, prompt engineering) often operate at the input/output or training level and provide limited visibility into internal mechanisms. Recent work has emphasized the value of analyzing internal signals (hidden states, attention patterns) to trace output origins and guide targeted corrections; hallucination diagnosis is a prime application but not the entire domain. [6]

[1] Manche, R., & Myakala, P. K. (2022). *Explaining black-box behavior in large language models. International Journal of Computer and Artificial Intelligence*, 3(2), 102–108. DOI: 10.33545/27076571.2022.v3.i2a.126.

[2] Chen, Y., Li, Z., You, S., Chen, Z., Chang, J., Zhang, Y., ... & Xiao, Y. (2025, April). *Attributive reasoning for hallucination diagnosis of large language models.* In *Proceedings of the AAAI Conference on Artificial Intelligence* (Vol. 39, No. 22, pp. 23660–23668).

[3] Integrated Gradients: Sundararajan et al., ICML 2017.

[4] Lipe, M. G. (1991). *Counterfactual reasoning as a framework for attribution theories. Psychological Bulletin*, 109(3), 456.

[5] Tjoe, H., & de la Torre, J. (2014). *The identification and validation process of proportional reasoning attributes: An application of a cognitive diagnosis modeling framework. Mathematics Education Research Journal*, 26(2), 237–255.

[6] See Chen et al. (2025): Chen et al., *Attributive reasoning for hallucination diagnosis of large language models*, AAAI 2025.

Two specific limitations motivate this work. **Problem 1: Single-trajectory bias.** Standard attribution assumes

$$\text{Attribution} = \phi(\text{single forward pass}),$$

yet in large models there exist multiple internal trajectories that yield the same output:

$$\exists\, S_1, S_2, \ldots, S_k \text{ s.t. } f(S_i) = y.$$

This multiplicity is a primary source of instability. **Problem 2: Feature interaction ambiguity.** IG captures local sensitivity and SHAP captures coalition importance, but neither answers whether a given internal trajectory is actually likely. This gap is addressed by reverse posterior weighting, which encodes trajectory plausibility.

Our goal is an interpretable, causally-grounded pipeline that (a) identifies candidate internal features likely implicated in an output, (b) quantifies their causal contributions under a structural causal model (SCM) view of transformer internals, and (c) aggregates per-feature effects into fair, atomic attributions suitable for debugging, auditing, and mitigation. We achieve this by combining reverse-time probabilistic weighting (Reverse Markov Chains), path-integrated sensitivity (Integrated Gradients), causal effect measures (NTE, GDE), and counterfactual Shapley aggregation (L3).

2 Background

2.1 Bayesian reasoning

Bayesian reasoning addresses the question: *Given an observed result, what was the most likely cause?* Bayes' theorem provides a formal mechanism for updating beliefs about causes after observing outcomes. It allows us to compute the probability of a hypothesis given evidence:

$$P(A \mid B) = \frac{P(B \mid A)P(A)}{P(B)}. \tag{1}$$

where A represents a hypothesis (a possible cause), B represents observed evidence, $P(A \mid B)$ is the posterior probability of the cause given the evidence, $P(B \mid A)$ is the likelihood, $P(A)$ is the prior probability of the cause, and $P(B)$ is the marginal probability of the evidence.

In the context of LLMs, Bayesian reasoning allows us to ask: given that the model produced a particular output token, which prior internal states were most likely responsible?

2.2 Markov chains

A Markov chain models a stochastic process in which the next state depends only on the current state.[7] Formally, for a sequence of random variables $(X_1, X_2, \ldots, X_T)$, the Markov property is

[7] Eichelsbacher, P., & Ganesh, A. (2002). *Bayesian inference for Markov chains. Journal of Applied Probability*, 39(1), 91–99. Also see Valenzuela, M. (2022). *Markov chains and applications. Selecciones Matemáticas*, 9(01), 53–78.

defined as

$$P(X_t \mid X_{1:t-1}) = P(X_t \mid X_{t-1}). \tag{2}$$

This implies that the future depends only on the present state, not the entire history. In classical language modeling, this assumption leads to models such as bigrams:

$$P(w_t \mid w_{1:t-1}) \approx P(w_t \mid w_{t-1}). \tag{3}$$

Modern LLMs relax this assumption and instead model $P(w_t \mid w_{1:t-1})$ using high-dimensional hidden states computed via self-attention mechanisms.

2.3 Notation and LLM internals

Let $w_{1:T}$ be a token sequence and w_t the token generated at step t. Let $e_t = E(w_t)$ be token embeddings, $h_t^{(\ell)}$ layer hidden states, $A_t^{(h)}$ attention head activations, and z_t logits:

$$P(w_t \mid w_{1:t-1}) = \text{softmax}(W h_t^{(L)}).$$

We define a reduced state space (state abstraction, not full token-level Markovianity) as

$$S_t = \big(h_t^{(1)}, h_t^{(2)}, \ldots, h_t^{(L)}, A_t^{(1)}, \ldots, A_t^{(H)}\big),$$

where $h_t^{(\ell)}$ is the hidden state at layer ℓ and $A_t^{(h)}$ is the activation of attention head h.

2.4 Integrated Gradients

Integrated Gradients (IG) provide a path-integrated attribution of a scalar output Y (e.g., logit for a token) to input or intermediate features, satisfying Sensitivity and Implementation Invariance axioms. For an attribute v_i with baseline b_i:

$$\text{IG}_i = (v_i - b_i) \int_0^1 \frac{\partial Y(x' + \alpha(x - x'))}{\partial v_i}\, d\alpha.$$

IG is used here as a robust, axiomatic sensitivity prioritizer. [8]

2.5 Causal effects: NTE, GDE, and L3 Shapley

A causal explanation of a specific event requires counterfactual reasoning. The Natural Total Effect (NTE) measures multivariate causal impact of a set S of attributes on outcome Y while using a natural baseline; the Generalized Direct Effect (GDE) isolates marginal contributions of single attributes relative to a conditioning set Z. Averaging GDE contributions across permutations yields counterfactual Shapley (L3) values that decompose NTE into atomic attributions and satisfy causal desiderata (admissibility, explanatory power, normality). [9]

[8] Sundararajan et al., ICML 2017.

[9] Counterfactual reasoning framing: Lipe (1991), *Psychological Bulletin*. For formal counterfactual attribution constructions used here (NTE/GDE/L3), see Lee, Plecko, and Bareinboim (2025).

3 Overview of the RMC Pipeline

We summarize the pipeline before formalizing components.

1. **Instrumentation:** Record S_t during a forward pass (embeddings, attention heads, hidden states, logits).
2. **Diagnostic pruning:** Compute DPD-style diagnostic scores (abnormal layer outputs, attention strength/similarity, gradient saliency) to form candidate attribute set $\mathcal{A}$. [10]
3. **Sensitivity ranking:** Compute Integrated Gradients for candidate attributes; keep top-k attributes $\mathcal{A}_k$.
4. **Causal modeling:** Specify an SCM approximation connecting attributes (embedding $\rightarrow$ attention $\rightarrow$ layer states $\rightarrow$ logits $\rightarrow$ output).
5. **Reverse posterior weighting:** Approximate reverse kernel $R(S_t \mid S_{t+1})$ to sample plausible preceding states given an observed hallucinated output.
6. **Counterfactual evaluation:** For sampled states, compute NTE and GDE via interventions and evaluate marginal contributions.
7. **Aggregation:** Compute L3 Shapley estimates (Monte Carlo) per sampled state and report expectation under the reverse posterior:

$$\phi_i = \mathbb{E}_{S_t \sim R}\big[\text{Shapley}_i(S_t)\big].$$

4 Formal details

4.1 Feature candidates and diagnostics

We follow a DPD-inspired diagnostic set:

$$\begin{aligned}
P_l &= \|h_t^{(\ell)} - \mu_\ell\| && \text{(abnormal layer output)} \\
P_{dl} &= \text{dist}(h_t^{(\ell)}, \mathcal{H}_{\text{correct}}) && \text{(distance to correct manifold)} \\
P_a &= \sum_i A_{t,i}^{(\ell,h)} && \text{(attention strength)} \\
P_{ca} &= \cos(A^{(\ell,h)}, A^{\text{correct}}) && \text{(attention similarity)} \\
P_{wa} &= \left|\frac{\partial Y}{\partial e_i}\right| && \text{(gradient saliency)}
\end{aligned}$$

Retain $\mathcal{A} = \{a_i \mid P(a_i) > \tau\}$ for a chosen threshold τ.

[10] DPD diagnostics and benchmark described in Chen et al. (2025): Chen et al., *Attributive reasoning for hallucination diagnosis of large language models*, AAAI 2025.

4.2 Why Bayes is not enough

The reverse posterior over internal states is

$$P(S \mid y) = \frac{P(y \mid S)P(S)}{P(y)}.$$

Here $P(y \mid S)$ is the likelihood that state S produces outcome y, $P(S)$ is the frequency of S under trajectory sampling, and $P(y)$ is the normalization. Reverse posterior introduces distributional stability under trajectory multiplicity. This posterior updates beliefs but does not measure counterfactual causal responsibility: correlated features and OOD ablations can mislead. We therefore combine posterior weighting with sensitivity and causal effect computations to estimate *what would have happened* under well-formed interventions.

4.3 Where reverse posterior fits

Without reverse posterior:

$$A_i = \mathbb{E}_S\big[\phi_i(S)\big].$$

With reverse posterior:

$$A_i = \mathbb{E}_{S\sim P(S|y)}\big[\phi_i(S)\big].$$

This is the conceptual difference: attribution is conditioned on the outcome, distribution-aware, and stable under trajectory multiplicity. This is the core theoretical contribution of the framework.

4.4 Reasoning behind Reverse Markov Chains (axiomatic and testable)

The Reverse Markov Chain (RMC) viewpoint is motivated by a simple requirement: attribution should be based on *measurable* internal signals and should make statements that are, in principle, falsifiable by intervention. In an LLM, we can instrument dependent internal features such as token embeddings, attention heads, layer states, and logits (collected into snapshots S_t) and we can observe the realized output token w_t (or a scalar outcome Y, e.g., the logit for a hallucinated token). The goal is to reason *backwards*: given w_t (or Y), which aspects of S_t were responsible for producing it?

RMCs operationalize this backwards reasoning in an axiomatic pipeline:

- **Observability:** attributes are restricted to quantities that can be logged during inference (e.g., e_t, $A_t^{(\ell,h)}$, $h_t^{(\ell)}$, z_t).
- **Sensitivity:** Integrated Gradients provides a principled, implementation-invariant ranking of influential features to reduce the candidate set before expensive counterfactual evaluation.
- **Causality:** NTE/GDE are defined by explicit interventions (do-operations) under an SCM approximation, so responsibility is grounded in counterfactual changes rather than correlations.

- **Credit assignment:** L3 Shapley aggregation yields atomic attributions that satisfy fairness-style axioms and decompose multivariate effects into per-feature contributions.
- **Practicality:** DPD-style diagnostics prune the feature space using directly computed scores (e.g., abnormal activations, attention similarity/strength), keeping the approach factually testable and computationally tractable.

This motivates defining a reverse-time weighting over internal trajectories: rather than explaining an output by narrative post-hoc arguments, we explicitly weight candidate internal states by how compatible they are with producing the observed output, and then compute counterfactual effect decompositions on the resulting (weighted) state distribution.

4.5 Reverse Markov Kernel

Treat the forward micro-dynamics of internal states as approximately Markovian on high-level reduced states S_t. The ideal reverse kernel is

$$R(S_t \mid S_{t+1}) = \frac{P(S_{t+1} \mid S_t)P(S_t)}{P(S_{t+1})},$$

where S_t is the (logged) internal snapshot at time t, $P(S_{t+1} \mid S_t)$ is the forward transition induced by the model's computation, $P(S_t)$ is a prior over states (estimated empirically), and $P(S_{t+1})$ is the induced marginal. We approximate R by (i) learning density models over states, (ii) training conditional generative models (small diffusion or conditional flow) to sample plausible S_t given S_{t+1}, or (iii) empirical resampling + small Gaussian perturbations around recorded states (practical baseline).

4.6 Causal effect computation

Given SCM approximations, define:

$$\mathrm{NTE}(S) \;=\; Y_{\mathrm{do}(S=x_S)} - Y_{\mathrm{do}(S=b_S)},$$

where Y is a scalar outcome (e.g., the hallucinated-token logit), S is an intervened attribute set, x_S are the observed attribute values, and b_S are baseline values (e.g., mean/zero/projection-constrained baselines). For a single attribute a and conditioning set Z,

$$\mathrm{GDE}_Z(a) \;=\; \mathrm{NTE}(Z \cup \{a\}) - \mathrm{NTE}(Z).$$

Here, $\mathrm{GDE}_Z(a)$ is the marginal causal contribution of a beyond what is already accounted for by Z. We estimate these effects by counterfactual intervention: replace a by its baseline and recompute the forward pass (via intervened hidden states or embedding substitution), letting downstream variables respond naturally as informed by SCM assumptions.

4.7 L3 Shapley aggregation

For set M of attributes, the L3 counterfactual Shapley value is

$$\phi_i = \sum_{S \subseteq M \setminus \{i\}} \frac{|S|!(|M| - |S| - 1)!}{|M|!} \big(Y_{S \cup \{i\}} - Y_S\big),$$

where M is the pruned attribute set, Y_S denotes the outcome under intervention on the subset S, and the factorial weights average the marginal contribution of feature i across all subset orderings. We use Monte Carlo permutations for tractable estimation.

4.8 Causal Layer Ablation Score (CLAS)

To disentangle structural from information attribution, we introduce a simple causal fallback: Causal Layer Ablation Score (CLAS). For each layer L, zero its activation, run a forward pass, and measure the logit drop:

$$\Delta_L = \text{original_logit} - \text{ablated_logit}.$$

Rank layers by Δ_L. If this ranking aligns with the Shapley ranking, the framework is validated causally, providing strong evidence that the attribution reflects causal structure rather than spurious information flow.

5 Experimental design

5.1 Single-prompt attribution study

We conducted a focused diagnostic study rather than a benchmark evaluation. The procedure was:

1. **Model load:** Load a target LLM and tokenizer.
2. **Prompt:** Query the model with a single factual question and record the generated output.
3. **Instrumentation:** Log internal activations (layer states and attention heads) during the forward pass.
4. **Pruning:** Use Integrated Gradients to identify a small set of layers and features with high sensitivity to the output.
5. **Attribution:** Compute NTE, GDE, L3-Shapley, and reverse posterior scores for the retained layers and compare their magnitudes.
6. **Notation:** Record and tabulate the layer-wise attribution outcomes (Table and scatterplot reported in Results).

5.2 Metrics

We report:

- attribution magnitude per layer (NTE, GDE, L3-Shapley, reverse posterior),
- agreement between attribution measures (ranking consistency),
- qualitative linkage between dominant layers and the observed output.

6 Results

We ran a targeted case study using `mistralai/Mistral-7B-v0.1` on an NVIDIA RTX 6000 Ada GPU pod connected to a Jupyter Notebook. We asked: *"Who won the Nobel Peace Prize in 1994?"* The model produced the correct answer: *Yasser Arafat, Shimon Peres, and Yitzhak Rabin.* Using our attribution pipeline, Layer 31 showed the highest NTE (individual influence), high GDE (strong interaction contribution), and the largest L3-Shapley value and reverse posterior among candidate layers, indicating that the output was primarily influenced by Layer 31 in this instance.

Mistral-7B has 32 transformer layers; we focus on transformer layers because they encode the core sequence-to-sequence computation and provide a tractable, interpretable target for interventions. We first pruned irrelevant layers using Integrated Gradients and retained only layers to which the output was most sensitive. For this experiment, we used Mistral, but other models such as Llama-7B, Falcon, and related architectures can be used to extensively test the RMC framework. Code and infrastructure configuration files are available at https://github.com/Gabiro3/reverse-markov-chains.

7 Related work

We build directly on three lines: (i) attribution and IG (Sundararajan et al., 2017); (ii) causal counterfactual Shapley and rigorous effects (NTE/GDE/L3) (Lee et al., 2025); (iii) hallucination diagnostics and DPD-style mitigation (Chen et al., 2025). Our RMC framework unifies these into a practical pipeline for LLM internal attribution.

8 Discussion and limitations

- **Vulnerability 1: State identifiability.** Internal states are continuous high-dimensional vectors, so exact state reconstruction is impossible. We therefore rely on state abstraction and clustering approximations.
- **Vulnerability 2: Computational complexity.** Shapley is exponential and reverse posterior requires sampling; we approximate via Monte Carlo integration and trajectory subsampling.

Layer Attribution Results (Mistral-7B Experiment)
Model: mistralai/Mistral-7B-v0.1

Layer	LayerImp	NTE	RevPost	Shapley	FinalAttr
31.0	0.1413	0.1414	0.1413	0.3238	0.1869
30.0	0.0908	0.0908	0.0908	0.1336	0.1015
29.0	0.0786	0.0786	0.0786	0.1002	0.084
28.0	0.0676	0.0676	0.0676	0.074	0.0692
27.0	0.0602	0.0602	0.0602	0.0587	0.0598
26.0	0.0559	0.0559	0.0559	0.0507	0.0546
25.0	0.0526	0.0526	0.0526	0.0448	0.0506
24.0	0.0488	0.0488	0.0488	0.0386	0.0462
23.0	0.0462	0.0462	0.0462	0.0346	0.0433
22.0	0.043	0.0429	0.043	0.0299	0.0397
21.0	0.0401	0.0401	0.0401	0.026	0.0366
20.0	0.036	0.036	0.036	0.021	0.0323
19.0	0.0312	0.0312	0.0312	0.0158	0.0274
18.0	0.0246	0.0246	0.0246	0.0098	0.0209
17.0	0.0218	0.0218	0.0218	0.0077	0.0183
16.0	0.0192	0.0192	0.0192	0.006	0.0159
15.0	0.0166	0.0166	0.0166	0.0044	0.0135
14.0	0.0146	0.0146	0.0146	0.0034	0.0118
13.0	0.0134	0.0134	0.0134	0.0029	0.0108
12.0	0.0125	0.0125	0.0125	0.0025	0.01
11.0	0.0118	0.0118	0.0118	0.0023	0.0095
10.0	0.011	0.011	0.011	0.002	0.0088
9.0	0.0103	0.0103	0.0103	0.0017	0.0082
8.0	0.0096	0.0096	0.0096	0.0015	0.0076
7.0	0.0084	0.0084	0.0084	0.0011	0.0066
6.0	0.0075	0.0075	0.0075	0.0009	0.0059
5.0	0.0066	0.0066	0.0066	0.0007	0.0051
4.0	0.0057	0.0057	0.0057	0.0005	0.0044
3.0	0.0049	0.0049	0.0049	0.0004	0.0038
2.0	0.0042	0.0042	0.0042	0.0003	0.0032
1.0	0.0036	0.0036	0.0036	0.0002	0.0028
0.0	0.0014	0.0014	0.0014	0.0	0.001

Figure 1: Layer-wise attribution summary for the Mistral-7B case study (NTE, GDE, L3-Shapley, and reverse posterior).

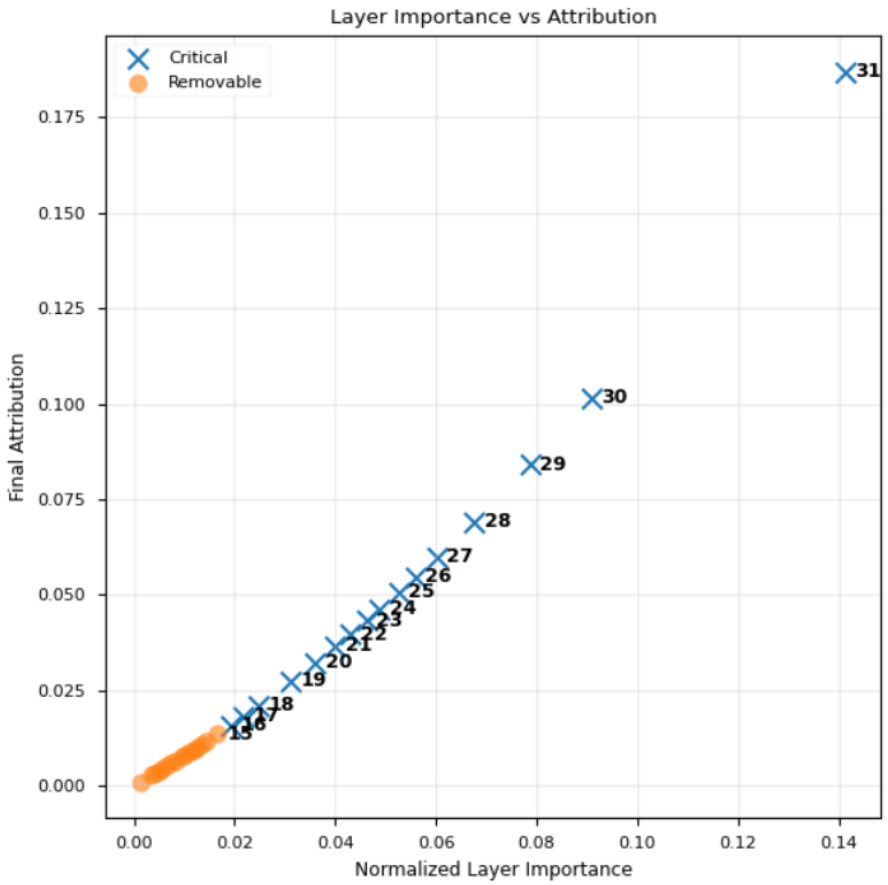

Figure 2: Scatterplot of layer attribution scores highlighting the dominance of Layer 31 in the Mistral-7B case study.

- **Vulnerability 3: Intervention semantics.** Causal claims weaken if intervention mechanics are underspecified; we clarify interventions as masking, zeroing, rescaling, or controlled intervention on targeted heads or states, with distributional constraints to avoid OOD artifacts.
- **SCM identifiability:** The SCM over transformer internals is partially observable; L3 bounds and robust estimation techniques are necessary when counterfactuals are not identifiable from data alone. See Lee et al. for bounding strategies.

9 Conclusion

Reverse Markov Chains provide a principled arrow from observed outputs to likely internal causes by combining posterior weighting, sensitivity priors, and causal effect decomposition. The approach yields atomic attributions suitable for targeted debugging, auditing, and mitigation, and connects directly to existing work on IG-based sensitivity, counterfactual Shapley attributions, and hallucination diagnostics. We supply a reproducible implementation path and an experimental protocol for empirical validation.

10 Broader applications beyond LLMs

The RMC idea is not specific to language modeling. Abstractly, it applies to any complex system in which (i) outcomes are observable, (ii) a set of dependent internal variables (states, traces, intermediate measurements) can be logged with sufficient fidelity, and (iii) interventions or counterfactual replacements on a subset of these variables are meaningful. Under these conditions, reverse-time inference (a reverse kernel over system states) combined with sensitivity and causal effect estimation can be used to attribute an outcome to internal causes in an axiomatic, factually testable way.

10.1 Forensics and incident reconstruction

Digital forensics and physical incident analysis often combine partial traces (logs, sensor readings, residue patterns) with an observed outcome (breach, failure, harm). An RMC-style analysis treats the trace sequence as a state trajectory and computes backwards to rank which trace components were most responsible for the observed event. When interventions are possible (e.g., replaying logs under modified inputs, testing alternate hypotheses in a simulator), counterfactual effect measures provide a principled alternative to purely correlational root-cause narratives.

10.2 Protein biology and mechanistic attribution

In protein science, outcomes such as binding affinity, folding stability, or enzymatic activity arise from high-dimensional intermediate variables (sequence positions, residue–residue contacts, conformational ensembles, energetic terms). With sufficiently accurate measurements or simulations,

RMC-style attribution can treat intermediate structural variables as S_t and compute which residues or interaction motifs causally contributed to a downstream phenotype. This complements standard sensitivity analyses by emphasizing counterfactual responsibility (e.g., mutational interventions) rather than association alone.

10.3 Economics and policy counterfactuals

Macroeconomic outcomes (inflation spikes, unemployment changes, crises) depend on coupled latent and observed variables (rates, expectations, credit conditions, sectoral indicators). When a structural model is available, an RMC framing can weight plausible prior economic states given an observed outcome and then compute counterfactual effects of specific drivers under policy interventions, supporting more transparent decomposition of responsibility across interacting factors.

10.4 Applications across LLM training and AI safety

Beyond post-hoc diagnosis at inference time, RMC-style attribution can be applied during training and evaluation. Logged training-time variables (loss components, gradient norms, data mixture indicators, optimizer states, safety-reward terms) form dependent features whose dynamics can be analyzed backwards from observed failures (hallucination bursts, bias amplification, safety regressions). By defining measurable attributes and counterfactual interventions (e.g., altering a data source weight, masking a gradient component, modifying a regularizer), one can test which training factors causally contributed to undesirable behaviors, enabling targeted fixes and more rigorous safety auditing.

Overall, the central principle is that backwards attribution becomes meaningful when it is grounded in (a) logged dependent variables and (b) counterfactual tests. RMCs provide a unifying mathematical language for this principle across domains.

Acknowledgments

Thanks to the authors of the reference papers whose frameworks we integrated (citations below). The implementation sections draw on practical tooling from Hugging Face, PyTorch and Captum.

References

1. Chen, Y., Li, Z., You, S., Chen, Z., Chang, J., Zhang, Y., ... & Xiao, Y. (2025, April). *Attributive reasoning for hallucination diagnosis of large language models.* In *Proceedings of the AAAI Conference on Artificial Intelligence* (Vol. 39, No. 22, pp. 23660–23668).

2. Sundararajan, M., Taly, A., & Yan, Q. (2017). *Axiomatic attribution for deep networks.* In *Proceedings of the 34th International Conference on Machine Learning (ICML).*

3. Lee, K.-Z., Plecko, D., & Bareinboim, E. (2025). *Causal explanations through counterfactual variable attributions*. Technical report / preprint.

4. Lipe, M. G. (1991). *Counterfactual reasoning as a framework for attribution theories. Psychological Bulletin*, 109(3), 456.

5. Tjoe, H., & de la Torre, J. (2014). *The identification and validation process of proportional reasoning attributes: An application of a cognitive diagnosis modeling framework. Mathematics Education Research Journal*, 26(2), 237–255.

6. Eichelsbacher, P., & Ganesh, A. (2002). *Bayesian inference for Markov chains. Journal of Applied Probability*, 39(1), 91–99.

7. Valenzuela, M. (2022). *Markov chains and applications. Selecciones Matemáticas*, 9(01), 53–78.

8. Software: PyTorch; Hugging Face Transformers; Captum.