Computation Graph Recovery from Chain-of-Thought

Anonymous ACL submission

Abstract

Frontier LLMs have made impressive progress on complex multi-step tasks by leveraging Chain-of-Thought (CoT) reasoning, yet we still lack a clear understanding as to whether the intermediate tokens generated as reasoning traces are causally faithful to the final answer i.e. whether models may produce spurious CoT that contains logical or factual errors while still giving the right answer. Using math and physics questions that require at least one implicit intermediate computation, we built an automated interpretability (auto-interp) framework that automatically recovers the *computation graph* of a model's CoT trace i.e. a directed acyclic graph (DAG) whose nodes are intermediate numerical values and whose edges represent genuine causal dependencies. Our pipeline generates token-length-aligned counterfactual trace pairs, performs token patching to produce causal influence heatmaps, and aggregates these into a per-trace DAG via Benjamini–Hochberg FDR-controlled edge selection. Because each edge in the recovered graph is defined by a causal intervention rather than textual co-occurrence, the graph directly reflects mechanistic dependencies in the model's computation. Critically, because models may solve the same problem via different algebraic paths, no single symbolic formula serves as universal ground truth. We compare the recovered graphs to annotated "ground-truth" reference-graphs as a measure of faithfulness.

1 Introduction

Chain-of-Thought (CoT) traces have become a defining feature of how frontier LLMs tackle complex tasks (Wei et al., 2022) across a wide range of reasoning-heavy tasks. As stronger models tend to produce much longer and more complex reasoning traces, a fundamental question about their faithfulness becomes increasingly urgent: do the intermediate tokens of CoT actually drive the internal computation mechanistically, or are they merely post-hoc scaffolding that accompanies, but does not causally determine the final answer? Existing CoT faithfulness studies test whether models' output generally depend on their CoT as a whole, but they often overlooked finer, token-level dependencies within the trace (Lanham et al., 2023; Turpin et al., 2023). Our work aims to identify these direct dependencies at the level of variables in the CoT by reconstructing a causal dependency graph among groups of tokens that denote numerical values.

As an example, consider the physics problem: "An object of m kg has E J of kinetic energy – how long does it take to travel d m?" Answering this question requires computing the bridge variable *velocity* v as an intermediate step by way of producing a CoT trace that contains v explicitly before computing the final answer. The central question is whether that written intermediate value is *causally necessary* for the downstream computation, and more generally whether the dependency structure visible in the text matches the dependency structure the model actually uses.

We formalize this question as computation graph recovery. Let a CoT trace contain numerical value nodes $V = \{v_0, v_1, \ldots, v_K\}$, where each node is a numeric span extracted from the trace. We define the *mechanistic computation graph* $G^{\text{patch}} = (V, E^{\text{patch}})$, where $(v_i, v_j) \in E^{\text{patch}}$ iff counterfactual patching tokens in v_i's span causally influences the generation of v_j. We further define the *annotated computation graph* $G^{\text{annotate}} = (V, E^{\text{annotate}})$, where $(v_i, v_j) \in E^{\text{annotate}}$ if and only if v_i is a direct symbolic input to the computation of v_j. Because a model may solve the same problem through different algebraic paths, the relevant reference graph is trace-specific rather than a single canonical formula graph.

We recover G^{patch} by generating token-length-aligned counterfactual trace pairs, patching one token position at a time, and measuring the resulting

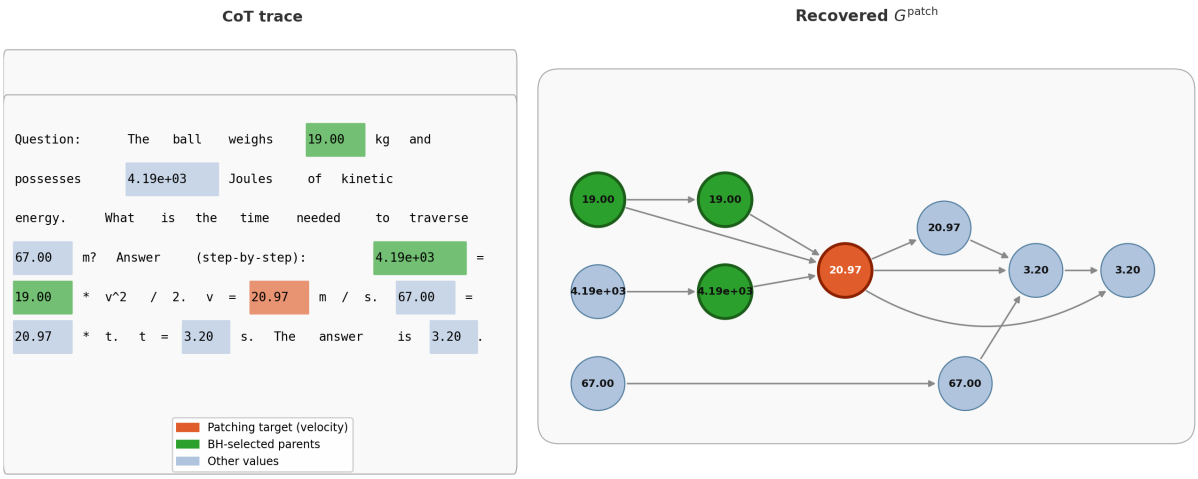

Figure 1: **Overview of our framework** on a representative instance. *Left:* A CoT trace in which intermediate numerical values are highlighted by role: orange is the patching target (velocity v), green its BH-selected causal parents (m, KE), and grey-blue all other values. *Right:* The recovered mechanistic graph G^{patch}. Green arrows from m and KE to v are the edges inferred by counterfactual token patching. Recovering such graphs – and comparing them to reference graphs from GPT-4o – is the main contribution of this work.

change in downstream output log-probabilities. Aggregating causal scores with Benjamini–Hochberg false-discovery-rate control selects statistically significant token positions, which are then mapped onto preceding value nodes to produce directed edges. We compare the resulting intervention-derived graph to the trace-conditioned reference graph G^{annotate} using edge precision, recall, F1, and transitive-closure agreement. Agreement between the two graphs indicates that the model's mechanistic dependencies match the written CoT; disagreement identifies either implicit causal dependencies absent from the text or written dependencies that are not load-bearing under intervention.

Our contribution. We introduce an automatic computation graph recovery pipeline that performs exhaustive single-position token-swap patching, constructs candidate causal graphs using FDR-controlled edge selection, and generates trace-conditioned frontier-model reference graphs for comparison.

2 Related Work

Chain-of-Thought faithfulness. CoT prompting substantially improves LLM performance on multi-step tasks (Wei et al., 2022), and has been extended through variants such as least-to-most prompting (Zhou et al., 2022), self-consistency sampling (Wang et al., 2022), and tree- and graph-of-thought search (Yao et al., 2023; Besta et al., 2024). Modern reasoning-specialised models (Guo et al., 2025) push these capabilities further, making the faithfulness question more pressing. Subsequent work has questioned how faithfully written reasoning reflects internal computation. Faithfulness testing by CoT truncation and deliberate mistake injection (Lanham et al., 2023) finds models are generally faithful on arithmetic, while biased-context experiments (Turpin et al., 2023) show models can produce plausible but mechanistically decoupled explanations. More recent analyses find CoT outputs frequently decoupled from final answers and propose constraints for alignment (Paul et al., 2024; Lyu et al., 2023). Even semantically empty filler tokens can provide the computation cycles needed for accuracy (Pfau et al., 2024), while structured reasoning emerges in base models without instruction tuning (Venhoff et al., 2025), and reasoning can be compressed into latent form without explicit CoT tokens (Deng et al., 2023) – suggesting that written tokens and mechanistic computation can be decoupled in both directions. Our work advances this line by measuring faithfulness *mechanistically* at the level of individual numerical value nodes rather than behaviorally, and by recovering the full causal graph structure rather than reporting a single aggregate faithfulness score.

Mechanistic interpretability and causal interventions. The mechanistic interpretability program (Elhage et al., 2021) seeks to reverse-engineer transformer computations by attributing effects to

specific circuit components. Activation patching – substituting hidden states from a source run into a base run – has been used to localize factual associations (Meng et al., 2022), to interpret arithmetic computations (Stolfo et al., 2023), and to trace multi-hop factual reasoning (Bronzini et al., 2024); however, subsequent work shows that localization-based edits do not always generalise (Hase et al., 2023). Causal abstraction (Geiger et al., 2021) provides formal grounding for such interventions by mapping model internals to structural causal models. Sparse autoencoders reveal monosemantic features in transformer residual streams (Bricken et al., 2023; Templeton et al., 2024; Miller et al., 2026), and recent circuit-tracing work uses attribution graphs to reveal full computational graphs inside language models (Ameisen et al., 2025). Grokking studies provide evidence that transformers learn structured internal algorithms (Power et al., 2022; Nanda et al., 2023), and probing work shows that LLMs represent spatial, temporal, and numerical quantities as structured linear features (Li et al., 2022; **?**; Zhu et al., 2025; Levy and Geva, 2025). Our pipeline adopts the same interventionist logic at the input-token level: rather than patching every layer-position activation, which is substantially more expensive at scale, we patch aligned counterfactual tokens and recover whole computation graphs rather than localising single facts. Token patching reduces the intervention grid from layer-token cells to token positions, yielding comparable graph-level behavior in downstream evaluations while cutting patching cost by over an order of magnitude (see Appendix 5), and targets a higher abstraction level than circuit tracing (Ameisen et al., 2025): token-span causal dependencies that scale to thousands of trace pairs across five models.

Causal graph learning and evaluation. Our evaluation borrows from the Bayesian network structure learning literature: we adopt the Structural Hamming Distance (SHD) (Tsamardinos et al., 2006) as a graph comparison metric, apply Benjamini–Hochberg FDR control (Benjamini and Hochberg, 1995) for edge selection, and interpret all token-patching interventions through Pearl's interventionist causality framework (Pearl, 2000).

Multi-hop and compositional reasoning. A related body of work examines whether transformers compose intermediate facts across reasoning hops (Yang et al., 2024; Biran et al., 2024; Li et al., 2024). While those works focus on accuracy-based evidence for multi-step composition, our focus is on whether the *written* intermediate values in a trace are causally load-bearing – a question that requires explicit graph recovery rather than answer-accuracy probing.

3 Method

3.1 Causal Interpretation of Token Patching

We justify our graph recovery procedure using the interventionist view of causality (Pearl, 2000). For a fixed model with parameters held constant, a CoT trace can be viewed as the output of a structural causal process over token positions: earlier tokens determine later hidden states and logits through the transformer's autoregressive generation (Vaswani et al., 2017), while later tokens cannot causally affect earlier predictions. Thus, token position imposes a natural partial order over possible causal dependencies.

Let $\mathbf{b} = (b_1, \ldots, b_N)$ be a base trace and $\mathbf{s} = (s_1, \ldots, s_N)$ a token-aligned counterfactual trace. For a target value node v_k, let p_k be the first token position at which the base and counterfactual spans for that value differ. For any earlier position $p < p_k$, token patching performs the intervention

$$\mathrm{do}(X_p = s_p) \tag{1}$$

on the base sequence while leaving all other input tokens and all model parameters fixed. Operationally, this creates a patched sequence $\mathbf{b}^{(p)}$ with $b_p^{(p)} = s_p$ and $b_q^{(p)} = b_q$ for $q \neq p$.

We measure the causal effect of position p on the generation of value v_k by the change it induces in the model's next-token distribution at p_k:

$$\Delta_{\text{src}}(p) = \log P_{\mathbf{b}^{(p)}}(s_{p_k}) - \log P_{\mathbf{b}}(s_{p_k}), \tag{2}$$

$$\Delta_{\text{base}}(p) = \log P_{\mathbf{b}^{(p)}}(b_{p_k}) - \log P_{\mathbf{b}}(b_{p_k}), \tag{3}$$

$$\Delta_{\text{diff}}(p) = \Delta_{\text{src}}(p) - \Delta_{\text{base}}(p) \tag{4}$$

This quantity is an individual-level causal contrast: it asks whether replacing one earlier token with its counterfactual value makes the model more likely to continue as the counterfactual trace rather than as the base trace. If $\Delta_{\text{diff}}(p)$ is statistically distinguishable from the background distribution of token effects, then token p is causally relevant for generating v_k under this intervention.

The value-level graph is obtained by lifting token-level effects to numerical spans. We add an edge (v_i, v_k) to G^{patch} when a significant causal effect on v_k is mapped to the earlier value v_i by

Domain	Bridge variable	Formula chain
VELOCITY	Velocity v	$T, m \to v \to t$
CURRENT	Current I	$P, R \to I \to Q$
RADIUS	Radius r	$A \to r \to C$
SIDE_LENGTH	Side length s	$V \to s \to$ SA
WAVELENGTH	Wavelength λ	$v_w, f \to \lambda \to d$
CROSS_SECTION	Cross-sec. area A	$r, v \to A \to$ vol
DISPLACEMENT	Displacement x	$k, F \to x \to E_p$
MARKET_CAP	Market cap M	price, $N_s \to M \to$ P/E

Table 1: Eight representative experimental domains, bridge variables, and symbolic formula chains out of 33 total. See Appendix 3.1 for the complete list.

the span-assignment rule described below. Because interventions are only considered for positions $p < p_k$, all recovered edges point forward in the generated trace, making G^{patch} a DAG. This graph should therefore be read as a causal graph over token-level information flow in the written trace, not as a full circuit-level structural model of the transformer's internal computation.

3.2 Computation Graph Recovery

Our pipeline consists of five stages: dataset generation, CoT trace collection, counterfactual pair construction, token patching, and graph construction. As formalized above, token patching is a causal intervention: given two forward passes that produce different numerical outputs, we substitute the token at a single position p from one run into the other and measure the resulting change in output log-probability, identifying which earlier tokens causally influence the downstream value. Figure 1 provides an overview.

Stage 1: Dataset generation. We design 33 problem families spanning mechanics, electromagnetism, optics, thermodynamics, geometry, and finance (Table 1; full list in Appendix 3.1). Each family has a "bridge" variable that must be computed as an intermediate step – e.g., in VELOCITY, kinetic energy and mass are given; velocity is the bridge; travel time is the target.

For each family we provide five surface-form templates, sample independent variables uniformly from a set range, and compute derived quantities exactly (n=10 examples per template). The bridge is sampled independently of other inputs to prevent spurious correlations.

Stage 2: CoT trace generation. We generate greedy-decoded (T=0) completions from each model. Traces with incorrect final answers are filtered out and removed from the pool.

Stage 3: Counterfactual pair construction. Causal patching requires counterfactual traces from which we extract tokens to patch. For these traces we are interested in the numerical variables, so we generate counterfactual traces that differ only in numerical variables. We also enforce that counterfactual traces have the same number of tokens as the original trace – pairs that violate this rule are discarded.

Counterfactual sampling. For each base trace, we independently resample all numerical input variables from the same distributions used in Stage 1. Derived variables (including the bridge) are recomputed exactly. To ensure counterfactual traces have an identical reasoning structure, we prompt an API model (using the same problem template) to rewrite the base trace's CoT with the new numbers, preserving all non-numerical wording and all reasoning steps verbatim.

Numeric span extraction. After alignment, we extract all numeric spans in both the source (counterfactual) and base (original) texts using a regular-expression pattern and map these values to corresponding tokens. For each matched value at position k in the trace, we record the span – the tokens that are a part of the value. These are used to identify the *scoring position* for patching: the first token index where source and base diverge within the span for value v_k.

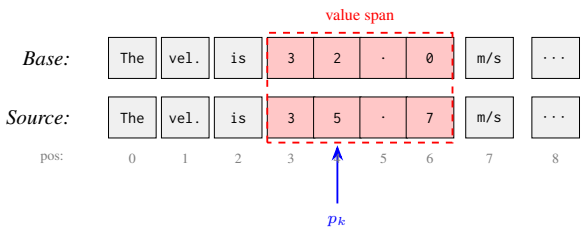

Figure 2: Token-level alignment for an example. Grey tokens are identical in both traces; red tokens encode different numerical values (`32.0` vs. `35.7`). p_k=4 is the first diverging token within the value span and serves as the scoring target for token patching.

Stage 4: Token patching. For each aligned pair and each intermediate value v_k (processed in reverse order, final answer first), we run the following procedure.

Let $\mathbf{s} = (s_1, \ldots, s_N)$ and $\mathbf{b} = (b_1, \ldots, b_N)$ denote the source (counterfactual) and base (original) token sequences respectively, both of length N. As described prior, we define the *scoring position* p_k

as the first position where $s_p \neq b_p$ within the span for value v_k.

The intervention is performed at the *input token sequence*, before the embedding lookup: for each token position p with $p < p_k$, we create a patched sequence $\mathbf{b}^{(p)}$ that is identical to the base sequence except that $b_p^{(p)} = s_p$. We patch exactly one position at a time, never cumulatively, and leave all model parameters unchanged. We then run a full forward pass on $\mathbf{b}^{(p)}$ and record the change in log-probabilities:

$$\Delta_{\text{src}}(p) = \log p_{\mathbf{b}^{(p)}}(s_{p_k}) - \log p_{\mathbf{b}}(s_{p_k}) \quad (5)$$

$$\Delta_{\text{base}}(p) = \log p_{\mathbf{b}^{(p)}}(b_{p_k}) - \log p_{\mathbf{b}}(b_{p_k}) \quad (6)$$

$$\Delta_{\text{diff}}(p) = \Delta_{\text{src}}(p) - \Delta_{\text{base}}(p) \quad (7)$$

The probabilities are read from the logits at position $p_k - 1$, i.e., the next-token distribution used to predict the first diverging token of value v_k. The *difference delta* Δ_{diff} captures simultaneously the increase in source-token probability and decrease in base-token probability.

Stage 5: Graph construction.

Statistical token selection. Among the token scores, we apply Benjamini–Hochberg (BH) FDR control (Benjamini and Hochberg, 1995) on robust z-scores computed from the median absolute deviation (MAD).

$$z_p = \frac{|\Delta_{\text{diff}}(p)| - \text{median}(\Delta_{\text{diff}})}{1.4826 \cdot \text{MAD}(\Delta_{\text{diff}})} \quad (8)$$

Two-sided normal p-values $p_p = \text{erfc}(z_p/\sqrt{2})$ are computed and the BH procedure at a particular q selects a set $\mathcal{P}_k$ of statistically significant token positions. A minimum of one token is always retained (falling back to the highest-scoring position).

Edge inference. For a child value v_k, we identify all tokens selected by FDR control and identify the tokens' corresponding values v_j as parents. Edge weight $w(v_j, v_k)$ is determined by taking the maximum causal score of all tokens in the span $\mathcal{P}_j$ of the parent value v_j.

$$w(v_j, v_k) = \max_{p \in \mathcal{P}_j} \Delta_{\text{diff}}(p) \quad (9)$$

Comparison: Frontier model graph annotation. Since the generated DAG is created via intervention, we treat it as an intervention-based estimate of how information is transferred through the CoT. We require annotation of a reference graph to compare the DAG to. To obtain a per-trace reference graph G^{annotate}, we prompt GPT-4o with the model's full CoT trace and ask it to label, for each numerical value $v_k \in V$, which other values in the trace were *directly* used in its computation. Critically, the annotator is constrained to draw edges only between values already identified in Stage 3's numeric span extraction, ensuring both graphs share the same node set of nodes V. The resulting annotated DAG G^{annotate} is compared to G^{patch} using edge precision, recall, F1, and transitive-closure agreement. Full prompt structure and output-parsing details are given in Appendix 2. Because the reference graphs are produced by an API model rather than human adjudicators, we interpret them as trace-conditioned reference annotations rather than gold labels. Appendix 4 reports a human annotation study evaluating the accuracy of LLM annotations.

3.3 Graph Faithfulness Interpretation

We interpret the comparison between G^{patch} and G^{annotate} as a direct measure of CoT faithfulness at the graph level. When the two graphs agree on an edge (v_i, v_j), it implies that the mechanistic causal dependency the model relies on is also explicitly expressed in the written CoT text – a faithful trace. When G^{patch} contains an edge absent from G^{annotate}, the model is mechanistically relying on a token that the trace does not identify as a direct input – a potential latent shortcut. Conversely, when G^{annotate} contains an edge absent from G^{patch}, the model writes an intermediate value but does not causally rely on it, suggesting redundant or decorative reasoning steps. This asymmetric analysis allows us to distinguish between models that reason faithfully and those that produce plausible-looking but mechanistically decoupled CoT.

4 Experiments

4.1 Dataset

We use the 33-family dataset described in §3.2 (Stage 1), yielding a total of $33 \times 5 \times 10 = 1{,}650$ base problems. After counterfactual rewriting, alignment, and numerical verification (Stages 2–3), the number of usable aligned pairs varies by model (see the "Pairs" column in Table 2).

4.2 Models

We evaluate our pipeline on a collection of open-weight base models spanning a range of scales: Qwen2.5-32B, Llama-3.1-70B, Gemma-4-31B,

Model	Pairs	Direct edges			Transitive closure		
		P	R	F1	P	R	F1
GPT-OSS-20B	377	0.533	0.790	0.617	0.846	0.795	0.797
Llama-3.1-70B	445	0.470	0.754	0.549	0.764	0.772	0.735
Mistral-Small-3.1-24B	821	0.422	0.763	0.529	0.692	0.763	0.694
Qwen2.5-32B	1254	0.491	0.576	0.529	0.836	0.552	0.659
Gemma-4-31B	1026	0.545	0.853	0.659	0.873	0.843	0.852

Table 2: Agreement between patching-derived graphs (G^{patch}) and frontier-model reference graphs (G^{annotate}) at the FDR threshold $q = 0.01$. "Pairs" = number of aligned trace pairs with both graph sources available. Direct-edge metrics compare immediate parent–child dependencies; transitive-closure metrics compare ancestor–descendant reachability. All metrics are micro-averaged over pairs.

Mistral-Small-3.1-24B, and GPT-OSS-20B. For frontier model graph annotation, we use GPT-4o.

4.3 Implementation Details

All patching experiments are training-free, use greedy decoding (temperature T=0), and require one patched forward pass per token (p) per scoring target p_k. Counterfactual rewrites are produced via the GPT-4o API; alignment and span extraction are deterministic post-processing steps applied before model evaluation. Additional implementation details and compute costs are reported in Appendix 1.

5 Results and Analysis

Agreement with Frontier-Model Reference Graphs. Table 2 reports end-to-end graph recovery results across models. For each model, we compare G^{patch} against the frontier-model reference graph G^{annotate} on the same node set. We report both direct-edge agreement and transitive-closure agreement. Direct-edge scores measure whether the two graphs identify the same immediate parent–child dependencies, while closure scores measure whether they agree on ancestor–descendant reachability even when the exact intermediate edge factorization differs.

Interpreting precision, recall, and F1. The direct-edge precision–recall split is the central signal in Table 2. High recall means that when the reference graph says a written value is an immediate input to a later value, token patching usually finds a causal dependency on some token in that parent span. In this sense, the models are not merely decorating their answers with inert intermediate numbers: most annotated reasoning steps are causally load-bearing. On the other hand, high precision indicates that most of the token patching identified causal parents align with the annotated graph. F1 summarizes this tension: it is highest when the recovered graph both covers the written computation and avoids adding many extra mechanistic dependencies.

Across all five models, direct-edge recall exceeds direct-edge precision, indicating that the dominant failure mode is over-inclusion rather than omission. This pattern is important for interpreting CoT faithfulness. A low-recall model would suggest post-hoc reasoning, where the model writes intermediate values but does not causally rely on them. Instead, we observe relatively high recall and moderate precision: the written values usually matter, but the mechanistic computation is less sparse than the textual derivation. Thus the CoT is faithful in the weaker sense that its stated intermediates are used, but not faithful in the stronger sense that the written graph fully specifies all causal parents.

Model-level differences. Gemma-4-31B shows the strongest overall agreement, with the highest direct-edge F1 (0.659), direct-edge recall (0.853), and closure F1 (0.852). This suggests that its traces expose a large fraction of the dependencies actually used by the model, while its additional patching edges mostly preserve the same ancestor–descendant structure as the reference graph. GPT-OSS-20B has the next-best direct-edge F1 (0.617) and similarly strong closure agreement (0.797), suggesting a comparable but slightly noisier alignment between written and mechanistic computation. Llama-3.1-70B and Mistral-Small-3.1-24B have similar recall to GPT-OSS-20B but lower precision, which points to a denser recovered graph: they recover many written dependencies, but patching also identifies more extra parents per target. Qwen2.5-32B is qualitatively different. It has relatively high precision but the lowest direct-edge recall (0.576) and closure recall (0.552), suggesting

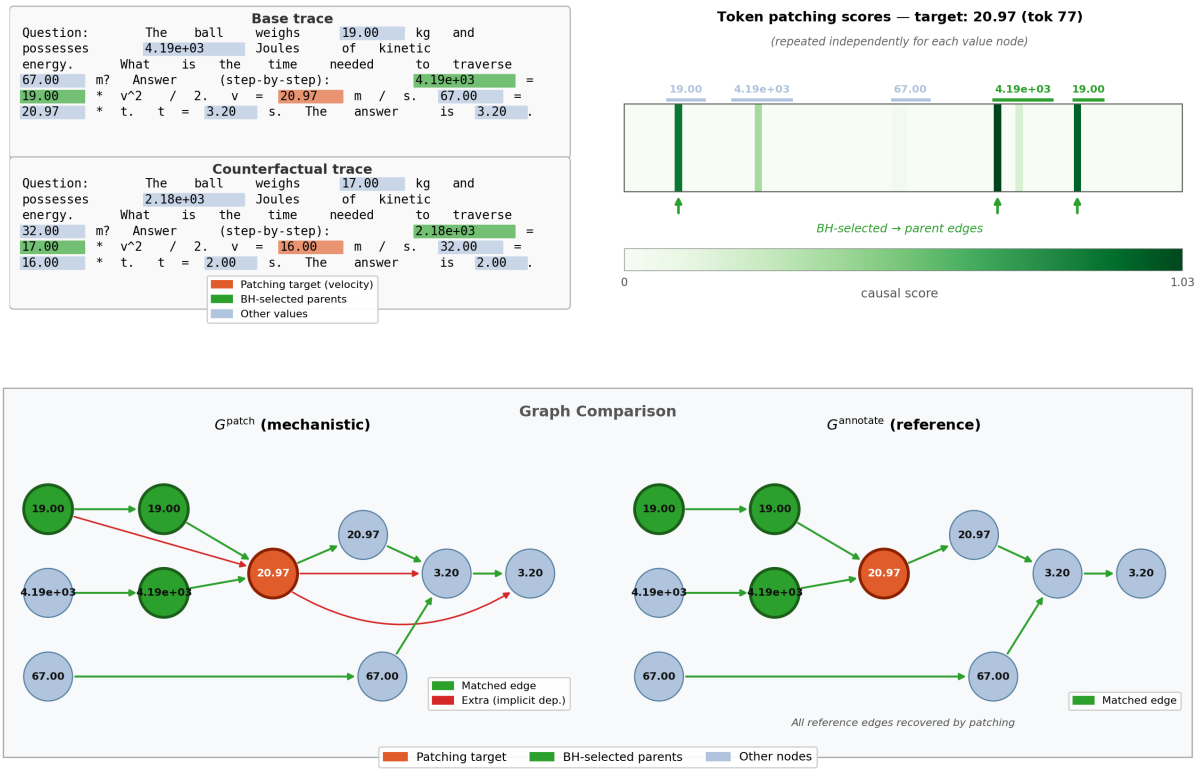

Figure 3: **Case study: Qwen2.5-32B, VELOCITY_FROM_KE, pair 20 (FDR q=0.05).** Values in the traces are coloured by role: orange = patching target (v), green = BH-selected parents, grey-blue = other values. **(A)** Base and counterfactual traces with numerical spans highlighted. **(B)** Per-position causal score Δ_{diff} for target v (20.97 m/s, truncation token 77). Each column is one token position; colour encodes score magnitude (darker green = higher causal influence). Brackets above mark the token spans of labelled values; green arrows below indicate the three BH-selected tokens (positions 7, 51, 62) at spans of m (CoT) and KE (CoT), which become parent edges. **(C)** Recovered G^{patch}: green edges appear in both graphs; red edges are extra implicit dependencies recovered by patching but absent from the textual reference. **(D)** Reference graph G^{annotate} from GPT-4o: green edges are shared; purple edges are reference-only (present in text but not recovered by patching).

that when Qwen edges are recovered they are often reference-consistent, but many annotated written dependencies do not appear as strong token-level causal parents under our intervention. One interpretation is that Qwen's successful traces rely more on recomputation from surrounding context or on a smaller subset of written values, rather than consistently reusing every explicit intermediate value.

Direct edges versus reachability. The gap between direct-edge and transitive-closure metrics is also informative. Direct-edge F1 ranges from 0.53 to 0.66, while closure F1 ranges from 0.66 to 0.85. This means many disagreements are not about whether the right information reaches the right downstream quantity, but about *where* the graph places the immediate edge. For example, the reference graph may treat $a \to b \to c$ as two direct symbolic steps, while token patching may find a direct sensitivity from a to c because the model attends back to the earlier copied value when generating c. Closure agreement counts both graphs as agreeing that a is an ancestor of c. The consistently higher closure F1 therefore suggests that the recovered graphs usually preserve the broad causal ordering of the computation even when they disagree with the annotator about the finest-grained edge factorization.

Relation to activation patching. In this section, we also perform activation patching as a complementary intervention to token patching, allowing us to test whether the graph-construction signals are consistent across intervention types. Appendix 5 compares both methods on 842 sampled pair/value records. Although token-position scores are only moderately correlated, value-level parent strengths are highly aligned: parent-value Pearson correlation is 0.938, top-3 parent Jaccard similarity is 0.908, and the strongest parent matches in 0.878 of

records. The gap between token-level and value-level agreement is expected: activation patching searches over hidden-state interventions across layers, while token patching changes the surface token supplied to the model. Thus the two methods need not rank the same individual token positions, especially within multi-token numbers or repeated value mentions. What matters for graph recovery is whether they identify the same *parent value*; on that abstraction, agreement is high. This supports our use of token patching as a scalable value-level intervention, while also clarifying its scope: it recovers which written CoT values are causally useful, not the internal circuit or layer at which the dependency is implemented.

Comparison to behavioural faithfulness measures. Behavioural studies test CoT faithfulness by truncating or corrupting reasoning and measuring the final accuracy drop (Lanham et al., 2023; Turpin et al., 2023). Our results are consistent with arithmetic traces being broadly useful, but add edge-level attribution: high recall shows that written intermediates are usually load-bearing, while imperfect precision shows that models also use implicit dependencies not named as direct textual inputs.

Alignment with recent faithfulness benchmarks. Our model-level pattern is directionally consistent with FaithCoT-Bench (Shen et al., 2026), an expert-annotated benchmark for instance-level CoT unfaithfulness. FaithCoT-Bench reports that GPT-4o-mini and Gemini-2.5-Flash produce the largest numbers of correct-faithful traces, while Llama-3.1-8B and Qwen2.5-7B exhibit higher unfaithfulness rates. We observe a similar ordering at the graph level: GPT-OSS-20B and Gemma-4-31B have the strongest agreement with annotated computation graphs, while Qwen2.5-32B and Llama-3.1-70B show weaker agreement, especially in direct-edge recall and F1. The match is not exact: FaithCoT-Bench studies different model sizes and versions, includes closed models, and labels instance-level faithfulness rather than value-level causal graphs. In particular, Gemma is not Gemini, although both are Google-developed model families. Still, the recurrence of a GPT/Gemini-or-Gemma versus Qwen/Llama contrast is suggestive. One possibility is that model-specific design and training choices – architecture details, pretraining mixtures, instruction tuning, or reasoning-style post-training – affect whether written CoT tokens are reused as computational state or merely accompany a computation carried by surrounding context. Our results contribute a complementary mechanistic signal: the same broad family-level contrast appears when faithfulness is measured by intervention-derived graph agreement rather than by human labels of entire reasoning traces.

6 Conclusion

We introduced a pipeline for automatically recovering the computation graph of a Chain-of-Thought trace, a directed acyclic graph whose edges are defined by counterfactual token patching rather than textual co-occurrence. Using 33 problem families spanning seven scientific and financial domains, we generated token-length-aligned counterfactual trace pairs, performed exhaustive token patching, applied Benjamini–Hochberg FDR control to select statistically significant causal token positions, and aggregated the results into per-trace DAGs. We compared the recovered graphs G^{patch} against per-trace annotated graphs G^{annotate} from GPT-4o, which provide a trace-conditioned reference that accounts for the model's actual algebraic path. Across all five models we consistently observe that recall substantially exceeds precision for direct edges: the patching-derived graphs recover most dependencies that the annotator identifies, but also include additional edges absent from the textual trace. Strikingly, transitive-closure F1 is much higher than direct-edge F1 across all models, indicating that even where the exact intermediate edge factorisation differs, the recovered graphs preserve the broader causal ordering – the right ancestors reach the right descendants. Together, these results suggest that explicit CoT tokens are largely, but not exclusively, causally load-bearing: the models do read their own intermediate values, but the mechanistic computation occasionally takes shortcuts not visible in the written text. This mechanistic window into CoT faithfulness provides a foundation for understanding and improving the reliability of chain-of-thought reasoning in large language models.

7 Limitations

Our analysis is restricted to templated numerical reasoning problems with explicitly identifiable scalar values. Real-world CoT traces often contain symbolic, relational, or natural-language subgoals whose boundaries are not captured by

regular-expression numeric span extraction; it is therefore unclear whether the precision–recall patterns we observe will generalise to open-ended reasoning (Wei et al., 2022; Yao et al., 2023). Even within numerical domains, the method assumes that the values relevant to the computation are written as surface-form spans that can be matched across a base and counterfactual trace. If a model performs a computation implicitly, uses verbal quantities rather than numerals, rounds aggressively, or changes notation between equivalent expressions, our span extractor can miss nodes or merge distinct quantities. The resulting graph should therefore be understood as a graph over observed numerical tokens in a trace, not a complete graph over every latent variable used internally by the model.

Our intervention is also deliberately coarse-grained. Counterfactual token patching cannot distinguish whether a causal effect is mediated by a specific internal circuit (Elhage et al., 2021), by lexical replacement, by attention to a copied value, or by downstream recomputation from surrounding context. Accordingly, G^{patch} should be interpreted as a causal graph over token-level information flow, not as a complete circuit-level explanation. Layerwise activation patching yields highly consistent graph-level results (Appendix 5), but the moderate token-level agreement between the two methods indicates that token patching is best suited for recovering value-level dependencies rather than localising fine-grained mechanisms. The single-token intervention design may also understate dependencies that only become visible when multiple tokens in a multi-token number are patched jointly.

Finally, our reference graphs are produced by GPT-4o rather than by full human adjudication (Turpin et al., 2023). While the human agreement study (Appendix 4) shows high fidelity (direct-edge F1 = 0.885), API-generated annotations may still miss implicit dependencies, over-regularise traces into canonical formula structures, or prefer a direct edge where a human would treat an intermediate textual value as load-bearing. For that reason, precision and recall against G^{annotate} should not be read as absolute truth measures. They are best interpreted as agreement with a trace-conditioned reference graph that is reliable enough to support aggregate comparisons, while individual edge disagreements may require manual inspection. The FDR threshold is also an engineering choice. Appendix 6 presents a sensitivity analysis across $q \in \{0.01, 0.02, 0.05, 0.1, 0.2, 0.3, 0.5\}$.

8 Ethical Considerations

This work aims to improve scientific understanding of when Chain-of-Thought traces reflect the computations that produce an answer. Our experiments use synthetic templated math, physics, and finance problems, and do not involve personal, private, or human-subject data. Traces are generated by open-weight models and annotated with an API model, so privacy risks are limited to standard handling of API requests, experiment logs, and generated outputs.

There are dual-use risks: methods for identifying causally important tokens can support auditing and debugging, but could also help reveal brittle reasoning patterns. We therefore frame the method as a tool for reliability evaluation, mechanistic diagnosis, and red-teaming, and report only aggregate graph-recovery methodology. Since our benchmark is synthetic and numerical, strong graph agreement should not be interpreted as evidence of broad safety, truthfulness, or rationale faithfulness in high-stakes domains.

The study also has compute costs. Although token patching is cheaper than exhaustive activation patching, the experiments required tens of H100-hours. We reduce this cost through training-free interventions, deterministic decoding, reusable aligned pairs, and CPU-bound graph construction after patching, and we report compute requirements to support future comparisons.

References

Emmanuel Ameisen, Jack Lindsey, Adam Pearce, Wes Gurnee, Nicholas L. Turner, Brian Chen, Craig Citro, David Abrahams, Shan Carter, Basil Hosmer, Jonathan Marcus, Michael Sklar, Adly Templeton, Trenton Bricken, Callum McDougall, Hoagy Cunningham, Thomas Henighan, Adam Jermyn, Andy Jones, Andrew Persic, Zhenyi Qi, T. Ben Thompson, Sam Zimmerman, Kelley Rivoire, Thomas Conerly, Chris Olah, and Joshua Batson. 2025. Circuit tracing: Revealing computational graphs in language models. *Transformer Circuits Thread*. 3

Yoav Benjamini and Yosef Hochberg. 1995. Controlling the false discovery rate: A practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society: Series B (Methodological)*, 57(1):289–300. 3, 5

Maciej Besta, Nils Blach, Ales Kubicek, Robert Gerstenberger, Michal Podstawski, Lukas Gianinazzi, Joanna Gajda, Tomasz Lehmann, Hubert Niewiadomski, Piotr Nyczyk, et al. 2024. Graph of thoughts:

Solving elaborate problems with large language models. In *Proceedings of the AAAI conference on artificial intelligence*, volume 38, pages 17682–17690. 2

Eden Biran, Daniela Gottesman, Sohee Yang, Mor Geva, and Amir Globerson. 2024. Hopping too late: Exploring the limitations of large language models on multi-hop queries. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 14113–14130. 3

Trenton Bricken, Adly Templeton, Joshua Batson, Brian Chen, Adam Jermyn, Tom Conerly, Nick Turner, Cem Anil, Carson Denison, Amanda Askell, Robert Lasenby, Yifan Wu, Shauna Kravec, Nicholas Schiefer, Tim Maxwell, Nicholas Joseph, Zac Hatfield-Dodds, Alex Tamkin, Karina Nguyen, Brayden McLean, Josiah E Burke, Tristan Hume, Shan Carter, Tom Henighan, and Christopher Olah. 2023. Towards monosemanticity: Decomposing language models with dictionary learning. *Transformer Circuits Thread*. 3

Martina Bronzini, Carlo Seganti, Gabriele Sarti, Lonneke van der Plas, Raffaella Bernardi, and Alberto Testoni. 2024. Unveiling factual recall behaviors of language models through knowledge neurons. *arXiv preprint arXiv:2408.05273*. 3

Yuntian Deng, Kiran Prasad, Roland Fernandez, Paul Smolensky, Vishrav Chaudhary, and Stuart Shieber. 2023. Implicit chain of thought reasoning via knowledge distillation. *arXiv preprint arXiv:2311.01460*. 2

Nelson Elhage, Neel Nanda, Catherine Olsson, Tom Henighan, Nicholas Joseph, Ben Mann, Amanda Askell, Yuntao Bai, Anna Chen, Tom Conerly, Nova DasSarma, Dawn Drain, Deep Ganguli, Zac Hatfield-Dodds, Danny Hernandez, Andy Jones, Jackson Kernion, Liane Lovitt, Kamal Ndousse, Dario Amodei, Tom Brown, Jack Clark, Jared Kaplan, Sam McCandlish, and Chris Olah. 2021. A mathematical framework for transformer circuits. *Transformer Circuits Thread*. 2, 9

Atticus Geiger, Hanson Lu, Thomas Icard, and Christopher Potts. 2021. Causal abstractions of neural networks. *Advances in neural information processing systems*, 34:9574–9586. 3

Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Peiyi Wang, Qihao Zhu, Runxin Xu, Ruoyu Zhang, Shirong Ma, Xiao Bi, et al. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*. 2

Peter Hase, Mohit Bansal, Been Kim, and Asma Ghandeharioun. 2023. Does localization inform editing? surprising differences in causality-based localization vs. knowledge editing in language models. *Advances in Neural Information Processing Systems*, 36:17643–17668. 3

Tamera Lanham, Anna Chen, Ansh Radhakrishnan, Benoit Steiner, Carson Denison, Danny Hernandez, Dustin Li, Esin Durmus, Evan Hubinger, Jackson Kernion, et al. 2023. Measuring faithfulness in chain-of-thought reasoning. *arXiv preprint arXiv:2307.13702*. 1, 2, 8

Amit Arnold Levy and Mor Geva. 2025. Language models encode numbers using digit representations in base 10. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 2: Short Papers)*, pages 385–395. 3

Kenneth Li, Aspen K Hopkins, David Bau, Fernanda Viégas, Hanspeter Pfister, and Martin Wattenberg. 2022. Emergent world representations: Exploring a sequence model trained on a synthetic task. *arXiv preprint arXiv:2210.13382*. 3

Zhaoyi Li, Gangwei Jiang, Hong Xie, Linqi Song, Defu Lian, and Ying Wei. 2024. Understanding and patching compositional reasoning in llms. In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 9668–9688. 3

Qing Lyu, Shreya Havaldar, Adam Stein, Li Zhang, Delip Rao, Eric Wong, Marianna Apidianaki, and Chris Callison-Burch. 2023. Faithful chain-of-thought reasoning. In *Proceedings of the 13th International Joint Conference on Natural Language Processing and the 3rd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 305–329. 2

Kevin Meng, David Bau, Alex Andonian, and Yonatan Belinkov. 2022. Locating and editing factual associations in gpt. *Advances in neural information processing systems*, 35:17359–17372. 3

Moritz Miller, Florent Draye, and Bernhard Schölkopf. 2026. Identifying intervenable and interpretable features via orthogonality regularization. *arXiv preprint arXiv:2602.04718*. 3

Neel Nanda, Lawrence Chan, Tom Lieberum, Jess Smith, and Jacob Steinhardt. 2023. Progress measures for grokking via mechanistic interpretability. *arXiv preprint arXiv:2301.05217*. 3

Debjit Paul, Robert West, Antoine Bosselut, and Boi Faltings. 2024. Making reasoning matter: Measuring and improving faithfulness of chain-of-thought reasoning. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 15012–15032. 2

Judea Pearl. 2000. *Causality: Models, Reasoning, and Inference*. Cambridge University Press. 3

Jacob Pfau, William Merrill, and Samuel R Bowman. 2024. Let's think dot by dot: Hidden computation in transformer language models. *arXiv preprint arXiv:2404.15758*. 2

Alethea Power, Yuri Burda, Harri Edwards, Igor Babuschkin, and Vedant Misra. 2022. Grokking: Generalization beyond overfitting on small algorithmic datasets. *arXiv preprint arXiv:2201.02177*. 3

Xu Shen, Song Wang, Zhen Tan, Laura Yao, Xinyu Zhao, Kaidi Xu, Xin Wang, and Tianlong Chen. 2026. Faithcot-bench: Benchmarking instance-level faithfulness of chain-of-thought reasoning. In *The Fourteenth International Conference on Learning Representations*. 8

Alessandro Stolfo, Yonatan Belinkov, and Mrinmaya Sachan. 2023. A mechanistic interpretation of arithmetic reasoning in language models using causal mediation analysis. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 7035–7052. 3

Adly Templeton, Tom Conerly, Jonathan Marcus, Jack Lindsey, Trenton Bricken, Brian Chen, Adam Pearce, Craig Citro, Emmanuel Ameisen, Andy Jones, Hoagy Cunningham, Nicholas L Turner, Callum McDougall, Monte MacDiarmid, C. Daniel Freeman, Theodore R. Sumers, Edward Rees, Joshua Batson, Adam Jermyn, Shan Carter, Chris Olah, and Tom Henighan. 2024. Scaling monosemanticity: Extracting interpretable features from claude 3 sonnet. *Transformer Circuits Thread*. 3

Ioannis Tsamardinos, Laura E Brown, and Constantin F Aliferis. 2006. The max-min hill-climbing bayesian network structure learning algorithm. *Machine learning*, 65(1):31–78. 3

Miles Turpin, Julian Michael, Ethan Perez, and Samuel Bowman. 2023. Language models don't always say what they think: Unfaithful explanations in chain-of-thought prompting. *Advances in Neural Information Processing Systems*, 36:74952–74965. 1, 2, 8, 9

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems*, 30. 3

Constantin Venhoff, Iván Arcuschin, Philip Torr, Arthur Conmy, and Neel Nanda. 2025. Base models know how to reason, thinking models learn when. *arXiv preprint arXiv:2510.07364*. 2

Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2022. Self-consistency improves chain of thought reasoning in language models. *arXiv preprint arXiv:2203.11171*. 2

Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837. 1, 2, 9

Sohee Yang, Elena Gribovskaya, Nora Kassner, Mor Geva, and Sebastian Riedel. 2024. Do large language models latently perform multi-hop reasoning? In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 10210–10229. 3

Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Tom Griffiths, Yuan Cao, and Karthik Narasimhan. 2023. Tree of thoughts: Deliberate problem solving with large language models. *Advances in neural information processing systems*, 36:11809–11822. 2, 9

Denny Zhou, Nathanael Schärli, Le Hou, Jason Wei, Nathan Scales, Xuezhi Wang, Dale Schuurmans, Claire Cui, Olivier Bousquet, Quoc Le, et al. 2022. Least-to-most prompting enables complex reasoning in large language models. *arXiv preprint arXiv:2205.10625*. 2

Fangwei Zhu, Damai Dai, and Zhifang Sui. 2025. Language models encode the value of numbers linearly. In *Proceedings of the 31st International Conference on Computational Linguistics*, pages 693–709, Abu Dhabi, UAE. Association for Computational Linguistics. 3

1 Implementation Details

All patching experiments are training-free and leave model parameters unchanged. For each target value v_k, token patching requires one patched forward pass for every earlier token position $p < p_k$. CoT traces are generated with greedy decoding (temperature T=0), and counterfactual rewrites are produced via the GPT-4o API. Alignment and numerical span extraction are deterministic post-processing steps applied before any model evaluation. Graph construction and span extraction are CPU-bound and negligible by comparison with model inference ($< 1\%$ of total runtime).

Compute budget. Experiments were distributed across clusters provided by Compute Canada and run on NVIDIA H100s. Token patching is the dominant cost because it requires one forward pass per token position per trace pair. Table 3 summarizes completed H100-hours by model. In total, patching across all five models consumed approximately 55 H100-hours, with full pipelines (trace generation, counterfactual rewriting, patching, and graph construction) totaling roughly 73 H100-hours.

Model	Patching H100-h	Total H100-h
GPT-OSS-20B	7.7	9.3
Llama-3.1-70B	11.0	22.0
Mistral-Small-3.1-24B	7.0	8.8
Qwen2.5-32B	20.9	22.1
Gemma-4-31B	8.2	10.3
Total	**54.8**	**72.5**

Table 3: Completed H100-hours by model. "Patching" counts only the token-patching stage; "Total" includes trace generation, counterfactual rewriting, patching, and graph construction.

2 Prompt for Graph Annotation

To establish ground-truth reference graphs G^{annotate}, we instructed frontier LLMs to label computational dependencies from CoT traces. We used the following prompt template:

```
You are an expert annotator
    building a DIRECTED ACYCLIC
    COMPUTATION
GRAPH over numeric values in a
    trace.
Each node is a labeled value (v0,
    v1, ...). Use ONLY provided
    nodes.

Edge semantics:
- source -> target means source
    value is directly used to
    compute target.
- Parents must occur earlier in
    token order than children.
- No self-loops, no cycles.
- The only nodes without parents
    should values in the prompt
  (Before "Answer (step-by-step):")
     --- all other values trace
    back
  to values in the prompt.
- The graph should be fully
    connected -- no isolated nodes.


Output STRICT JSON only:
{
  "root": "v0",
  "edges": [
    {"source": "vA", "target": "vB
    "}
  ]
}
PAIR_ID: {pair_id}
ALLOWED_NODE_IDS: {node_ids}

NODE_TABLE:
{node_table}

SOURCE_TRACE_WITH_LABELS:
{annotated_text}

SOURCE_TRACE_RAW:
{source_text}
```

Annotation is zero-shot with temperature 0. The system prompt instructs the model to return strict JSON. The user prompt supplies: (i) the list of allowed node IDs; (ii) a node table giving each value's text and its token/character span in the trace; and (iii) the full CoT trace with node labels inserted inline. Edge semantics are stated explicitly: `source` $\rightarrow$ `target` means the source value is directly used to compute the target; parents must precede children in token order; self-loops and cycles are forbidden; prompt-given values may be roots; and the output must be a JSON object with a `root` field and an `edges` list of `{"source": "vA", "target": "vB"}` entries. During post-processing, we discard malformed edges, edges referencing nodes outside V, self-loops, duplicate edges, and temporally invalid edges whose source does not precede its target in token order. If the API response is not parseable as JSON, the pair is marked as an annotation failure and excluded from graph-agreement evaluation.

3 Dataset Problem Families

The complete set of specific prompts used for generating the base traces across all 33 problem families is listed below in Section 3.2.

3.1 Summary of Bridge Variables

Table 4 lists all 33 problem families used in our experiments, grouped by bridge variable category.

3.2 Full List of Templates

`velocity_from_ke`

Moving mass with kinetic energy → velocity (hidden) → travel time

- A **{m}** kg **{obj}** has **{ke}** Joules of kinetic energy. How long does it take to travel **{d}** m?
- Given a **{m}** kg **{obj}** with **{ke}** Joules of kinetic energy, calculate the duration required to cover a distance of **{d}** m.
- The **{obj}** weighs **{m}** kg and possesses **{ke}** Joules of kinetic energy. What is the time needed to traverse **{d}** m?
- Mass: **{m}** kg, Kinetic Energy: **{ke}** Joules. Determine the time interval necessary for this **{obj}** to displace **{d}** m.
- Consider a **{m}** kg **{obj}** with **{ke}** Joules of kinetic energy. Find the number of seconds needed to move **{d}** m.

`velocity_from_ke_momentum`

KE to Momentum: "An object of mass m has Kinetic Energy K. What is its momentum?"

- An **{obj}** of mass **{m}** kg has Kinetic Energy **{ke}** Joules. What is its momentum in kg·m/s?
- Given a **{m}** kg **{obj}** with **{ke}** J of kinetic energy, calculate the momentum.
- The **{obj}** weighs **{m}** kg and has **{ke}** Joules of kinetic energy. Determine its momentum.
- Mass: **{m}** kg, KE: **{ke}** J. Find the momentum of this **{obj}** in kg·m/s.
- Consider a **{m}** kg **{obj}** possessing **{ke}** Joules of kinetic energy. What is its momentum?

`velocity_from_momentum`

Momentum to Displacement: "An object of mass m has momentum p. How far does it travel in time t?"

- An **{obj}** of mass **{m}** kg has momentum **{p}** kg·m/s. How far does it travel in **{t}** seconds?
- Given a **{m}** kg **{obj}** with momentum **{p}** kg·m/s, calculate the distance traveled in **{t}** seconds.
- The **{obj}** weighs **{m}** kg and has momentum **{p}** kg·m/s. Determine the distance covered in **{t}** s.
- Mass: **{m}** kg, Momentum: **{p}** kg·m/s. Find the displacement of this **{obj}** after **{t}** seconds.
- Consider a **{m}** kg **{obj}** with momentum **{p}** kg·m/s. How many meters does it move in **{t}** seconds?

`velocity_from_energy_conservation`

Height to Speed (Conservation): "A ball of mass m is dropped from height h. What is its speed just before impact?"

- A **{obj}** of mass **{m}** kg is dropped from height **{h}** meters. What is its speed in m/s just before impact?
- Given a **{m}** kg **{obj}** falling from **{h}** m, calculate the velocity just before it hits the ground.
- The **{obj}** weighs **{m}** kg and is dropped from **{h}** meters. Determine its impact speed in m/s.
- Mass: **{m}** kg, Height: **{h}** m. Find the final velocity of this **{obj}** just before landing.
- Consider a **{m}** kg **{obj}** dropped from **{h}** meters. What is its speed upon impact?

`current_from_power`

Resistor dissipating heat with power and resistance → current (hidden) → charge transferred

- A **{obj}** has **{r}** ohms of resistance and dissipates **{p}** watts of power. How much charge flows through it after **{t}** seconds?

Category	Domain key	Bridge	Formula chain
Velocity	VELOCITY_FROM_KE	v	$T, m \to v \to t$
	VELOCITY_FROM_KE_MOMENTUM	v	$T, m \to v \to p$
	VELOCITY_FROM_MOMENTUM	v	$p, m \to v \to \text{KE}$
	VELOCITY_FROM_ENERGY_CONSERVATION	v	$E, m \to v \to t$
	VELOCITY_FROM_DISTANCE_TIME	v	$d, t \to v \to t_2$
Current	CURRENT_FROM_POWER	I	$P, R \to I \to Q$
	CURRENT_FROM_VOLTAGE	I	$V, R \to I \to P$
	CURRENT_FROM_CHARGE_TIME	I	$q, t \to I \to q_2$
Acceleration	ACCELERATION_FORCE_TO_DISTANCE	a	$F, m \to a \to d$
	ACCELERATION_VELOCITY_TO_FORCE	a	$\Delta v, t \to a \to F$
	ACCELERATION_SPRING	a	$k, x, m \to a \to F$
	ACCELERATION_FROM_VELOCITY_CHANGE	a	$u, v, t \to a \to F$
Radius	RADIUS_FROM_AREA	r	$A \to r \to C$
	RADIUS_FROM_SPHERE_VOLUME	r	$V \to r \to A$
	RADIUS_FROM_CYLINDER_VOLUME	r	$V, h \to r \to A$
	RADIUS_FROM_CIRCUMFERENCE	r	$C \to r \to A$
Side length	SIDE_LENGTH_FROM_VOLUME	s	$V \to s \to \text{SA}$
	SIDE_LENGTH_FROM_AREA	s	$A \to s \to V$
	SIDE_LENGTH_FROM_PERIMETER	s	$P \to s \to A$
Wavelength	WAVELENGTH_FROM_SPEED	λ	$v_w, f \to \lambda \to d$
Cross-section	CROSS_SECTION_FROM_FLOW	A	$r, v \to A \to \text{vol}$
Displacement	DISPLACEMENT_FROM_SPRING	x	$k, F \to x \to E_p$
Market cap	MARKET_CAP_FROM_SHARES	M	price, $N_s \to M \to$ P/E
Moles	MOLES_IDEAL_GAS_TO_MASS	n	$P, V, T \to n \to m$
	MOLES_MASS_TO_PRESSURE	n	$m, M_r \to n \to P$
Energy	ENERGY_POWER_TO_HEIGHT	E	$P, t \to E \to h$
	ENERGY_FRICTION_TO_HEAT	W	$f, d \to W \to \Delta T$
Force	FORCE_PRESSURE_TO_ACCELERATION	F	$P, A \to F \to a$
	FORCE_FROM_MASS_ACCELERATION	F	$m, a \to F \to J$
Frequency	FREQUENCY_PENDULUM_TO_GRAVITY	f	$L, T \to f \to g$
	FREQUENCY_FROM_PERIOD	f	$T \to f \to N$
Angle	ANGLE_RAMP_TO_ACCELERATION	$\sin\theta$	$h, L \to \sin\theta \to a$
	ANGLE_COMPONENTS_TO_MAGNITUDE	F_{tot}	$F_x, F_y \to F_{\text{tot}} \to a$

Table 4: All 33 problem families grouped by bridge variable category. Each family has 5 surface-form templates × 10 examples = 50 instances. Formula chains use the notation inputs → bridge → target.

- Given a **{obj}** with **{r}** ohms of resistance and **{p}** watts of power output, calculate the total charge in Coulombs that passes through over **{t}** seconds.
- The **{obj}** dissipates **{p}** watts across a resistance of **{r}** ohms. Determine the magnitude of charge transferred during a **{t}** second interval.
- Resistance: **{r}** ohms, Power: **{p}** watts. Find the net charge flow accumulated in this **{obj}** after **{t}** seconds.
- Consider a **{obj}** rated at **{r}** ohms and **{p}** watts. How many Coulombs of charge travel through this component in **{t}** seconds?

`current_from_voltage`

Ohm's Law to Power: "A resistor R is connected to a voltage V. How much power is dissipated?"

- A **{obj}** with **{r}** ohms of resistance is connected to a voltage of **{v}** volts. How much power is dissipated in watts?
- Given a **{r}** ohm **{obj}** connected to **{v}** V, calculate the power dissipation.
- The **{obj}** has **{r}** ohms and is connected to **{v}** volts. Determine the power dissipated in watts.
- Resistance: **{r}** Ω, Voltage: **{v}** V. Find the power dissipated by this **{obj}**.
- Consider a **{r}** ohm **{obj}** with **{v}** volts applied. What is the power dissipation in watts?

`acceleration_force_to_distance`

Force to Distance: "A block of mass m is pushed from rest by a net force F. How far does it travel in t seconds?"

- A **{obj}** of mass **{m}** kg is pushed from rest by a net force of **{f}** Newtons. How far does it travel in **{t}** seconds?
- Given a **{m}** kg **{obj}** accelerated by **{f}** N from rest, calculate the distance traveled in **{t}** seconds.
- The **{obj}** weighs **{m}** kg and experiences a force of **{f}** N starting from rest. Determine the displacement after **{t}** seconds.
- Mass: **{m}** kg, Force: **{f}** N, Initial velocity: 0. Find the distance this **{obj}** travels in **{t}** seconds.
- Consider a **{m}** kg **{obj}** pushed by **{f}** Newtons from rest. How many meters does it move in **{t}** seconds?

`acceleration_velocity_to_force`

Velocity to Force: "An object accelerates from rest to velocity v over distance d. If mass is m, what was the net force?"

- An **{obj}** of mass **{m}** kg accelerates from rest to **{v}** m/s over a distance of **{d}** meters. What was the net force applied in Newtons?
- Given a **{m}** kg **{obj}** that reaches **{v}** m/s from rest after traveling **{d}** m, calculate the net force.
- The **{obj}** weighs **{m}** kg and accelerates from 0 to **{v}** m/s over **{d}** meters. Determine the applied force in N.
- Mass: **{m}** kg, Final velocity: **{v}** m/s, Distance: **{d}** m. Find the net force on this **{obj}** starting from rest.
- Consider a **{m}** kg **{obj}** that accelerates to **{v}** m/s over **{d}** meters from rest. What force was applied?

`acceleration_spring`

Spring to Acceleration: "A spring with constant k compressed by x launches a ball of mass m. What is initial acceleration?"

- A spring with constant **{k}** N/m is compressed by **{x}** meters. If it launches a **{obj}** of mass **{m}** kg, what is the initial acceleration in m/s^2?
- Given a spring (k = **{k}** N/m) compressed **{x}** m that releases a **{m}** kg **{obj}**, calculate the initial acceleration.
- The spring has stiffness **{k}** N/m and is compressed **{x}** meters. It launches a **{m}** kg **{obj}**. Determine the initial acceleration.
- Spring constant: **{k}** N/m, Compression: **{x}** m, Mass: **{m}** kg. Find the initial acceleration of this **{obj}**.
- Consider a spring (k = **{k}** N/m) compressed **{x}** m launching a **{m}** kg **{obj}**. What is the initial acceleration?

`radius_from_area`

Circle area → radius (hidden) → circumference

- A **{shape}** has an area of **{area}** square meters. What is its circumference in meters?
- Given a **{shape}** with a surface area of **{area}** square meters, calculate the perimeter.
- The **{shape}** covers **{area}** square meters of area. Determine the length around its edge in meters.
- Area: **{area}** square meters. Find the distance around the boundary of this **{shape}**.
- Consider a **{shape}** occupying **{area}** square meters. How many meters is the path around it?

`radius_from_sphere_volume`

Volume to Surface Area: "A sphere has volume V. What is its surface area?"

- A **{obj}** has a volume of **{vol}** cubic meters. What is its surface area in square meters?
- Given a **{obj}** with volume **{vol}** m^3, calculate the surface area.
- The **{obj}** has a volume of **{vol}** cubic meters. Determine its surface area in m^2.
- Volume: **{vol}** m^3. Find the surface area of this **{obj}**.
- Consider a **{obj}** with **{vol}** cubic meters of volume. What is its surface area?

`radius_from_cylinder_volume`

Cylinder Volume to Diameter: "A cylinder with height h has volume V. What is its diameter?"

- A **{obj}** with height **{h}** cm has a volume of **{vol}** cubic centimeters. What is its diameter in cm?
- Given a **{obj}** of height **{h}** cm and volume **{vol}** cm^3, calculate the diameter.
- The **{obj}** is **{h}** cm tall and has a volume of **{vol}** cubic cm. Determine its diameter.
- Height: **{h}** cm, Volume: **{vol}** cm^3. Find the diameter of this **{obj}**.
- Consider a **{obj}** with height **{h}** cm and volume **{vol}** cm^3. What is the diameter?

`side_length_from_volume`

Cube volume → side length (hidden) → surface area

- A **{obj}** has a volume of **{vol}** cubic centimeters. What is its total surface area in square centimeters?
- Given a **{obj}** that contains **{vol}** cubic centimeters of space, calculate the surface area.
- The **{obj}** has an internal capacity of **{vol}** cubic centimeters. Determine its exterior surface area.
- Volume: **{vol}** cubic centimeters. Find the total area of all faces of this **{obj}** in square centimeters.
- Consider a **{obj}** with **{vol}** cubic centimeters of volume. How much surface area does it have?

`side_length_from_area`

Area to Diagonal: "A square has area A. What is the length of its diagonal?"

- A **{obj}** has an area of **{area}** square centimeters. What is the length of its diagonal in cm?
- Given a **{obj}** with area **{area}** cm^2, calculate the diagonal length.
- The **{obj}** covers **{area}** square centimeters. Determine the diagonal length in cm.
- Area: **{area}** cm^2. Find the diagonal of this **{obj}**.
- Consider a **{obj}** with **{area}** cm^2 of area. What is its diagonal length?

`wavelength_from_speed`

Wave with speed and frequency → wavelength (hidden) → distance between n crests

- A **{wave}** travels at **{speed}** m/s with a frequency of **{freq}** Hz. What is the distance between **{n}** consecutive crests?
- Given a **{wave}** moving at **{speed}** m/s and oscillating at **{freq}** Hz, calculate the span between **{n}** adjacent wave peaks.
- The **{wave}** has a speed of **{speed}** m/s and frequency of **{freq}** Hz. Determine the separation of **{n}** successive crests in meters.
- Speed: **{speed}** m/s, Frequency: **{freq}** Hz. Find the distance covered by **{n}** complete wavelengths of this **{wave}**.
- Consider a **{wave}** at **{freq}** Hz traveling at **{speed}** m/s. How many meters separate the first crest from the **{n}**th crest?

`cross_section_from_flow`

Water through cylindrical pipe with radius and speed → cross-sectional area (hidden) → volume

- Water flows through a **{conduit}** with a radius of **{r}** cm at a speed of **{v}** cm/s. How much water (in cubic cm) passes through after **{t}** seconds?

- Given a **{conduit}** with radius **{r}** cm where water moves at **{v}** cm/s, calculate the volume of water in cubic centimeters that flows through for **{t}** seconds.
- The **{conduit}** has a **{r}** cm radius and water travels at **{v}** cm/s. Determine the total volume discharged in **{t}** seconds.
- Radius: **{r}** cm, Flow speed: **{v}** cm/s. Find the cubic centimeters of water transported through this **{conduit}** after **{t}** seconds.
- Consider a **{conduit}** of **{r}** cm radius carrying water at **{v}** cm/s. What volume in cubic cm flows through in **{t}** seconds?

`displacement_from_spring`

Spring with spring constant and force → displacement (hidden) → potential energy

- A **{spring}** with spring constant **{k}** N/m has **{f}** Newtons of force applied to it. What is the elastic potential energy stored in Joules?
- Given a **{spring}** with k = **{k}** N/m stretched by a force of **{f}** N, calculate the potential energy.
- The **{spring}** has stiffness **{k}** N/m and is under **{f}** N of force. Determine the stored elastic energy in Joules.
- Spring constant: **{k}** N/m, Applied force: **{f}** N. Find the potential energy stored in this **{spring}**.
- Consider a **{spring}** with spring constant **{k}** N/m pulled with **{f}** Newtons. How much elastic potential energy does it contain in Joules?

`market_cap_from_shares`

Company with share price and shares → market cap (hidden) → P/E ratio

- A **{company}** has a share price of $**{price}** and **{shares}** million shares outstanding, with annual net income of $**{income}** million. What is its P/E ratio?
- Given a **{company}** trading at $**{price}** per share with **{shares}** million shares and $**{income}** million in net income, calculate the price-to-earnings ratio.
- The **{company}** has **{shares}** million shares at $**{price}** each and annual earnings of $**{income}** million. Determine the P/E ratio.
- Share price: $**{price}**, Shares outstanding: **{shares}** million, Net income: $**{income}** million. Find the P/E ratio for this **{company}**.
- Consider a **{company}** with **{shares}** million shares worth $**{price}** each that earns $**{income}** million annually. What is the price-to-earnings ratio?

`moles_ideal_gas_to_mass`

Ideal Gas to Mass: "A tank of volume V holds O2 at pressure P and temperature T. What is the mass?"

- A tank of volume **{v}** liters holds **{gas_name}** (**{formula}**) at pressure **{p}** atm and temperature **{t}** K. What is the mass of the gas in grams?
- Given a **{v}** L container of **{gas_name}** (**{formula}**) at **{p}** atm and **{t}** K, calculate the mass in grams.
- The tank contains **{v}** liters of **{gas_name}** (**{formula}**) at **{p}** atmospheres and **{t}** Kelvin. Determine the mass.
- Volume: **{v}** L, Gas: **{gas_name}** (**{formula}**), Pressure: **{p}** atm, Temperature: **{t}** K. Find the mass in grams.
- Consider **{v}** liters of **{gas_name}** (**{formula}**) at **{p}** atm and **{t}** K. What is the mass of the gas?

`moles_mass_to_pressure`

Mass to Pressure: "x grams of Helium in container V at temperature T. What is the pressure?"

- **{mass}** grams of **{gas_name}** (**{formula}**) are placed in a container of volume **{v}** liters at temperature **{t}** K. What is the pressure in atm?
- Given **{mass}** g of **{gas_name}** (**{formula}**) in **{v}** L at **{t}** K, calculate the pressure.
- The container holds **{mass}** grams of **{gas_name}** (**{formula}**) with volume **{v}** liters and temperature **{t}** Kelvin. Determine the pressure.

- Mass: **{mass}** g, Gas: **{gas_name}** (**{formula}**), Volume: **{v}** L, Temperature: **{t}** K. Find the pressure in atmospheres.
- Consider **{mass}** grams of **{gas_name}** (**{formula}**) in **{v}** liters at **{t}** K. What is the pressure?

`energy_power_to_height`

Power to Climb: "A motor uses P watts to lift mass m. How high can it lift the mass in time t?"

- A motor uses **{p}** watts to lift a **{obj}** of **{m}** kg. How high in meters can it lift the **{obj}** in **{t}** seconds?
- Given a motor rated at **{p}** W lifting a **{m}** kg **{obj}**, calculate the height reached in **{t}** seconds.
- The motor delivers **{p}** watts of power to lift a **{m}** kg **{obj}**. Determine the height in **{t}** seconds.
- Power: **{p}** W, Mass: **{m}** kg, Time: **{t}** s. Find the height this **{obj}** can be lifted.
- Consider a motor using **{p}** watts to lift a **{m}** kg **{obj}**. What height is reached after **{t}** seconds?

`energy_friction_to_heat`

Friction to Heat: "Block mass m slides distance d against friction f. If all energy goes into heat (specific heat c),

- A **{material}** **{obj}** of mass **{m}** kg slides **{d}** meters against a friction force of **{f}** N. If all energy goes into heating the **{obj}** (specific heat **{c}** J/(kg·K)), how much does its temperature rise in Kelvin?
- Given a **{m}** kg **{material}** **{obj}** sliding **{d}** m with friction force **{f}** N, calculate the temperature rise. Specific heat: **{c}** J/(kg·K).
- The **{material}** **{obj}** (**{m}** kg) experiences **{f}** N of friction over **{d}** meters. With specific heat **{c}** J/(kg·K), determine the temperature increase.
- Mass: **{m}** kg, Material: **{material}**, Friction: **{f}** N, Distance: **{d}** m, Specific heat: **{c}** J/(kg·K). Find the temperature rise of this **{obj}**.
- Consider a **{m}** kg **{material}** **{obj}** sliding **{d}** m against **{f}** N friction. With c = **{c}** J/(kg·K), what is the temperature increase?

`force_pressure_to_acceleration`

Pressure to Acceleration: "A piston of area A is pushed by pressure P. If it moves mass m, what is acceleration?"

- A **{obj}** with cross-sectional area **{a}** cm^2 is pushed by a pressure of **{p}** Pa. If it moves a mass of **{m}** kg, what is the acceleration in m/s^2?
- Given a **{obj}** of area **{a}** cm^2 under pressure **{p}** Pa pushing mass **{m}** kg, calculate the acceleration.
- The **{obj}** has area **{a}** cm^2 and experiences pressure **{p}** Pa. It accelerates a **{m}** kg mass. Determine the acceleration.
- Area: **{a}** cm^2, Pressure: **{p}** Pa, Mass: **{m}** kg. Find the acceleration produced by this **{obj}**.
- Consider a **{obj}** with **{a}** cm^2 area at **{p}** Pa moving **{m}** kg. What is the acceleration?

`force_electric_field_to_acceleration`

Electric Field to Acceleration: "A particle of charge q and mass m is in electric field E. What is acceleration?"

- A **{particle}** (charge **{q_str}** C, mass **{m_str}** kg) is placed in an electric field of **{e}** N/C. What is its acceleration in m/s^2?
- Given a **{particle}** with q = **{q_str}** C and m = **{m_str}** kg in field E = **{e}** N/C, calculate the acceleration.
- The **{particle}** has charge **{q_str}** C, mass **{m_str}** kg, and is in an **{e}** N/C electric field. Determine the acceleration.
- Charge: **{q_str}** C, Mass: **{m_str}** kg, Field: **{e}** N/C. Find the acceleration of this **{particle}**.
- Consider a **{particle}** (q = **{q_str}** C, m = **{m_str}** kg) in E = **{e}** N/C. What is the acceleration?

`frequency_pendulum_to_gravity`

Pendulum to Gravity: "A pendulum of length L has period T. What is gravitational acceleration g?"

- A **{obj}** with length **{l}** meters has a period of **{t}** seconds. What is the gravitational acceleration g in m/s^2?
- Given a **{obj}** of length **{l}** m oscillating with period **{t}** s, calculate the gravitational acceleration.
- The **{obj}** is **{l}** meters long and completes one swing in **{t}** seconds. Determine g in m/s^2.
- Length: **{l}** m, Period: **{t}** s. Find the gravitational acceleration for this **{obj}**.
- Consider a **{obj}** with **{l}** m length and **{t}** s period. What is g?

`angle_ramp_to_acceleration`

Ramp to Acceleration: "A block slides down a frictionless ramp of height h and length L. What is acceleration?"

- A **{obj}** slides down a frictionless ramp with height **{h}** meters and length **{l}** meters. What is its acceleration down the ramp in m/s^2?
- Given a **{obj}** on a frictionless ramp (height **{h}** m, length **{l}** m), calculate the acceleration.
- The **{obj}** is on a ramp that rises **{h}** meters over **{l}** meters of length. Determine the acceleration down the slope.
- Height: **{h}** m, Length: **{l}** m. Find the acceleration of this **{obj}** sliding down the frictionless ramp.
- Consider a **{obj}** on a frictionless ramp with **{h}** m height and **{l}** m length. What is the acceleration?

`angle_components_to_magnitude`

Components to Magnitude: "A force has x-component Fx and y-component Fy. What acceleration does it give mass m?"

- A force vector has x-component **{fx}** N and y-component **{fy}** N. What acceleration does it give a mass of **{m}** kg in m/s^2?
- Given a force with Fx = **{fx}** N and Fy = **{fy}** N acting on **{m}** kg, calculate the magnitude of acceleration.
- The force has components **{fx}** N (x-direction) and **{fy}** N (y-direction). Determine the acceleration of a **{m}** kg mass.
- Fx: **{fx}** N, Fy: **{fy}** N, Mass: **{m}** kg. Find the magnitude of acceleration.
- Consider a force with components **{fx}** N and **{fy}** N on a **{m}** kg object. What is the acceleration magnitude?

`velocity_from_distance_time`

Distance and time imply velocity, which is then used to extrapolate a different segment of travel.

- A **{obj}** covers **{d}** meters in **{t}** seconds. If it keeps the same pace, how long will it take to cover another **{d2}** meters?
- Given that a **{obj}** goes **{d}** m in **{t}** s, estimate the time needed to travel an additional **{d2}** meters.
- The **{obj}** moves **{d}** meters in **{t}** seconds. How many seconds are required for it to go **{d2}** more meters at the same speed?
- Distance: **{d}** m, Time: **{t}** s. At this pace, what is the travel time for a further **{d2}** m for this **{obj}**?
- Consider a **{obj}** that covers **{d}** meters in **{t}** seconds. How long does it need to travel **{d2}** extra meters?

`current_from_charge_time`

Charge and time imply current, which is then used to extrapolate a later charge transfer.

- A **{obj}** transfers **{q}** coulombs in **{t}** seconds. If the current stays constant, how much charge moves in the next **{t2}** seconds?
- Given a **{obj}** carrying **{q}** C over **{t}** s, estimate the charge moved during an additional **{t2}** seconds.
- The **{obj}** moves **{q}** coulombs in **{t}** seconds. How much charge passes through it over another **{t2}** seconds?

- Charge: **{q}** C, Time: **{t}** s. At the same rate, find the charge flow in the next **{t2}** seconds for this **{obj}**.

- Consider a **{obj}** with **{q}** coulombs over **{t}** seconds. How much charge will flow in a further **{t2}** seconds?

`force_from_mass_acceleration`

Mass and acceleration imply force, which is then used to compute impulse over time.

- A **{obj}** of mass **{m}** kg accelerates at **{a}** m/s^2 for **{t}** seconds. How much momentum change does it accumulate?

- Given a **{m}** kg **{obj}** with acceleration **{a}** m/s^2 over **{t}** s, determine the total impulse delivered.

- The **{obj}** has mass **{m}** kg and acceleration **{a}** m/s^2. What impulse is delivered after **{t}** seconds?

- Mass: **{m}** kg, Acceleration: **{a}** m/s^2. Find the impulse gained by this **{obj}** over **{t}** seconds.

- Consider a **{m}** kg **{obj}** accelerating at **{a}** m/s^2. Calculate the impulse after **{t}** seconds.

`acceleration_from_velocity_change`

Velocity change over time implies acceleration, which is then used to compute force.

- A **{obj}** changes speed from **{u}** m/s to **{v}** m/s in **{t}** seconds. What acceleration did it experience?

- Given a **{m}** kg **{obj}** going from **{u}** m/s to **{v}** m/s in **{t}** s, determine its acceleration.

- The **{obj}** speeds up from **{u}** m/s to **{v}** m/s over **{t}** seconds. What is the acceleration of the motion?

- Initial speed: **{u}** m/s, Final speed: **{v}** m/s, Time: **{t}** s. Find the acceleration of this **{obj}**.

- Consider a **{m}** kg **{obj}** whose speed changes from **{u}** m/s to **{v}** m/s in **{t}** seconds. Calculate the acceleration.

`radius_from_circumference`

Circumference implies radius, which is then used to compute area.

- A circular track has circumference **{c}** meters. What area would a filled disk with that same boundary cover?

- Given a circle with circumference **{c}** m, calculate the area enclosed by the same ring.

- The circumference of a round table is **{c}** meters. Determine the tabletop area.

- Circumference: **{c}** m. Find the area enclosed by the circle.

- Consider a circle with circumference **{c}** meters. What is the area inside it?

`side_length_from_perimeter`

Perimeter implies side length, which is then used to compute area.

- A square fence has perimeter **{p}** meters. What area of ground does it enclose?

- Given a square with perimeter **{p}** m, determine the land area inside it.

- The square's perimeter is **{p}** meters. Calculate the enclosed area.

- Perimeter: **{p}** m. Find the area inside the square.

- Consider a square with perimeter **{p}** meters. How many square meters are inside?

`frequency_from_period`

Period implies frequency, which is then used to count oscillations over time.

- A metronome ticks once every **{period}** seconds. How many ticks will it make in **{t}** seconds?

- Given a **{obj}** period of **{period}** s, calculate how many cycles occur in **{t}** seconds.

- The **{obj}** repeats every **{period}** seconds. Determine the number of repetitions over **{t}** seconds.

- Period: **{period}** s. Find the total number of oscillations in **{t}** seconds for this **{obj}**.

- Consider a **{obj}** with period **{period}** seconds. How many complete cycles happen in **{t}** seconds?

4 Human–LLM Graph Agreement Study

We evaluated the reliability of the LLM-generated reference graphs by comparing them against a human annotation sample. We sampled 100 generated graphs, approximately balanced across the three main evaluated models. The annotator was shown the model trace, the numeric value nodes, and the LLM graph, and then wrote a complete human graph in the `ADDED_EDGES` section of the annotation file. Nine examples were left blank because the underlying LLM response was incoherent; these were discarded before scoring.

Agreement was computed by treating the LLM graph as the prediction and the human graph as the reference. We report direct-edge precision, recall, and F1, as well as the same metrics after transitive closure of both graphs. Closure agreement measures whether the two graphs imply the same ancestor–descendant dependencies even when they differ in intermediate edge placement.

Across the 91 retained annotations, the LLM graphs contained 757 direct edges and the human graphs contained 787 direct edges. Of these, 683 edges appeared in both graphs, 74 LLM edges were not present in the human graph, and 104 human edges were missing from the LLM graph. The resulting micro-averaged direct-edge F1 was 0.885; transitive-closure F1 was similar at 0.879. That the closure score does not improve over the direct score indicates that most disagreements are not merely due to harmless differences in graph granularity – when the two annotators disagree on an edge, the disagreement typically represents a genuine structural difference rather than one annotator inserting an intermediate node that the other skipped.

Graphs	Dir. P	Dir. R	Dir. F1	Clos. P	Clos. R	Clos. F1
91	0.902	0.868	0.885	0.922	0.839	0.879

Table 5: Human–LLM reference graph agreement across 91 annotated graphs. "Dir." = direct-edge metrics; "Clos." = transitive-closure metrics. All metrics are micro-averaged.

The high direct-edge F1 (0.885) indicates that the LLM-generated reference graphs are largely consistent with human judgment, as further discussed in Appendix 4, and the modest gap between direct and closure F1 (0.879 vs. 0.885) confirms that the disagreements are structurally meaningful rather than trivial edge-granularity differences. These results support treating G^{annotate} as a reliable trace-conditioned reference for the graph-agreement analyses reported in the main paper.

5 Token Patching vs. Activation Patching

To justify the use of token patching in place of the more expensive layerwise activation patching, we ran both methods on 842 randomly sampled pair/-value records. Layerwise activation scores were aggregated across layers with `max_abs` before comparison. We compare the resulting token-position score vectors and the derived parent-value causal-strength vectors using Pearson and Spearman correlation, top-10 token overlap, and the match rate of the highest-scoring parent value.

n	Tok. r (Pearson)	Top-10 overlap	Par. r (Pearson)	Top-3 par. Jaccard	Top-par. match
842	0.559	0.592	0.938	0.908	0.878

Table 6: Comparison of token patching and layerwise activation patching (`max_abs` aggregation) on 842 randomly sampled pair/value records. "Tok. r" = Pearson correlation of per-position causal scores; "Top-10 overlap" = fraction of the top-10 token positions shared; "Par. r" = Pearson correlation of parent-value causal strengths; "Top-3 par. Jaccard" = Jaccard similarity of the top-3 parent values; "Top-par. match" = fraction of records where both methods agree on the single strongest parent.

Token-level score vectors show moderate correlation with activation-patching scores (Pearson ≈ 0.56, Spearman ≈ 0.63), reflecting that the two methods operate at different granularities. However, when scores are aggregated to the value level – which is the level at which graph edges are decided – agreement is substantially higher: parent-value Pearson correlations exceed 0.93, top-3 parent Jaccard similarity is 0.91, and the strongest-parent match rate is 0.88. This indicates that token patching and activation patching identify the same dominant causal parents in the vast majority of cases, justifying the use of the cheaper token-patching intervention for full-scale graph recovery.

6 FDR Threshold Sensitivity

All results in the main paper use the BH FDR threshold $q{=}0.01$. To assess sensitivity to this choice, we re-ran graph construction at $q \in \{0.01, 0.02, 0.05, 0.1, 0.2, 0.3, 0.5\}$ and evaluated direct-edge and transitive-closure agreement

against the frontier-model reference graphs. Figure 4 plots direct-edge precision, recall, and F1 as a function of q for all five models.

The precision–recall trade-off is consistent across all five models: tightening q retains fewer edges, increasing precision at the cost of recall, while loosening q recovers more edges at the cost of including noisy ones. F1 is broadly stable between $q{=}0.01$ and $q{=}0.05$ for most models (variation < 3 percentage points), and declines more sharply beyond $q{=}0.1$ where recall gains can no longer compensate for precision loss. These results confirm that the main reported numbers are not particularly sensitive to the exact choice of FDR threshold within a reasonable range, and that $q{=}0.01$ represents a conservative operating point that favours precision.

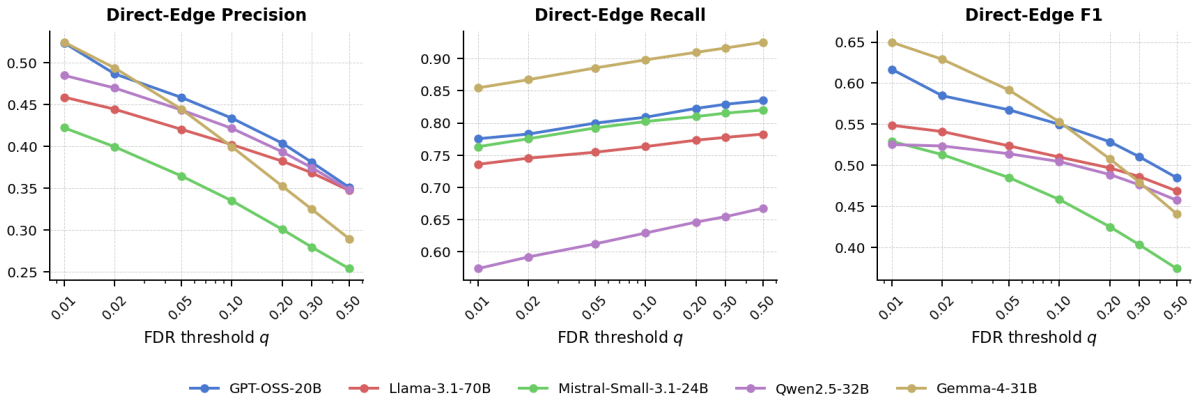

Figure 4: Direct-edge precision, recall, and F1 as a function of the BH FDR threshold q across all five models. As q increases, more token positions pass the significance test, so more edges are proposed: precision falls monotonically while recall rises. F1 peaks near q=0.01–0.02 for most models and declines at higher thresholds, confirming that q=0.01 (the value used throughout the paper) is a reasonable operating point. Gemma-4-31B (sand) achieves the highest F1 across all thresholds; Mistral-Small-3.1-24B (green) is the most sensitive to loosening the threshold.