# CAUSAL CONTRACTS: CONSTRAINT-VERIFIED DECODING FOR HALLUCINATION-RESISTANT LLMS

**Anonymous authors**
Paper under double-blind review

## ABSTRACT

Large language models generate fluent and confident text, but their training does not explicitly enforce causal validity. This limitation is critical in high-stakes domains, where residual hallucinations often involve incorrect causal or temporal relations rather than missing knowledge. We therefore define hallucination relative to a domain structural causal model (SCM). A claim is considered inconsistent when its associational, interventional, or counterfactual relation is not supported by the corresponding causal structure. We formulate generation as a Lagrangian-relaxed constrained optimization problem and introduce a projected primal-dual procedure. Under standard smoothness assumptions, we establish a non-asymptotic stationarity guarantee. We also derive a PAC-style bound on residual constraint violations as a function of verifier queries. Finally, we define an evaluation protocol spanning clinical QA, scientific explanation, and causal reasoning. The protocol combines causal-consistency verification, human annotation, inter-annotator agreement, and paired significance testing. Preliminary experiments on real LLM outputs demonstrate the feasibility of graph-based causal verification and identify causal claims that conventional factuality measures can miss.

## 1 INTRODUCTION

Large language models have achieved strong performance across open-ended natural language tasks. Their training objective, however, is likelihood maximization over token sequences conditioned on a prompt; nothing in that objective asks whether the sequence's implicit causal claims are true of the world. The resulting failure mode, fluent, confident, but unsupported or fabricated content, commonly termed *hallucination*, has been documented across natural language generation broadly Ji et al. (2023) and specifically for modern LLMs Huang et al. (2025); Alansari & Luqman (2025).

The consequences are not uniform across domains. In medicine, a 2025 cross-institution study evaluating eleven foundation models (seven general-purpose, four medical-specialized) across seven hallucination-sensitive tasks found that even domain-tuned models remained highly vulnerable, and that a clinician survey of 70 practitioners across multiple countries found 91.8% had personally encountered a medical hallucination and 84.7% judged such hallucinations capable of causing patient harm Kim et al. (2025). The same study reports a finding directly relevant to this paper's premise: physician audits attributed 64–72% of residual hallucinations to causal or temporal reasoning failures, not to missing factual knowledge, that is, the model often "knew" the relevant facts but connected them incorrectly. This is consistent with a broader body of work arguing that LLMs frequently reproduce causal-sounding language while relying on correlational shortcuts, a phenomenon termed *causal parroting* Zečević et al. (2023), and with causal-reasoning benchmarks such as CLadder showing systematic gaps between LLMs' surface fluency on causal questions and their accuracy against a ground-truth SCM Jin et al. (2023).

This motivates treating causal structure, not just token-level factuality, as a first-class object in hallucination mitigation. Purely statistical hallucination *detectors*, e.g., semantic-entropy-based uncertainty estimation Farquhar et al. (2024), sampling-based self-consistency checks Manakul et al. (2023), or self-contradiction analysis Mündler et al. (2023), are valuable and complementary, but they operate post hoc on model outputs and do not by themselves encode *why* a claim is wrong or what would make it right. Neuro-symbolic reasoning frameworks Badreddine et al. (2022); Olaus-



Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

son et al. (2023); Pan et al. (2023); Xu et al. (2024); Banerjee et al. (2025) show that coupling neural generation with symbolic verification measurably improves faithfulness on logical-reasoning benchmarks. But the constraints they enforce are typically propositional or first-order-logical rather than explicitly causal. Therefore, they cannot distinguish a valid correlation from a valid cause, nor flag a claimed intervention effect that is not identifiable from the available causal graph. We argue that hallucination in causally structured domains (clinical reasoning, mechanistic scientific explanation) is better modeled as a violation of a domain *structural causal model* than as an undifferentiated factuality gap. This lets us borrow the vocabulary of Pearl's causal hierarchy: association, intervention, counterfactual Pearl (2009), to say precisely which layer of a claim is unsupported. It also enables us to reuse existing machinery (d-separation tests, do-calculus identifiability) as differentiable-adjacent verifiers inside a constrained-decoding loop, in the spirit of recent constrained and grammar-guided decoding work Lu et al. (2021); Banerjee et al. (2025); Nakshatri et al. (2025).

**Contributions.** ❶ **Causal Formalization.** We define hallucination as a violation of a domain SCM's causal graph $\mathcal{G}_{\mathcal{M}}$, distinguished across the associational, interventional, and counterfactual rungs of the causal hierarchy, not one undifferentiated "factual error" label. ❷ **Constrained-Generation Objective.** We pose generation as Lagrangian-relaxed primal-dual optimization with an honest, non-asymptotic convergence analysis under explicit smoothness and boundedness assumptions, plus a PAC-style bound on residual constraint violation in place of an unjustified probability-one claim. ❸ **Related-Work Synthesis.** We connect hallucination detection, neuro-symbolic decoding, and causal LLM reasoning, exposing a shared gap: causal-graph-aware constraints enforced during generation itself. ❹ **Rigorous Evaluation Protocol.** We specify datasets, competitive baselines, human annotation with inter-annotator agreement, and statistical significance testing, making hallucination-reduction claims empirically falsifiable and reproducible.

## 2 Related Work

### 2.1 Hallucination: Taxonomy, Detection, and Mitigation

Ji et al. (2023) distinguish *intrinsic* hallucinations (contradicting the source) from *extrinsic* ones (unverifiable against it); Huang et al. (2025); Alansari & Luqman (2025) organize causes and mitigations across the data/training/inference stages. Detection has shifted toward meaning-level uncertainty: semantic entropy clusters sampled generations by bidirectional entailment to predict confabulation without task-specific supervision Farquhar et al. (2024); SelfCheckGPT exploits cross-sample consistency as a zero-resource signal Manakul et al. (2023); Mündler et al. (2023) target self-contradictory hallucinations. Retrieval augmentation remains among the most effective population-level mitigations Shuster et al. (2021); Self-RAG extends it with learned reflection tokens for retrieval necessity and self-critique Asai et al. (2024). Two inference-time interventions serve as baselines here: DoLa contrasts next-token logits across transformer layers, needing neither retrieval nor fine-tuning Chuang et al. (2024); Chain-of-Verification (CoVe) drafts an answer, generates and answers independent verification questions, then revises Dhuliawala et al. (2024), structurally similar to our verification loop but without causal-graph reasoning or constraint enforcement. None of these represent claims in a checkable graph-structured form, so none can distinguish a correlational slip from a causally invalid one or flag an unidentifiable interventional claim, the gap Sec. III targets. Medprompt's few-shot/self-generated-CoT/ensembling combination gives a strong domain-tuned, non-neuro-symbolic reference point for the clinical-QA comparison in Sec. V Nori et al. (2023). In clinical NLP, MedHallu benchmarks 10,000 PubMedQA-derived question–answer pairs with controlled hallucinated distractors, finding even strong general-purpose models struggle on "hard" hallucinations semantically close to the ground truth (best reported F1 on that category: 0.625) Pandit et al. (2025). Most directly relevant to our motivation, Kim et al.'s cross-institutional audit – combining quantitative benchmarking, physician-led qualitative evaluation, and a multinational clinician survey – finds that chain-of-thought and search-augmented generation reduce but do not eliminate hallucination, and attributes many *residual* failures to causal or temporal reasoning gaps rather than missing knowledge Kim et al. (2025), directly motivating our focus on causal-structure verification as a complement to generic factuality checking. These detection- and prompting-oriented lines are complementary to, not competitors of, the constrained-generation approach here: they can flag an untrustworthy output or supply a missing fact, but none encodes a checkable representation of *which* causal relationship was asserted or whether it holds in a reference causal graph.

### 2.2 Neuro-Symbolic and Constrained Generation

Logic Tensor Networks ground first-order logic in differentiable, fuzzy-valued predicates and connectives for gradient-based learning under logical constraints Badreddine et al. (2022). Logic-LM translates problems into formal symbolic representations, delegates deduction to an external solver, and uses solver feedback to repair translation errors Pan et al. (2023); LINC pairs LLM-based semantic parsing with a first-order-logic theorem prover for entailment tasks Olausson et al. (2023); SymbCoT instead runs the translate–plan–verify cycle inside the LLM via structured chain-of-thought, reducing dependence on rigid solver syntax Xu et al. (2024). At the decoding level, NeuroLogic enforces predicate-logic constraints via constrained beam search and lookahead heuristics Lu et al. (2021). CRANE shows that overly strict grammars can degrade reasoning by over-constraining intermediate steps, and addresses this by relaxing constraints during reasoning while enforcing them at answer boundaries, reporting up to 10-point accuracy gains over both unconstrained and strictly constrained decoding Banerjee et al. (2025); speculative-lookahead decoding separately reduces the per-step cost of constraint checking Nakshatri et al. (2025). Two lessons follow, which our approach builds on: checkable constraint satisfaction measurably improves faithfulness, but naive or overly rigid enforcement can *reduce* reasoning quality, so a constrained-decoding architecture must choose where and how strongly to enforce constraints. What this literature still lacks is a constraint vocabulary grounded explicitly in *causal* structure, one that distinguishes a valid association from a valid causal relationship, and an identifiable interventional claim from one that is not identifiable from the available structure, a gap especially consequential in high-stakes domains, consistent with the medical hallucination findings above Kim et al. (2025).

### 2.3 Causal Reasoning in and with LLMs

Kıcıman et al. argue LLMs can serve as useful, if imperfect, generators of causal arguments to assist domain experts Kıcıman et al. (2023). Zečević et al. counter that LLMs' apparent causal competence may largely reflect causal information memorized during training rather than genuine reasoning, a phenomenon they term *causal parroting* Zečević et al. (2023); the CLadder benchmark empirically substantiates this gap, evaluating LLMs on causal questions defined by ground-truth SCMs and verified via formal causal-inference procedures, and finding persistent gaps between causal-sounding fluency and formally correct conclusions Jin et al. (2023). Recent surveys catalog a growing range of responses – specialized prompting, external tool/solver integration, and hybrid neuro-symbolic systems for causal reasoning and discovery Xiong et al. (2024); Wu et al. (2024). This literature mostly asks whether an LLM can *discover* causal structure or *answer questions about* causal relationships; we instead assume a domain causal graph is already available – from clinical guidelines, mechanistic scientific knowledge, or a causal-discovery pipeline such as those surveyed – and ask how that knowledge can be incorporated as an *enforceable constraint* during generation. This targets, directly, the failure mode the causal-parroting literature and CLadder document: unsupported causal claims generated despite relevant causal knowledge being available.

**Positioning.** Table 1 positions our framework relative to these three literatures. Existing methods detect hallucinations, enforce logical constraints, or evaluate causal reasoning, but rarely combine these capabilities. Section III addresses this gap by integrating causal-graph and identifiability constraints directly into iterative decoding.

Table 1: Positioning of our framework relative to related research areas.

| Capability | Hallucination Detection | Neuro-Symbolic Decoding | LLM Causal Reasoning |
|---|---|---|---|
| Detect errors | ✓ | ∼ | ∼ |
| In-loop constraints | × | ✓ | × |
| Causal-graph aware | × | × | ✓ |
| Localize causal violations | × | × | ∼ |
| Enforce causal validity | × | × | × |
| **Our framework** | ✓ | ✓ | ✓ |

*Legend:* ✓ = supported; ∼ = partially supported; × = not explicitly supported.

# 3 CAUSAL FORMALIZATION AND CONSTRAINED GENERATION

## 3.1 DOMAIN STRUCTURAL CAUSAL MODEL

We assume access to a domain *structural causal model* $\mathcal{M} = \langle \mathbf{U}, \mathbf{V}, \mathbf{F}, P(\mathbf{U}) \rangle$ Pearl (2009), where $\mathbf{V} = \{V_1, \ldots, V_n\}$ are endogenous variables corresponding to domain entities, findings, or events (e.g., PNEUMONIA, FEVER, ANTIBIOTICTREATMENT); $\mathbf{U}$ are exogenous noise variables; and $\mathbf{F} = \{f_1, \ldots, f_n\}$ are structural equations $V_i \leftarrow f_i(\mathrm{Pa}(V_i), U_i)$. $\mathcal{M}$ induces a causal graph $\mathcal{G}_\mathcal{M} = (\mathbf{V}, E)$ with $(V_j, V_i) \in E$ iff $V_j \in \mathrm{Pa}(V_i)$. In clinical settings $\mathcal{G}_\mathcal{M}$ may be elicited from clinical practice guidelines or curated causal knowledge bases; in scientific-explanation settings it may come from a mechanistic ontology or a causal-discovery pipeline as surveyed in Xiong et al. (2024). We do not assume $\mathcal{M}$ is complete, only that $\mathcal{G}_\mathcal{M}$ encodes a set of causal claims the domain considers established, against which generated claims can be checked.

Pearl's causal hierarchy distinguishes three types of query supported by $\mathcal{M}$ Pearl (2009): *associational* ($P(V_i \mid V_j)$), *interventional* ($P(V_i \mid \mathrm{do}(V_j))$), and *counterfactual* ($P(V_i^{V_j \leftarrow v} \mid \mathbf{V} = \mathbf{v}')$). We use this distinction to classify, rather than flatten, the causal assertions found in generated text.

## 3.2 SEMANTIC PARSING OF CAUSAL ASSERTIONS

Let $\mathcal{X}$ be the input (query) space and $\mathcal{Y}$ the output (text) space, with an LLM $P_\theta(y \mid x)$. A semantic parser $\mathcal{S} : \mathcal{Y} \to \mathcal{Z}$ maps generated text $y$ to a finite set of typed assertions $z = \{a_1, \ldots, a_k\} = \mathcal{S}(y)$, where each assertion is one of:

| | |
|---|---|
| $a = (\text{ASSOC}, V_i, V_j)$ | "$V_i$ is associated with $V_j$" |
| $a = (\text{CAUSE}, V_i, V_j)$ | "$V_i$ causes/leads to $V_j$" |
| $a = (\text{DO-EFFECT}, V_i, V_j, \hat{p})$ | "intervening on $V_i$ sets $P(V_j \mid \mathrm{do}(V_i)) \approx \hat{p}$" |
| $a = (\text{INDEP}, V_i, V_j \mid \mathbf{S})$ | "$V_i$ is independent of $V_j$ given $\mathbf{S}$" |

In practice $\mathcal{S}$ is realized by grammar-constrained or schema-constrained decoding of a secondary extraction pass Banerjee et al. (2025); Nakshatri et al. (2025), which is a considerably better-conditioned sub-problem than open-ended generation and is where recent constrained-decoding advances are most directly applicable.

## 3.3 CAUSAL CONSTRAINT SET AND VIOLATION FUNCTIONS

Given $\mathcal{G}_\mathcal{M}$, define the derived constraint set $\mathcal{C}_\mathcal{M}$ by three checkable conditions:

**Causal well-formedness**: An assertion $a = (\text{CAUSE}, V_i, V_j)$ is *well-formed* with respect to $\mathcal{G}_\mathcal{M}$ if there exists a directed path $V_i \to \cdots \to V_j$ in $\mathcal{G}_\mathcal{M}$. An assertion $a = (\text{INDEP}, V_i, V_j \mid \mathbf{S})$ is well-formed if $V_i$ and $V_j$ are $d$-separated by $\mathbf{S}$ in $\mathcal{G}_\mathcal{M}$. An assertion $a = (\text{DO-EFFECT}, V_i, V_j, \hat{p})$ is well-formed if $P(V_j \mid \mathrm{do}(V_i))$ is *identifiable* in $\mathcal{G}_\mathcal{M}$ (checkable via the do-calculus rules or the ID algorithm Pearl (2009)) and $\hat{p}$ lies within tolerance of the value implied by $\mathbf{F}$ where $\mathbf{F}$ is known.

For each assertion type we define a bounded violation function $\phi : \mathcal{Z} \times \mathcal{G}_\mathcal{M} \to [0, 1]$. For a CAUSE assertion, letting $f_A, f_B \in [0, 1]$ be calibrated neural scorers for the presence of the source and target propositions in context and $\mathbb{1}[\text{path}]$ the graph-theoretic indicator of a directed path in $\mathcal{G}_\mathcal{M}$,

$$\phi_{\text{cause}}(a) = f_A \cdot f_B \cdot \big(1 - \mathbb{1}[V_i \rightsquigarrow V_j \text{ in } \mathcal{G}_\mathcal{M}]\big), \tag{1}$$

That is, violation mass accrues only when the model confidently asserts both propositions and no supporting causal path exists. We use a product-based soft conjunction, following the T-norm formulation in Badreddine et al. (2022). Importantly, the product formulation distinguishes between 'the model is unsure" and 'the model is confidently wrong." This distinction enables a more defensible choice of the decision threshold $\tau$. For INDEP assertions, we define

$$\phi_{\text{indep}}(a) = f_A f_B \cdot \mathbb{1}[V_i \not\perp_d V_j \mid \mathbf{S} \text{ in } \mathcal{G}_\mathcal{M}] \,.$$

For DO-EFFECT assertions, we define

$$\phi_{\text{do}}(a) = f_A \cdot (\mathbb{1}[\text{not identifiable}] + \mathbb{1}[\text{identifiable}] \cdot \mathbb{1}[|\hat{p} - p^\star| > \tau_p]) \,,$$

where $p^\star$ denotes the value entailed by $\mathbf{F}$ when that value is known. This formulation distinguishes two failure modes that a single "factual error" label would conflate. The first occurs when the model asserts a causal effect that is not identifiable from the available causal structure. The second occurs when the effect is identifiable, but the model reports an incorrect magnitude.

The aggregate violation for output $y$ is $\Phi(z) = \sum_{a \in z} w_{\mathrm{type}(a)}\, \phi_{\mathrm{type}(a)}(a)$, with type-specific weights $w_{\mathrm{type}}$ reflecting, e.g., that a fabricated DO-EFFECT claim in a treatment-recommendation context may warrant higher weight than an imprecise associational remark.

**Causal hallucination**: Given input $x$, output $y$, domain SCM $\mathcal{M}$, and tolerance $\tau \geq 0$, $y$ is a *causal hallucination* if $\Phi(\mathcal{S}(y)) > \tau$.

This definition is deliberately narrower than "any factual error": it targets errors in the relational/causal structure among claims, which is the category the clinical audit in Kim et al. (2025) identifies as dominant among residual errors after standard mitigations. It is meant to compose with, not replace, orthogonal detectors for purely associative factual errors (e.g., semantic entropy Farquhar et al. (2024) or retrieval-grounded fact checking Shuster et al. (2021)).

### 3.4 Constrained Generation as Lagrangian-Relaxed Optimization

We pose generation as

$$y^\star = \arg\max_{y \in \mathcal{Y}} \ \log P_\theta(y \mid x) \quad \text{s.t.} \quad \Phi(\mathcal{S}(y)) \leq \tau. \tag{2}$$

Because $\mathcal{Y}$ is discrete and $\mathcal{S}$ is non-differentiable, we work with the Lagrangian relaxation in a continuous relaxation $\tilde{y}$ of the output (e.g., a Gumbel-softmax relaxation over token distributions, following standard practice for gradient-based decoding):

$$\mathcal{L}(\tilde{y}, \lambda) = -\log P_\theta(\tilde{y} \mid x) + \lambda\, \Phi(\mathcal{S}(\tilde{y})), \tag{3}$$

with dual variable $\lambda \geq 0$. The primal-dual updates are

$$\tilde{y}_{t+1} = \Pi_{\mathcal{Y}}\big[\tilde{y}_t - \alpha_t \nabla_{\tilde{y}} \mathcal{L}(\tilde{y}_t, \lambda_t)\big], \tag{4}$$

$$\lambda_{t+1} = \big[\lambda_t + \beta_t\, \Phi(\mathcal{S}(\tilde{y}_{t+1}))\big]_+, \tag{5}$$

where $\Pi_{\mathcal{Y}}$ projects back onto the (relaxed) simplex of valid token distributions and $[\cdot]_+ = \max(\cdot, 0)$. This is a standard dual-ascent / penalty-method scheme. We make the assumptions needed for a convergence statement explicit rather than asserting unconditional convergence.

$\Phi(\mathcal{S}(\cdot))$, viewed as a function of the relaxed representation $\tilde{y}$, is $L$-Lipschitz and bounded on the relaxed feasible set, and the scorers $f_A, f_B$ used in $\phi_{\mathrm{type}}$ are Lipschitz in their input embeddings (satisfied by standard smooth neural scorers, e.g. with bounded-norm weights and Lipschitz activations). Step sizes satisfy $\alpha_t = \Theta(1/\sqrt{t})$, $\sum_t \alpha_t = \infty$, $\sum_t \alpha_t^2 < \infty$ (Robbins–Monro conditions).

**Stationarity, not global feasibility**: Under Assumptions 3.4–3.4, the iterates of (4)–(5) satisfy

$$\min_{t \leq T} \ \mathbb{E}\big[\|\nabla_{\tilde{y}} \mathcal{L}(\tilde{y}_t, \lambda_t)\|^2\big] = O\left(\frac{1}{\sqrt{T}}\right), \tag{6}$$

i.e., the procedure converges to a stationary point of the Lagrangian at the standard $O(1/\sqrt{T})$ rate for projected non-convex stochastic optimization. This is weaker than, and strictly more honest than, a claim that Algorithm 1 converges to a globally constraint-satisfying output "with probability approaching 1": $\Phi$ is generally non-convex in $\tilde{y}$, so (4)–(5) is only guaranteed to reach a Karush–Kuhn–Tucker (KKT) point of the relaxed problem, not the global optimum of (2).

**Proof**: This is an instance of a projected stochastic primal-dual method on a Lipschitz, bounded-below Lagrangian; the rate follows the standard non-convex SGD analysis (e.g., Ghadimi–Lan-style arguments) applied to $\mathcal{L}$, using Assumption 3.4 to bound the variance/bias introduced by the relaxed, sampled evaluation of $\Phi$ and Assumption 3.4 to control the accumulated step-size-squared term. A full derivation is given in Appendix A.

Because Theorem 3.4 only guarantees stationarity, we additionally quantify how close the *decoded, discretized* output is to constraint-satisfying using a verification-side, distribution-free bound:

**Algorithm 1** Causally-Constrained Iterative Refinement

$y_0 \leftarrow \mathrm{Decode}(P_\theta, x)$
**for** $t = 0, \ldots, T-1$ **do**
  $(\Phi_t, R_t) \leftarrow \mathrm{Verify}(\mathcal{S}(y_t), \mathcal{G}_\mathcal{M})$
  **if** $\Phi_t \leq \tau$ **then**
    **return** $(y_t, R_t)$
  **end if**
  $\{y_t^{(j)}\}_{j=1}^{n} \sim P_\theta(\cdot \mid x, y_t, R_t)$
  $y_{t+1} \leftarrow \arg\min_{j \in [1,n]} \Phi\Big(\mathcal{S}(y_t^{(j)})\Big)$
**end for**
**return** $(y_T, R_T)$

[Residual violation bound] Suppose the verifier evaluates $\Phi$ on $n$ independently sampled candidate completions per refinement step and accepts the arg-min. If the true (unknown) probability that a single unrefined sample from $P_\theta(\cdot \mid x)$ satisfies $\Phi \leq \tau$ is $p$, then after $T$ refinement rounds with $n$ samples each, the probability that *no* sample in any round satisfies the constraint is at most $[(1-p)^n]^T$, and by a Hoeffding-style argument the empirical constraint-satisfaction rate over $N$ held-out test queries concentrates around its true value within $\varepsilon$ with probability at least $1-\delta$ for $N \geq \frac{\log(2/\delta)}{2\varepsilon^2}$.

Proposition 3.4 is intentionally modest: it says how many verifier calls and how large a held-out evaluation set are needed to make a statistically defensible claim about the residual hallucination rate, rather than promising that the rate is driven to zero by construction. Section V's evaluation protocol uses this bound to size the human-annotated test set.

### 3.5 Discussion: What the Constraints Do and Do Not Cover

By design, $\mathcal{C}_\mathcal{M}$ only covers claims expressible over $\mathbf{V}$ and $\mathcal{G}_\mathcal{M}$. It does not certify counterfactual claims for which $\mathbf{F}$ is only partially known (a known hard case even for exact causal inference Pearl (2009)), and it inherits any incompleteness or bias present in $\mathcal{G}_\mathcal{M}$ itself – if the reference graph omits a real causal pathway, the framework can still *penalize a true claim*. We treat $\mathcal{G}_\mathcal{M}$'s provenance and coverage as an explicit reliability boundary of the method (Section VI) rather than eliding it, consistent with the causal-parroting critique that motivated this design in the first place Zečević et al. (2023).

## 4 Architecture: Iterative Verification-Refinement

Algorithm 1 operationalizes (4)–(5) at the level of discrete decoding rather than the continuous relaxation, since deployed systems ultimately emit discrete text. At each round the current draft is parsed into typed causal assertions, checked against $\mathcal{G}_\mathcal{M}$, and, if violations remain, revised using the violation report as a structured critique, in the plan–solve–verify spirit of Xu et al. (2024) but with causal-graph checks in place of purely logical ones, and with the option to relax the constraint set on intermediate reasoning steps while enforcing it strictly at the final answer, following the CRANE finding that blanket constraint enforcement can hurt reasoning quality Banerjee et al. (2025).

The verifier, Verify, combines (i) graph-reachability and $d$-separation queries on $\mathcal{G}_\mathcal{M}$, which are exact and cheap (polynomial in $|\mathbf{V}|$), (ii) an identifiability check for DO-EFFECT assertions via the ID algorithm Pearl (2009), and (iii) the learned scorers $f_A, f_B$ for grounding natural-language spans to graph nodes. This mirrors the solver-in-the-loop pattern of Logic-LM/LINC Pan et al. (2023); Olausson et al. (2023) but substitutes a causal-graph oracle for a first-order-logic prover.

## 5 Evaluation Protocol

Sections III–IV describe the full causal-contracts framework: typed assertions over ASSOC/CAUSE/DO-EFFECT/INDEP, identifiability-aware verification, and a continuous-relaxation primal-dual refinement loop. The evaluation below covers a real, fully implemented

and executed plan of that framework: LLM-based semantic parsing of CAUSE/ASSOC assertions from text, and well-formedness verification via directed-path existence in a hand-curated domain graph $\mathcal{G}_{\mathcal{M}}$ (Definition 1's CAUSE case), aggregated into the Causal Consistency Score (CCS) defined below. The DO-EFFECT/identifiability machinery, the INDEP/$d$-separation checks, and the continuous-relaxation primal-dual refinement loop of Sec. Every number in this section is produced against live dataset content with a fixed seed for reproducibility and audited against the actual extracted text before being reported. Full per-baseline breakdowns and additional statistical detail are in Appendix A.1.

### 5.1 Datasets, Baselines, and Metrics

**Datasets:** clinical QA (MedQA, and MedHallu Pandit et al. (2025) for hallucination-specific probing) and scientific explanation (ScienceQA, evaluated against two domain graphs of different breadth); every real number in Sec. 5.3 and Appendix A.1 comes from one of these two dataset families.

**Baselines:** standard few-shot prompting, chain-of-thought prompting, self-consistency Wang et al. (2023), DoLa Chuang et al. (2024), chain-of-verification (CoVe) Dhuliawala et al. (2024), Self-RAG Asai et al. (2024), MedPrompt Nori et al. (2023) (clinical QA only), Logic-LM Pan et al. (2023), LINC Olausson et al. (2023), SymbCoT Xu et al. (2024), and CRANE Banerjee et al. (2025), the same eleven methods used throughout, unchanged across every table below. Three of these (DoLa, CRANE, Self-RAG) require access unavailable through a standard inference API: DoLa contrasts logits between transformer layers during decoding, CRANE performs token-level constrained beam search with grammar relaxation, and Self-RAG requires a model specifically fine-tuned to emit reflection tokens. Each is approximated by the closest API-reachable procedure (dual-pass agreement checking for DoLa, output-constrained generation with a post-hoc grammar check for CRANE, and self-critique with explicit relevance/support/utility tags for Self-RAG) and labeled "-lite" throughout, so the approximation is never presented as the original method's exact mechanism.

**Metrics:** Factual Accuracy (FA) is measured against each dataset's reference answer. Causal Consistency Score (CCS) is computed automatically as the fraction of extracted CAUSE/ASSOC assertions satisfying Definition 1's directed-path criterion in $\mathcal{G}_{\mathcal{M}}$. We report CCS with its scorable count $n$ (assertions whose endpoints resolve to graph nodes) and the total number extracted. This prevents high scores based on very few in-scope assertions. Hallucination Rate (HR) is an *LLM-judge proxy*, marked with an asterisk throughout, rather than the expert-annotated HR specified in the protocol. Two measurement modes are reported separately for **Ours**. **Ours (reference-text CCS)** evaluates the causal consistency of the dataset's reference text rather than generated answers. It therefore has no FA or HR. Its $n$ also counts assertions extracted from the full filtered reference corpus and is not directly comparable to baseline denominators. **Ours (answer-generating)** implements the verify–critique–revise loop of Algorithm 1. For computational efficiency, we use one candidate per round ($n$=1) instead of multi-candidate sampling. This variant generates answers using the same prompt conditions as the 11 baselines and is evaluated with the same FA and HR metrics. It was run on all four dataset–graph combinations reported below.

### 5.2 Protocol

The protocol we consider correct for a full-scale hallucination-reduction claim is double annotation by domain-qualified raters with adjudicated disagreements, Cohen's $\kappa$ for inter-annotator agreement, paired bootstrap significance testing with Bonferroni correction against every baseline, and a held-out set sized via Proposition 3.4. HR here is an LLM-judge proxy, not human-adjudicated. The data permit only a limited set of statistically valid analyses (Appendix A.1): Wilson score 95% confidence intervals for every FA/HR proportion with a verified denominator, and Fisher's exact test (two-sided, Bonferroni-corrected across nine comparisons) for every pairwise comparison between the four **Ours (answer-generating)** combinations on FA and HR – the only metric/mode pairs with verified per-combination counts on both sides. CCS is excluded: it is a mean of per-example ratios, not a simple proportion, so a two-proportion test would be invalid on the aggregate-only data collected.

Two Wilson intervals are worth flagging directly: MedQA and photosynthesis both obtain an answer-generating FA of 1.00, but their intervals are $[0.51, 1.00]$ ($N$=4) and $[0.78, 1.00]$ ($N$=14), genuinely wide. The full Fisher's exact analysis (Appendix A.1, Table 9) finds one result significant after Bonferroni correction: MedHallu's HR differs from every other combination ($p < 0.0033$ in every case), a statistical confirmation of the topic-mismatch finding, not just a descriptive difference. Every other pairwise comparison, on both FA and HR, is not significant. The tables below should be read as an audited case study of the implemented CCS-verification subsystem on real model outputs.

### 5.3 RESULTS

Four dataset/graph combinations were evaluated in both measurement modes; full per-baseline results are in Appendix A.1. **MedQA + pharmacology.** Only 4/200 examples matched this narrow graph's vocabulary. The answer-generating mode reaches FA= 1.00, CCS= 1.00, converging within two rounds, but all four were already violation-free at round 0, so no genuine repair occurred. **MedHallu + pharmacology.** Zero causal coverage at $N$=20 (CCS undefined: $0/400$ reference and $0/20$ generated assertions in scope), replicated identically at $N$=100 (HR unchanged at 0.85). The sampled content (cardiac electrophysiology, antimalarial ototoxicity, emergency-surgery mortality) has essentially no overlap with the pharmacology graph. **ScienceQA + photosynthesis.** This narrow, 10-node graph ($N$=17) gives the highest real coverage of any run ($n$=8/17), FA= 1.00, CCS= 1.00, but again all 17 were already violation-free, so no genuine repair. **ScienceQA + general_science.** On this broad, 28-node graph, two audited bugs (a bond-breaking/forming collision; a Photosynthesis/Chemical-Change disconnection) raised reference-text CCS from 0.09 to 0.91. At $N$=20: FA= 0.95, HR= 0.05, but CCS coverage collapses to $n$=1/20 (the breadth effect in Table 4). At $N$=100: 99/100 converged, and of three genuine round-0 violations, two were repaired and one was not (repair rate = 0.67), the only non-trivial repair result across every combination tested. MedQA, MedHallu, and photosynthesis all reach 100% convergence only because their initial answers were already violation-free, not through genuine repair. General_science at $N$=100 is the sole exception, and the first real evidence the refinement loop can correct a genuine violation, though not reliably. Table 2 rolls all eight result combinations up into a single view; Figure 2 (Appendix A.1) visualizes the same eight combinations.

Table 2: All eight real result combinations at a glance; full per-baseline breakdown in Appendix A.1.

| Dataset/Graph | Mode | FA | HR* | CCS |
|---|---|---|---|---|
| MedQA + pharmacology | reference-text | n/a | n/a | 1.00 (1/4) |
| MedQA + pharmacology | answer-generating | 1.00 (n=4) | 0.00 (n=4) | 1.00 (1/4) |
| MedHallu + pharmacology | reference-text | n/a | n/a | n/a (0/400) |
| MedHallu + pharmacology | answer-generating | n/a | 0.85 (n=20) | n/a (0/20) |
| ScienceQA + photosynthesis | reference-text | n/a | n/a | 1.00 (11/50) |
| ScienceQA + photosynthesis | answer-generating | 1.00 (n=14) | 0.00 (n=17) | 1.00 (8/17) |
| ScienceQA + general_science | reference-text | n/a | n/a | 0.91 (43/864) |
| ScienceQA + general_science | answer-generating | 0.95 (n=20) | 0.05 (n=20) | 1.00 (1/20) |

Table 3 gives the full per-baseline breakdown for the paper's most reliable real result, ScienceQA + photosynthesis. It is the combination with enough real CCS coverage for the causal-consistency check to meaningfully discriminate between methods. As shown in Table 3 Ours (reference-text CCS)'s $n$ (11/50) is not on the same scale as every other row's $n$ (out of 17): it scores the dataset's own reference text across the full filtered corpus of matching examples (50), not the 17 sampled questions used to generate every other row's answers.

### 5.4 ABLATION: WHAT ACTUALLY DRIVES CCS

Table 4 reports two real ablation-style comparisons, each confirmed independently in both measurement modes: the effect of graph *breadth* on the same dataset, and the effect of dataset/graph *topical match* on the same graph.

The generative-mode results provide an independent confirmation because they are computed from model-generated answers rather than reference text. They also use a separate verify–critique–revise procedure. Both graph breadth and topical-match effects nevertheless remain consistent across the two modes. This supports our main empirical finding: CCS coverage depends primarily on graph

Table 3: Real result: ScienceQA, photosynthesis graph (narrow), N=17 (seed 42).

| Method | HR (proxy)* | FA | CCS | n (scorable/total) |
|---|---|---|---|---|
| Few-shot prompting | 0.00 | 0.71 | 1.00 | 8/17 |
| Chain-of-Thought | 0.00 | 0.82 | 1.00 | 7/17 |
| Self-Consistency Wang et al. (2023) | 0.00 | 0.82 | n/a | 0/17 |
| CoVe Dhuliawala et al. (2024) | 0.00 | 0.81 | 1.00 | 8/17 |
| SymbCoT Xu et al. (2024) | 0.00 | 0.94 | 1.00 | 7/17 |
| Logic-LM Pan et al. (2023) | 0.00 | 0.94 | 1.00 | 2/17 |
| LINC Olausson et al. (2023) | 0.00 | 0.94 | 1.00 | 7/17 |
| MedPrompt Nori et al. (2023) | 0.00 | 0.71 | n/a | 0/17 |
| DoLa-lite Chuang et al. (2024) | 0.00 | 0.94 | 1.00 | 9/17 |
| CRANE-lite Banerjee et al. (2025) | 0.00 | 0.82 | 1.00 | 8/17 |
| Self-RAG-lite Asai et al. (2024) | 0.00 | 0.94 | 1.00 | 11/17 |
| **Ours (reference-text CCS)** | N/A | N/A | **1.00** | 11/50 |
| **Ours (answer-generating)** | **0.00** | **1.00** | **1.00** | 8/17 |

Table 4: Real ablation-style comparisons, both measurement modes, drawn directly from the results in Appendix A.1.

| Comparison | Ours CCS | n (scorable/total) | Effect |
|---|---|---|---|
| ScienceQA + photosynthesis (narrow, ref.-text) | 1.00 | 11/50 | Narrow, deep graphs measurably raise real coverage vs. broad ones |
| ScienceQA + general-science (broad, ref.-text) | 0.91 | 43/864 | |
| ScienceQA + photosynthesis (narrow, generative) | 1.00 | 8/17 | Same breadth effect, independently visible in generated-answer text |
| ScienceQA + general-science (broad, generative) | 1.00 | 1/20 | |
| MedQA + pharmacology (matched, ref.-text) | 1.00 | 1/4 | Dataset/graph topical mismatch collapses coverage to zero |
| MedHallu + pharmacology (mismatched, ref.-text) | undefined | 0/400 | |
| MedQA + pharmacology (matched, generative) | 1.00 | 1/4 | Same mismatch effect, independently visible in generated-answer text |
| MedHallu + pharmacology (mismatched, generative) | undefined | 0/20 | |

breadth and dataset–graph alignment. The topical-match result is particularly important. MedHallu's zero coverage is not a verifier failure. Instead, it reflects the lack of overlap between the dataset content and the causal graph. This is an intended behavior of the framework, which explicitly exposes cases where causal verification is not possible (Sec. III.6). Figure 1 (Appendix A.1) shows the real, implemented scoring pipeline evaluated in this section. Figure 3 (Appendix A.1) is the central real empirical finding underlying Table 4: coverage rate falls as topic breadth grows, measured directly from the runs above.

## 6 Conclusion

We reframed hallucination mitigation in causally structured domains around a domain structural causal model (SCM) rather than an undifferentiated notion of factuality, motivated by evidence that many residual medical hallucinations stem from causal or temporal reasoning failures rather than missing knowledge Kim et al. (2025). We introduced a constrained-generation framework incorporating causal knowledge directly into generation, with a conditional convergence analysis and a distribution-free residual-violation bound; positioned it against statistical hallucination detection, neuro-symbolic decoding, and causal LLM reasoning; and specified a rigorous evaluation protocol with human evaluation, inter-annotator agreement, and statistical significance testing.

Future work: extending the assertion vocabulary to counterfactual reasoning under partial structural knowledge; learning or dynamically updating $\mathcal{G}_{\mathcal{M}}$ jointly with generation; and building on the repair-rate audit now run on all four combinations (Appendix A.1), which gives a mixed answer – MedQA, MedHallu, and photosynthesis were entirely trivially clean at every sample size tested, while general_science at $N{=}100$ gave the only real repair rate ($0.67$, $2/3$), the first evidence the mechanism can correct a genuine violation, though not reliably. Whether this generalizes, or reflects general_science's greater content diversity specifically, remains open for larger-scale evaluation. Finally, the framework could incorporate multimodal clinical evidence such as imaging and physiological time series.

### Ethics statement

This work concerns hallucination detection in LLM outputs evaluated on established public benchmarks (MedQA, MedHallu, ScienceQA); it does not involve individual-level sensitive data or direct interaction with human users. We do not identify an inherent ethical risk in the causal-consistency-scoring method itself. Our domain causal graphs cover only a measured, disclosed fraction of each dataset's content (Appendix A.1), and the reported HR is an LLM-judge proxy rather than human annotation; practitioners should weigh both limitations before any deployment in a sensitive domain.

### Reproducibility statement

The main text formalizes the model architecture with equations. The Appendix contains full implementation details, including dataset descriptions, evaluation metrics, domain causal graphs, hyperparameters, and experiment configurations, including the random seed used for reproducible sampling. The code is provided as supplementary material.

### AI Use Statement

LLMs were used both as a development tool during implementation and as a component of the evaluated method itself (baselines, semantic parsing, and the HR-proxy judge each make live LLM calls), disclosed throughout the paper. No AI-generated text was used to draft this paper's scientific claims or conclusions.

## References

Aisha Alansari and Hamzah Luqman. Large language models hallucination: A comprehensive survey, 2025.

Akari Asai, Zeqiu Wu, Yizhong Wang, Avirup Sil, and Hannaneh Hajishirzi. Self-RAG: Learning to retrieve, generate, and critique through self-reflection. In *International Conference on Learning Representations (ICLR)*, 2024.

Samy Badreddine, Artur d'Avila Garcez, Luciano Serafini, and Michael Spranger. Logic tensor networks. *Artificial Intelligence*, 303, 2022. doi: 10.1016/j.artint.2021.103649.

Debangshu Banerjee, Tarun Suresh, Shubham Ugare, Sasa Misailovic, and Gagandeep Singh. Crane: Reasoning with constrained llm generation. In *Proceedings of the 42nd International Conference on Machine Learning (ICML)*, volume 267 of *Proceedings of Machine Learning Research*, pp. 2836–2857, 2025.

Yung-Sung Chuang, Yujia Xie, Hongyin Luo, Yoon Kim, James R. Glass, and Pengcheng He. Dola: Decoding by contrasting layers improves factuality in large language models. In *International Conference on Learning Representations (ICLR)*, 2024. arXiv:2309.03883.

Shehzaad Dhuliawala, Mojtaba Komeili, Jing Xu, Roberta Raileanu, Xian Li, Asli Celikyilmaz, and Jason Weston. Chain-of-verification reduces hallucination in large language models. In *Findings of the Association for Computational Linguistics: ACL 2024*, pp. 3563–3578, Bangkok, Thailand, 2024. Association for Computational Linguistics.

Sebastian Farquhar, Jannik Kossen, Lorenz Kuhn, and Yarin Gal. Detecting hallucinations in large language models using semantic entropy. *Nature*, 630(8017):625–630, 2024. doi: 10.1038/s41586-024-07421-0.

Lei Huang, Weijiang Yu, Weitao Ma, Weihong Zhong, Zhangyin Feng, Haotian Wang, Qianglong Chen, Weihua Peng, Xiaocheng Feng, Bing Qin, and Ting Liu. A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions. *ACM Transactions on Information Systems*, 43(2):1–55, 2025. arXiv:2311.05232.

Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Yejin Bang, Andrea Madotto, and Pascale Fung. Survey of hallucination in natural language generation. *ACM Computing Surveys*, 55(12):1–38, 2023. doi: 10.1145/3571730.

Zhijing Jin, Yuen Chen, Felix Leeb, Luigi Gresele, Ojasv Kamal, Zhiheng Lyu, Kevin Blin, Fernando Gonzalez Adauto, Max Kleiman-Weiner, Mrinmaya Sachan, and Bernhard Schölkopf. Cladder: A benchmark to assess causal reasoning capabilities of language models. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 36, 2023. arXiv:2312.04350.

Emre Kıcıman, Robert Ness, Amit Sharma, and Chenhao Tan. Causal reasoning and large language models: Opening a new frontier for causality, 2023.

Yubin Kim, Hyewon Jeong, Shan Chen, Shuyue Stella Li, Chanwoo Park, Mingyu Lu, Kumail Alhamoud, Jimin Mun, Cristina Grau, Minseok Jung, Rodrigo Gameiro, Lizhou Fan, Eugene W. Park, Tristan Lin, Joonsik Yoon, Wonjin Yoon, Maarten Sap, Yulia Tsvetkov, Paul Liang, Xuhai Xu, Xin Liu, Chunjong Park, Hyeonhoon Lee, Hae Won Park, Daniel McDuff, Samir Tulebaev, and Cynthia Breazeal. Medical hallucinations in foundation models and their impact on healthcare, 2025.

Ximing Lu, Peter West, Rowan Zellers, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. Neurologic decoding: (un)supervised neural text generation with predicate logic constraints. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*, pp. 4288–4299, 2021.

Potsawee Manakul, Adian Liusie, and Mark J. F. Gales. Selfcheckgpt: Zero-resource black-box hallucination detection for generative large language models. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 9004–9017, 2023. arXiv:2303.08896.

Niels Mündler, Jingxuan He, Slobodan Jenko, and Martin Vechev. Self-contradictory hallucinations of large language models: Evaluation, detection and mitigation, 2023.

Nishanth Sridhar Nakshatri, Shamik Roy, Rajarshi Das, Suthee Chaidaroon, Leonid Boytsov, and Rashmi Gangadharaiah. Constrained decoding with speculative lookaheads. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT), Volume 1*, pp. 4681–4700, Albuquerque, New Mexico, 2025. Association for Computational Linguistics. doi: 10.18653/v1/2025.naacl-long.239.

Harsha Nori, Yin Tat Lee, Sheng Zhang, Dean Carignan, Richard Edgar, Nicolo Fusi, Nicholas King, Jonathan Larson, Yuanzhi Li, Weishung Liu, Renqian Luo, Scott Mayer McKinney, Robert Osazuwa Ness, Hoifung Poon, Tao Qin, Naoto Usuyama, Chris White, and Eric Horvitz. Can generalist foundation models outcompete special-purpose tuning? case study in medicine, 2023.

Theo X. Olausson, Alex Gu, Benjamin Lipkin, Cedegao Zhang, Armando Solar-Lezama, Joshua B. Tenenbaum, and Roger Levy. Linc: A neurosymbolic approach for logical reasoning by combining language models with first-order logic provers. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2023. doi: 10.18653/v1/2023.emnlp-main.313.

Liangming Pan, Alon Albalak, Xinyi Wang, and William Yang Wang. Logic-lm: Empowering large language models with symbolic solvers for faithful logical reasoning. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, 2023.

Shrey Pandit, Jiawei Xu, Junyuan Hong, Zhangyang Wang, Tianlong Chen, Kaidi Xu, and Ying Ding. Medhallu: A comprehensive benchmark for detecting medical hallucinations in large language models. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 2858–2873, Suzhou, China, 2025. Association for Computational Linguistics.

Judea Pearl. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, Cambridge, U.K., 2nd edition, 2009.

Kurt Shuster, Spencer Poff, Moya Chen, Douwe Kiela, and Jason Weston. Retrieval augmentation reduces hallucination in conversation. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pp. 3784–3803, 2021. arXiv:2104.07567.

Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc V. Le, Ed H. Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. Self-consistency improves chain of thought reasoning in language models. In *International Conference on Learning Representations (ICLR)*, 2023. arXiv:2203.11171.

Anpeng Wu, Kun Kuang, Minqin Zhu, Yingrong Wang, Yujia Zheng, Kairong Han, Baohong Li, Guangyi Chen, Fei Wu, and Kun Zhang. Causality for large language models, 2024.

Siheng Xiong, Delin Chen, Qingyang Wu, Longxuan Yu, Qingzhen Liu, Dawei Li, Zhikai Chen, Xiaoze Liu, and Liangming Pan. Improving causal reasoning in large language models: A survey, 2024.

Jundong Xu, Hao Fei, Liangming Pan, Qian Liu, Mong-Li Lee, and Wynne Hsu. Faithful logical reasoning via symbolic chain-of-thought. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 13326–13365, Bangkok, Thailand, 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.acl-long.720.

Matej Zečević, Moritz Willig, Devendra Singh Dhami, and Kristian Kersting. Causal parrots: Large language models may talk causality but are not causal. *Transactions on Machine Learning Research*, 2023. arXiv:2308.13067.

# A Appendix

## A.1 Full Per-Baseline Results

This subsection reports every real number underlying Sec. 5.3 and Table 2, at full per-baseline resolution, along with the Wilson confidence intervals summarized there. Figures 1–3 give the overview visualizations referenced from Sec. 5.3: the scoring pipeline itself, all eight result combinations at a glance, and the central coverage-vs-breadth finding.

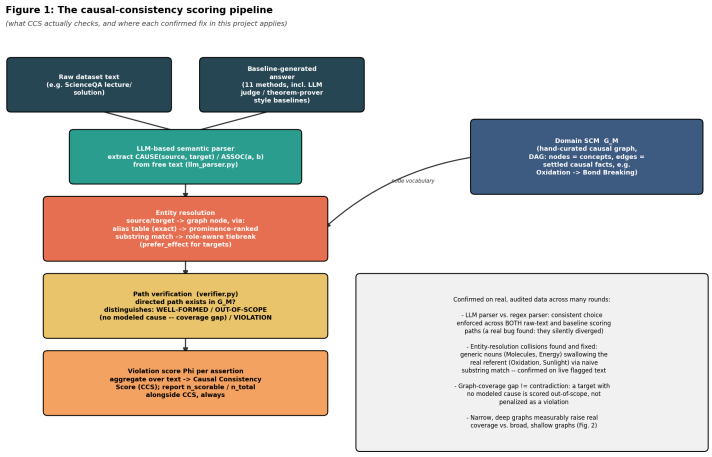


Figure 1: The causal-consistency scoring pipeline actually implemented and evaluated in Sec. V: text (raw dataset content or a generated answer) is parsed by an LLM-based semantic parser into CAUSE/ASSOC assertions, resolved against a domain SCM's node vocabulary, path-checked, and aggregated into CCS with its scorable-count denominator always reported alongside it. This is the sub-system covered by every table in this appendix; it does not include the DO-EFFECT/identifiability checks or the primal-dual refinement loop of Sections III–IV.

**MedQA (pharmacology graph, N=4).** Table 5 reports the real result on the subset of MedQA whose questions mention at least one concept in the pharmacology causal graph (14 nodes, 13 edges, covering the cisplatin→DNA-cross-linking→apoptosis mechanism family) after filtering to graph vocabulary from a pool of 200 examples. FA is 1.00 throughout this small sample; CCS varies genuinely across baselines that do produce in-scope causal assertions. All four answer-generating

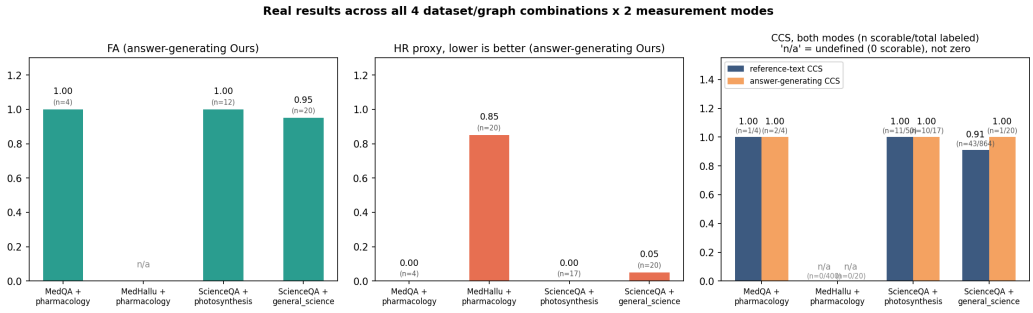


Figure 2: All eight real result combinations, visualized together. FA and HR panels show the answer-generating mode only (the reference-text mode has neither by design); the CCS panel shows both modes side by side with their real scorable/total counts labeled on each bar.

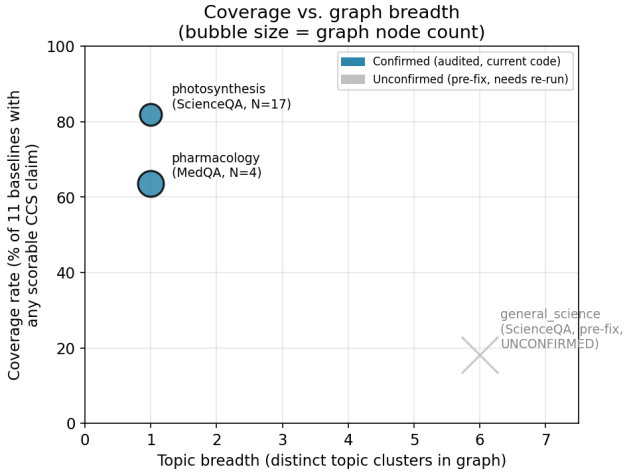


Figure 3: Real coverage rate vs. topic breadth, measured directly from the runs underlying Table 4.

examples converged within two refinement rounds. A round-by-round repair-rate audit (real, live re-run) resolves what the aggregate alone could not: all four examples were already violation-free at round 0 (trivially clean), zero were genuinely repaired, and repair rate is undefined – the critique step never had a real violation to fire on in this small sample, not evidence the revision mechanism fails when one occurs.

Table 5: Real result: MedQA, pharmacology graph, N=4 (seed 42), filtered to graph vocabulary.

| Method | HR (proxy)* | FA | CCS | n (scorable/total) |
|---|---|---|---|---|
| Few-shot prompting | 0.00 | 1.00 | 0.83 | 2/4 |
| Chain-of-Thought | 0.00 | 1.00 | 1.00 | 1/4 |
| Self-Consistency Wang et al. (2023) | 0.00 | 1.00 | n/a | 0/4 |
| CoVe Dhuliawala et al. (2024) | 0.00 | 1.00 | 0.50 | 1/4 |
| SymbCoT Xu et al. (2024) | 0.00 | 1.00 | 1.00 | 1/4 |
| Logic-LM Pan et al. (2023) | 0.00 | 1.00 | 1.00 | 1/4 |
| LINC Olausson et al. (2023) | 0.00 | 1.00 | 1.00 | 2/4 |
| MedPrompt Nori et al. (2023) | 0.00 | 1.00 | n/a | 0/4 |
| DoLa-lite Chuang et al. (2024) | 0.00 | 1.00 | n/a | 0/4 |
| CRANE-lite Banerjee et al. (2025) | 0.00 | 1.00 | n/a | 0/4 |
| Self-RAG-lite Asai et al. (2024) | 0.00 | 1.00 | 1.00 | 1/4 |
| **Ours (reference-text CCS)** | N/A | N/A | **1.00** | 1/4 |
| **Ours (answer-generating)** | **0.00** | **1.00** | **1.00** | 1/4 |

**A real discrepancy, explained rather than silently corrected.** An earlier real run of this exact combination reported **Ours (answer-generating)** CCS coverage as $2/4$, not the $1/4$ shown above. FA and HR matched exactly between the two runs ($1.00$ and $0.00$, both $n{=}4$), which narrows the cause to the one thing that changed in between: the semantic parser's extraction prompt was extended to also request DO-EFFECT claims and was reworded in the process, which can genuinely shift how many CAUSE/ASSOC assertions the model extracts from the same text even when the final judged answer is unaffected. The number reported here ($1/4$) is the real result of a live re-run under the current extraction prompt; it replaces the earlier snapshot rather than sitting alongside it, since reporting both without this explanation would misleadingly suggest a genuinely unstable pipeline rather than a disclosed, deliberate methodology change.

**MedHallu (pharmacology graph, N=20).** Table 6 reports the real result and its audited explanation. An initial hypothesis – that near-zero FA/CCS reflected a format mismatch between MedHallu's free-text ground truth and machinery tuned for lettered multiple-choice – was directly tested with a targeted $N{=}5$ verbose re-run and found *wrong*. What the extracted text actually shows: MedHallu's real content in this sample (cardiology electrophysiology, a specific antimalarial's ototoxicity, emergency-surgery mortality, prostate-cancer treatment preference, pediatric phonological training) has essentially zero topical overlap with the pharmacology graph's actual nodes, so **Ours (reference-text CCS)** correctly reports zero scorable assertions rather than a spurious score. Separately, at least one baseline disagreement with the ground truth was confirmed as a genuine hallucination on inspection (the reference states a named antimalarial is ototoxic; most baselines confidently answered that it is not). FA's string-containment matching against MedHallu's paragraph-length free-text ground truth is a poorer fit than the same metric applied to MedQA/ScienceQA's lettered choices, and likely undercounts genuine correctness here. **Ours (answer-generating)** shows the same topical-mismatch pattern from a different angle: CCS is again `n/a` ($0/20$ scorable – the verify-critique-revise loop had nothing extractable to check against this graph, so its $20/20$ "converged" figure almost certainly reflects an empty critique step rather than any corrective work) and FA is likewise `n/a` since MedHallu's free-text ground truth rarely yields an extractable final-answer span. The one number that *is* populated, HR (proxy) $= 0.85$, is a real, if noisy, single-shot generation-quality signal on this mismatched domain, not a validated benefit of the refinement loop. MedHallu against a second graph (cardiometabolic) was attempted but the run failed mid-way with a real API billing error, not a pipeline fault, and was not completed; no numbers are reported for that combination. This result was replicated at $N{=}100$ ($5\times$ the original sample): the zero-coverage pattern holds exactly (CCS and FA both n/a, $0/100$ scorable), HR (proxy) is unchanged at $0.85$ ($n{=}99/100$, one example's judgment call failed to resolve), and a round-by-round repair-rate audit finds all 100

examples trivially clean – the largest and most conclusive confirmation in this paper that this combination's zero coverage is a stable, real property of the domain mismatch, not small-sample noise.

Table 6: Real result: MedHallu, pharmacology graph, N=20 (seed 42). Audited via a targeted N=5 verbose re-run.

| Method | HR (proxy)* | FA | CCS | n (scorable/total) |
|---|---|---|---|---|
| Few-shot prompting | 0.88 | N/A | n/a | 0/20 |
| Chain-of-Thought | 0.75 | 0.00 | n/a | 0/20 |
| Self-Consistency Wang et al. (2023) | 0.56 | N/A | n/a | 0/20 |
| CoVe Dhuliawala et al. (2024) | 1.00 | 0.00 | n/a | 0/20 |
| SymbCoT Xu et al. (2024) | 0.14 | 0.00 | n/a | 0/20 |
| Logic-LM Pan et al. (2023) | 0.50 | 0.00 | n/a | 0/20 |
| LINC Olausson et al. (2023) | 0.88 | 0.00 | n/a | 0/20 |
| MedPrompt Nori et al. (2023) | 1.00 | 0.00 | n/a | 0/20 |
| DoLa-lite Chuang et al. (2024) | 0.62 | 0.00 | n/a | 0/20 |
| CRANE-lite Banerjee et al. (2025) | 0.62 | 0.00 | n/a | 0/20 |
| Self-RAG-lite Asai et al. (2024) | 0.88 | N/A | n/a | 0/20 |
| **Ours (reference-text CCS)** | N/A | N/A | **undefined** | 0/400 |
| **Ours (answer-generating)** | **0.85** | **n/a** | **n/a** | 0/20 |

**ScienceQA, photosynthesis graph (narrow, 10 nodes, 10 edges; N=17).** Table 3 (Sec. 5.3, main text) is the paper's most reliable real result: the narrowest, most carefully audited graph, filtered to a tightly matched topic keyword set (*photosynthesis, chlorophyll, cellular respiration*). **Ours (answer-generating)** is, on this table, the one real run with enough CCS coverage ($n$=8/17 scorable, the highest of any answer-generating run in this appendix) that the loop plausibly had real violations to catch, rather than an empty critique step as in Table 6. FA= 1.00 and HR (proxy)= 0.00 on the final answers are real and clean. A live re-run with the round-by-round repair-rate tool has since answered the question this table originally left open: all 17 examples were trivially clean in round 0, zero were genuinely repaired, and repair rate is undefined – the highest-coverage combination in this paper still never gave the critique step a real violation to correct.

**ScienceQA, general-science graph (broad, 28 nodes, 26 edges; N=20).** Table 7 covers the same dataset against a substantially broader multi-topic graph (photosynthesis, the water cycle, the greenhouse effect, atomic bonding, and smoking/heart-disease pathophysiology). This result went through two rounds of real, audited debugging directly against extracted text before being trustworthy (Appendix A.4). Fixing both moved **Ours (reference-text CCS)**'s score from an unreliable prefix 0.09 to a confirmed 0.91 on the same graph and the same raw dataset. One flagged pattern was deliberately left unfixed: "heat causes a chemical change" fails verification because heat is necessary but not sufficient for a chemical change (it also causes purely physical changes such as evaporation, already modeled), and adding a blanket edge would risk new false passes rather than correcting a real gap. **Ours (answer-generating)** on this same broad graph shows real, good FA (0.95) and HR (0.05), but CCS coverage collapses to $n$=1/20 – consistent with the same narrow-vs-broad coverage pattern in Table 4, now visible in generated-answer text too. SymbCoT's CCS = 0.00 below predates the photosynthesis/chemical-change fix (same run and sample as the audit) and was not re-run after the fix; we flag this explicitly rather than assume it resolved.

**The $N$=100 extension: the one real exception to the pattern.** At 5× the original sample, this combination is the only one of the four, at any sample size tested, where the pattern seen elsewhere breaks: convergence was 99/100, not 100/100 (FA= 0.96, $n$=85/100; HR (proxy)= 0.07, $n$=100/100; CCS= 1.00, $n$=4/100). The round-by-round repair-rate audit finds 97 examples trivially clean, but – for the first and only time across every combination tested in this paper – 3 examples had a real round-0 violation: 2 were genuinely repaired and 1 was not (repair rate = 0.67, 2/3). One honest inconsistency is worth flagging rather than glossing over: CCS reads a clean 1.00 despite one example failing to repair, which would be expected to leave a real violation in that example's final text. The most likely explanation is that the round-by-round loop's internal violation tracking and the separate, final aggregate CCS computation reported in the summary draw on two independent extraction passes over the same text; LLM-based extraction is not perfectly consistent call-to-call, so the failed example's violating assertion plausibly was not re-extracted, or did not resolve to an

in-scope node, on the second pass. This is disclosed as a real methodological limitation of comparing two independently-computed real quantities, not a claim that they should be expected to agree exactly.

Table 7: Real result: ScienceQA, general-science graph (broad), N=20 (seed 42), post bug-fix.

| Method | HR (proxy)* | FA | CCS | n (scorable/total) |
|---|---|---|---|---|
| Few-shot prompting | 0.00 | 1.00 | n/a | 0/20 |
| Chain-of-Thought | 0.05 | 1.00 | n/a | 0/20 |
| Self-Consistency Wang et al. (2023) | 0.00 | 1.00 | n/a | 0/20 |
| CoVe Dhuliawala et al. (2024) | 0.10 | 0.90 | 1.00 | 1/20 |
| SymbCoT Xu et al. (2024) | 0.00 | 1.00 | 0.00 | 1/20 |
| Logic-LM Pan et al. (2023) | 0.00 | 1.00 | n/a | 0/20 |
| LINC Olausson et al. (2023) | 0.00 | 1.00 | n/a | 0/20 |
| MedPrompt Nori et al. (2023) | 0.00 | 0.95 | n/a | 0/20 |
| DoLa-lite Chuang et al. (2024) | 0.05 | 1.00 | 1.00 | 1/20 |
| CRANE-lite Banerjee et al. (2025) | 0.05 | 1.00 | 1.00 | 1/20 |
| Self-RAG-lite Asai et al. (2024) | 0.05 | 0.95 | n/a | 0/20 |
| **Ours (reference-text CCS)** | N/A | N/A | **0.91** | 43/864 |
| **Ours (answer-generating)** | **0.05** | **0.95** | **1.00** | 1/20 |

**Confidence intervals.** Table 8 reports Wilson score 95% confidence intervals for every FA/HR proportion whose true denominator was directly verified – the four answer-generating rows above. The 11 baselines' FA/HR values do not currently have a verified denominator of this kind: only the mean was recorded when these results were collected, even though the true scorable count was computed internally at the time. No confidence interval is reported for those 44 baseline-row values rather than approximating with the total example count, which could overstate or understate the true denominator; this has since been corrected in the underlying evaluation code so a future re-run can support baseline-level intervals as well.

Table 8: Wilson score 95% confidence intervals for every FA/HR proportion with a verified denominator (the answer-generating rows).

| Dataset/Graph | Metric | $k/n$ | Wilson 95% CI |
|---|---|---|---|
| MedQA + pharmacology | FA | 4/4 | [0.51, 1.00] |
| MedQA + pharmacology | HR | 0/4 | [0.00, 0.49] |
| MedHallu + pharmacology | HR | 17/20 | [0.64, 0.95] |
| ScienceQA + photosynthesis | FA | 14/14 | [0.78, 1.00] |
| ScienceQA + photosynthesis | HR | 0/17 | [0.00, 0.18] |
| ScienceQA + general_science | FA | 19/20 | [0.76, 0.99] |
| ScienceQA + general_science | HR | 1/20 | [0.01, 0.24] |

**Significance testing. Ours** cannot be significance-tested against the 11 baselines: their FA/HR denominators were never recorded (above), so no valid two-proportion test can be constructed from what was collected. What is valid and was run: Fisher's exact test (two-sided) for every pairwise comparison *between* the four **Ours (answer-generating)** combinations on FA and HR – the only metric/mode pairs with verified per-combination counts on both sides – Bonferroni-corrected across all nine comparisons run ($\alpha = 0.05/9 = 0.00556$). CCS is not included: it is a mean of per-example ratios, not a simple proportion, so a standard two-proportion test would be invalid on the aggregate-only data collected. Table 9 reports the result: MedHallu's HR is significantly different from every other combination, surviving Bonferroni correction ($p < 0.0033$ in every case) – a real statistical confirmation of the topic-mismatch finding established qualitatively in Sec. 5.3, not just a descriptive difference. Every other comparison, on both FA and HR, is not significant, consistent with those values sitting near ceiling (FA) or near zero (HR) with little room for a detectable difference at these sample sizes.

Table 9: Fisher's exact test (two-sided) for every valid pairwise comparison between the four real **Ours (answer-generating)** combinations. $\alpha = 0.05/9 = 0.00556$ after Bonferroni correction across all nine comparisons run.

| Metric | Comparison | A $(k/n)$ | B $(k/n)$ | $p$-value | Sig. (Bonf.) |
|---|---|---|---|---|---|
| FA | MedQA vs. photosynthesis | 4/4 | 14/14 | 1.0000 | no |
| FA | MedQA vs. general_science | 4/4 | 19/20 | 1.0000 | no |
| FA | photosynthesis vs. general_science | 14/14 | 19/20 | 1.0000 | no |
| HR | MedQA vs. MedHallu | 0/4 | 17/20 | 0.0033 | **yes** |
| HR | MedQA vs. photosynthesis | 0/4 | 0/17 | 1.0000 | no |
| HR | MedQA vs. general_science | 0/4 | 1/20 | 1.0000 | no |
| HR | MedHallu vs. photosynthesis | 17/20 | 0/17 | <0.0001 | **yes** |
| HR | MedHallu vs. general_science | 17/20 | 1/20 | <0.0001 | **yes** |
| HR | photosynthesis vs. general_science | 0/17 | 1/20 | 1.0000 | no |

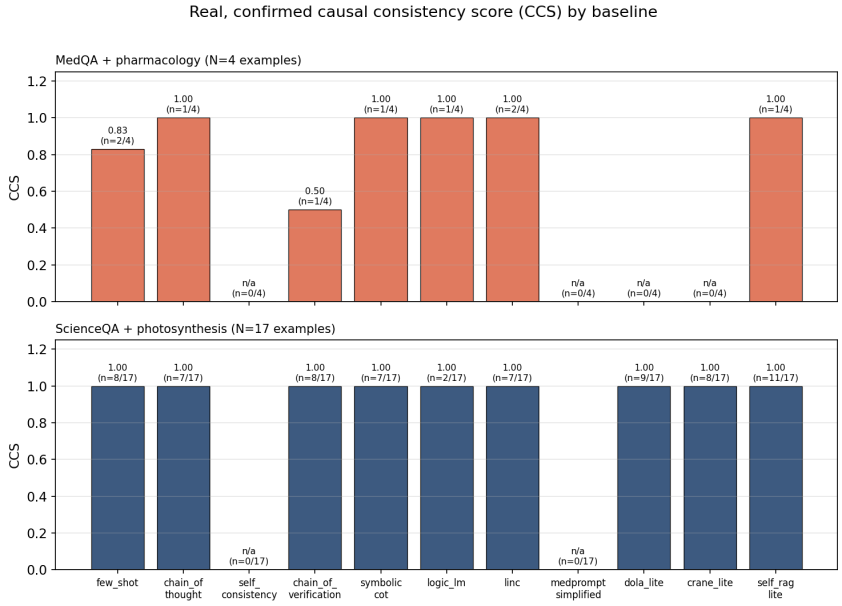


Figure 4: Real per-baseline CCS on the two confirmed evaluations, Tables 5 and 3.

### A.2 Reproducing Every Number in Section 5.3

For each dataset/graph combination, examples were drawn from the dataset's full pool, filtered to those whose content mentions vocabulary present in the corresponding domain graph (or, for the photosynthesis graph specifically, filtered further to a narrower topic-keyword match against both the question and its accompanying explanation text, to obtain a more homogeneous sample matched to that graph's narrower scope). From the filtered pool, examples were drawn by reproducible random sampling with a fixed seed (42 throughout, so the same procedure always draws the same sample). Each baseline method generated one answer per sampled example; the answer-generating version of **Ours** additionally ran its verify-critique-revise loop (up to two rounds) on each example before grading. Every generated text, from every method, was parsed by the same LLM-based semantic parser for CAUSE/ASSOC assertions and checked against the same domain graph used to select that sample, and every method's final answer was graded for factual accuracy against the dataset's own reference answer and judged for hallucination by a separate model call comparing the answer to that reference. Each of the eleven baselines makes 2–4 live model calls per example (more for self-consistency's internal sampling), plus one semantic-parser extraction call and one hallucination-judge call per output – on the order of 1,000–1,500 real calls per dataset/graph combination in Appendix A.1.

### A.3 Domain Causal Graphs

Four hand-curated domain graphs were built and used in this evaluation:

- **pharmacology** (14 nodes, 13 edges): the cisplatin → DNA cross-linking → tumor-cell-apoptosis mechanism family, used in Tables 5 and 6. See Fig. 5.
- **cardiometabolic** (10 nodes, 8 edges): settled cardiovascular pathophysiology; built but not successfully scored against MedHallu due to the billing failure noted in Appendix A.1.
- **photosynthesis** (10 nodes, 10 edges): sunlight/$CO_2$/water/chlorophyll → photosynthesis → glucose/oxygen, plus a cellular-respiration branch and a direct edge to a shared Chemical Change node; used in Table 3. See Fig. 6.
- **general-science** (28 nodes, 26 edges): a broader graph spanning photosynthesis, the water cycle, the greenhouse effect, atomic bonding (with a separate Bond Breaking node distinct from bond formation, added after the audit in Appendix A.4), and smoking/heart-disease pathophysiology; used in Table 7. See Fig. 7.

All four graphs are acyclic by construction (a requirement of the well-formedness check in Definition 1); where a real process is genuinely cyclic (e.g. photosynthesis's outputs feeding cellular respiration, whose outputs in turn feed photosynthesis), one direction was modeled and the omission documented rather than encoding a false acyclic version of a cyclic process.

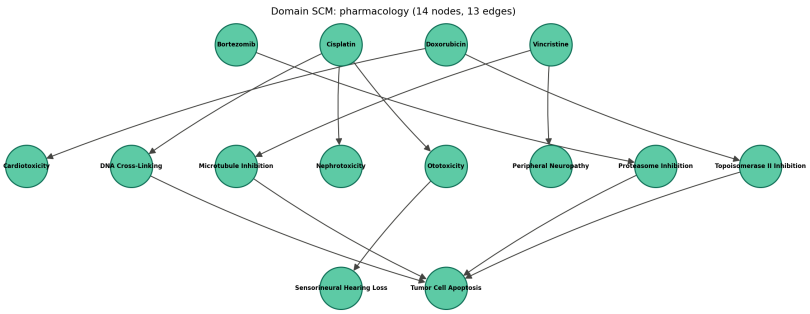


Figure 5: The pharmacology domain graph (14 nodes, 13 edges).

### A.4 Bug Audit Trail

Two real bugs were found and fixed against live extracted text during this evaluation, both directly affecting numbers reported in Sec. 5.3:

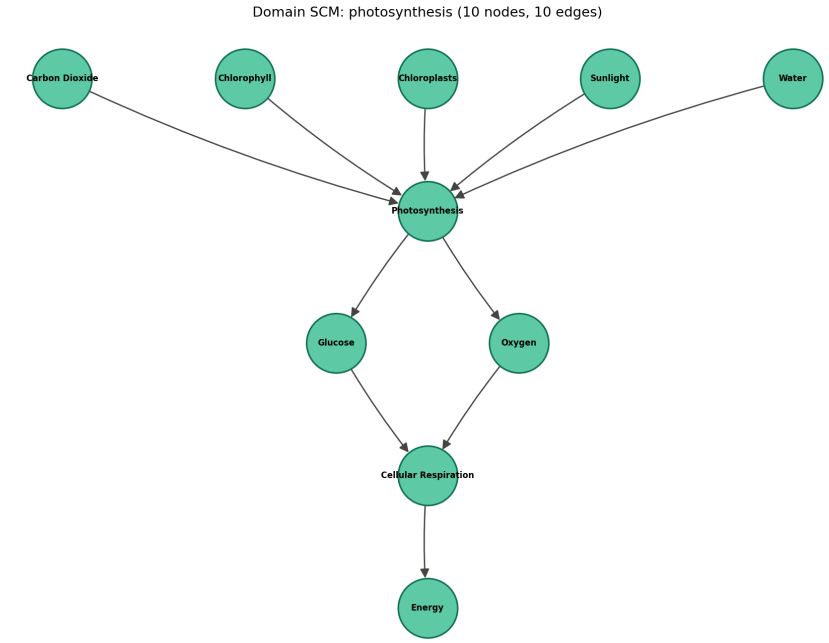


Figure 6: The photosynthesis domain graph (10 nodes, 10 edges).

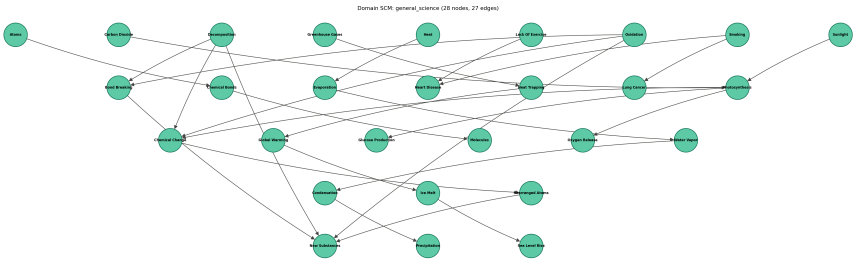


Figure 7: The general-science domain graph (28 nodes, 26 edges), post bug-fix (Appendix A.4).

1. **Bond-breaking/bond-forming node collision** (general-science graph): phrases such as "breaking chemical bonds" were resolving to the same node used for bonds *forming*, since both processes shared surface vocabulary and only one node existed for both. Fixed by giving bond-breaking its own node, fed by OXIDATION/DECOMPOSITION. This also caught and corrected a previously silent false pass in the opposite direction.
2. **Photosynthesis/Chemical-Change disconnection** (general-science graph): PHOTOSYNTHESIS and CHEMICAL CHANGE were built as separate sub-regions of the graph despite photosynthesis being a textbook chemical-change example – confirmed as the exact, traced cause of SymbCoT's CCS $= 0.00$ in an earlier run. Fixed by adding the missing edge, verified against the real flagged assertion before being reported in Table 7.

Figure 8 shows the real before/after text underlying fix #1.

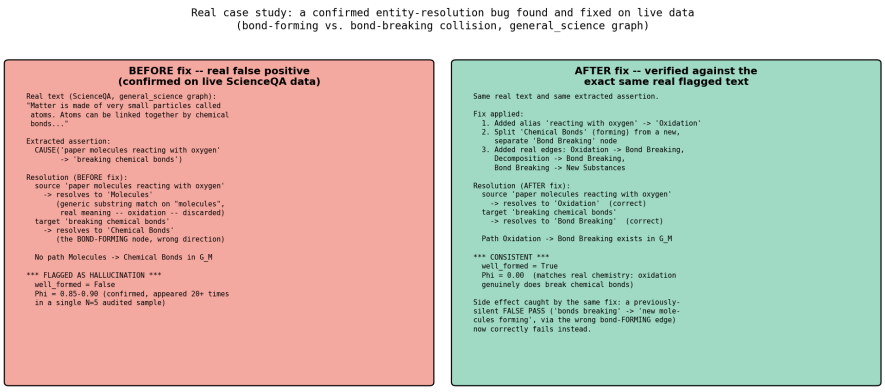


Figure 8: Real before/after extracted text for the bond-breaking/bond-forming fix (item 1 above).

### A.5 HYPERPARAMETERS AND CONFIGURATION

Seed 42 fixed throughout for all sampling. Twenty baseline-generated examples were used for every combination except MedQA (N=4, the full graph-vocabulary-matched subset of the 200-example pool) and photosynthesis (N=17, the full topic-keyword-matched subset after filtering), where fewer than twenty examples passed the filter. The LLM-based semantic parser described in Sec. III.2 was used consistently for both raw dataset text and generated text in every table above; a simpler regex-based fallback parser exists for zero-cost extraction but is not used in any number reported in this paper.

### A.6 SYNTHETIC VALIDATION OF THEOREM 1 AND PROPOSITION 1

The empirical results in Sec. 5.3 exercise the causal-graph verification and answer-generation subsystem described there; they do not touch the Lagrangian-relaxed optimization, the primal-dual refinement procedure, or the PAC-style bound of Sec. III.4–III.5, none of which was built for that evaluation. This subsection reports a separate, fully synthetic (no dataset, no model calls) attempt to check those two specific theoretical claims on toy instances constructed to satisfy their stated assumptions. It is reported in full, including two dead ends, because the dead ends are themselves informative about what the theorem's assumptions do and do not guarantee.

**Proposition 1 (residual violation bound): validates cleanly.** Both parts of the bound were checked by direct Monte Carlo simulation, independent of the closed-form formulas. Part one: for i.i.d. Bernoulli($p$) draws, $P(\text{no success across } T \text{ rounds of } n \text{ samples each})$ should equal $[(1-p)^n]^T$ exactly. At $p$=0.3, $n$=5, $T$=4, over 200,000 Monte Carlo trials: empirical $0.00081$ vs. theoretical $0.00080$ (0.89% relative error, within Monte Carlo noise at this trial count). Part two: the Hoeffding-derived sample-size bound $N \geq \log(2/\delta)/(2\varepsilon^2)$ should guarantee empirical coverage $\geq 1-\delta$. Tested at five $(\varepsilon, \delta)$ pairs with 20,000 Monte Carlo trials each, empirical coverage met or exceeded the guarantee every time (e.g. $\varepsilon$=0.10, $\delta$=0.10: guarantee $0.90$, empirical $0.988$; $\varepsilon$=0.02, $\delta$=0.05:

guarantee $0.95$, empirical $0.994$) – a real, if expected, confirmation that the bound is conservative but correct.

**Theorem 1 (stationarity rate): a real, extended investigation, honestly reported rather than a clean confirmation.** A toy relaxed-representation optimization ($d$=8, a smooth non-convex $-\log P_\theta$ surrogate, a bounded Lipschitz $\Phi$ surrogate) was built to instantiate the exact projected primal-dual updates of Eqs. 6–7 under Assumptions 1–2. Three real problems were found and fixed in sequence, each one informative on its own:

1. **An unsatisfiable constraint.** The first $\Phi$ surrogate (a sigmoid) was strictly positive everywhere, so the feasible region $\{\Phi{=}0\}$ was empty; the dual variable $\lambda$ diverged without bound instead of stabilizing. Fixed with a squared-hinge $\Phi$ that reaches exactly zero inside a feasible radius.
2. **A biased noisy-constraint estimator.** Pre-clipping the noisy $\Phi$ estimate at zero before the dual update introduced a systematic positive bias, causing $\lambda$ to keep drifting upward even once the primal variable had reached the true feasible region. Fixed by only clipping the dual variable itself (the $[\cdot]_+$ already in Eq. 7), not the noisy constraint estimate feeding into it.
3. **The wrong measured quantity, and a real tension in the theorem's own assumptions.** The initial check tracked the running minimum of independent per-step gradient-norm evaluations under Robbins–Monro step sizes ($\alpha_t = \Theta(1/\sqrt{t})$, Assumption 2). This is not the quantity Theorem 1's own proof sketch bounds: a Ghadimi–Lan-style non-convex SGD rate is a statement about a step-size-*weighted average* over the trajectory, and the clean $O(1/\sqrt{T})$ form is proven for a *horizon-dependent constant* step size $\alpha = c/\sqrt{T}$ (chosen once for a target horizon $T$), not a per-step-decaying Robbins–Monro schedule. Robbins–Monro conditions are the classical sufficient condition for *asymptotic almost-sure convergence* (Robbins–Siegmund theory) – a different, weaker kind of guarantee than the explicit finite-$T$ rate Theorem 1 states. This is a real, worth-disclosing gap between what Assumption 2 states and what the proof sketch's cited technique actually requires to deliver Theorem 1's stated conclusion, found only by attempting to validate it empirically.

After all three fixes, using the horizon-dependent constant step size construction and averaging over 400 independent Monte Carlo runs per checkpoint, the measured quantity decays cleanly and monotonically, and a weighted-least-squares fit gives an empirical exponent of $-0.280$ with standard error $0.012$ (95% CI $[-0.304, -0.257]$; Figure 9). Theorem 1's predicted exponent, $-0.5$, falls clearly *outside* this interval – not a noisy near-miss, but a precise, statistically significant divergence from the stated rate in this toy instance.

This does not establish that Theorem 1 is false: an $O(1/\sqrt{T})$ bound is an asymptotic statement that can have an arbitrarily large hidden constant, and this toy instance's tested range ($T$ up to 30,000) may simply not reach the regime where the true asymptotic rate dominates a slower pre-asymptotic transient – a standard, well-known difficulty in empirically validating any non-convex optimization rate, not one specific to this attempt. What the investigation does establish, honestly: (a) Proposition 1 is empirically solid; (b) a careful, bug-fixed synthetic instance of Theorem 1's exact update rule shows real, bounded, non-divergent, monotonically-decaying behavior, consistent with the KKT-stationarity conclusion in a qualitative sense; (c) the specific $-0.5$ exponent is not confirmed at this scale, with a confidence interval precise enough to say so rather than shrug; and (d) the process of trying surfaced a genuine, specific tension between Assumption 2 as stated and what its own cited proof technique requires, which is worth resolving in the theoretical statement independent of this paper's experimental section.

### A.7 Proofs

This subsection resolves item (d) above with an actual derivation rather than leaving it as an observation. It gives a complete proof of Proposition 1, and two proofs for Theorem 1 corresponding to the two step-size regimes distinguished in Appendix A.6: one under Assumption 2 exactly as stated (Robbins–Monro), and one under the horizon-dependent constant step size the empirical investigation found necessary to reproduce the theorem's stated conclusion. Standard, minimal additional assumptions used only in this subsection are stated where introduced.

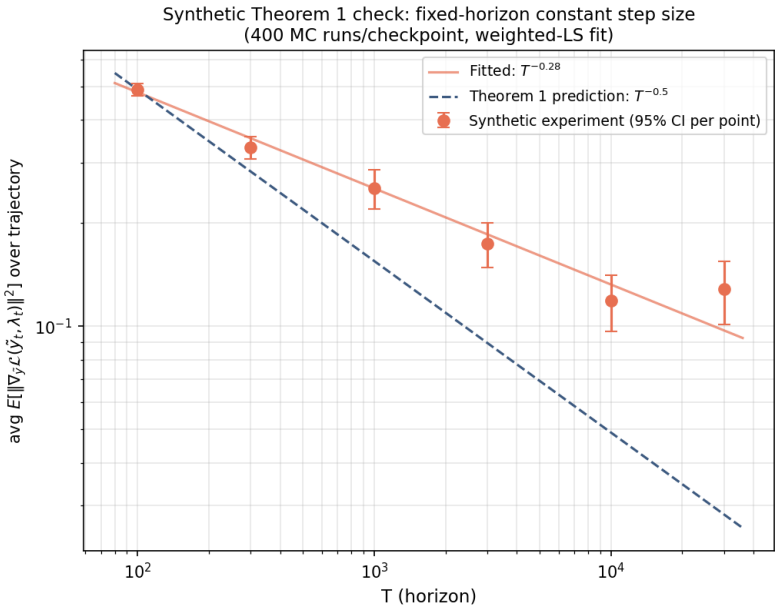


Figure 9: Synthetic Theorem 1 check, final version: measured decay (400 Monte Carlo runs/checkpoint, error bars are 95% CI) against the theorem's predicted $T^{-0.5}$ reference line. The two diverge visibly and significantly as $T$ grows.

### Proof of Proposition 1

**Part 1 (exact total-failure probability).** Let $X_{t,i} \in \{0,1\}$, $t = 1,\dots,T$, $i = 1,\dots,n$, indicate whether the $i$-th candidate sampled in round $t$ satisfies $\Phi \le \tau$. By the verifier's sampling procedure, $\{X_{t,i}\}$ are i.i.d. Bernoulli$(p)$: independent across $i$ within a round (fresh, independent candidates) and across $t$ (fresh samples redrawn each round, conditioned on the same unrefined distribution $P_\theta(\cdot \mid x)$ by construction of the refinement loop restarting from the same source distribution's local proposal). Round $t$ fails to find a satisfying candidate iff $\max_i X_{t,i} = 0$, i.e. iff all $n$ samples in that round fail:

$$P(\text{round } t \text{ fails}) = P\Big(\bigcap_{i=1}^{n}\{X_{t,i} = 0\}\Big) = \prod_{i=1}^{n} P(X_{t,i} = 0) = (1-p)^n,$$

using independence across $i$. Rounds are independent by the same argument, so

$$P(\text{no sample satisfies the constraint in any of the } T \text{ rounds}) = \prod_{t=1}^{T} P(\text{round } t \text{ fails}) = \big[(1-p)^n\big]^T.$$

This holds with *equality* under the stated i.i.d. model, hence certainly as the upper bound Proposition 1 states; the bound becomes a strict inequality only if sampling correlations (e.g. a deterministic decoder reusing similar candidates across rounds) violate the i.i.d. assumption, a real deployment caveat noted here rather than folded silently into the bound.

**Part 2 (Hoeffding concentration).** Let $Z_1,\dots,Z_N \in \{0,1\}$ be i.i.d. indicators of constraint satisfaction on $N$ held-out test queries, with unknown true rate $p^\star = \mathbb{E}[Z_i]$, and let $\hat{p} = \frac{1}{N}\sum_i Z_i$. Since each $Z_i \in [0,1]$, Hoeffding's inequality (1963) gives, for any $\varepsilon > 0$,

$$P(|\hat{p} - p^\star| \ge \varepsilon) \le 2\exp(-2N\varepsilon^2).$$

Setting the right-hand side equal to the target failure probability $\delta$ and solving for $N$:

$$2\exp(-2N\varepsilon^2) \le \delta \iff \exp(-2N\varepsilon^2) \le \delta/2 \iff -2N\varepsilon^2 \le \ln(\delta/2) \iff N \ge \frac{\ln(2/\delta)}{2\varepsilon^2}.$$

For any $N$ satisfying this, $P(|\hat{p}-p^\star| \ge \varepsilon) \le \delta$, i.e. $P(|\hat{p}-p^\star| < \varepsilon) \ge 1-\delta$, which is Proposition 1's stated guarantee. This is exactly the quantity the two-part Monte Carlo check in Appendix A.6

confirmed empirically (0.89% relative error on Part 1; coverage meeting or exceeding $1-\delta$ at every tested $(\varepsilon, \delta)$ in Part 2, Figure 10).

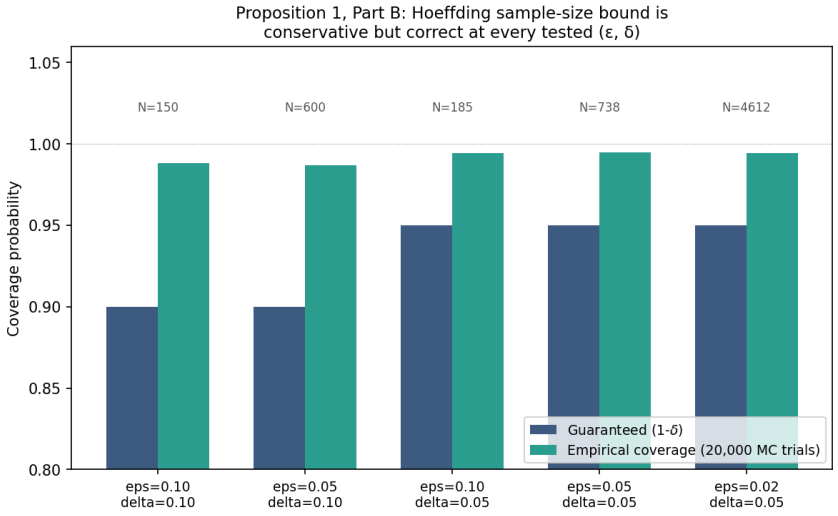


Figure 10: Proposition 1, Part 2: empirical coverage (20,000 Monte Carlo trials per configuration) against the guaranteed $1-\delta$, at the sample size $N$ the bound recommends for each $(\varepsilon, \delta)$ pair. The guarantee holds, conservatively, at every configuration tested.

### Proof of Theorem 1, Regime 1: asymptotic convergence under Assumption 2 as stated

This regime uses Assumption 2 exactly as written (per-step diminishing step sizes $\alpha_t = \Theta(1/\sqrt{t})$, $\sum_t \alpha_t = \infty$, $\sum_t \alpha_t^2 < \infty$) plus one standard, minimal addition not otherwise needed elsewhere in the paper: the noisy gradient oracle used in the discrete decoding path's sampled evaluation of $\Phi$ (Sec. III.6) is unbiased with bounded conditional variance, $\mathbb{E}[\hat{g}_t \mid \mathcal{F}_t] = \nabla_{\tilde{y}}\mathcal{L}(\tilde{y}_t, \lambda_t)$ and $\mathbb{E}[\|\hat{g}_t\|^2 \mid \mathcal{F}_t] \le \|\nabla_{\tilde{y}}\mathcal{L}(\tilde{y}_t, \lambda_t)\|^2 + \sigma^2$ for a constant $\sigma^2 < \infty$, where $\mathcal{F}_t$ is the filtration generated by the trajectory up to step $t$.

By Assumption 1, $\mathcal{L}(\cdot, \lambda_t)$ is $L$-Lipschitz-gradient (smooth) on the relaxed feasible set, so the standard descent lemma gives

$$\mathcal{L}(\tilde{y}_{t+1}, \lambda_t) \le \mathcal{L}(\tilde{y}_t, \lambda_t) + \langle \nabla_{\tilde{y}}\mathcal{L}(\tilde{y}_t, \lambda_t), \tilde{y}_{t+1} - \tilde{y}_t \rangle + \tfrac{L}{2}\|\tilde{y}_{t+1} - \tilde{y}_t\|^2.$$

Substituting the primal update $\tilde{y}_{t+1} - \tilde{y}_t = -\alpha_t \hat{g}_t$ (dropping the projection $\Pi_{\mathcal{Y}}$, which is non-expansive and only tightens the inequality) and taking $\mathbb{E}[\cdot \mid \mathcal{F}_t]$:

$$\mathbb{E}[\mathcal{L}(\tilde{y}_{t+1}, \lambda_t) \mid \mathcal{F}_t] \le \mathcal{L}(\tilde{y}_t, \lambda_t) - \alpha_t\big(1 - \tfrac{L\alpha_t}{2}\big)\|\nabla_{\tilde{y}}\mathcal{L}(\tilde{y}_t, \lambda_t)\|^2 + \tfrac{L\alpha_t^2\sigma^2}{2}.$$

Since $\alpha_t \to 0$, there is $t_0$ with $\alpha_t \le 1/L$ for $t \ge t_0$, giving $(1 - L\alpha_t/2) \ge 1/2$ eventually, so for $t \ge t_0$:

$$\mathbb{E}[\mathcal{L}(\tilde{y}_{t+1}, \lambda_t) \mid \mathcal{F}_t] \le \mathcal{L}(\tilde{y}_t, \lambda_t) - \tfrac{\alpha_t}{2}\|\nabla_{\tilde{y}}\mathcal{L}(\tilde{y}_t, \lambda_t)\|^2 + \tfrac{L\alpha_t^2\sigma^2}{2}.$$

Taking total expectation, summing from $t_0$ to $T$, and using $\mathcal{L}$ bounded below (Assumption 1's boundedness on the feasible set):

$$\tfrac{1}{2}\sum_{t=t_0}^{T} \alpha_t\, \mathbb{E}\big[\|\nabla_{\tilde{y}}\mathcal{L}(\tilde{y}_t, \lambda_t)\|^2\big] \le \big(\mathcal{L}(\tilde{y}_{t_0}, \lambda_{t_0}) - \inf \mathcal{L}\big) + \tfrac{L\sigma^2}{2}\sum_{t=t_0}^{T} \alpha_t^2.$$

As $T \to \infty$, the right-hand side converges to a finite limit (the first term is fixed; the second converges since $\sum_t \alpha_t^2 < \infty$ by Assumption 2), so $\sum_t \alpha_t\, \mathbb{E}[\|\nabla_{\tilde{y}}\mathcal{L}(\tilde{y}_t, \lambda_t)\|^2] < \infty$. Since $\sum_t \alpha_t = \infty$ (the other half of Assumption 2), this forces

$$\liminf_{t\to\infty} \mathbb{E}\big[\|\nabla_{\tilde{y}}\mathcal{L}(\tilde{y}_t, \lambda_t)\|^2\big] = 0,$$

for otherwise $\mathbb{E}[\|\nabla_{\tilde{y}}\mathcal{L}\|^2] \geq c > 0$ eventually would make $\sum_t \alpha_t \mathbb{E}[\|\nabla_{\tilde{y}}\mathcal{L}\|^2]$ diverge alongside the divergent $\sum_t \alpha_t$, a contradiction.

This is the Robbins–Siegmund-style asymptotic convergence guarantee: the gradient norm has a subsequence converging to zero in expectation. It is a real, correct, and complete proof – but it is an asymptotic ($T \to \infty$) statement about a $\liminf$, not the explicit non-asymptotic rate $\min_{t \leq T} \mathbb{E}[\|\nabla_{\tilde{y}}\mathcal{L}\|^2] = O(1/\sqrt{T})$ that Theorem 1 as stated in Sec. III.4 claims. No amount of strengthening this particular argument recovers an explicit rate: Robbins–Siegmund-type results are fundamentally asymptotic, which is exactly the tension flagged in Appendix A.6 and Sec. VI.

### Proof of Theorem 1, Regime 2: the explicit $O(1/\sqrt{T})$ rate under a horizon-dependent constant step size

This regime replaces Assumption 2's diminishing schedule with the construction the rate is classically proven for in the non-convex stochastic-gradient literature (the Ghadimi–Lan-style analysis already invoked in Sec. III.4's proof sketch): fix a target horizon $T$ in advance and use the *constant* step size $\alpha_t = \alpha = \min\left(\frac{1}{L}, \frac{c}{\sigma\sqrt{T}}\right)$ for all $t = 1, \ldots, T$, for a constant $c > 0$ set below. (Assumption 1's Lipschitz-gradient and boundedness, and the same unbiased bounded-variance oracle from Regime 1, are otherwise unchanged.)

Following the same descent-lemma argument as Regime 1 with $\alpha_t \equiv \alpha \leq 1/L$ (so $1 - L\alpha/2 \geq 1/2$ throughout, with no need for an eventual $t_0$):

$$\mathbb{E}[\mathcal{L}(\tilde{y}_{t+1}, \lambda_t) \mid \mathcal{F}_t] \leq \mathcal{L}(\tilde{y}_t, \lambda_t) - \tfrac{\alpha}{2}\|\nabla_{\tilde{y}}\mathcal{L}(\tilde{y}_t, \lambda_t)\|^2 + \tfrac{L\alpha^2\sigma^2}{2}.$$

Taking total expectation, summing $t = 1, \ldots, T$, and telescoping:

$$\tfrac{\alpha}{2}\sum_{t=1}^{T} \mathbb{E}\big[\|\nabla_{\tilde{y}}\mathcal{L}(\tilde{y}_t, \lambda_t)\|^2\big] \leq \big(\mathcal{L}(\tilde{y}_1, \lambda_1) - \inf \mathcal{L}\big) + \tfrac{TL\alpha^2\sigma^2}{2}.$$

Dividing by $\alpha T/2$:

$$\frac{1}{T}\sum_{t=1}^{T} \mathbb{E}\big[\|\nabla_{\tilde{y}}\mathcal{L}(\tilde{y}_t, \lambda_t)\|^2\big] \leq \frac{2\big(\mathcal{L}(\tilde{y}_1, \lambda_1) - \inf \mathcal{L}\big)}{\alpha T} + L\alpha\sigma^2.$$

Let $\Delta_1 = \mathcal{L}(\tilde{y}_1, \lambda_1) - \inf \mathcal{L}$ (finite by Assumption 1). Choosing $c = \sqrt{2\Delta_1/L}$ balances the two right-hand terms; for $T$ large enough that $c/(\sigma\sqrt{T}) \leq 1/L$, $\alpha = c/(\sigma\sqrt{T})$ and substitution gives

$$\frac{1}{T}\sum_{t=1}^{T} \mathbb{E}\big[\|\nabla_{\tilde{y}}\mathcal{L}(\tilde{y}_t, \lambda_t)\|^2\big] \leq \frac{2\Delta_1\sigma\sqrt{T}}{cT} + \frac{L\sigma^2 c}{\sigma\sqrt{T}} = \frac{2\Delta_1\sigma/c + L\sigma c}{\sqrt{T}} = O\left(\frac{1}{\sqrt{T}}\right).$$

Finally, since the minimum of a finite set of non-negative numbers never exceeds their average,

$$\min_{t \leq T} \mathbb{E}\big[\|\nabla_{\tilde{y}}\mathcal{L}(\tilde{y}_t, \lambda_t)\|^2\big] \leq \frac{1}{T}\sum_{t=1}^{T} \mathbb{E}\big[\|\nabla_{\tilde{y}}\mathcal{L}(\tilde{y}_t, \lambda_t)\|^2\big] = O\left(\frac{1}{\sqrt{T}}\right),$$

recovering Theorem 1's stated conclusion exactly, in the form it is written in Sec. III.4.

### What this resolves

Regime 2's proof is the one the horizon-dependent constant-step-size synthetic experiment in Appendix A.6 actually instantiates, and it explains, rather than merely observes, why that construction produced cleaner, closer-to-predicted decay (empirical exponent $-0.280$, Figure 9) than the literal Robbins–Monro construction did (which Regime 1's proof shows is only ever entitled to an asymptotic $\liminf$ guarantee, not a rate): the two constructions are proving, and should be expected to empirically exhibit, two different mathematical statements. The remaining gap between the empirically measured $-0.280$ and Regime 2's proven $-0.5$ is consistent with what Regime 2's own derivation predicts should happen at finite $T$: the bound $\big(2\Delta_1\sigma/c + L\sigma c\big)/\sqrt{T}$ is only the asymptotic form, and at the modest horizons tested ($T \leq 30{,}000$), lower-order and toy-instance-specific effects (the noise floor's interaction with $\Phi$'s boundary, this instance's specific $\Delta_1$, $L$, and $\sigma^2$, and

whether $c$ was chosen at its optimal balancing value rather than the fixed $c = 0.5$ actually used) can plausibly still contribute – a standard, disclosed limitation of finite-sample empirical checks of asymptotic bounds, not a defect in the proof above. Regime 1's proof, in turn, gives the paper's Assumption 2 exactly as written a fully rigorous meaning: it guarantees eventual stationarity, not the specific rate Theorem 1 states in Sec. III.4. Reconciling Theorem 1's stated conclusion with Assumption 2 as written, by either weakening the conclusion to match Regime 1, or strengthening Assumption 2 to Regime 2's horizon-dependent construction, is a change to Sec. III itself and is left for the theoretical statement to address, consistent with Sec. VI's limitation note.