

Causal World Models for a Typed-Anchor Temporal Knowledge Graph

Executive synthesis

A correlational temporal graph answers questions of the form “what tended to happen after what.” A causal world model answers questions of the form “what mechanism changed, under which intervention, and would the downstream outcome still have occurred under a different action or context.” In the SCM/do-calculus tradition, the required architectural shift is not merely a new edge-labeling rule. It is a change in graph semantics: nodes must denote endogenous variables in a generative process, edges must denote parent sets in structural assignments, hidden context/noise must be represented explicitly, and task executions must be modeled as interventions that modify one or more mechanisms rather than as ordinary events in the same log. SCMs are designed precisely to support intervention distributions and counterfactuals, not just observational prediction. ¹

For the typed-anchor setting, the practical implication is that “anchor selected pointing at frame Y” should usually be split into at least three variables: a latent or observed world state, a quality or task-success variable for the frame, and a selection/attention variable generated by the policy or anchoring subsystem. Otherwise the graph conflates “this frame was good,” “this frame was chosen,” and “this frame was chosen because it was good,” which are distinct causal claims. A causal graph for this system must therefore separate state, action, outcome, and measurement/selection processes. ²

The answer to whether execution records are interventional data is conditional. They can be treated as interventional only when the execution event corresponds to an actual manipulation of a target mechanism and the logging setup makes the action exogenous enough for the question being asked. If execution is chosen by a policy using hidden information, the resulting log is observational from the learner’s perspective and can remain confounded. In sequential settings, identification from logged actions typically requires a sequential ignorability condition or an equivalent no-unobserved-confounding assumption. ³

Architectural shift from temporal graph to causal world model

In an SCM, each variable is generated by a structural assignment $X_j := f_j(\text{Pa}_j, N_j)$, where N_j is exogenous noise and the set of mechanisms is assumed to be modular or autonomous. Interventions are represented by replacing one or more structural assignments while leaving the others unchanged. This is the formal change that turns a graph from an event-sequence representation into a causal model. The independence-of-mechanisms view is central: the system is modeled as autonomous modules whose conditional mechanisms can be changed separately. ⁴

For a temporal knowledge graph, the most direct causal upgrade is a dynamic SCM or interventional dynamic Bayesian network over time slices. A minimal abstraction for the typed-anchor case is

$$S_{t+1} = f_S(S_t, T_t, C_t, U_t^S), \quad Q_t = f_Q(S_t, T_t, C_t, U_t^Q), \quad A_t = f_A(S_t, Q_t, \pi_t, U_t^A),$$

where S_t is system/world state, T_t is task execution, C_t is regime/context, Q_t is frame quality or task-relevant outcome, and A_t is anchor selection. In this factorization, the anchor points to evidence about Q_t or S_t ; it is not itself the causal fact. This formulation matches the SCM view of sequential decision making, in which actions are explicit endogenous variables and interventional queries such as $P(S_{t+1} \mid S_t, do(T_t = \tau))$ are well-defined. ⁵

The architectural change is therefore fourfold. First, replace “happened-before” edges with mechanism-parent edges. Second, add exogenous/context variables so that the graph can represent common causes and regime shifts instead of forcing all dependence through observed events. Third, separate outcome variables from selection variables, so that anchor selection is modeled as a policy or measurement process rather than as direct proof of outcome quality. Fourth, add an intervention layer that records which task was executed, what variables or mechanisms it was intended to affect, and whether the intervention is hard, soft, or unknown-target. This last distinction matters because many real execution events alter mechanisms without severing all incoming dependencies, which is closer to soft intervention than to a hard do-operation. ⁶

Causal discovery on an event-sourced graph

For temporally indexed variables, classical causal discovery algorithms can be applied after a lag expansion or time-slice unrolling, but their assumptions differ substantially. The PC algorithm is constraint-based: it starts from a complete undirected graph and removes edges using conditional independence tests, then orients what it can into a CPDAG. It is suited to sparse graphs and can distinguish direct from indirect relationships when its assumptions hold, but it remains observational unless intervention contexts are added, and it is sensitive to hidden confounding and test error. ⁷

GES is score-based and searches directly over equivalence classes rather than individual DAGs. In the observational setting it searches over DAG equivalence classes; in the interventional setting, GIES refines this to interventional equivalence classes. This matters in the present context because interventional data can make the equivalence classes strictly finer than under observational data alone, which means more edge directions become identifiable. ⁸

NOTEARS reframes DAG learning as continuous optimization over weighted adjacency matrices using a smooth, exact acyclicity constraint. For temporal data, DYNOTEARS extends this idea to learn intra-slice and inter-slice dependencies jointly and reports applicability at roughly $d \approx 100$ variables, but it assumes a fixed DBN structure through time. That assumption is directly relevant for a temporal graph whose causal structure may drift. ⁹

DCDI is the closest match among the listed methods when execution logs can be treated as interventions. It explicitly supports perfect, imperfect, and unknown-target interventions, and identifies the interventional Markov equivalence class under regularity assumptions. The paper also notes that when every node is individually targeted, the true graph becomes isolated in its equivalence class. For execution-heavy systems with known or inferable task targets, this is the strongest of the listed options. ¹⁰

For asynchronous event streams rather than regularly sampled time series, point-process or event-sequence methods are often a better fit than lagged DAG learners. Recent examples include CASCADE for

event sequences, Structural Hawkes models for event logs, and MOCHA, which models multi-order dynamic causality in temporal point processes using time-varying DAGs with sparsity and acyclicity constraints. These methods are especially relevant when your graph is event-sourced rather than frame-synchronous.

11

Task execution records should therefore be treated as **candidate intervention contexts**, not automatically as interventions. They can count as interventional data if the task truly manipulates a mechanism and either the task choice is randomized or the policy parents of the task variable are sufficiently observed to support sequential ignorability. If task choice depends on hidden state, the log is observational with respect to the learner, even though an action occurred in the environment. A useful compromise is to represent execution explicitly as its own node and let the model infer whether it acts as a hard intervention, a soft intervention, or merely an observed treatment assignment.

12

Observational and interventional identification

Identification from observational data asks whether a causal query such as $P(Y \mid do(X = x))$ can be uniquely expressed using the pre-intervention distribution and a causal graph. In the SCM framework, this is the role of do-calculus and graphical identification results: if the do-expression can be converted into a do-free expression under the assumed graph, the effect is identifiable from passive data. If not, more assumptions or more data types are needed.

13

In the typed-anchor context, observational data alone often leaves several explanations compatible with the same event history: a task may improve frame quality, a hidden context may cause both the task and the quality, or the anchoring policy may systematically attend to already-good frames. Observational discovery also tends to stop at a Markov equivalence class. This is why merely learning that “task X and anchor A often precede frame Y” does not answer the user’s actual question of whether executing X caused Y to be high quality.

14

Interventional data changes that situation in two ways. First, it can orient edges that are indistinguishable from observational data; Hauser and Bühlmann show that interventional equivalence classes are finer than observational ones, and DCDI proves recovery of the interventional equivalence class under known or even unknown targets. Second, it licenses direct estimation of action effects under explicit intervention semantics rather than only conditional prediction. In sequential SCMs, this is the difference between $P(S_{t+1} \mid S_t, T_t)$ and $P(S_{t+1} \mid S_t, do(T_t = \tau))$.

15

Treating task execution as an explicit intervention node gives the system three concrete gains. It enables effect estimation for planning, because the model can ask what happens under alternative task choices rather than only what happened under the logged policy. It enables counterfactual analysis, because SCM-based sequential models can ask whether the next state or outcome would have differed had a different action been taken. It also localizes confounding: once action choice is explicit, the system can ask whether hidden or observed variables were parents of the action, which is exactly the distinction between observational treatment assignment and intervention.

16

Recent literature on scalable and changing causal world models

The 2023–2025 literature does not yet converge on a single scalable causal world-model architecture for large, nonstationary agentic systems. The strongest high-confidence pattern is instead a set of recurring design principles: factorize the world into sparse modules or objects, learn local causal structure rather than dense global structure, use interventions or regime changes to improve identifiability, and allow multiple graphs or changing mechanisms rather than assuming one fixed global DAG. The theoretical backdrop is an ICLR 2024 result showing a formal link between causality and robust generalization, arguing that agents that robustly solve shifted decision tasks must learn a causal model of the data-generating process. At the same time, the new CausalDynamics benchmark argues that much of the discovery literature still lags behind the complexity of realistic dynamical systems. ¹⁷

On the world-model side, a 2023 explainable RL framework learns a sparse causal world model from environment transitions specifically so explanations and decisions use the same structure. In 2024, OOCMD extends causal dynamics models to large-scale object-oriented environments by sharing causalities and parameters across objects of the same class, explicitly targeting “a multitude of objects” and varying object counts. SPARTAN then pushes the same trend further with an object-factored sparse transformer that learns context-dependent local interaction graphs and adapts to sparse interventions with unknown targets. A 2025 object-centric extraction approach, COMET, combines object-centric state extraction and symbolic regression to derive interpretable causal world models from sensory inputs. ¹⁸

On the changing-structure side, the literature increasingly rejects the assumption of a single fixed causal graph. CtrlINS identifies distribution shifts and latent temporal factors under nonstationary sparse transitions. SAM addresses “superposed causal relationships” by inferring per-trajectory or per-episode causal graphs rather than one global graph, precisely because a global graph becomes over-dense when multiple mechanisms are mixed. The 2025 “Meta-Causal Graph” work goes further and represents the world as a collection of causal subgraphs indexed by latent meta-state, discovered through curiosity-driven interventions rather than passive observation alone. ¹⁹

For event-sequence and industrial settings, recent work is similar in spirit but uses temporal event models instead of object-centric simulators. RealTCD frames temporal causal discovery for industrial scenarios using interventional data, while CASCADE and MOCHA target event sequences directly, with MOCHA introducing a time-varying DAG over temporal point processes to handle evolving, multi-order causal structure. These methods are relevant if the typed-anchor graph is closer to an event stream than to a fixed-vector state sequence. ²⁰

The main limitation of this recent literature is scope. Most methods achieve scalability by assuming object factorization, local sparsity, or a small number of latent regimes. Many demonstrations remain on synthetic, object-centric, or moderate-scale domains rather than on large partially observed agentic systems with millions of heterogeneous events. The field is moving toward mixture-of-graphs and regime-conditioned world models, but end-to-end causal world models that are simultaneously scalable, nonstationary, partially observed, and counterfactually faithful remain an open problem. ²¹

Representing causal uncertainty

A posterior over graph structures is principled, but it is not sufficient for the belief state you described. Recent work on uncertain causal graphs makes this explicit: observational data are often compatible with multiple graphs, and methods that commit to a single structure become overconfident, but the Bayesian object of interest is not only a posterior over graphs. It is a posterior over graphs **and** functional mechanisms, because downstream interventional distributions depend on both. Dhir et al. argue that even given a posterior over plausible graphs, one still needs posteriors over the mechanisms to estimate intervention distributions; this is exactly why they target the intervention posterior directly. ²²

In your context, there is a further layer: intervention-target uncertainty. BaCaDI models unknown interventions as a joint inference problem over graph structure, intervention targets, and intervention-effect parameters, and reports that learning from multiple interventional datasets is inference over multiple mutilated graphs linked through an unobserved underlying graph. That is much closer to the typed-task setting than a pure structure posterior, because the system may be uncertain not only about whether an edge exists, but about **which mechanism a task perturbs** and whether the perturbation is hard or soft. ²³

A richer belief state should therefore maintain at least five coupled uncertainties: graph structure, functional mechanisms/effect sizes, intervention targets and intervention type, latent confounders or omitted variables, and regime/meta-state when structure can change over time. A concise way to write the target object is a posterior over a dynamic SCM family,

$$p(G_t, f_t, I_t, C_t, U_t \mid D),$$

together with posterior predictive answers to task-relevant interventional queries such as $p(Q_t \mid do(T_t = \tau), H_t)$ and counterfactual queries such as $p(Q_t^{T_t \leftarrow \tau'} \mid H_t)$. This is more useful for decision making than a posterior over graph topology alone. ²⁴

The decision-theoretic literature also shows that structural uncertainty matters only when it changes what the system should do. A 2025 analysis of structural uncertainty in causal decision making shows that Bayesian averaging over structures is most valuable when structural uncertainty is moderate to high, competing structures imply meaningfully different causal effects, and the loss function is sensitive to those differences. For an anchor-selection system, that means uncertainty should be propagated at least to the posterior over task values and intervention outcomes, not just stored as a background graph uncertainty score. ²⁵

Design blueprint for the typed-anchor system

A practical design consistent with the literature is to split the current temporal graph into a **causal state layer** and an **anchor/measurement layer**. The causal state layer contains task executions, latent or observed state factors, frame-quality variables, and regime/context variables with explicit structural assignments. The anchor layer stores what was pointed to, selected, or observed. Anchors should attach to variables or evidence statements; they should not themselves stand in for the causal claim. This separation is the smallest architectural change that preserves your current event-sourced logging while making SCM semantics possible. ²

A workable schema is:

- **State variables:** pre-task and post-task factors that can actually change. ²⁶
- **Task variables:** execution event, parameters, and intended targets. Model these as explicit treatments/interventions, not just events. ²⁷
- **Outcome variables:** frame quality, task success, downstream utility, and any validation signals. ²⁸
- **Selection variables:** why an anchor was selected, including policy state and attention or salience features. This is where “chosen because good” must be distinguished from “chosen because policy preferred it.” ²⁹
- **Regime variables:** context, domain, or latent meta-state for cases where the causal graph changes over time. ³⁰

For discovery and learning, a hybrid pipeline is more aligned with the literature than a single method. Use lag-expanded or DBN-style discovery for coarse temporal skeletons when variables are synchronized; use interventional discovery such as DCDI or GIES when task targets are known or inferable; use nonstationary methods such as CtrlINS or regime-conditioned graph models when structure drifts; and use event-sequence models such as CASCADE or point-process approaches when the log is asynchronous and sparse in continuous time. ³¹

The belief state should then expose causal queries directly. For example: “What is the posterior effect of executing task X on frame quality?” “What is the probability that anchor A would still have been selected if task X had not run?” “Under which latent regimes does task X improve quality?” This is the operational difference between a correlational temporal graph and a causal world model. ³²

Open questions and limitations

The main unresolved issue is intervention semantics. If a task record means “the agent attempted task X,” that is not automatically the same as “the mechanism targeted by X was successfully intervened upon.” Many real systems require explicit modeling of failed, partial, delayed, or unknown-target interventions. DCDI and BaCaDI support imperfect and unknown-target interventions, but deploying that idea in production depends on having reliable task metadata. ²⁷

The second unresolved issue is hidden confounding in adaptive execution logs. If task choice depends on hidden context, then historical execution traces remain observational and can bias both effect estimates and discovery. The sequential-decision literature is explicit that off-policy evaluation can be heavily biased under unobserved confounding and requires additional assumptions or robustness analysis. ³³

The third limitation is scale under nonstationarity. Recent work has improved scalability through sparsity, object factorization, and mixture-of-graphs representations, but the literature still lacks a broadly validated causal world model that is simultaneously large-scale, partially observed, asynchronous, and dynamically changing in structure. For systems like a typed-anchor temporal knowledge graph, some architectural choices will still need to be engineering judgments rather than consequences of a settled theory. ³⁴

3 12 33 <https://arxiv.org/abs/2106.14421>

<https://arxiv.org/abs/2106.14421>

5 16 26 https://proceedings.neurips.cc/paper_files/paper/2023/file/09ae6beae5f1ff38f05c05979097ea0f-Paper-Conference.pdf

https://proceedings.neurips.cc/paper_files/paper/2023/file/09ae6beae5f1ff38f05c05979097ea0f-Paper-Conference.pdf

7 <https://www.stat.cmu.edu/~brian/peem/brainstorming%20talk%20I%20gave%20in%20networkshop/PC%20algorithm/Sirtes-Glymour-Scheines%20Causation-Prediction-Search%20SPICPA-2v1.pdf>

<https://www.stat.cmu.edu/~brian/peem/brainstorming%20talk%20I%20gave%20in%20networkshop/PC%20algorithm/Sirtes-Glymour-Scheines%20Causation-Prediction-Search%20SPICPA-2v1.pdf>

8 <https://www.jmlr.org/papers/volume3/chickering02b/chickering02b.pdf>

<https://www.jmlr.org/papers/volume3/chickering02b/chickering02b.pdf>

9 <https://papers.neurips.cc/paper/8157-dags-with-no-tears-continuous-optimization-for-structure-learning.pdf>

<https://papers.neurips.cc/paper/8157-dags-with-no-tears-continuous-optimization-for-structure-learning.pdf>

10 27 https://proceedings.neurips.cc/paper_files/paper/2020/file/f8b7aa3a0d349d9562b424160ad18612-Paper.pdf

https://proceedings.neurips.cc/paper_files/paper/2020/file/f8b7aa3a0d349d9562b424160ad18612-Paper.pdf

11 https://proceedings.neurips.cc/paper_files/paper/2024/file/2b22bacd7ad8677f4837b28a11fe496f-Paper-Conference.pdf

https://proceedings.neurips.cc/paper_files/paper/2024/file/2b22bacd7ad8677f4837b28a11fe496f-Paper-Conference.pdf

14 15 <https://www.jmlr.org/papers/volume13/hauser12a/hauser12a.pdf>

<https://www.jmlr.org/papers/volume13/hauser12a/hauser12a.pdf>

17 https://proceedings.iclr.cc/paper_files/paper/2024/file/44a2b9f7bf9aec3f1fa333ad875b0ee0-Paper-Conference.pdf

https://proceedings.iclr.cc/paper_files/paper/2024/file/44a2b9f7bf9aec3f1fa333ad875b0ee0-Paper-Conference.pdf

18 <https://www.ijcai.org/proceedings/2023/0505.pdf>

<https://www.ijcai.org/proceedings/2023/0505.pdf>

19 30 https://proceedings.neurips.cc/paper_files/paper/2024/file/8cef4e4bcb85f7d4a1005a9db018d6b6-Paper-Conference.pdf

https://proceedings.neurips.cc/paper_files/paper/2024/file/8cef4e4bcb85f7d4a1005a9db018d6b6-Paper-Conference.pdf

20 <https://mn.cs.tsinghua.edu.cn/xinwang/PDF/papers/>

[2024_RealTCD%20Temporal%20Causal%20Discovery%20from%20Interventional%20Data%20with%20Large%20Language%20M](https://mn.cs.tsinghua.edu.cn/xinwang/PDF/papers/)

<https://mn.cs.tsinghua.edu.cn/xinwang/PDF/papers/>

[2024_RealTCD%20Temporal%20Causal%20Discovery%20from%20Interventional%20Data%20with%20Large%20Language%20Model.pdf](https://mn.cs.tsinghua.edu.cn/xinwang/PDF/papers/)

21 34 <https://arxiv.org/html/2505.16620v2>

<https://arxiv.org/html/2505.16620v2>

22 24 <https://openreview.net/pdf?id=MYzGsDPtHG>

<https://openreview.net/pdf?id=MYzGsDPtHG>

23 <https://proceedings.mlr.press/v206/hagele23a/hagele23a.pdf>

<https://proceedings.mlr.press/v206/hagele23a/hagele23a.pdf>

²⁵ <https://arxiv.org/pdf/2507.23495>

<https://arxiv.org/pdf/2507.23495>

²⁹ <https://proceedings.neurips.cc/paper/2020/file/da21bae82c02d1e2b8168d57cd3fbab7-Paper.pdf>

<https://proceedings.neurips.cc/paper/2020/file/da21bae82c02d1e2b8168d57cd3fbab7-Paper.pdf>

³¹ <https://proceedings.mlr.press/v108/pamfil20a/pamfil20a.pdf>

<https://proceedings.mlr.press/v108/pamfil20a/pamfil20a.pdf>