

DeepSWIP: Quotient-WMC Counterfactuals for Neural Probabilistic Logic Programs

Saimun Habib

University of Edinburgh

S.HABIB-1@SMS.ED.AC.UK

Vaishak Belle

University of Edinburgh

VBELLE@ED.AC.UK

Fengxiang He

University of Edinburgh

FHE@ED.AC.UK

Abstract

Neurosymbolic systems such as DeepProbLog combine neural perception with probabilistic logic, but standard inference is associational. Counterfactual reasoning additionally requires a causal semantics for interventions and evidence. We introduce DeepSWIP, a single-world counterfactual semantics for DeepProbLog programs. Using *neural materialization*, we reduce fixed-context neural predicates to ordinary ProbLog choices, apply Single World Intervention Programs (SWIPs), and compute counterfactuals by weighted model counting (WMC) over a single transformed program. Under finite grounding and unique-supported-model assumptions, DeepSWIP is exact relative to the learned materialized FCM. The standard quotient-WMC form of ProbLog conditionals identifies active neural probabilities and explains intervention cleaning, calibration sensitivity, and rare-evidence instability. Experiments on MPI3D confirm the transformation against a DeepTwin construction against 12,000 queries, as predicted and a $2.14\times$ inference speedup from avoiding the Twin’s endogenous duplication. A SUMO HOV experiment shows that neural calibration degradation biases plug-in estimates, while a correctly scoped randomized-policy AIPW estimator removes most first-order bias for population mean and ATE estimands. Code is at https://github.com/saibib/deep_SWIP

1. Introduction

In DeepProbLog, Scallop, and related probabilistic logic programming (PLP) systems, neural modules map raw inputs to distributions over symbolic atoms, and symbolic inference composes these uncertain atoms through structured logical rules (Manhaeve et al., 2018; Li et al., 2023). This architecture is effective for prediction under perceptual uncertainty, but prediction is not intervention. A model that answers $\mathbb{P}(Y \mid X = x)$ cannot answer what would have happened under $\text{do}(X = x)$, nor what would have happened under $\text{do}(X = x)$ given contradictory factual evidence.

Counterfactual reasoning is central to explanation, robustness, fairness, and decision making (Pearl, 2009; Halpern, 2016). In PLPs, counterfactuals have been given causal semantics by translating programs into functional causal models (FCMs) and evaluating counterfactual queries through Twin Networks (Kiesel et al., 2023). The Twin construction is semantically convenient. It duplicates the endogenous structure into factual and counterfactual copies while sharing exogenous choices. However, duplication is expensive for relational programs and awkward for neural predicates Habib et al. (2026). In a visual DeepProbLog program,

should the convolutional predicate be evaluated twice, once per world, or shared as an exogenous symbolic uncertainty? More importantly, duplication obscures the single-world semantics of interventions that motivates Single World Intervention Graphs (SWIGs) and creates unnecessary (and unverifiable) independence assumptions (Richardson and Robins, 2013, 2023; Habib et al., 2026).

The SWIP methodology of Habib et al. (2026) translates SWIG-style intervention surgery into ProbLog program rewriting. The methodology defines causal semantics for removing clauses defining intervened atoms, redirecting downstream uses to fixed intervention atoms, and querying the resulting single-world program to produce programs equivalent to the interventional and counterfactual distribution specified by the FCM. DeepSWIP extends this idea to DeepProbLog. Once neural predicates are materialized into probabilistic choices, counterfactual inference becomes a quotient of weighted model counts,

$$Q(p) = \mathbb{P}(Y_x = y \mid E = e) = \frac{\Phi_{Y,E}^x(p)}{\Phi_E^x(p)}, \quad (1)$$

where p are materialized neural probabilities and Φ are multilinear WMC polynomials. This algebraic representation is the paper’s organizing principle and explains why intervened mechanisms have zero derivative, why calibration error rather than classification accuracy is the relevant neural failure mode, and why rare evidence causes structural instability through the denominator Φ_E^x .

Empirically, we evaluate DeepSWIP on two complementary settings. MPI3D tests the full neurosymbolic pipeline with neural predicates, symbolic rules, exact factor-level counterfactual ground truth, and a Twin-style comparison. The SUMO traffic experiments tests the statistical side of the framework with imperfect neural traffic-state estimates, rare-evidence amplification, and the distinction between individual plug-in counterfactual prediction and population-level AIPW/DML estimation.

We make four contributions: (i) neural materialization for fixed-context DeepProbLog counterfactuals; (ii) a correctness result lifting SWIP to the learned materialized FCM; (iii) a quotient-WMC analysis of active neural predicates, intervention cleaning, calibration sensitivity, and rare-evidence instability; and (iv) MPI3D and SUMO experiments separating exact symbolic inference from statistical error in learned predicates.

2. Background and Related Work

Probabilistic logic and counterfactuals. A ground ProbLog program consists of probabilistic facts and deterministic rules. Under distribution semantics, probabilistic facts define independent exogenous choices and the logic program maps each exogenous assignment to a unique symbolic world, provided the program has unique supported models (De Raedt et al., 2007). Kiesel et al. (2023) show how to interpret such programs as FCMs and answer counterfactual queries by a Twin Network transformation, proving consistency with Pearl-style counterfactuals and CP-logic (Vennekens et al., 2009). Habib et al. (2026) instead proposes single-world intervention programs, following the SWIG idea that an intervention splits a node into a random factual part and a fixed intervention part (Richardson and Robins, 2013). Our work assumes these ordinary-ProbLog results and asks how they extend when some exogenous symbolic choices are produced by neural predicates.

Neural probabilistic logic and statistical correction. DeepProbLog introduces neural predicates of the form $\text{nn}(m, x, y)$, where a neural model m maps input x to a distribution over symbolic values y (Manhaeve et al., 2018). Inference is typically reduced to arithmetic circuits or WMC after evaluating neural outputs. Scallop and related systems similarly reduce neural-symbolic inference to provenance circuits or weighted logical inference (Li et al., 2023; Valentin et al., 2025). This circuit view is unified by algebraic model counting, which casts WMC, gradients, and related tasks as semiring evaluations over a shared circuit (Kimmig et al., 2017; Maene and De Raedt, 2025); we use this perspective in Section 5. DeepSWIP is compatible with such solvers but our contribution does not focus on them. Instead, we treat neural predicates as probability estimators inside a causal logic program and this separation lets us analyze how neural nuisance error propagates through symbolic counterfactuals. To create causal estimands with desirable properties, we implement Double/debiased machine learning (DML) (Chernozhukov et al., 2018; Robins et al., 1994). We do not claim universal DML for arbitrary logic queries but we prove a plug-in sensitivity result for general DeepSWIP queries and use AIPW only in experiments where the target is a randomized binary-treatment population estimand.

3. Preliminaries

Let $\mathcal{P} = (\mathcal{F}, \mathcal{R}, \mathcal{N})$ be a finite ground DeepProbLog program, where \mathcal{F} contains ordinary probabilistic choices, \mathcal{R} deterministic rules, and \mathcal{N} neural predicates. Let \mathcal{B} be the finite Herbrand base induced by the query, evidence, intervention, and neural contexts under consideration. We write $\mathcal{E} \subseteq \mathcal{B}$ for the external atoms; ordinary probabilistic choices from \mathcal{F} together with the categorical choices obtained by neural-materializing predicates in \mathcal{N} . The internal atoms are $\mathcal{I} = \mathcal{B} \setminus \mathcal{E}$ and are defined by the rules in \mathcal{R} . For an assignment $u \in \{0, 1\}^{\mathcal{E}}$ of external atoms, write

$$\text{SM}_{\mathcal{R}}(u) = \{v \in \{0, 1\}^{\mathcal{I}} : u \cup v \text{ is a supported model of } \mathcal{R}\}.$$

For an intervention $X = x$, let \mathcal{R}_x denote the rule set obtained by deleting all rules whose head is in X and adding the corresponding constant assignments $X := x$.

We assume for the query, evidence, intervention, and neural contexts under consideration, the grounding of \mathcal{P} is finite. Moreover, each neural predicate $k \in \mathcal{N}$ has a finite output domain $\mathcal{Y}_k = \{y_{k,1}, \dots, y_{k,r_k}\}$, and for fixed parameters η_k and context c_k returns a categorical distribution

$$p_{k,j}(\eta_k, c_k) = \mathbb{P}_{\eta_k}(Y_k = y_{k,j} \mid c_k), \quad \sum_{j=1}^{r_k} p_{k,j}(\eta_k, c_k) = 1. \quad (2)$$

We write $p(\eta, c)$ for all ordinary and neural-materialized weights. We also require $\mathbb{P}_{M_{\eta,c}}(E = e) > 0$ for every conditional query $\mathbb{P}(Y_x = y \mid E = e)$, so that the WMC quotient is defined.

The semantics of \mathcal{P} define a probability distribution over the possible worlds of the assignments of probabilistic facts and neural-materialized choices as mutually independent external choices. Thus, for fixed (η, c) , every assignment $u \in \{0, 1\}^{\mathcal{E}}$ has probability

$$\mathbb{P}_{\eta,c}(u) = \prod_{a \in u} p_a(\eta, c) \prod_{a \in \mathcal{E} \setminus u} (1 - p_a(\eta, c)),$$

with annotated-disjunction choices interpreted as one-hot categorical choices.

Finally, we assume unique supported models before and after intervention:

$$|\text{SM}_{\mathcal{R}}(u)| = 1, \quad |\text{SM}_{\mathcal{R}_x}(u)| = 1 \quad \text{for all } u \in \{0, 1\}^{\mathcal{E}}.$$

Equivalently, the materialized program induces a well-defined FCM. Acyclic rule dependency is sufficient but not necessary.

4. The DeepSWIP Transformation

The purpose of DeepSWIP is to separate neural perception from symbolic counterfactual inference. A DeepProbLog neural predicate is not intervened on as a neural network, rather, for fixed parameters and fixed input context, it induces a finite distribution over symbolic values. We therefore freeze the neural computation and expose only the induced symbolic uncertainty to the logic program. This turns a neural probabilistic logic program into an ordinary ProbLog program, after which standard causal program transformations can be applied.

Definition 1 (Neural materialization) *Fix neural parameters η and contexts c . The materialized program $\mathcal{P}_{\text{mat}}(\eta, c)$ is the ordinary ProbLog program obtained by replacing each neural predicate $\text{nn}(m_k, c_k, Y_k)$ by the categorical probabilistic choice*

$$p_{k,1} :: a_k(y_{k,1}); \dots ; p_{k,r_k} :: a_k(y_{k,r_k}),$$

where $p_{k,j} = p_{k,j}(\eta_k, c_k)$. The deterministic rule set \mathcal{R} is unchanged.

Theorem 2 (Materialization equivalence) *Fix η and c . For every ground formula φ over the symbolic atoms of \mathcal{P} ,*

$$\mathbb{P}_{\mathcal{P}}(\varphi \mid \eta, c) = \mathbb{P}_{\mathcal{P}_{\text{mat}}(\eta, c)}(\varphi).$$

Materialization gives the object on which intervention surgery is well-defined. The intervention itself is a program transformation that removes the mechanisms defining intervened atoms, redirects downstream rules to fixed intervention atoms, and leaves all non-intervened mechanisms unchanged. Thus, DeepSWIP inherits the single-world interpretation of SWIGs while retaining the ordinary distribution semantics of the materialized ProbLog program.

Theorem 3 (DeepSWIP correctness against FCM) *Assume $\mathcal{P}_{\text{mat}}(\eta, c)$ has unique supported models, and let $M_{\eta, c}$ be its induced FCM. Let $\mathcal{S}_x(\mathcal{P}_{\text{mat}})$ be the SWIP transformation under $\text{do}(X = x)$. For any query $Y = y$ and evidence $E = e$ with $\mathbb{P}_{M_{\eta, c}}(E = e) > 0$,*

$$\mathbb{P}_{\text{DeepSWIP}}(Y_x = y \mid E = e; \eta, c) = \mathbb{P}_{M_{\eta, c}}(Y_x = y \mid E = e).$$

The theorem is conditional on the learned/materialized FCM. It says that symbolic counterfactual inference is exact given the materialized neural probabilities. It does not say that η is the true data-generating parameter. Algorithm 1 details the implementation of the program transformation and counterfactual querying.

We give a minimal example that illustrates the complete DeepSWIP pipeline. Consider a DeepProbLog traffic-monitoring fragment in which a neural predicate classifies an upstream sensor image into a latent traffic state, and symbolic rules derive congestion and delay:

Algorithm 1 DeepSWIP transformation

Input: Ground DeepProbLog program $P = (F, R, N)$; neural parameters η ; contexts c ; intervention $do(X = x)$; evidence $E = e$; query $Y = y$

Output: Counterfactual probability $P(Y_x = y \mid E = e; \eta, c)$

$P_{\text{mat}} \leftarrow F \cup R$ **foreach** $nn(m_k, c_k, Y_k) \in N$ **do**
 | Evaluate $m_k(c_k)$ and obtain $p_{k,j} = P_{\eta_k}(Y_k = y_{k,j} \mid c_k)$ for $j = 1, \dots, r_k$ Replace $nn(m_k, c_k, Y_k)$ in P_{mat} by
 | $p_{k,1} :: a_k(y_{k,1}); \dots; p_{k,r_k} :: a_k(y_{k,r_k})$
end

$S_x \leftarrow$ probabilistic choices of P_{mat} **foreach** clause $C = (h \leftarrow b_1, \dots, b_\ell)$ in R **do**
 | **if** h defines an intervened atom in X **then**
 | | discard C
 | **else**
 | | replace each body occurrence of X_i by $X_{i,\text{fix}}(x_i)$ add the rewritten clause to S_x
 | **end**
end

foreach intervention assignment $X_i = x_i$ **do**
 | add the deterministic fixed assertion $X_{i,\text{fix}}(x_i)$ to S_x
end

Let Δ_x be the Boolean theory induced by the transformed program $\Phi_{Y,E}^x(p) \leftarrow \text{WMC}(\Delta_x \wedge Y = y \wedge E = e; p)$
 $\Phi_E^x(p) \leftarrow \text{WMC}(\Delta_x \wedge E = e; p)$ **return** $\Phi_{Y,E}^x(p) / \Phi_E^x(p)$

```

nn(state_net, img42, [free,queued])::traffic_state(S).

0.20::rain.
0.10::roadworks.

incident :- rain, roadworks.
congested :- traffic_state(queued).
congested :- incident.
delayed  :- congested.
reroute  :- incident.
    
```

For fixed neural parameters and input image `img42`, neural materialization evaluates the network once and replaces the neural predicate with its softmax output as an annotated disjunction. For example, if the network returns probability 0.30 for `free` and 0.70 for `queued`, the materialized program is now an ordinary ProbLog program containing `0.30::traffic_state(free); 0.70::traffic_state(queued)`.

Now, suppose the factual evidence is that congestion was observed while it was raining, but the counterfactual query asks what the delay would have been under $do(\text{congested} = \text{false})$. The counterfactual program is now:

```

0.30::traffic_state(free); 0.70::traffic_state(queued).

0.20::rain.
0.10::roadworks.

incident :- rain, roadworks.

congested_obs :- traffic_state(queued).
congested_obs :- incident.

0.0::congested_fix_true.
delayed :- congested_fix_true.
    
```

```

reroute :- incident.

evidence(congested_obs, true).
evidence(rain, true).
query(delayed).

```

The causal pathways between `traffic_state(queued) → congested_obs` and `incident → congested_obs` remain in the denominator of the WMC quotient. They perform the abduction step by reweighting the exogenous explanations of the factual congestion, including the neural-materialized traffic-state choice. Only the outgoing use of `congested` in the descendant rule for `delayed` is redirected to the fixed false intervention.

The counterfactual probability is then computed as the corresponding WMC quotient with numerator additionally conjoining `delayed` and denominator conditioning on `congested_obs` and `rain`. As desired, the query evaluates to zero because `delayed` cannot be derived through the intervened congestion mechanism, even though factual evidence says congestion occurred.

By contrast, $\mathbb{P}_{\text{DeepSWIP}}(\text{reroute} \mid \text{congested}, \text{do}(\text{congested} = \text{false}))$ is still governed by `reroute :- incident`. Thus DeepSWIP does not erase factual causes of the intervened variable, instead it only severs the intervened atom’s defining equations and redirects downstream uses to the fixed value. In general, we can bound how much the materialized program changes.

Proposition 4 (Program growth) *Let \mathcal{P}_{mat} have m probabilistic choices and $|\mathcal{R}|$ deterministic clauses. For an intervention on k ground atoms, $\mathcal{S}_x(\mathcal{P}_{\text{mat}})$ contains at most $m + k$ probabilistic choices if fixed atoms are encoded as degenerate facts, and m otherwise. Its number of clauses is at most $|\mathcal{R}| + k$ before pruning.*

Proposition 4 is a syntactic statement. WMC can still be exponential in circuit size or treewidth.

5. Quotient-WMC Semantics and Consequences

The quotient (4) is an instance of *algebraic model counting* (AMC), which generalizes weighted model counting from the probability semiring to arbitrary commutative semirings, unifying a range of inference tasks as one circuit evaluation (Kimmig et al., 2017). Both $\Phi_{Y,E}^x$ and Φ_E^x are AMC evaluations over the probability semiring on the single SWIP-transformed circuit Δ^x , so the counterfactual conditional is a ratio of two AMC values on the same circuit. This view is more than terminological. The local sensitivities of Theorems 8 and 9 are exactly the quantities computed by AMC over the *gradient semiring*, the construction used to train neural predicates in DeepProbLog, and recently shown to unify gradient and model-counting computations in a single algebraic framework (Maene and De Raedt, 2025).

Consequently ∇Q is not only characterized analytically by (6) but is computable by one gradient-semiring pass over Δ^x . Intervention cleaning (Theorem 6) is likewise algebraic with SWIP removing the intervened literals from the labeling, so their semiring weights never enter the count, and the active neural set is the support of the labeling on the pruned circuit. We read the results below as a quotient-AMC account of neurosymbolic counterfactuals and the causal SWIP transformation and the counterfactual reading of the quotient are our additions on top of the AMC machinery.

Let Δ_x be the Boolean theory induced by the SWIP-transformed materialized program. For a world ω over probabilistic choices p_1, \dots, p_m , define

$$w_p(\omega) = \prod_{i:\omega_i=1} p_i \prod_{i:\omega_i=0} (1 - p_i), \quad \text{WMC}(\Gamma; p) = \sum_{\omega \models \Gamma} w_p(\omega). \quad (3)$$

The previous section establishes that DeepSWIP produces the correct transformed program. After materialization and intervention surgery, neural dependence enters only through the weights p , while the logic determines which assignments satisfy query and evidence. Thus a counterfactual conditional is a quotient of two weighted model counts. This quotient is the central algebraic object of the paper.

Theorem 5 (Quotient-polynomial WMC semantics) *For any query $Y = y$ and evidence $E = e$ with $\text{WMC}(\Delta_x \wedge E = e; p) > 0$,*

$$Q(p) = \mathbb{P}_{\Delta_x}(Y = y \mid E = e) = \frac{\Phi_{Y,E}^x(p)}{\Phi_E^x(p)}, \quad (4)$$

where

$$\Phi_{Y,E}^x(p) = \text{WMC}(\Delta_x \wedge Y = y \wedge E = e; p), \quad \Phi_E^x(p) = \text{WMC}(\Delta_x \wedge E = e; p).$$

Both $\Phi_{Y,E}^x$ and Φ_E^x are multilinear polynomials in the materialized probabilities.

The remaining results are consequences of this quotient representation. The transformation changes which weights appear in the numerator and denominator and this gives a clean lens to view intervention, pruning, calibration, and rare event stability through.

Theorem 6 (Intervention cleaning) *If a materialized probability p_k appears only in clauses defining an intervened atom removed by SWIP, then Q does not depend on p_k . Hence $\partial Q / \partial p_k = 0$ wherever the derivative is defined.*

This formalizes the intuition that intervened mechanisms are no longer active sources of uncertainty for the counterfactual query. Neural errors in a mechanism that has been surgically replaced cannot affect the resulting quotient. In effect, we should only consider the active neural set.

Proposition 7 (Active neural set) *Let $Q(p) = \Phi_{Y,E}^x(p) / \Phi_E^x(p)$ be the DeepSWIP quotient. If a materialized neural probability p_k does not occur in either $\Phi_{Y,E}^x$ or Φ_E^x after SWIP and relevance pruning, then*

$$\partial Q / \partial p_k = 0.$$

Thus each query induces an active set of neural probabilities and only probabilities that survive intervention surgery and appear in the numerator or denominator can influence the answer. It follows that the query is sensitive to the materialized neural probabilities and the logical rules and we can decompose this sensitivity accordingly

Theorem 8 (Local calibration sensitivity) *Let p^* be ideal materialized probabilities and assume $\Phi_E^x(p^*) > 0$. For $\hat{p} = p^* + \varepsilon$ sufficiently close to p^* ,*

$$Q(\hat{p}) - Q(p^*) = \nabla Q(p^*)^\top \varepsilon + O(\|\varepsilon\|^2), \quad (5)$$

where for each active coordinate

$$\frac{\partial Q}{\partial p_k}(p) = \frac{1}{\Phi_E^x(p)} \left[\frac{\partial \Phi_{Y,E}^x}{\partial p_k}(p) - Q(p) \frac{\partial \Phi_E^x}{\partial p_k}(p) \right]. \quad (6)$$

The leading error term is a product of neural probability error and logical sensitivity. This is why calibration, rather than classification accuracy alone, is the relevant neural property for counterfactual reliability. A classifier may have high accuracy but poor probability calibration and such calibration error can still be amplified by the symbolic quotient. This also reveals a structural instability: when the evidence probability is small, small neural calibration errors may induce large counterfactual errors.

Theorem 9 (Rare-evidence instability) *Let $\delta = \Phi_E^x(p) > 0$. If $|\partial \Phi_{Y,E}^x / \partial p_k| \leq C_1$ and $|\partial \Phi_E^x / \partial p_k| \leq C_2$ in a neighborhood of p , then*

$$\left| \frac{\partial Q}{\partial p_k}(p) \right| \leq \frac{C_1 + C_2}{\delta}.$$

Thus local sensitivity can grow as $1/\mathbb{P}(E = e)$ and diverge.

Rare-evidence instability is a property of the counterfactual functional itself. It is not caused by a particular neural architecture or optimizer. Better neural training can reduce $\|\hat{p} - p^*\|$, but it cannot remove the denominator $\Phi_E^x(p)$ from the counterfactual query. This motivates the use of double machine learning in the context of neurosymbolic counterfactuals.

Proposition 10 (Cross-fitting and scoped DML) *Let A be the active set after intervention cleaning and pruning. For any plug-in estimator \bar{p} of p^* ,*

$$\mathbb{E}[Q(\bar{p}) - Q(p^*)] = \nabla_A Q(p^*)^\top \mathbb{E}[\bar{p}_A - p_A^*] + O(\mathbb{E}\|\bar{p}_A - p_A^*\|^2).$$

Therefore calibration or cross-fitting can reduce first-order plug-in bias only to the extent that it reduces active probability bias. This is not a universal DML guarantee. AIPW/DML requires a target estimand, valid treatment/propensity structure, and a Neyman-orthogonal score. The proposition separates two claims that should not be conflated.

Full proofs are provided in Appendix A. The main point for experiments is that DeepSWIP has two separable error sources: exact symbolic inference relative to p , and statistical/perceptual error in estimating p .

6. Experiments

Our experiments are designed around the theory rather than breadth of applications. MPI3D tests visual neural materialization, the speedup from the single-world transformation, pruning, and intervention cleaning, using DeepSWIP–DeepTwin agreement as an implementation correctness check and factor labels as counterfactual ground truth. SUMO HOV tests calibration-sensitive plug-in error and scoped AIPW correction.

6.1. MPI3D visual counterfactuals

Setup. MPI3D-toy is factorized by color, shape, size, camera, background, and pose (Gondal et al., 2019). We train neural predicates for shape, size, and color, materialize and use symbolic rules for `can_roll`, `can_stack`, `stable`, and `risky_on_shelf`. Counterfactuals intervene on shape or size and compare DeepSWIP with a DeepTwin construction over the same materialized program. Factor labels supply ground-truth potential outcomes.

Results. With 8,368 training and 2,000 test images, neural predicate accuracies were 0.889 for shape, 0.987 for size, and 1.000 for color. Table 1 reports across 500 evaluation images, six shape interventions, and four queries (12,000 ProbLog calls), DeepSWIP and DeepTwin agreed on 100% of queries. DeepSWIP averaged 3.57 ms/query and DeepTwin 7.65 ms/query, a 2.14 \times speedup. Ground-truth RMSE was zero for intervention-determined queries and small for queries depending on non-intervened predicates: 0.058 for `can_stack` and 0.082 for `risky_on_shelf`. Pruning under `do(shape=sphere)` reduced statement counts from 18 to 5, 8, 11, and 7 for the four queries, pruning color in all cases.

DeepSWIP and DeepTwin are both exact relative to the materialized FCM (Theorem 3), so we use their agreement as a correctness check on the transformation rather than as evidence for the semantics. Across 500 evaluation images, six shape interventions, and four queries (12,000 ProbLog calls) they agreed on 100% of queries, exactly as the equivalence theorem predicts, and any disagreement would have indicated a transformation or implementation bug. The substantive results are the speedup and the propagation of neural error. Avoiding the Twin’s endogenous duplication (Theorem 4) reduced inference from 7.65 to 3.57 ms/query, a 2.14 \times speedup on this program; ?? shows the ratio widening with the number of endogenous groundings, as the duplication argument predicts. Because DeepSWIP is exact, the only error against factor ground truth comes from neural calibration on the *active* predicates: RMSE was zero for intervention-determined queries (`can_roll`, `stable`), empirically exhibiting intervention cleaning, and small for queries depending on non-intervened predicates (0.058 for `can_stack`, 0.082 for `risky_on_shelf`). Pruning under `do(shape=sphere)` reduced statement counts from 18 to 5, 8, 11, and 7, removing color in all cases. Neural predicate accuracies were 0.889 (shape), 0.987 (size), and 1.000 (color).

Table 1: MPI3D visual counterfactuals.

Component	Metric	Value	Comparison	Value
Shape predicate	Accuracy	0.889	Calibration gap	-0.046
Size predicate	Accuracy	0.987	Calibration gap	-0.026
Colour predicate	Accuracy	1.000	Calibration gap	-0.001
DeepSWIP vs DeepTwin	Agreement	100.0%	Queries	12,000
Inference time	DeepSWIP	3.57 ms	DeepTwin	7.65 ms
Runtime ratio	Twin/SWIP	2.14 \times		
<code>can_roll</code>	RMSE	0.000	Bias	0.000
<code>can_stack</code>	RMSE	0.058	Bias	0.001
<code>risky_on_shelf</code>	RMSE	0.082	Bias	-0.003
<code>stable</code>	RMSE	0.000	Bias	0.000

6.2. HOV randomized-policy calibration and DML stress test

Setup. The HOV experiment uses paired SUMO simulations of restricted and open-lane policies, yielding potential outcomes $Y_i(0)$ and $Y_i(1)$ for throughput (Lopez et al., 2018). We construct a randomized-policy observation process with $T_i \sim \text{Bernoulli}(0.5)$ and target $\psi_1 = \mathbb{E}[Y(1)]$ and $\tau = \mathbb{E}[Y(1) - Y(0)]$. A neural traffic-state predicate is trained under increasing sensor noise. The plug-in DeepSWIP nuisance model maps state probabilities to state-conditional throughput means, while AIPW uses known propensity $g = 0.5$ and fold-specific observed-only nuisance estimates. AIPW is evaluated as a population estimator, not an individual counterfactual RMSE.

The HOV experiment uses paired SUMO simulations of restricted and open-lane policies, yielding potential outcomes $Y_i(0)$ and $Y_i(1)$ for throughput (Lopez et al., 2018). To test valid AIPW/DML, we construct a randomized-policy observation process: $T_i \sim \text{Bernoulli}(0.5)$ and $Y_i = T_i Y_i(1) + (1 - T_i) Y_i(0)$. The target estimands are $\psi_1 = \mathbb{E}[Y(1)]$ and $\tau = \mathbb{E}[Y(1) - Y(0)]$. A neural traffic-state predicate is trained under increasing sensor noise. The plug-in DeepSWIP nuisance model maps state probabilities to state-conditional throughput means. AIPW is evaluated as a population estimator, not an individual counterfactual RMSE.

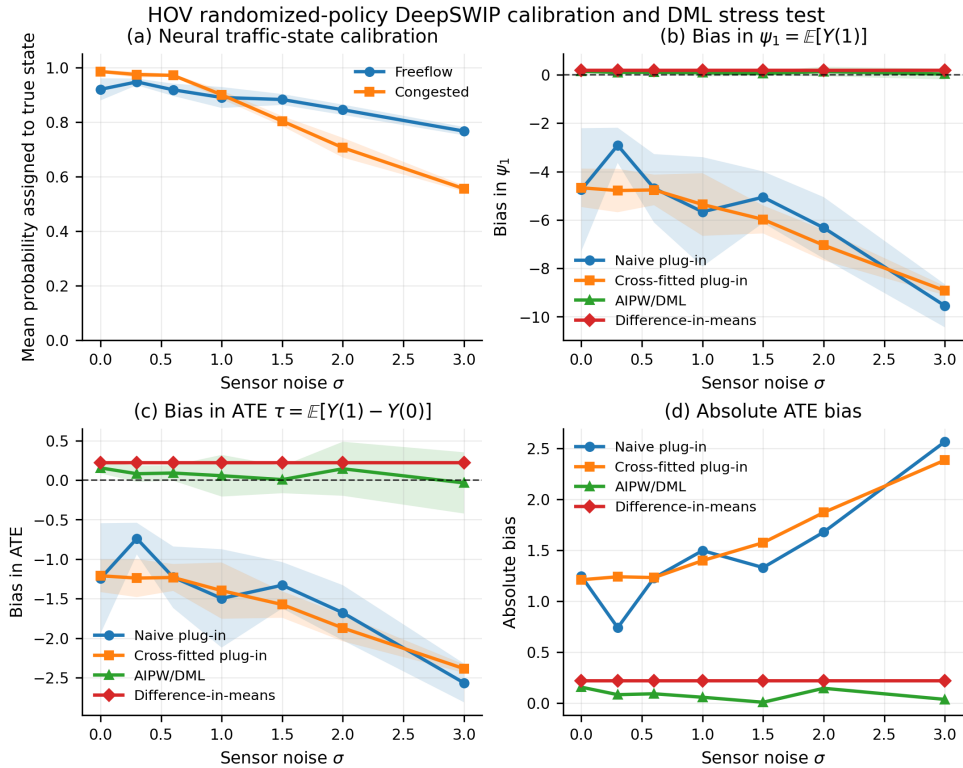


Figure 1: **HOV randomized-policy calibration and DML stress test.**

Results. The dataset has $N = 5000$ episodes: 4247 freeflow and 753 congested. The true values are $\mathbb{E}[Y(1)] = 248.221$, $\mathbb{E}[Y(0)] = 184.893$, and $\tau = 63.329$. Figure 1 shows, as sensor noise increases, congested calibration falls from 0.985 to 0.556. Naive plug-in bias worsens for both ψ_1 (-4.75 to -9.55) and ATE (-1.24 to -2.57), while cross-fitted plug-in remains

biased. AIPW/DML stays near zero bias across noise levels, as does difference-in-means under randomized treatment.

7. Discussion and Limitations

DeepSWIP’s exactness is exactness relative to the learned materialized FCM. If a neural predicate is miscalibrated, the symbolic counterfactual can be exactly wrong with respect to the external data-generating process. This is not a defect of SWIP; it is the distinction between causal-symbolic inference and statistical estimation. The quotient derivative in [Theorem 8](#) makes this distinction apparent as only active, non-intervened neural probabilities matter, and their errors are amplified by the logic and by $1/\Phi_E$.

The DML result is deliberately scoped. The HOV AIPW experiment is valid because we construct a randomized binary treatment with known propensity and evaluate population means/ATEs. This does not imply that arbitrary DeepSWIP counterfactual queries have an orthogonal score, nor that AIPW estimates individual counterfactual outcomes. For general path-specific, continuous, or relational counterfactuals, deriving the efficient influence function is a separate statistical problem. Similarly, approximate WMC methods may accelerate inference, but approximation error in a quotient requires denominator-aware control; we leave approximate counterfactual WMC to future work.

Finally, our MPI3D rules are intentionally simple. The experiment is not a claim about state-of-the-art vision or physics. Its role is to breakdown the NeSy counterfactuals as neural predicates are materialized, symbolic interventions are performed in a single world, and WMC computes the resulting quotient exactly.

8. Conclusion

We presented DeepSWIP, a framework for counterfactual reasoning in neural probabilistic logic programs. The method evaluates neural predicates once, materializes them as ordinary probabilistic choices, applies a single-world intervention transformation, and computes the result by WMC. This yields exact counterfactuals relative to the learned materialized FCM and exposes the counterfactual query as a quotient of WMC polynomials. The quotient view provides a common explanation for program growth, pruning, intervention cleaning, calibration sensitivity, rare-evidence instability, and the limited role of DML. Experiments on MPI3D and SUMO support these claims while keeping the statistical scope explicit. DeepSWIP is therefore best understood not as a universal causal discovery method, but as a precise counterfactual inference layer for NeSy programs with explicit causal structure and learned probabilistic mechanisms.

Acknowledgments

Funding for this research was provided by NERC through an E4 DTP studentship (NE/S007407/1).

References

- Victor Chernozhukov, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, Whitney Newey, and James Robins. Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal*, 21(1):C1–C68, 2018.
- Luc De Raedt, Angelika Kimmig, and Hannu Toivonen. Problog: A probabilistic prolog and its application in link discovery. *Proceedings of the International Joint Conference on Artificial Intelligence*, 2007.
- Muhammad Waleed Gondal, Manuel Wuthrich, Djordje Miladinovic, et al. On the transfer of inductive bias from simulation to the real world: A new disentanglement dataset. In *Advances in Neural Information Processing Systems*, 2019.
- Saimun Habib, Vaishak Belle, and Fengxiang He. Efficient counterfactual reasoning in problog via single world intervention programs. *arXiv preprint arXiv:2603.20505*, 2026.
- Joseph Y. Halpern. *Actual Causality*. MIT Press, 2016.
- Rafael Kiesel, Kilian Ruckschloss, and Felix Weitkamper. “what if?” in probabilistic logic programming. *arXiv preprint arXiv:2305.15318*, 2023.
- Angelika Kimmig, Guy Van den Broeck, and Luc De Raedt. Algebraic model counting. *Journal of Applied Logic*, 22:46–62, 2017.
- Jiani Li, Ruitao Zhu, Yewen Cheng, et al. Scallop: A language for neurosymbolic programming. In *Proceedings of the ACM SIGPLAN Conference on Programming Language Design and Implementation*, 2023.
- Pablo Alvarez Lopez, Michael Behrisch, Laura Bieker-Walz, et al. Microscopic traffic simulation using sumo. In *IEEE International Conference on Intelligent Transportation Systems*, 2018.
- Jaron Maene and Luc De Raedt. The gradient of algebraic model counting. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2025. arXiv:2502.18406.
- Robin Manhaeve, Sebastijan Dumančić, Angelika Kimmig, Thomas Demeester, and Luc De Raedt. Deepproblog: Neural probabilistic logic programming. In *Advances in Neural Information Processing Systems*, 2018.
- Judea Pearl. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, 2 edition, 2009.
- Thomas S. Richardson and James M. Robins. Single world intervention graphs: A unification of the counterfactual and graphical approaches to causality. *Working paper*, 2013.
- Thomas S. Richardson and James M. Robins. Single world intervention graphs. In *Probabilistic and Causal Inference: The Works of Judea Pearl*. ACM Books, 2023.

James M. Robins, Andrea Rotnitzky, and Lue Ping Zhao. Estimation of regression coefficients when some regressors are not always observed. *Journal of the American Statistical Association*, 89(427):846–866, 1994.

Thomas Jean-Michel Valentin, Luisa Sophie Werner, Pierre Genevès, and Nabil Layaïda. On scaling neurosymbolic programming through guided logical inference. In *arXiv preprint arXiv:2501.18202*, 2025.

Joost Vennekens, Marc Denecker, and Maurice Bruynooghe. Cp-logic: A language of causal probabilistic events and its relation to logic programming. *Theory and Practice of Logic Programming*, 9(3):245–308, 2009.

Appendix A. Proofs

A.1. Proof of [Theorem 2](#)

Fix η and c . Each neural predicate k induces a categorical distribution over a finite symbolic choice $U_k \in \mathcal{Y}_k$ with probabilities $p_{k,j}(\eta_k, c_k)$. Materialization replaces this random choice by an ordinary probabilistic choice with exactly the same probabilities. Ordinary probabilistic facts are unchanged. Therefore every complete assignment u to all ordinary and neural-materialized choices has the same probability before and after materialization:

$$\mathbb{P}_{\mathcal{P}}(U = u \mid \eta, c) = \mathbb{P}_{\mathcal{P}_{\text{mat}}(\eta, c)}(U = u).$$

For a fixed u , the deterministic rules are identical in \mathcal{P} and \mathcal{P}_{mat} . By unique supported models, the closure of the rules under u is unique and the same in both programs. Let Ω_φ be the set of assignments whose closure satisfies φ . Then

$$\mathbb{P}_{\mathcal{P}}(\varphi \mid \eta, c) = \sum_{u \in \Omega_\varphi} \mathbb{P}_{\mathcal{P}}(u \mid \eta, c) = \sum_{u \in \Omega_\varphi} \mathbb{P}_{\mathcal{P}_{\text{mat}}}(u) = \mathbb{P}_{\mathcal{P}_{\text{mat}}}(\varphi).$$

A.2. Proof of [Theorem 3](#)

By [Theorem 2](#), the fixed-parameter DeepProbLog program and \mathcal{P}_{mat} induce the same distribution over symbolic worlds. Since \mathcal{P}_{mat} has unique supported models, it induces an FCM $M_{\eta, c}$ whose exogenous variables are the probabilistic choices and whose endogenous structural assignments are the deterministic closures of the rules. The SWIP transformation removes clauses defining X , redirects downstream body literals to fixed atoms, and asserts the intervention values. This is exactly the program-level implementation of replacing the structural assignments for X by constants while leaving other mechanisms unchanged. The ordinary-ProbLog SWIP correctness theorem of [Habib et al. \(2026\)](#), applied to \mathcal{P}_{mat} , gives equality between the probability computed on $\mathcal{S}_x(\mathcal{P}_{\text{mat}})$ and the FCM counterfactual in $M_{\eta, c}$. The condition $\mathbb{P}(E = e) > 0$ ensures the conditional probability is defined.

A.3. Proof of Theorem 4

Materialization yields m probabilistic choices. SWIP does not duplicate any retained rule. For each intervened atom, all clauses defining that atom are removed; all remaining clauses are either unchanged or have body literals renamed to fixed atoms. This operation cannot increase the number of retained original clauses. SWIP then adds at most one fixed assertion per intervened atom. If these assertions are deterministic facts they add no probabilistic choices; if encoded as degenerate probabilistic facts they add at most k . Hence the bounds $m+k$ and $|\mathcal{R}|+k$. A Twin construction creates factual and counterfactual endogenous copies, giving two copies of each endogenous rule before intervention-specific deletions. Exogenous choices may be shared, but endogenous symbolic structure is duplicated.

A.4. Proof of Theorem 5

Conditional probability gives

$$\mathbb{P}_{\Delta_x}(Y = y \mid E = e) = \frac{\mathbb{P}_{\Delta_x}(Y = y, E = e)}{\mathbb{P}_{\Delta_x}(E = e)}.$$

Under WMC, these probabilities are exactly $\text{WMC}(\Delta_x \wedge Y = y \wedge E = e; p)$ and $\text{WMC}(\Delta_x \wedge E = e; p)$. For any formula Γ ,

$$\text{WMC}(\Gamma; p) = \sum_{\omega \models \Gamma} \prod_{i: \omega_i=1} p_i \prod_{i: \omega_i=0} (1 - p_i).$$

Each summand contains each p_i at most once, either as p_i or $1 - p_i$. Expanding yields a multilinear polynomial. Thus both numerator and denominator are multilinear polynomials and the query is their quotient.

A.5. Proof of Theorem 6

If p_k appears only in clauses defining an intervened atom, those clauses are removed by SWIP. Therefore the transformed theory Δ_x contains no dependence of Y or E on the Boolean choice weighted by p_k . In the WMC sum, assignments to that variable factor as $p_k + (1 - p_k) = 1$ in both numerator and denominator. Thus neither polynomial depends on p_k , and the quotient derivative is zero.

A.6. Proof of Theorem 8

By Theorem 5, $Q(p) = N(p)/D(p)$ with $N = \Phi_{Y,E}^x$ and $D = \Phi_E^x$. Since N and D are polynomials and $D(p^*) > 0$, Q is differentiable in a neighborhood of p^* . Taylor's theorem gives (5). The quotient rule gives

$$\frac{\partial Q}{\partial p_k} = \frac{D \partial N / \partial p_k - N \partial D / \partial p_k}{D^2} = \frac{1}{D} \left[\frac{\partial N}{\partial p_k} - Q \frac{\partial D}{\partial p_k} \right],$$

which is (6).

A.7. Proof of [Theorem 9](#)

Using (6) and $0 \leq Q(p) \leq 1$,

$$\left| \frac{\partial Q}{\partial p_k} \right| \leq \frac{1}{\Phi_E^x(p)} \left(\left| \frac{\partial \Phi_{Y,E}^x}{\partial p_k} \right| + |Q(p)| \left| \frac{\partial \Phi_E^x}{\partial p_k} \right| \right) \leq \frac{C_1 + C_2}{\delta}.$$

Thus the derivative bound diverges as $\delta \rightarrow 0$.

A.8. Proof of [Theorem 10](#)

Partition worlds as $(\omega_A, \omega_{\bar{A}})$, where $A = A_x(Y, E)$. By definition, variables outside A have no directed path to Y or E in the transformed dependency graph. Under the stated constraint-separation assumption, changing $\omega_{\bar{A}}$ cannot affect satisfaction of Y , E , or the retained clauses. Hence

$$\Phi_{Y,E}^x(p) = \sum_{\omega_A} \mathbb{I}\{\omega_A \models \Delta_x^A \wedge Y \wedge E\} w_{p_A}(\omega_A) \sum_{\omega_{\bar{A}}} w_{p_{\bar{A}}}(\omega_{\bar{A}}).$$

The inner sum is one because it sums the probability mass of independent removed choices. Therefore $\Phi_{Y,E}^x(p) = \Phi_{Y,E}^{x,A}(p_A)$. The same argument gives $\Phi_E^x(p) = \Phi_E^{x,A}(p_A)$, and taking quotients proves the claim. Now, apply [Theorem 8](#) to $\bar{p} = p^* + \varepsilon$. Coordinates outside A are absent or have zero derivative by [Theorem 6](#) and the above. Hence

$$Q(\bar{p}) - Q(p^*) = \nabla_A Q(p^*)^\top (\bar{p}_A - p_A^*) + O(\|\bar{p}_A - p_A^*\|^2).$$

Taking expectations yields the stated expansion. The final statement follows from the definition of DML: cross-fitting controls nuisance overfitting, but first-order bias is removed only when the estimator is built from a Neyman-orthogonal score for a valid target estimand.

Appendix B. Experimental DeepProbLog Programs

This appendix gives the DeepProbLog-style program templates used in the experiments. The actual experiments first evaluate the neural predicates at a fixed input context and then replace them by ordinary probabilistic choices before applying the DeepSWIP transformation. Thus the code below specifies the neural-symbolic model before materialization; each experimental query is then evaluated on the corresponding materialized and SWIP-transformed ProbLog program.

B.1. MPI3D visual counterfactual program

The MPI3D experiment uses neural predicates for visual factors and deterministic rules for physical predicates. For each image, the neural predicates return distributions over shape, size, and color. Counterfactual interventions are then applied to materialized shape or size atoms.

Listing 1: DeepProbLog template for the MPI3D visual counterfactual experiment.

```
% Neural visual predicates.
% For a fixed image I, these are evaluated once and materialized.
```

```

nn(shape_net, I, Shape, [cone,cube,cylinder,hexagonal,pyramid,sphere]) :: shape(I,
    Shape).
nn(size_net, I, Size, [small,large]) :: size(I, Size).
nn(color_net, I, Color, [red,green,blue,white,brown,olive]) :: color(I, Color).

% Shape-derived physical properties.
rollable(I) :- shape(I, sphere).
rollable(I) :- shape(I, cylinder).

flat_base(I) :- shape(I, cube).
flat_base(I) :- shape(I, cylinder).
flat_base(I) :- shape(I, hexagonal).

pointed(I) :- shape(I, cone).
pointed(I) :- shape(I, pyramid).

% Size-derived physical properties.
heavy(I) :- size(I, large).
light(I) :- size(I, small).

% Symbolic task predicates.
can_roll(I) :- rollable(I).

can_stack(I) :- flat_base(I).
can_stack(I) :- light(I), \+ pointed(I).

stable(I) :- flat_base(I), heavy(I).
stable(I) :- flat_base(I), light(I).

risky_on_shelf(I) :- rollable(I).
risky_on_shelf(I) :- pointed(I), heavy(I).

% Example queries used after materialization and SWIP transformation.
query(can_roll(I)).
query(can_stack(I)).
query(stable(I)).
query(risky_on_shelf(I)).

```

For example, the intervention $do(\text{shape}(I, \text{sphere}))$ removes the materialized mechanism defining the shape of image I , redirects downstream uses of shape to the fixed intervention atom, and then evaluates one of the four queries above. The same materialized program is also used to construct the DeepTwin baseline, so that disagreement would indicate an error in the counterfactual transformation rather than in the neural visual classifiers.

B.2. SUMO HOV traffic program

The SUMO HOV experiment uses a neural traffic-state predicate learned from simulated sensor features. The symbolic program maps traffic state and lane policy to coarse throughput regimes. The randomized-policy AIPW analysis is not itself a ProbLog program: it is a population estimator applied around the DeepSWIP plug-in nuisance model. The program below is the symbolic nuisance model whose materialized probabilities are perturbed by sensor noise.

Listing 2: DeepProbLog template for the SUMO HOV traffic experiment.

```

% Neural traffic-state predicate.
% Sensor features Z are produced by SUMO loop-detector counts and speeds.
nn(traffic_state_net, Z, State, [freeflow,congested]) :: traffic_state(Episode, State).

% Policy atoms. In the randomized-policy experiment, treatment is assigned
% outside the program and then materialized as one of these facts.
hov_restricted(Episode) :- policy(Episode, restricted).
hov_open(Episode)      :- policy(Episode, open).

% State indicators.
freeflow(Episode) :- traffic_state(Episode, freeflow).
congested(Episode) :- traffic_state(Episode, congested).

% Coarse symbolic throughput regimes.
high_throughput(Episode) :- freeflow(Episode), hov_open(Episode).
high_throughput(Episode) :- freeflow(Episode), hov_restricted(Episode).

medium_throughput(Episode) :- congested(Episode), hov_open(Episode).
low_throughput(Episode)    :- congested(Episode), hov_restricted(Episode).

% Binary summaries used by the plug-in DeepSWIP nuisance model.
good_outcome(Episode) :- high_throughput(Episode).
good_outcome(Episode) :- medium_throughput(Episode).

% Example evidence and query forms.
% The experiments instantiate these per episode after materialization.
query(good_outcome(Episode)).
query(high_throughput(Episode)).
query(medium_throughput(Episode)).
query(low_throughput(Episode)).

```

For the HOV counterfactual query, the intervention $do(\text{policy} = \text{open})$ replaces the policy mechanism with the fixed open-lane policy. Conditioning evidence such as $\text{congested}(i)$ is evaluated in the factual world, while the query is evaluated in the SWIP-transformed program. The AIPW/DML correction reported in the experiment is then applied only to the constructed randomized-policy population estimands, not to arbitrary individual counterfactual queries.

B.3. Materialized form used by DeepSWIP

After neural evaluation, each neural predicate is replaced by an ordinary probabilistic choice. For example, an MPI3D image with shape probabilities (0.01, 0.02, 0.05, 0.03, 0.04, 0.85) is represented as:

Listing 3: Example materialized MPI3D neural predicate.

```

0.01::shape(i, cone);
0.02::shape(i, cube);
0.05::shape(i, cylinder);
0.03::shape(i, hexagonal);
0.04::shape(i, pyramid);
0.85::shape(i, sphere).

```

Similarly, a SUMO episode with predicted traffic-state probabilities (0.84, 0.16) is represented as:

Listing 4: Example materialized SUMO neural predicate.

```
0.84::traffic_state(e, freeflow);
0.16::traffic_state(e, congested).
```

DeepSWIP is applied only after this materialization step. This is why the counterfactual solver operates on ordinary ProbLog programs even though the original experimental models contain neural predicates.

The AIPW estimator uses known propensity $g = 0.5$ and fold-specific observed-only nuisance estimates:

$$\hat{\psi}_1^{\text{AIPW}} = \frac{1}{n} \sum_{i=1}^n \left[\hat{\mu}_1^{(-i)}(W_i) + \frac{T_i}{g} \{Y_i - \hat{\mu}_1^{(-i)}(W_i)\} \right], \quad (7)$$

with the analogous expression for ψ_0 and $\hat{\tau} = \hat{\psi}_1 - \hat{\psi}_0$.

Appendix C. Additional Experimental Details

Table 2: MPI3D visual counterfactuals. SWIP and Twin agree exactly over 12,000 materialized visual-symbolic counterfactual queries. RMSE is against factor-ground-truth potential outcomes.

Query	SWIP=Twin	Bias	RMSE	Pruning
can_roll	100%	+0.0000	0.0000	18 → 5
can_stack	100%	+0.0013	0.0580	18 → 8
risky_on_shelf	100%	-0.0026	0.0821	18 → 11
stable	100%	+0.0000	0.0000	18 → 7

Table 3: SUMO HOV corrected DML/AIPW stress test.

σ	Cal(ff)	Cal(cg)	Naive ψ_1	CV ψ_1	AIPW ψ_1	Naive ATE	AIPW ATE
0.0	0.920	0.985	-4.75	-4.67	+0.14	-1.24	+0.15
1.0	0.890	0.900	-5.66	-5.36	+0.08	-1.50	+0.06
2.0	0.845	0.706	-6.32	-7.04	+0.13	-1.68	+0.14
3.0	0.767	0.556	-9.55	-8.92	+0.03	-2.57	-0.03