PCIWM: A UNIFIED DRIVING WORLD STATE FOR PREDICTION AND PLANNING

Anonymous authors
Paper under double-blind review

ABSTRACT

We introduce PCIWM, a driving world model that explicitly constructs physical (P), causal-semantic (C), interaction (I), and ego (E) states and maps them jointly into one world-state representation. A geometry-aware composer, shared transition under external ego-motion inputs, and typed decoders turn this definition into candidate-specific predicted futures. We evaluate the architecture through two connected questions: what it predicts, and what decisions those predictions enable. Controlled nuScenes studies test factor inclusion, composition, relation identity, and transition sharing. Historical controls show relation-prediction and long-horizon stability gains; common-budget and three-seed comparisons reveal observable-dependent tradeoffs and remaining single-step prediction errors. On NAVSIM, predicted traffic improves a component readout over static traffic, and a state-motivated candidate library improves coverage and selected utility over equal-count alternatives. Raw consequences yield a positive but uncertain five-seed selection gain. A further control tests the incremental use of later raw-state updates: repeating each candidate's first predicted state outperforms trajectory-only inputs in an exploratory three-seed comparison, while the full multi-step sequence performs worse than that repetition. These cached, annotation-assisted development studies connect an explicit unified-state architecture to controlled prediction-to-decision evidence, rather than equating predictive fidelity with planning utility.

1 INTRODUCTION

PCIE is our definition of a unified driving world state. We explicitly construct physical context (P), causal-semantic context (C), interaction structure (I), and ego state (E), and map them jointly into a shared world-model representation. The starting question is what a driving world state must preserve to support both future prediction and decisions about ego motion. Geometry, traffic semantics, participant relationships, and ego motion describe complementary aspects of the same evolving world; our formulation makes their construction, integration, and evolution explicit.

This leads to two connected commitments: *define the world state explicitly, and assess its usefulness through downstream decision quality alongside prediction accuracy*. Prediction errors diagnose the captured observables; planning tests whether predicted consequences help distinguish available actions. We operationalize this established motivation through controlled state, intervention, and readout comparisons.

PCIWM implements the state definition as

$$W_t = (P_t, C_t, I_t, E_t), \qquad Z_t = \mathcal{M}_\theta(W_t). \tag{1}$$

Here W_t is the typed world state and Z_t is its shared latent representation. A geometry-aware composer aligns the four factors with shared slots. A common transition evolves those slots under an external ego-motion input, and typed decoders recover the predicted state. Repeating this process under alternative motions produces candidate-specific futures for planning. State definition, coupled dynamics, explicit intervention, alternative futures, and decision readout thus form one computational chain (Fig. 1).

PCIE was formulated and experimentally instantiated before our documented planning-metric analysis. Its later correspondence with safety, geometric compliance, and progress criteria guided braking

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.

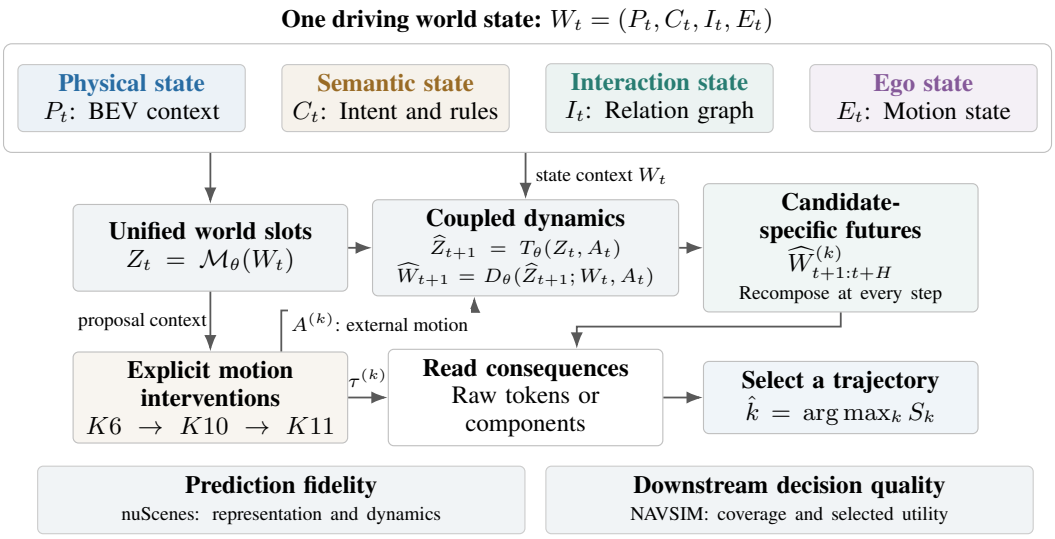

Figure 1: **Explicit PCIE states, one unified representation, two tests of value.** Physical, semantic, interaction, and ego states jointly define W_t and are mapped into its shared representation Z_t. Coupled dynamics under external motion inputs produce candidate-specific futures. nuScenes tests prediction fidelity; NAVSIM tests how predicted consequences inform decisions through coverage and selection. These are complementary evaluations on their respective datasets.

and geometric extensions of the proposal set. This provides a concrete use of the state definition in intervention design; the experiments separately measure what the resulting candidate library covers and what a selector can recover.

We make two contributions. **First, an explicitly structured world-model architecture.** PCIWM constructs the four PCIE states, maps them into one representation, and learns their action-conditioned evolution. Factor, composition, identity, and transition controls test whether these choices improve specific predicted observables and their stability under recursive feedback. **Second, a controlled prediction-to-decision study.** We distinguish candidate coverage, access to predicted consequences, and the extra value of their multi-step evolution. Provider and input controls measure selected utility through two readout interfaces, connecting the proposed state construction to decisions without treating predictive accuracy as sufficient.

nuScenes tests representation and dynamics, including nine model variants and three-seed factor controls; NAVSIM tests downstream decision utility. Section 4 specifies the cached, annotation-assisted development setting. The evidence characterizes where this architecture helps, its observable-specific tradeoffs, and remaining prediction and selection gaps.

2 Related Work

Latent dynamics support planning and policy learning through imagined futures (Hafner et al., 2019; 2020). In driving, MILE learns action-conditioned latent dynamics together with an imitation policy (Hu et al., 2022); OccWorld jointly models occupancy and ego tokens (Zheng et al., 2024); and DriveWorld uses static and dynamic scene structure for representation pretraining and downstream tasks (Min et al., 2024). These approaches already incorporate spatial structure, semantics, or ego behavior. Interaction Networks establish a broader precedent for explicitly modeling entities and relations (Battaglia et al., 2016). GraphAD connects ego, agents, and map elements through geometry-guided interaction graphs for prediction and planning (Zhang et al., 2025).

Recent driving models connect future prediction to different planning interfaces. LAW uses action-conditioned future latent supervision to improve driving representations (Li et al., 2025a). World4Drive predicts intention-conditioned physical latents and scores candidate intentions (Zheng et al., 2025); WoTE evaluates candidate trajectories through predicted BEV futures and a learned re-

ward model (Li et al., 2025b). GraphWorld constructs an interaction-aware world state to condition long-horizon planning, emphasizing latent refinement rather than explicit multi-step autoregressive rollout (Song et al., 2026). Thus, action conditioning, structured states, and future-based candidate evaluation each have prior art.

PCIWM's architectural focus is the explicit state definition $W_t = (P_t, C_t, I_t, E_t)$ with $Z_t = \mathcal{M}(W_t)$: physical context, semantic context, interaction structure, and ego state are factors of one driving world state, while actions remain a separate transition input. Geometry-aware joint mapping, shared transition, and typed recovery implement this definition. The accompanying evaluation asks both whether the state supports prediction and whether reading its predicted consequences improves decisions. Our contribution is this specific construction and its controlled study. The semantic factor C does not identify causal variables (Schölkopf et al., 2021); our counterfactual futures denote model predictions under alternative ego inputs, without establishing real-world intervention effects. The MILE-style and OccWorld-style experiments are internal adaptations under a shared protocol, rather than reproductions of the original systems.

3 A Unified Driving World State

3.1 Design rationale

Evaluating a braking alternative involves ego motion, nearby participants and relations, surrounding geometry, and traffic conditions. PCIE preserves these complementary observables explicitly, maps them jointly so dynamics can use their dependencies, and recovers typed targets for evaluation. Factor-inclusion and composition controls test this design hypothesis; the factors may overlap and need not be statistically independent.

3.2 Explicit construction of the four state factors

The native state $W_t = (P_t, C_t, I_t, E_t)$ preserves each factor's structure before their joint mapping into Z_t (Eq. 1). P_t is a cached temporal BEV feature grid. C_t is the causal-semantic context: grounded actor semantics including intent, traffic-rule, affordance, risk, and right-of-way attributes. I_t is a directed interaction graph with participant identities, node geometry, relation attributes, and validity masks. E_t describes ego motion, including speed, velocity, and yaw rate. Here "causal-semantic" denotes context relevant to driving decisions, not identified causal variables. The evaluated BEV grids are obtained with a BEVFormer front end (Li et al., 2022); language-derived fields use Qwen2.5-VL (Bai et al., 2025). The native interaction construction uses annotation-derived actors and geometric rules, with a GraphAD-style graph encoder (Zhang et al., 2025); map nodes supported by a generic builder are not assumed present in the evaluated nuScenes inputs.

Input contract. In all new nuScenes architecture controls and three-seed factor repetitions, the current C intent is recomputed from current and previous annotations; used timestamps are at or before the anchor. Other semantic fields are preserved. The inspected nuScenes I/E motion estimates retain adjacent-frame dependencies. The formal NAVHARD adapter constructs inputs from four observed frames, ending at the current frame; its future participants are predicted by PCIWM. Historical dynamics and NAVSIM development caches have separately recorded input lineages (Appendix A). Annotation assistance, observation timing, and future supervision are distinct properties of these inputs.

3.3 Mapping PCIE into one shared representation

Each enabled factor is encoded as tokens with masks and spatial coordinates. Factor-specific cross-attention aligns these tokens with learned world slots, using geometry embeddings in both queries and source tokens. At layer ℓ, source tokens x_{si} with geometry q_{si} and mask m_{si} form $k_{si}^{\ell} = x_{si} + \gamma_s^{\ell}(q_{si})$. Each factor/layer has its own scalar gate:

$$\bar{k}_s^{\ell} = \frac{\sum_i m_{si} k_{si}^{\ell}}{\max(\sum_i m_{si}, 1)}, \qquad g_s^{\ell} = \sigma\Big(\mathrm{MLP}_s^{\ell}(\bar{k}_s^{\ell})\Big). \tag{2}$$

The MLP is Linear(d, d)–SiLU–Linear$(d, 1)$. Its per-sample gate is broadcast over slots and channels. With U_s^ℓ the LayerNorm of the factor's attention output, composition uses

$$\widetilde{Z}^\ell = Z^\ell + \frac{\sum_s g_s^\ell U_s^\ell}{\max(\sum_s g_s^\ell, \epsilon)}, \qquad Z^{\ell+1} = \widetilde{Z}^\ell + \mathrm{FFN}_\ell(\widetilde{Z}^\ell). \tag{3}$$

Both sums include only enabled factors; $\epsilon = 10^{-6}$. Final LayerNorm gives $Z_t = \mathcal{M}_\theta(W_t)$. Removing optional factors yields PE, PCE, and PIE controls. A monolithic-composition control concatenates the same native tokens before shared attention. It retains their tokenization, geometry, and graph encoders, isolating the composition mechanism rather than removing all structure.

3.4 Controlled transition and persistent relations

We keep the proposed motion increment external to the state so that the same world can be evaluated under alternative ego motions:

$$\widehat{Z}_{t+1} = T_\theta(Z_t, A_t), \qquad \widehat{W}_{t+1} = D_\theta(\widehat{Z}_{t+1}; W_t, A_t). \tag{4}$$

An adapter maps A_t to control tokens attended by the transition transformer. Structured interaction decoding additionally uses current graph context and action, as indicated explicitly in Eq. 4. Our shared transition is compared with factor-specific transitions that each receive the same already-fused slots and action. This tests parameter sharing after composition, not independent access to isolated factor inputs.

Nodes are aligned by persistent participant identity. Baseline controls identify edges by endpoint–relation keys (u, v, r). The persistent-identity variant instead uses the ordered endpoint pair (u, v); relation type becomes a changing label rather than part of identity. A relation change can consequently occupy the same edge slot. Existence masks distinguish invalid or absent elements. The rollout evaluations freeze the initial support and report future-element coverage separately, so persistent identity is not evidence of unrestricted node or edge birth prediction.

3.5 Learning state dynamics

Training combines decoded-state supervision with a latent target:

$$\mathcal{L} = \sum_{s \in \{P,C,I,E\}} \lambda_s \mathcal{L}_s + \lambda_Z \|\widehat{Z}_{t+1} - \mathrm{sg}[\mathcal{M}_\theta(W_{t+1})]\|_2^2 + \lambda_{\Delta I} \mathcal{L}_{\Delta I}. \tag{5}$$

Physical and ego predictions use Smooth-L1 losses; semantic attributes use masked absolute error. Graph supervision covers node and edge attributes, existence, and type with their applicable masks. The optional final term emphasizes relation changes. Target slots are composed without gradients; factor-specific transitions receive the same latent target. Additional configured auxiliary terms and weights are specified with each protocol.

For multi-step training, a scheduled selector chooses the next recorded state or the decoded prediction for feedback. Predicted feedback is detached, and each transition is differentiated separately; this is not full backpropagation through the rollout. Teacher-forced, autoregressive, and persistence evaluations expose different error sources. Recorded ego motion conditions dynamics evaluation; this measures prediction along observed actions, not alternative-action ground-truth accuracy.

3.6 Explicit motion alternatives

The implemented action is local incremental motion $A_t = (\Delta x_t, \Delta y_t, \Delta \psi_t)$, derived from successive ego poses for training and integrated for planning. It is not a calibrated throttle or steering command. A proposal head reads current world slots and ego history to generate six action sequences, K6. Recursively applying Eq. 4 produces deterministic candidate-specific futures; diversity comes from different actions rather than stochastic world sampling.

Semantic-guided expansion adds four fixed transformations: two time-decayed motion sequences and two variants scaling lateral and yaw increments. These produce K10; adding constant velocity produces K11. This is a fixed library motivated by braking and geometric failure modes, not an online factor-wise intervention policy. Candidate-coverage oracles are kept separate from achievable selection scores.

Table 1: **Distinct questions and evaluation contracts.** A shared conceptual pipeline does not imply a single checkpoint lineage. All cohorts are development analyses.

Study	Data and training	Horizon	Question
Architecture controls	nuScenes: 1,484 pairs; nine models, seed 7, four epochs	149 chains; H6 ($\sim$3 s)	Transition sharing and composition
C-factor repeats	Same pairs and endpoint; four models, seeds 7/8/9	Same 149 chains; H6	Incremental C effects across fitted seeds
Raw consequences	NAVSIM: 419/186 tokens; two arms, five seeds; temporal control on original three	K11, H8 (4 s)	State access and multi-step increment
Component readout	NAVSIM: 419/186 tokens; three providers, three seeds	K11, H8	Predicted traffic vs. static/CV
Candidate expansion	482 evaluation tokens; six frozen selectors	Three K11 pools	Coverage and selected utility
Branch selection	NAVHARD: 225 linked pairs; 450 support entries	Official two-stage evaluator	Full-pool vs. strict Top-3 search

3.7 Reading consequences for planning

On NAVSIM, we examine two planning interfaces for the predicted state. The **raw-consequence selector** encodes each predicted P/C/I/E factor and the unified latent into five typed tokens per step; the latter is a derived representation, not a fifth state factor. A candidate-relative world stream subtracts the across-candidate mean, while a separate trajectory stream encodes actions and integrated poses. Their scores are combined. The original soft-ranking and world-contrast objectives train this readout using candidate PDM labels. An independently trained trajectory-only control zeros all five consequence-token types during training and evaluation, preserving the architecture and trajectory inputs. A second control repeats each candidate's first predicted raw state across the horizon, testing later state evolution while retaining its initial prediction and full candidate trajectory.

The **component readout** instead evaluates decoded participant tracks and candidate geometry with explicit planning criteria, followed by the original residual ranker. Predicted, static, and constant-velocity providers recompute the component features, aggregate prior, and fallback selection; no predicted-provider prior is retained in the controls. Within each interface, all input controls share its world-conditioned proposals. Their controls therefore test the incremental use of consequences within a fixed candidate pool, rather than an end-to-end system without a world model.

4 Experimental Design

We organize the evidence by question, using nuScenes (Caesar et al., 2020) for representation and dynamics, and NAVSIM (Dauner et al., 2024) for planning. Table 1 separates the experimental contracts. The common dynamics study uses 1,484 selected training pairs, 765 development pairs, and 149 continuous evaluation chains from 30 scenes. All nine models are evaluated at epoch 3: a fixed four-epoch prefix of a twelve-epoch schedule, seed 7, and 92 training batches. The four factor-inclusion models are also repeated with seeds 8 and 9 on the same pool and fixed endpoint. This matches exposure, not convergence or active capacity. Historical studies retain their separate endpoints in Appendix B. The reference architecture is shared-transition PCIE with 128 width-256 slots, three composer and six transition layers. The historical NAVSIM world weights and readout caches are separate branches, not a single end-to-end checkpoint shared with these nuScenes controls (Appendix A).

All studies use cached, annotation-assisted state inputs. The common dynamics study repairs the C intent field to use history; inspected nuScenes motion fields retain adjacent-frame dependencies. The formal NAVHARD adapter uses observed history, while historical NAVSIM development studies retain their own cache lineages. We evaluate the representation and readout interfaces under those

Table 2: AR velocity-L2 change from adding C at both fixed endpoints (m/s).

Comparison	Endpoint	Seed 7	Seed 8	Seed 9	Mean ± SD
PCE − PE	H6	−0.1486	−0.0579	+0.0234	-0.0610 ± 0.0860
PCE − PE	Mean H1–H6	−0.0571	−0.0069	+0.0086	-0.0185 ± 0.0343
PCIE − PIE	H6	−0.1550	+0.3018	+0.0136	$+0.0535 \pm 0.2310$
PCIE − PIE	Mean H1–H6	−0.0891	+0.2412	+0.0384	$+0.0635 \pm 0.1666$

All three paired seeds are retained; negative favors C. Mean and sample SD describe these fixed fitted instances, without a training-seed population inference.

inputs, without claiming a verified sensor-only policy. The cohorts are historical development data. NAVSIM train/development tokens and scene IDs are disjoint in the 419/186 study, but 66 recording groups overlap. The full field-level contract, sampling rules, model lineages, and hash-indexed artifacts accompany the code.

We report raw candidate PDM and NAVHARD combined EPDMS on a 0–100 scale; they are different quantities. Candidate oracle denotes the best labeled trajectory in a fixed pool, not a selector supplied with ground-truth future worlds. Multi-seed tables use mean ± sample SD. Recording-group bootstrap intervals are post-hoc, marginal 95% percentile intervals conditional on the trained instances. They do not remove adaptive development or train/recording overlap. Fixed terminal checkpoints are primary; development-best endpoints are exploratory. The new first-step raw-state control reuses the original three readout seeds; its comparison with trajectory-only was added after observing its results and is separately marked exploratory.

5 State Representation and Long-Horizon Dynamics

We test the state design through factor inclusion, composition and transition controls: which information survives prediction, and how errors accumulate when the unified state evolves recursively.

Shared evolution can improve long-horizon stability. A historical same-data, twelve-epoch comparison gives H5 ego AR/TF error amplification of 2.7454 for shared dynamics versus 3.4282 for independent transitions (Appendix B). The independent branches receive the same composed state; this tests transition coupling after fusion. The new seed-7, four-epoch study exposes complementary tradeoffs (Table 3): independent transitions and monolithic composition improve speed and velocity-vector error, while shared PCIE has lower yaw-rate and lateral-velocity errors. Thus coupling benefits some long-horizon observables without uniformly dominating either control. The internally adapted MILE-style and OccWorld-style controls have larger common P/E errors here; these are not comparisons with their original architectures, data or converged training.

Explicit factors change what the state preserves. Across seeds 7/8/9, adding C to PIE lowers lateral-velocity and both BEV-feature errors in all three seeds at H6 and the six-step mean, but raises mean speed and longitudinal-velocity errors. The primary velocity-vector endpoint is also mixed: PCE improves over PE by 0.0610 m/s at H6 in the mean, with two seeds improving; PCIE worsens over PIE by 0.0535 m/s, with only seed 7 improving. Six-step means follow the same pattern (Table 2). These repetitions identify observable-specific effects, not stable incremental C benefit for PIE or necessity of every factor. The favorable seed-7 contrasts and their conditional scene intervals remain separately reported (Appendix D).

Structured identity matters for relation dynamics. Changing the edge key from endpoint–relation to ordered endpoints in the historical controlled experiment reduces terminal validation loss by 6.47% and graph-edge loss by 14.60%. Persistent participants therefore provide a useful anchor for changing relations. Separate full214 H5 experiments test feedback training and changed-edge weighting; the latter improves changed-relation F1 while sometimes harming ordinary type classification (Appendix B).

Localizing predictive limitations. The C-associated velocity disadvantage already appears under TF: three-seed PCIE–PIE six-step error is $+0.0738$ m/s, with conditional scene interval

Table 3: **Nine-model three-second AR evaluation: seed 7, fixed epoch 3.** Lower errors are better. H6 velocity L2 and mean velocity L2 use m/s; yaw-rate MAE uses rad/s; BEV MAE is in feature units. Parameter counts include disabled modules.

Model	Params (M)	H6 vel. L2	Mean vel. L2	H6 yaw rate	H6 BEV
PE	26.728	0.8187	0.5543	0.0539	0.2874
PCE	26.728	0.6701	0.4972	0.0379	0.2799
PIE	26.728	0.7723	0.5725	0.0391	0.2802
PCIE	26.728	0.6173	0.4834	0.0369	0.2797
Independent transition	46.487	0.5499	0.4245	0.0382	0.2794
Monolithic composition	22.365	0.4691	0.4413	0.0461	0.2813
Ordered edge identity	26.728	0.7789	0.5980	0.0383	0.2804
MILE-style	18.403	1.4336	1.4313	0.0413	0.2966
OccWorld-style	19.952	2.1036	2.1478	0.0735	0.2876
Persistence	—	1.0313	0.6091	0.0549	0.2483

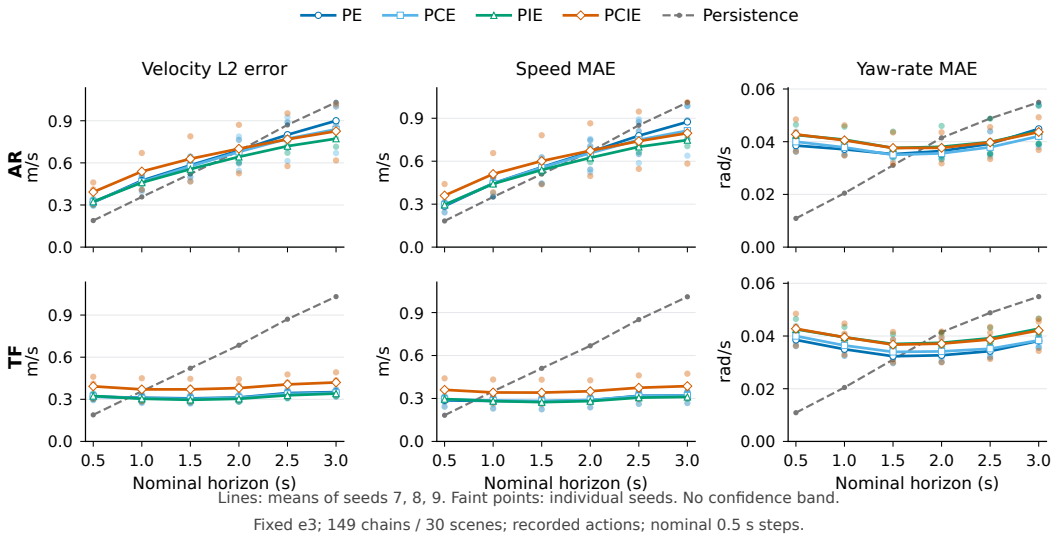

Figure 2: **Factor inclusion across three fitted training seeds.** Fixed epoch 3; 149 common chains from 30 scenes. Lines are means of seeds 7/8/9, faint points individual seeds; persistence is shared. AR appears above TF, with no confidence band. Time is nominal at 2 Hz (actual H6 mean: 2.990 s). All seven readouts and paired differences appear in Appendix B.4.

$[0.0444, 0.0993]$. The AR–TF difference does not consistently grow from the first to last three steps. This rules out accumulated feedback as its sole explanation, without separating C input from auxiliary supervision. At seed 7, copying the true preceding P/E state at each TF step beats all nine models' mean velocity error (PCIE .3038 versus .2044 m/s) and their BEV errors. This last-observation baseline is distinct from initial-state persistence: the three-seed AR six-step velocity mean is also higher than initial persistence (.6424 versus .6091 m/s), despite a lower H6 error (.8257 versus 1.0313). Frozen current-latent decoding also remains above that baseline; it is an off-label readout, not an information-loss decomposition (Appendix E.2).

Action shuffling/zeroing changes the historical curriculum model's ego prediction by .0808/.7619 MAE across all 765 pairs, establishing functional dependence rather than unexecuted-action accuracy. Relation dynamics remain a bottleneck: PCIE's H6 AR changed-relation F1 is zero, versus about .805 under TF, on initial tracked support. At seed 7, the preceding checkpoint has H6 velocity error 1.1260 m/s, falling to .6173 at epoch 3, while BEV improves by only .57%. This finite-budget result shows continuing motion improvement, not a converged capacity limit or birth prediction (Appendix E.5).

Table 4: **Consequence readouts under input controls.** PDM $\times 100$, mean $\pm$ sample SD at fixed terminal checkpoints. The first panel retains all five raw-readout seeds; the temporal control uses the original three. Compare rows within each panel; cache lineage and training budget differ. Gap is candidate oracle minus selection.

Readout	Input	Selected ↑	Gap ↓
Raw consequence tokens: four epochs; candidate oracle 94.66			
Original two-stream	Trajectory only	57.58 ± 2.93	37.08 ± 2.93
Original two-stream	Trajectory + future	60.93 ± 2.59	33.73 ± 2.59
Raw temporal control: same three seeds; four epochs; oracle 94.66			
Original two-stream	Trajectory only	58.12 ± 3.27	36.54 ± 3.27
Original two-stream	Full multi-step future	61.27 ± 0.86	33.39 ± 0.86
Original two-stream	First predicted state repeated	65.07 ± 4.02	29.59 ± 4.02
Future-derived components: 60 epochs; candidate oracle 94.15			
Original residual ranker	Static participants	84.62 ± 0.51	9.52 ± 0.51
Original residual ranker	Constant-velocity participants	84.49 ± 0.37	9.66 ± 0.37
Original residual ranker	Predicted participants	87.46 ± 0.12	6.69 ± 0.12

6 Evaluating Futures through Decisions

We test the world state's decision value through consequence-provider controls, candidate coverage and selected-plan quality. All pools already contain world-conditioned proposals.

Predicted traffic can improve a component readout. Replacing predicted traffic with static or constant-velocity participants recomputes each candidate's planning components, prior and fallback while retaining the residual-ranker architecture. At the fixed 60-epoch endpoint, predicted traffic yields 87.46 PDM, versus 84.62/84.49 for static/CV, with all three seed differences favoring prediction (Table 4). Recording-paired differences are +2.83 [0.79, 5.10] over static and +2.97 [-0.69, 6.81] over CV: the latter retains uncertainty.

Raw consequences give a positive but uncertain selection gain. The original two-stream architecture and trajectory-only control share paired initialization, labels and sample orders, with all five consequence token types zeroed in the control. At fixed epoch 3, five seeds yield 60.93 versus 57.58: four gains and one loss (-0.83 points), averaging +3.35 [-0.36, 7.10]. Against always-CV (60.08), the difference is +0.85 [-6.45, 7.95]. Two seeds were appended after observing the original three; all five remain reported. The oracle is 94.66. Component and raw panels have different cache lineages and budgets, so their scores do not isolate architecture quality.

Predicted-state access and temporal evolution are distinct. On the original three seeds, an additional control repeats each candidate's first predicted raw state across H8, retaining actions, trajectories, and the first state's embedded trajectory proposals. Full multi-step input scores 61.27 versus 65.07 for this control: all three differences are negative, averaging -3.80 [-8.98, 1.20]. Its subsequently requested comparison with trajectory-only is exploratory: $+6.95$ [1.11, 13.09], positive in all three seeds. Thus predicted-state access has positive evidence here, while additional raw-state evolution has no demonstrated gain. This is not a current-state or world-free control; it retains future information. The negative primary comparison and exploratory positive comparison are reported together (Appendix E.1).

Intervention design improves coverage and selected utility. The fixed braking and lateral/yaw library expands K6 to K10; CV adds K11. On 482 tokens, three eleven-candidate pools share K7 and vary four additions, with eleven distinct pose hashes per token. The semantic-guided pool improves oracle PDM by 3.28/3.60 points over generic lateral/speed pools, reducing zero-oracle cases to 2.07% from 5.81%/6.64%. It recovers 27 of 37 K7 zero-score cases, versus 9/5.

With frozen selectors, guided-pool selected utility improves by 3.46–4.28 points across two input variants and three seeds. All four marginal recording-bootstrap intervals are positive, while oracle–selection-gap differences remain uncertain. These selectors were trained on the guided pool: this

Table 5: **Equal-count candidate expansion.** 482 development tokens; K11 in each pool. Selected PDM uses frozen guided-trained selectors with current-state or kinematic inputs (three seeds each). These are separate auxiliary selectors, not the readouts in Table 4.

Expansion	Oracle ↑	Zero oracle ↓	Current selected ↑	Kinematic selected ↑
Generic lateral	89.495	5.81%	72.177 ± 0.909	72.151 ± 1.552
Generic speed	89.174	6.64%	72.438 ± 1.408	72.800 ± 1.246
Semantic-guided	92.775	2.07%	75.897 ± 0.783	76.428 ± 2.857

is pool transfer, not independently retrained pool fairness. Amplitudes are unmatched, and the geometry-only matching attempt failed fixed acceptance checks. The finding supports the tested expansion design as a whole, not an isolated semantic advantage.

Selection must preserve useful candidate support. The historical NAVHARD comparison fixes each of two scoring rules while retaining K11 or pruning to strict Top-3. Full-pool combined EPDMS is 53.66/54.17, falling to 49.16/48.97 after pruning. The full-pool winner is removed in 95/225 and 84/225 linked pairs; when retained, its selection is unchanged. Thus candidate support and search affect downstream outcomes without changing the predictor. This uses the official evaluator locally with benchmark-aware linked-branch search, not blind leaderboard evaluation. Appendix C retains the full decomposition, oracle bounds and mixed auxiliary readout results.

7 Discussion and Limitations

The explicit state definition makes representation, transition and readout choices testable. Semantic inclusion changes observables differently, and its velocity tradeoff already appears under TF. Last-observation baselines expose single-step errors alongside unstable AR relations. Planning controls separate useful predicted-state access from additional temporal evolution: the first-step control supplies the former while outperforming the full sequence in these fits. This identifies a concrete bottleneck for improving the interface, without claiming a mechanism or a universal temporal disadvantage.

Recorded targets cannot validate unexecuted-action counterfactuals. Limited training budgets and internal style baselines restrict claims about converged superiority. Annotation-assisted fields, historical cache differences, recording overlap, and repeated development analysis limit deployment and generalization claims. The reported confidence intervals are conditional descriptions of these studies; all adverse outcomes remain in the complete result tables. The first-step versus trajectory-only contrast is explicitly result-aware exploration, not a replacement for the primary temporal control. The nuScenes controls do not measure individual factors' effects on NAVSIM PDM.

8 Conclusion

PCIWM explicitly constructs physical, causal-semantic, interaction, and ego states and maps them into one shared world-model representation. Coupled dynamics under external motion interventions turn that state into alternative predicted futures for planning. Controlled studies test both their predictive fidelity and decision utility, revealing useful consequence readouts alongside factor tradeoffs, unstable relations, and limited temporal increments. The contribution is a concrete unified-state architecture and its controlled prediction-to-decision study.

Reproducibility Statement

The accompanying anonymous code archive contains the model and original planning implementations, experiment wrappers, sampled token manifests, configuration and dependency records, saved results, and result-replay scripts. Main dynamics and readout metrics can be reaggregated from saved per-example outputs without model inference. Other historical summaries are identified as saved-result extraction rather than a new evaluator run. Full retraining and inference additionally require the relevant licensed datasets, front ends, generated caches, and pretrained checkpoint inputs specified in the archive. The archive does not bundle raw datasets or trained weight files. Appendix A documents artifact lineages and reproduction levels.

AI Usage Statement

AI assistants were used to assist code implementation, experiment orchestration and verification, analysis of saved results, literature organization, and manuscript drafting and editing. Model-generated suggestions were checked against available source code and experimental artifacts. The authors are responsible for the claims, correctness, attribution, and final submission.

References

Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibo Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, Humen Zhong, Yuanzhi Zhu, Mingkun Yang, Zhaohai Li, Jianqiang Wan, Pengfei Wang, Wei Ding, Zheren Fu, Yiheng Xu, Jiabo Ye, Xi Zhang, Tianbao Xie, Zesen Cheng, Hang Zhang, Zhibo Yang, Haiyang Xu, and Junyang Lin. Qwen2.5-VL technical report. *arXiv preprint arXiv:2502.13923*, 2025. URL https://arxiv.org/abs/2502.13923.

Peter Battaglia, Razvan Pascanu, Matthew Lai, Danilo Jimenez Rezende, and Koray Kavukcuoglu. Interaction networks for learning about objects, relations and physics. In *Advances in Neural Information Processing Systems*, volume 29, 2016. URL https://proceedings.neurips.cc/paper/2016/hash/3147da8ab4a0437c15ef51a5cc7f2dc4-Abstract.html.

Holger Caesar, Varun Bankiti, Alex H. Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuScenes: A multimodal dataset for autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020. URL https://openaccess.thecvf.com/content_CVPR_2020/html/Caesar_nuScenes_A_Multimodal_Dataset_for_Autonomous_Driving_CVPR_2020_paper.html.

Wei Cao, Marcel Hallgarten, Tianyu Li, Daniel Dauner, Xunjiang Gu, Caojun Wang, Yakov Miron, Marco Aiello, Hongyang Li, Igor Gilitschenski, Boris Ivanovic, Marco Pavone, Andreas Geiger, and Kashyap Chitta. Pseudo-simulation for autonomous driving. In *Proceedings of The 9th Conference on Robot Learning*, volume 305 of *Proceedings of Machine Learning Research*, pp. 4709–4722, 2025. URL https://proceedings.mlr.press/v305/cao25a.html.

Daniel Dauner, Marcel Hallgarten, Andreas Geiger, and Kashyap Chitta. Parting with misconceptions about learning-based vehicle motion planning. In *Proceedings of The 7th Conference on Robot Learning*, volume 229 of *Proceedings of Machine Learning Research*, pp. 1268–1281, 2023. URL https://proceedings.mlr.press/v229/dauner23a.html.

Daniel Dauner, Marcel Hallgarten, Tianyu Li, Xinshuo Weng, Zhiyu Huang, Zetong Yang, Hongyang Li, Igor Gilitschenski, Boris Ivanovic, Marco Pavone, Andreas Geiger, and Kashyap Chitta. NAVSIM: Data-driven non-reactive autonomous vehicle simulation and benchmarking. In *Advances in Neural Information Processing Systems*, volume 37, pp. 28706–28719, 2024. doi: 10.52202/079017-0902. URL https://proceedings.neurips.cc/paper_files/paper/2024/hash/32768f7faf1995026ef9821c696f3404-Abstract-Datasets_and_Benchmarks_Track.html.

Danijar Hafner, Timothy Lillicrap, Ian Fischer, Ruben Villegas, David Ha, Honglak Lee, and James Davidson. Learning latent dynamics for planning from pixels. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 2555–2565, 2019. URL https://proceedings.mlr.press/v97/hafner19a.html.

Danijar Hafner, Timothy Lillicrap, Jimmy Ba, and Mohammad Norouzi. Dream to control: Learning behaviors by latent imagination. In *International Conference on Learning Representations*, 2020. URL https://arxiv.org/abs/1912.01603.

Anthony Hu, Gianluca Corrado, Nicolas Griffiths, Zachary Murez, Corina Gurau, Hudson Yeo, Alex Kendall, Roberto Cipolla, and Jamie Shotton. Model-based imitation learning for urban driving. In *Advances in Neural Information Processing Systems*, volume 35, pp. 20703–20716, 2022. doi: 10.52202/068431-1505. URL https://proceedings.neurips.cc/paper_files/paper/2022/hash/827cb489449ea216e4a257c47e407d18-Abstract-Conference.html.

Yingyan Li, Lue Fan, Jiawei He, Yuqi Wang, Yuntao Chen, Zhaoxiang Zhang, and Tieniu Tan. Enhancing end-to-end autonomous driving with latent world model. In *International Conference on Learning Representations*, 2025a. URL https://openreview.net/forum?id=fd2u60ryG0.

Yingyan Li, Yuqi Wang, Yang Liu, Jiawei He, Lue Fan, and Zhaoxiang Zhang. End-to-end driving with online trajectory evaluation via BEV world model. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 27137–27146, 2025b. URL https://openaccess.thecvf.com/content/ICCV2025/html/Li_End-to-End_Driving_with_Online_Trajectory_Evaluation_via_BEV_World_Model_ICCV_2025_paper.html.

Zhiqi Li, Wenhai Wang, Hongyang Li, Enze Xie, Chonghao Sima, Tong Lu, Yu Qiao, and Jifeng Dai. BEVFormer: Learning bird's-eye-view representation from multi-camera images via spatiotemporal transformers. In *European Conference on Computer Vision*, pp. 1–18, 2022. URL https://arxiv.org/abs/2203.17270.

Chen Min, Dawei Zhao, Liang Xiao, Jian Zhao, Xinli Xu, Zheng Zhu, Lei Jin, Jianshu Li, Yulan Guo, Junliang Xing, Liping Jing, Yiming Nie, and Bin Dai. DriveWorld: 4D pre-trained scene understanding via world models for autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 15522–15533, 2024. URL https://openaccess.thecvf.com/content/CVPR2024/html/Min_DriveWorld_4D_Pre-trained_Scene_Understanding_via_World_Models_for_Autonomous_CVPR_2024_paper.html.

Bernhard Schölkopf, Francesco Locatello, Stefan Bauer, Nan Rosemary Ke, Nal Kalchbrenner, Anirudh Goyal, and Yoshua Bengio. Toward causal representation learning. *Proceedings of the IEEE*, 109(5):612–634, 2021. URL https://arxiv.org/abs/2102.11107.

Ziying Song, Caiyan Jia, Lin Liu, Lei Yang, Shengkai Zhang, Feiyang Jia, Fengda Zhao, Peiliang Wu, Shaoqing Xu, Chen Lv, and Yadan Luo. GraphWorld: Long-horizon planning with world models for end-to-end autonomous driving. *arXiv preprint arXiv:2606.16274*, 2026. URL https://arxiv.org/abs/2606.16274v1.

Yunpeng Zhang, Deheng Qian, Ding Li, Yifeng Pan, Yong Chen, Zhenbao Liang, Zhiyao Zhang, Yingzong Liu, Jianhui Mei, Maolei Fu, Yun Ye, Zhujin Liang, Yi Shan, and Dalong Du. GraphAD: Interaction scene graph for end-to-end autonomous driving. In *Proceedings of the Thirty-Fourth International Joint Conference on Artificial Intelligence*, pp. 2422–2430, 2025. doi: 10.24963/ijcai.2025/270. URL https://www.ijcai.org/proceedings/2025/270.

Wenzhao Zheng, Weiliang Chen, Yuanhui Huang, Borui Zhang, Yueqi Duan, and Jiwen Lu. OccWorld: Learning a 3D occupancy world model for autonomous driving. In *European Conference on Computer Vision*, 2024. URL https://www.ecva.net/papers/eccv_2024/papers_ECCV/papers/02024.pdf.

Yupeng Zheng, Pengxuan Yang, Zebin Xing, Qichao Zhang, Yuhang Zheng, Yinfeng Gao, Pengfei Li, Teng Zhang, Zhongpu Xia, Peng Jia, XianPeng Lang, and Dongbin Zhao. World4Drive: End-to-end autonomous driving via intention-aware physical latent world model. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 28632–28642, 2025. URL https://openaccess.thecvf.com/content/ICCV2025/html/Zheng_World4Drive_End-to-End_Autonomous_Driving_via_Intention-aware_Physical_Latent_World_Model_ICCV_2025_paper.html.

A PROTOCOLS, INPUT LINEAGE, AND REPRODUCTION

A.1 EVIDENCE ORGANIZATION

Each result belongs to a specified state contract, candidate pool, checkpoint, and evaluator. nuScenes evaluates representation and dynamics; NAVSIM evaluates planning. We retain resolved configurations, membership manifests, source and checkpoint identifiers, and numerical outputs. Tables generated from archived results preserve their original denominators and endpoint rules. Recoverable original artifacts take precedence over retrospective summaries. Where a historical launch or generation-time binding is absent, we record that field as unknown rather than infer it from a directory name or file modification time.

A.2 WHAT IS AVAILABLE TO THE MODEL

The evaluated native states are annotation-assisted. Physical features come from cached temporal BEV representations. Actor identity, geometry, and motion support semantic grounding and interaction construction. In the historical nuScenes construction, current semantic intent used adjacent annotations including a future frame. The new nuScenes comparison replaces this field with an audited history-based rule: 19,737 files covering 372,858 actor records were checked, with 25,459 intents changed and other semantic fields preserved. The inspected VLM prompt did not consume the replaced intent field.

This is a scoped nuScenes repair. Its inspected ego and interaction writers also use adjacent-frame motion estimates. Future state targets, evaluator-derived training labels, and proposed ego actions have separate roles from initial-state fields. Detailed field definitions accompany the dynamics appendix.

NAVSIM observation timing. The formal NAVHARD input adapter takes four observed frames with the current frame last, clears its successor link, and computes ego yaw rate from the current and preceding frames. This is the inspected history-only input path; future participants used for selection come from PCIWM rollout. Current annotation assistance and benchmark-aware search are separate limitations, not evidence that this adapter observes future ground truth.

The historical 419/186 development studies retain different initial caches. The raw-consequence cache generator rolls out the current state under candidate actions; future poses and boxes are used only for detached supervision costs, not as rollout inputs. A check of one corrected historical ego-cache anchor found a stored yaw rate matching the forward difference over the next 0.500 s, rather than the backward difference. Thus the history-only NAVHARD contract cannot be assigned to every historical cache. For historical C, replaying the saved writer at the same anchor reproduced all five cached intents; two actor paths read the successor, but removing it changed none of these five labels. The historical generation command is not bound by this check, so it does not establish generation-time availability or future dependence of those labels. We retain these field- and cache-specific distinctions rather than classify all NAVSIM inputs by one timing label.

A.3 TRAINING BUDGETS AND ENDPOINTS

The new nine-model dynamics comparison uses a fixed 1,484-pair subset of 14,837 training pairs, seed 7, and four epochs. All models are evaluated at epoch 3 on the same 149 six-transition chains from 30 development scenes; the original development manifest contains 765 pairs. The optimizer follows the first four epochs of the original 12-epoch schedule. The separate six-epoch PE/PCE extension does not replace the common-budget rows. Historical 12-epoch factor comparisons and the six-epoch multi-step recipe retain their own contracts; the latter uses 4,291 H5 training sequences and 214 development chains.

The four factor-inclusion models were subsequently repeated with training seeds 8 and 9, after inspection of seed 7. All eight runs and both fixed velocity endpoints were specified before observing their outcomes. They retain the same 1,484-pair pool, history-based C intent, 765-pair development manifest, four epochs and original cosine prefix. Each epoch uses 23 full batches of 64 with drop-last, two loader workers and prefetch factor 1; 92 denotes training batches, not a guarantee of successful AMP updates. Separate bounded epoch processes restore the preceding checkpoint,

including optimizer, scheduler, scaler and RNG state. H6 evaluation retains RNG seed 7 and the same 149 chains. All eight fixed-e3 evaluations completed; the four models across three seeds yield 32,184 saved chain–horizon–mode records. No run or endpoint is omitted for its direction.

A configuration-parser mismatch in an external audit rejected eight successful epoch-zero runs. A recorded continuation adopted those checkpoints once and completed epochs 1–3 with the original loader; training source, configuration bytes, sample pool and budget were unchanged. Both audit records and checkpoint identities are retained, without counting epoch zero twice.

For NAVSIM, the component readout uses 419 training and 186 development tokens, three providers, seeds 1701–1703, and 60 epochs: two minibatch updates per epoch, 120 updates per fit. Its fixed terminal epoch 59 and development-selected endpoint are both retained. The raw-consequence selector uses the same memberships and two input arms, with four epochs and 1,676 single-scene updates per fit. Its original seeds 1701–1703 were subsequently supplemented by seeds 1704–1705 after inspecting the first three outcomes. Both extra seeds and all four paired fits were specified before their results were observed; all five seeds, including the negative paired outcome, enter the final fixed-epoch mean. Its consequence encoder and scorer start from scratch; an archived architecture configuration does not warm-start their weights. Candidate trajectories, labels, initial parameters, and realized training orders are checked across the paired input arms. Fixed epoch 3 is distinguished from development-best selection. Both use K11 and eight half-second motion steps. The first-predicted-state persistence control adds three fits with seeds 1701–1703, the same membership, architecture, initialization and four-epoch budget. The full-sequence comparison was fixed before its outcomes; the first-state versus trajectory-only contrast was requested after observing them and is marked exploratory (Appendix E.1).

The auxiliary summary-readout analysis uses an explicitly reduced 438-token sample: 332 training, 39 development, 48 selector-test, and 19 historical-report tokens, with three sources, three seeds, and 32 epochs. Its altered provider/support contract and sample size prevent direct full-library score comparisons. The candidate-transfer analysis reuses frozen selectors; it introduces no additional selector training.

A.4 Development exposure and statistical units

The shared 419/186 NAVSIM memberships have no token or scene overlap, but share 66 recording groups. These groups contain 252 training and 172 development tokens. The complete development set spans 74 recordings, 112 logs, and 124 scenes. It remains a historical development set. For the component comparison, paired bootstrap draws resample whole recordings, retain all their tokens and multiplicities, and average the same three trained seeds. The 10,000 shared draws yield marginal percentile intervals conditional on those models; they are not confidence intervals over a population of training seeds. Development-best intervals do not correct model-selection bias. Across-seed dispersion is sample standard deviation, not a dataset confidence interval. The five-seed raw-consequence analysis follows the same grouping method, first averaging the five fitted models per token and then resampling whole recordings. All 10 fits, 40 per-epoch prediction files and 7,440 prediction rows were replayed. A separately specified scene-group analysis is reported as sensitivity analysis; it does not replace the recording-group endpoint.

Candidate-transfer evaluation separates 482 selector-test tokens in 16 recordings from the 186-token historical report set. The former comes from the historical training library; isolation from upstream world-model training or candidate-design decisions is not established. All reused selectors were trained on the guided pool. Equal candidate count therefore does not imply matched training exposure or perturbation amplitude.

A.5 Cache and evaluator contracts

The raw-consequence experiment uses the original epoch-5 NAVSIM world checkpoint and legacy K6/add-on/CV caches. Its name does not imply that these futures were regenerated with the later repaired provider. The component comparison retains the original mixed legacy-training and corrected-development lineage. Each provider recomputes components, aggregate priors, and fallback choices on common candidates and labels. Stored poses, source entries, and available hashes are verified; historical K6 labels lack a complete generation-time trajectory-hash binding.

A candidate oracle selects the highest evaluator label within the fixed pool, using information unavailable to the deployed selector. It is not a ground-truth-future input fed through that selector. A matched ground-truth-future-input experiment for the three-provider component comparison has not been verified. Similarly, action perturbation measures functional dependence without supplying unobserved intervention outcomes.

Raw candidate PDM and NAVHARD combined EPDMS are reported separately. Zero-rate analyses identify their thresholds: 10^{-8} for the residual ranker and candidate-transfer analysis, and 10^{-12} for the new raw-consequence input study. Each comparison retains its native threshold. NAVHARD results use the official evaluator locally but include benchmark-aware batch search and repeated development feedback. They are transductive development measurements, not leaderboard submissions or blind independent benchmark estimates.

A.6 Reproduction interfaces

Source profiles separate the latest implementation, the exact reduced-data dynamics implementation, and the frozen legacy planning implementation. They are named `source/latest`, `source/h1_reduced`, and `source/legacy`, respectively. Experiment adapters identify which profile they import. Portable templates and runtime materialization maintain separate hashes from the executed originals. CPU result replay verifies saved selections and regenerates tables; it does not rerun world-model inference. Training and cache generation additionally require licensed datasets, external front ends and devkits, and the corresponding checkpoints. Recorded environment versions are observations, not a complete dependency lock. Raw caches and pretrained weights are not bundled, so checkpoint hashes alone do not provide offline pretrained inference.

B State contracts and complete dynamics evidence

B.1 What the state variables contain

PCIE defines the structured native state $W_t = (P_t, C_t, I_t, E_t)$. The composer maps it to shared slots $Z_t = \mathcal{M}(W_t)$; an external motion input conditions their transition. The interaction decoder also receives the current structural context and action. Thus the operational interface is $\widehat{W}_{t+1} = F(W_t, A_t)$, while the native tuple and its latent representation are distinct objects. Table 6 lists the implemented fields rather than all optional fields supported by the general software interface.

The native C encoder uses deterministic one-hot records, followed by a trainable mapper. The VLM prompt asks for visual/relational semantic attributes; it does not predict the grounding intent. The inspected historical writer derives that intent partly from future annotations and passes it into the current input. The new past-intent transformation rewrote 19,737 cache files, covering 372,858 actor records and changing 25,459 intents. Independent byte-hash, non-intent-field and timestamp checks passed over the complete transformation. All other semantic fields and record order were preserved. This repairs the specified C route, not the entire observation pipeline. The inspected E and I writers retain their stated annotation/adjacent-frame dependencies.

For NAVSIM, the input contract must additionally be tied to the relevant cache generation: corrected coordinates do not imply past-only fields. The synthetic NAVHARD adaptation has a current-only raw-context exception; it should not be conflated with all training and internal-validation caches. These distinctions concern input availability, not whether the saved experiments can be reproduced.

Table 6: Native state fields and their verified input-time boundaries.

Factor	Implemented fields	Origin and time boundary
P	Temporal BEV feature, $256 \times 200 \times 200$ in the documented native configuration; patch size 8.	Cached BEVFormer representation. Feature MAE is not occupancy or semantic-map accuracy. Optional occupancy fields in the general interface do not establish their use in these runs.
C	32-dimensional native record: intent 9, rule 7, affordance 5, risk 3, right-of-way 4, confidence 1, padding 3. Grounding xy is separate.	Qwen predicts rule/affordance/risk/right-of-way; identity, geometry and intent come from grounding. Historical intent uses adjacent annotations including the future. The new controls replace only intent with the audited past-intent rule.
I	Agent identity and type; raw node 18 = category 10 + $x, y, v_x, v_y, \psi, l, w,$ is_ego; edge 9 = relative position/velocity, distance, bearing sine/cosine, TTC and closest-approach distance.	Annotation-derived agents and geometry/TTC graph. The inspected writer uses adjacent ego frames and SDK object velocity. Its native call supplies agents, not map nodes, although the general builder supports maps.
E	$[0, 0, 0, \text{speed}, v_x, v_y, \text{yawrate}, 0]$. Global pose is separately cached for pair construction.	The inspected nuScenes writer prioritizes current-to-next pose differencing. This audit does not establish a past-only ego interface. Zero position targets do not define a predicted-trajectory ADE.
A	Local incremental motion $(\Delta x, \Delta y, \Delta\psi)$, represented separately from PCIE.	Training/evaluation actions are obtained from consecutive ego poses. Planning integrates proposed increments. These are motion proxies, not calibrated actuator commands or randomized interventions.

The table identifies implemented routes, not an assertion that every historical cache was generated by the same immutable writer snapshot. Source hashes, line ranges, run configuration hashes and result pointers accompany the numerical release. The past-intent repair does not retroactively repair all historical C caches or verify sensor-only P/I/E inputs.

B.2 A COMMON BUDGET FOR THE REPRESENTATION AND ARCHITECTURE CONTROLS

The new nuScenes comparison uses the same 1,484 training pairs, selected by a fixed label-independent hash rule with selection seed 1701; the 765-pair development set is unchanged. All nine models start with training seed 7 and are compared at checkpoint epoch 3. Four training epochs use 23 complete batches per epoch, or 92 training batches. PE/PCE later extensions are not used in this common comparison.

The shared optimizer recipe is AdamW, learning rate 10^{-4}, minimum rate 10^{-6}, weight decay 10^{-4}, gradient clipping 1, batch 64 and AMP. The original cosine schedule retains $T_{\max} = 12$ during this four-epoch prefix. Runtime loading uses two workers and prefetch factor 1. Loss weights are 1 for BEV and .2 for token and ego; enabled C/I branches also use .2, while disabled branches use zero. MILE-style retains its 64-dimensional stochastic component and KL weight .01; OccWorld-style retains its 512-code codebook, commitment .25 and VQ weight .1. Their objectives are consequently not identical aggregate quantities.

All native PCIE controls use width 256, 128 world slots, eight attention heads, three composer layers and six dynamics layers. The interaction encoder has three layers, raw-node dimension 18, embedding dimension 64 and edge dimension 9. Table 7 identifies the architectural changes and capacities. Independent transitions still read the same already fused world slots. Monolithic composition removes per-factor alignment branches and gates, while retaining tokenization and the graph encoder. The style baselines are internally adapted models evaluated through the project interface, not full reproductions of the original MILE or OccWorld systems. Total parameter equality in the factor-ablation rows includes inactive modules; other architecture rows are not capacity-matched.

Table 7: Common H1 training controls and model capacity.

Model	C	I	Edge ID	Composer	Transition	Params (M)
PE	No	No	relation	struct.	unified	26.728
PCE	Yes	No	relation	struct.	unified	26.728
PIE	No	Yes	relation	struct.	unified	26.728
PCIE	Yes	Yes	relation	struct.	unified	26.728
Independent	Yes	Yes	relation	struct.	branch	46.487
Monolithic	Yes	Yes	relation	mono.	unified	22.365
E1 identity	Yes	Yes	ordered	struct.	unified	26.728
MILE-style	Yes	Yes	relation	n/a	n/a	18.403
OccWorld-style	Yes	Yes	relation	n/a	n/a	19.952

All runs use the same 1,484 selected training pairs and 765 development pairs, seed 7, batch 64, AMP, and a fixed epoch-3 checkpoint (four-epoch prefix; 92 training batches). The cosine schedule retains $T_{\max} = 12$, rather than restarting at four epochs. Total parameters include disabled modules; they are not an active-capacity match. Relation identity means (u, v, r) and ordered identity means (u, v).

B.3 Long-horizon behavior along recorded actions

All nine epoch-3 models are evaluated on the same 149 connected chains from 30 development scenes. Each chain supplies six recorded motion increments. AR recursively feeds predictions back, TF uses recorded current states at each step, and persistence copies the initial state. There are 2,682 chain–horizon–mode records per model. Actual H6 elapsed time has mean 2.9903 s and range 2.7996–3.1502 s; "3 s" is nominal. This is recorded-action-conditioned state prediction, not closed-loop policy driving or a ground-truth alternative-action test.

Tables 8–11 retain all seven common physical/ego readouts at H6 and over all six horizons. Every entry is a chain-weighted mean; the six-step aggregate is an equal-weight mean of the six horizon means. Speed and planar velocity errors are computed from their native components. The zero-valued position fields are not reported as ADE/FDE. Appendix D adds paired scene intervals conditional on the completed seed-7 models; these do not estimate training-seed uncertainty.

For seed 7, adding C improves all seven common AR readouts at H6 and in the six-step mean for both controlled factor pairs. This observation is specific to that seed, the stated budget and corrected-intent route; the repetition below does not establish a stable C benefit. It does not imply that unified or structured PCIE wins every architecture comparison: branch transitions and monolithic composition can have lower speed/velocity errors, while the trade-off differs for yaw rate. Persistence has lower BEV feature error than the learned forecasts throughout these horizons, which limits claims about the physical prediction objective.

Table 8: H6 AR: all seven common readouts (lower is better).

Model	Speed	Vel. L2	Yaw rate	v_x	v_y	BEV MAE	BEV SL1
PE	0.7519	0.8187	0.0539	0.8032	0.0638	0.2874	0.0719
PCE	0.6385	0.6701	0.0379	0.6530	0.0516	0.2799	0.0688
PIE	0.7542	0.7723	0.0391	0.7540	0.0565	0.2802	0.0690
PCIE	0.5832	0.6173	0.0369	0.6002	0.0535	0.2797	0.0688
Independent	0.5190	0.5499	0.0382	0.5294	0.0546	0.2794	0.0686
Monolithic	0.4357	0.4691	0.0461	0.4481	0.0627	0.2813	0.0693
E1 identity	0.7818	0.7789	0.0383	0.7708	0.0577	0.2804	0.0690
MILE-style	1.3823	1.4336	0.0413	1.4112	0.1139	0.2966	0.0759
OccWorld-style	2.0393	2.1036	0.0735	2.0917	0.1199	0.2876	0.0718
Persistence	1.0103	1.0313	0.0549	1.0107	0.0754	0.2483	0.0636

149 chains from 30 scenes; one fixed seed and four training epochs. Speed, velocity L2, v_x and v_y use m/s; yaw rate uses rad/s; BEV values are feature-space quantities. SL1 denotes Smooth L1. The six-step mean averages the six horizon means; it is not a trajectory ADE. Persistence copies the initial state and is shared across models.

Table 9: Equal-weight mean over H1–H6, AR: all seven common readouts (lower is better).

Model	Speed	Vel. L2	Yaw rate	v_x	v_y	BEV MAE	BEV SL1
PE	0.4970	0.5543	0.0402	0.5397	0.0523	0.2819	0.0696
PCE	0.4742	0.4972	0.0353	0.4814	0.0481	0.2800	0.0689
PIE	0.5649	0.5725	0.0371	0.5576	0.0512	0.2801	0.0689
PCIE	0.4547	0.4834	0.0342	0.4679	0.0486	0.2798	0.0688
Independent	0.4028	0.4245	0.0358	0.4076	0.0492	0.2794	0.0686
Monolithic	0.4012	0.4413	0.0425	0.4236	0.0581	0.2814	0.0693
E1 identity	0.6004	0.5980	0.0350	0.5891	0.0540	0.2804	0.0690
MILE-style	1.3795	1.4313	0.0392	1.4131	0.1111	0.2965	0.0757
OccWorld-style	2.0862	2.1478	0.0726	2.1362	0.1169	0.2897	0.0728
Persistence	0.5954	0.6091	0.0346	0.5961	0.0501	0.2198	0.0525

149 chains from 30 scenes; one fixed seed and four training epochs. Speed, velocity L2, v_x and v_y use m/s; yaw rate uses rad/s; BEV values are feature-space quantities. SL1 denotes Smooth L1. The six-step mean averages the six horizon means; it is not a trajectory ADE. Persistence copies the initial state and is shared across models.

Table 10: H6 TF: all seven common readouts (lower is better).

Model	Speed	Vel. L2	Yaw rate	v_x	v_y	BEV MAE	BEV SL1
PE	0.2679	0.3230	0.0355	0.3007	0.0531	0.2794	0.0685
PCE	0.3002	0.3237	0.0369	0.3004	0.0531	0.2793	0.0685
PIE	0.2983	0.3184	0.0404	0.2975	0.0523	0.2801	0.0689
PCIE	0.3038	0.3326	0.0344	0.3100	0.0507	0.2793	0.0686
Independent	0.2994	0.3200	0.0381	0.2995	0.0527	0.2788	0.0683
Monolithic	0.4309	0.4621	0.0484	0.4420	0.0584	0.2812	0.0693
E1 identity	0.3910	0.4069	0.0351	0.3917	0.0563	0.2798	0.0687
MILE-style	1.2337	1.2875	0.0407	1.2641	0.1107	0.2948	0.0751
OccWorld-style	1.8798	1.9097	0.0847	1.8959	0.1409	0.2982	0.0769
Persistence	1.0103	1.0313	0.0549	1.0107	0.0754	0.2483	0.0636

149 chains from 30 scenes; one fixed seed and four training epochs. Speed, velocity L2, v_x and v_y use m/s; yaw rate uses rad/s; BEV values are feature-space quantities. SL1 denotes Smooth L1. The six-step mean averages the six horizon means; it is not a trajectory ADE. Persistence copies the initial state and is shared across models.

Table 11: Equal-weight mean over H1–H6, TF: all seven common readouts (lower is better).

Model	Speed	Vel. L2	Yaw rate	v_x	v_y	BEV MAE	BEV SL1
PE	0.2440	0.2946	0.0326	0.2760	0.0488	0.2795	0.0685
PCE	0.2747	0.2925	0.0343	0.2728	0.0492	0.2797	0.0687
PIE	0.2878	0.2960	0.0377	0.2785	0.0486	0.2802	0.0689
PCIE	0.2806	0.3038	0.0325	0.2855	0.0467	0.2796	0.0687
Independent	0.2779	0.2930	0.0352	0.2757	0.0482	0.2790	0.0684
Monolithic	0.3969	0.4357	0.0451	0.4184	0.0559	0.2813	0.0693
E1 identity	0.3538	0.3673	0.0329	0.3533	0.0531	0.2800	0.0688
MILE-style	1.2782	1.3318	0.0385	1.3126	0.1093	0.2951	0.0752
OccWorld-style	1.9657	1.9996	0.0816	1.9865	0.1350	0.2989	0.0772
Persistence	0.5954	0.6091	0.0346	0.5961	0.0501	0.2198	0.0525

149 chains from 30 scenes; one fixed seed and four training epochs. Speed, velocity L2, v_x and v_y use m/s; yaw rate uses rad/s; BEV values are feature-space quantities. SL1 denotes Smooth L1. The six-step mean averages the six horizon means; it is not a trajectory ADE. Persistence copies the initial state and is shared across models.

B.4 C-FACTOR REPETITION ACROSS TRAINING SEEDS

PE/PCE/PIE/PCIE are compared across seeds 7, 8 and 9 at the same fixed epoch 3, using the same corrected-intent inputs, training pool, schedule and 149 evaluation chains. These repetitions are distinct from the single-seed nine-model architecture comparison above. Every new run completed the prescribed four epochs and H6 evaluation; no partial run is substituted.

The paired difference is with-C minus without-C, so negative error differences favor C. PCE minus PE has mean H6 velocity-L2 difference -0.0610 ± 0.0860 m/s and six-step difference -0.0185 ± 0.0343 m/s; two of three fitted seeds improve at each endpoint. PCIE minus PIE instead has differences $+0.0535 \pm 0.2310$ and $+0.0635 \pm 0.1666$ m/s, with improvement only for seed 7 at both endpoints. The favorable seed-7 contrasts therefore do not establish stable incremental C utility for PIE under this budget. These are means and sample SDs of three paired fitted instances, not population confidence intervals or converged comparisons.

The four common-readout tables report all seven P/E metrics, with model means and SDs on separate rows. The following eight paired tables retain every fitted seed, both factor pairs, both endpoints and the AR and TF modes, including adverse and zero differences. Signs use unrounded values; small changes can round to zero. Persistence is one shared deterministic baseline; its zero differences are summarized below and retained in the data files. All six horizons remain in the accompanying 1,512-row CSV. This repetition does not add C/I-head or graph-support evaluations to the seed-7 evidence. The scene-bootstrap intervals in Appendix D remain conditional on seed 7 and must not be applied to these three-seed means.

Table 12: AR, H6: all common factor-comparison readouts.

Model	Speed	Vel. L2	Yaw rate	v_x	v_y	BEV MAE	BEV SL1
PE mean	0.8743	0.9000	0.0449	0.8879	0.0642	0.2844	0.0706
PE SD	0.1170	0.0928	0.0080	0.0969	0.0118	0.0027	0.0012
PCE mean	0.8131	0.8390	0.0421	0.8240	0.0623	0.2852	0.0711
PCE SD	0.1757	0.1778	0.0036	0.1789	0.0094	0.0088	0.0039
PIE mean	0.7476	0.7723	0.0440	0.7550	0.0678	0.2839	0.0704
PIE SD	0.0395	0.0571	0.0083	0.0590	0.0120	0.0054	0.0022
PCIE mean	0.7960	0.8257	0.0437	0.8139	0.0591	0.2825	0.0699
PCIE SD	0.2141	0.2004	0.0063	0.2053	0.0074	0.0036	0.0015
Persistence	1.0103	1.0313	0.0549	1.0107	0.0754	0.2483	0.0636

Speed, velocity L2, v_x, and v_y are in m/s; yaw rate is in rad/s; BEV errors are feature-space values. SL1 denotes Smooth L1. Lower is better. Three training seeds (7, 8, 9), fixed epoch 3, 149 common chains from 30 scenes; mean $\pm$ sample SD is descriptive, not a confidence interval. Means and sample SDs occupy separate rows to retain readable precision. Persistence repeats the initial state and is shared across models and seeds.

Table 13: AR, Equal-weight mean H1–H6: all common factor-comparison readouts.

Model	Speed	Vel. L2	Yaw rate	v_x	v_y	BEV MAE	BEV SL1
PE mean	0.6011	0.6272	0.0386	0.6145	0.0556	0.2815	0.0694
PE SD	0.0922	0.0692	0.0016	0.0714	0.0100	0.0004	0.0002
PCE mean	0.5856	0.6087	0.0381	0.5950	0.0544	0.2821	0.0697
PCE SD	0.1002	0.1031	0.0024	0.1046	0.0068	0.0033	0.0014
PIE mean	0.5582	0.5789	0.0405	0.5634	0.0596	0.2825	0.0698
PIE SD	0.0187	0.0303	0.0060	0.0311	0.0097	0.0031	0.0012
PCIE mean	0.6137	0.6424	0.0403	0.6306	0.0529	0.2818	0.0695
PCIE SD	0.1647	0.1552	0.0060	0.1584	0.0068	0.0023	0.0009
Persistence	0.5954	0.6091	0.0346	0.5961	0.0501	0.2198	0.0525

Speed, velocity L2, v_x, and v_y are in m/s; yaw rate is in rad/s; BEV errors are feature-space values. SL1 denotes Smooth L1. Lower is better. Three training seeds (7, 8, 9), fixed epoch 3, 149 common chains from 30 scenes; mean $\pm$ sample SD is descriptive, not a confidence interval. Means and sample SDs occupy separate rows to retain readable precision. Persistence repeats the initial state and is shared across models and seeds.

Table 14: TF, H6: all common factor-comparison readouts.

Model	Speed	Vel. L2	Yaw rate	v_x	v_y	BEV MAE	BEV SL1
PE mean	0.3205	0.3513	0.0382	0.3323	0.0556	0.2797	0.0687
PE SD	0.0456	0.0252	0.0024	0.0282	0.0068	0.0003	0.0002
PCE mean	0.3203	0.3464	0.0384	0.3260	0.0556	0.2794	0.0686
PCE SD	0.0208	0.0206	0.0036	0.0234	0.0059	0.0002	0.0001
PIE mean	0.3108	0.3408	0.0428	0.3183	0.0593	0.2811	0.0692
PIE SD	0.0118	0.0195	0.0033	0.0180	0.0088	0.0009	0.0003
PCIE mean	0.3862	0.4201	0.0421	0.4016	0.0547	0.2805	0.0690
PCIE SD	0.0848	0.0810	0.0067	0.0853	0.0049	0.0011	0.0004
Persistence	1.0103	1.0313	0.0549	1.0107	0.0754	0.2483	0.0636

Speed, velocity L2, v_x, and v_y are in m/s; yaw rate is in rad/s; BEV errors are feature-space values. SL1 denotes Smooth L1. Lower is better. Three training seeds (7, 8, 9), fixed epoch 3, 149 common chains from 30 scenes; mean $\pm$ sample SD is descriptive, not a confidence interval. Means and sample SDs occupy separate rows to retain readable precision. Persistence repeats the initial state and is shared across models and seeds.

Table 15: TF, Equal-weight mean H1–H6: all common factor-comparison readouts.

Model	Speed	Vel. L2	Yaw rate	v_x	v_y	BEV MAE	BEV SL1
PE mean	0.2961	0.3242	0.0352	0.3078	0.0513	0.2797	0.0687
PE SD	0.0452	0.0260	0.0022	0.0281	0.0066	0.0003	0.0001
PCE mean	0.2978	0.3210	0.0363	0.3039	0.0511	0.2798	0.0687
PCE SD	0.0233	0.0257	0.0037	0.0280	0.0058	0.0003	0.0001
PIE mean	0.2917	0.3162	0.0397	0.2968	0.0548	0.2812	0.0693
PIE SD	0.0060	0.0174	0.0033	0.0159	0.0078	0.0009	0.0003
PCIE mean	0.3591	0.3900	0.0395	0.3744	0.0503	0.2807	0.0691
PCIE SD	0.0820	0.0799	0.0062	0.0827	0.0049	0.0011	0.0004
Persistence	0.5954	0.6091	0.0346	0.5961	0.0501	0.2198	0.0525

Speed, velocity L2, v_x, and v_y are in m/s; yaw rate is in rad/s; BEV errors are feature-space values. SL1 denotes Smooth L1. Lower is better. Three training seeds (7, 8, 9), fixed epoch 3, 149 common chains from 30 scenes; mean $\pm$ sample SD is descriptive, not a confidence interval. Means and sample SDs occupy separate rows to retain readable precision. Persistence repeats the initial state and is shared across models and seeds.

Table 16: PCE $-$ PE, AR, H6: every paired seed.

Seed	Speed	Vel. L2	Yaw rate	v_x	v_y	BEV MAE	BEV SL1
7	-0.1134	-0.1486	-0.0160	-0.1502	-0.0121	-0.0075	-0.0031
8	-0.0750	-0.0579	$+0.0024$	-0.0577	-0.0102	-0.0031	-0.0012
9	$+0.0048$	$+0.0234$	$+0.0052$	$+0.0162$	$+0.0167$	$+0.0131$	$+0.0060$
Mean	-0.0612	-0.0610	-0.0028	-0.0639	-0.0019	$+0.0008$	$+0.0005$
Sample SD	0.0603	0.0860	0.0115	0.0834	0.0161	0.0108	0.0048

Differences are with-C minus without-C; negative favors C. Speed, velocity L2, v_x, and v_y are in m/s; yaw rate is in rad/s; BEV errors are feature-space values. SL1 denotes Smooth L1. Lower is better. Three training seeds (7, 8, 9), fixed epoch 3, 149 common chains from 30 scenes; mean $\pm$ sample SD is descriptive, not a confidence interval.

Table 17: PCE − PE, AR, Equal-weight mean H1–H6: every paired seed.

Seed	Speed	Vel. L2	Yaw rate	v_x	v_y	BEV MAE	BEV SL1
7	−0.0228	−0.0571	−0.0049	−0.0584	−0.0042	−0.0019	−0.0007
8	−0.0199	−0.0069	+0.0011	−0.0055	−0.0053	−0.0014	−0.0006
9	−0.0040	+0.0086	+0.0023	+0.0054	+0.0059	+0.0048	+0.0021
Mean	−0.0156	−0.0185	−0.0005	−0.0195	−0.0012	+0.0005	+0.0003
Sample SD	0.0101	0.0343	0.0038	0.0341	0.0062	0.0037	0.0016

Differences are with-C minus without-C; negative favors C. Speed, velocity L2, v_x, and v_y are in m/s; yaw rate is in rad/s; BEV errors are feature-space values. SL1 denotes Smooth L1. Lower is better. Three training seeds (7, 8, 9), fixed epoch 3, 149 common chains from 30 scenes; mean ± sample SD is descriptive, not a confidence interval.

Table 18: PCIE − PIE, AR, H6: every paired seed.

Seed	Speed	Vel. L2	Yaw rate	v_x	v_y	BEV MAE	BEV SL1
7	−0.1710	−0.1550	−0.0023	−0.1538	−0.0030	−0.0005	−0.0002
8	+0.3061	+0.3018	−0.0087	+0.3132	−0.0128	−0.0001	+0.0000
9	+0.0102	+0.0136	+0.0100	+0.0174	−0.0103	−0.0036	−0.0015
Mean	+0.0484	+0.0535	−0.0003	+0.0589	−0.0087	−0.0014	−0.0006
Sample SD	0.2408	0.2310	0.0095	0.2363	0.0051	0.0019	0.0008

Differences are with-C minus without-C; negative favors C. Speed, velocity L2, v_x, and v_y are in m/s; yaw rate is in rad/s; BEV errors are feature-space values. SL1 denotes Smooth L1. Lower is better. Three training seeds (7, 8, 9), fixed epoch 3, 149 common chains from 30 scenes; mean ± sample SD is descriptive, not a confidence interval.

Table 19: PCIE − PIE, AR, Equal-weight mean H1–H6: every paired seed.

Seed	Speed	Vel. L2	Yaw rate	v_x	v_y	BEV MAE	BEV SL1
7	−0.1102	−0.0891	−0.0030	−0.0896	−0.0025	−0.0004	−0.0002
8	+0.2465	+0.2412	−0.0066	+0.2488	−0.0094	−0.0001	−0.0001
9	+0.0302	+0.0384	+0.0092	+0.0426	−0.0083	−0.0017	−0.0006
Mean	+0.0555	+0.0635	−0.0001	+0.0672	−0.0068	−0.0007	−0.0003
Sample SD	0.1797	0.1666	0.0083	0.1705	0.0037	0.0009	0.0003

Differences are with-C minus without-C; negative favors C. Speed, velocity L2, v_x, and v_y are in m/s; yaw rate is in rad/s; BEV errors are feature-space values. SL1 denotes Smooth L1. Lower is better. Three training seeds (7, 8, 9), fixed epoch 3, 149 common chains from 30 scenes; mean ± sample SD is descriptive, not a confidence interval.

Table 20: PCE − PE, TF, H6: every paired seed.

Seed	Speed	Vel. L2	Yaw rate	v_x	v_y	BEV MAE	BEV SL1
7	+0.0323	+0.0007	+0.0013	−0.0002	+0.0000	−0.0001	+0.0000
8	−0.0057	+0.0040	+0.0023	+0.0047	−0.0009	−0.0008	−0.0003
9	−0.0273	−0.0197	−0.0030	−0.0234	+0.0010	+0.0001	+0.0000
Mean	−0.0002	−0.0050	+0.0002	−0.0063	+0.0000	−0.0003	−0.0001
Sample SD	0.0302	0.0128	0.0028	0.0150	0.0009	0.0004	0.0002

Differences are with-C minus without-C; negative favors C. Speed, velocity L2, v_x, and v_y are in m/s; yaw rate is in rad/s; BEV errors are feature-space values. SL1 denotes Smooth L1. Lower is better. Three training seeds (7, 8, 9), fixed epoch 3, 149 common chains from 30 scenes; mean ± sample SD is descriptive, not a confidence interval.

Table 21: PCE − PE, TF, Equal-weight mean H1–H6: every paired seed.

Seed	Speed	Vel. L2	Yaw rate	v_x	v_y	BEV MAE	BEV SL1
7	+0.0307	−0.0021	+0.0017	−0.0031	+0.0005	+0.0002	+0.0001
8	−0.0027	+0.0079	+0.0037	+0.0088	−0.0012	−0.0004	−0.0002
9	−0.0229	−0.0155	−0.0018	−0.0175	+0.0002	+0.0003	+0.0001
Mean	+0.0017	−0.0032	+0.0012	−0.0039	−0.0002	+0.0000	+0.0000
Sample SD	0.0271	0.0117	0.0028	0.0132	0.0009	0.0004	0.0002

Differences are with-C minus without-C; negative favors C. Speed, velocity L2, v_x, and v_y are in m/s; yaw rate is in rad/s; BEV errors are feature-space values. SL1 denotes Smooth L1. Lower is better. Three training seeds (7, 8, 9), fixed epoch 3, 149 common chains from 30 scenes; mean ± sample SD is descriptive, not a confidence interval.

Table 22: PCIE − PIE, TF, H6: every paired seed.

Seed	Speed	Vel. L2	Yaw rate	v_x	v_y	BEV MAE	BEV SL1
7	+0.0055	+0.0142	−0.0060	+0.0125	−0.0017	−0.0008	−0.0003
8	+0.1513	+0.1417	−0.0011	+0.1514	−0.0091	−0.0007	−0.0003
9	+0.0691	+0.0819	+0.0052	+0.0860	−0.0029	−0.0004	−0.0001
Mean	+0.0753	+0.0793	−0.0007	+0.0833	−0.0046	−0.0006	−0.0002
Sample SD	0.0731	0.0638	0.0056	0.0695	0.0040	0.0002	0.0001

Differences are with-C minus without-C; negative favors C. Speed, velocity L2, v_x, and v_y are in m/s; yaw rate is in rad/s; BEV errors are feature-space values. SL1 denotes Smooth L1. Lower is better. Three training seeds (7, 8, 9), fixed epoch 3, 149 common chains from 30 scenes; mean ± sample SD is descriptive, not a confidence interval.

Table 23: PCIE − PIE, TF, Equal-weight mean H1–H6: every paired seed.

Seed	Speed	Vel. L2	Yaw rate	v_x	v_y	BEV MAE	BEV SL1
7	−0.0073	+0.0078	−0.0052	+0.0071	−0.0019	−0.0006	−0.0003
8	+0.1456	+0.1352	−0.0020	+0.1438	−0.0077	−0.0005	−0.0002
9	+0.0639	+0.0785	+0.0064	+0.0821	−0.0042	−0.0003	−0.0001
Mean	+0.0674	+0.0738	−0.0003	+0.0776	−0.0046	−0.0005	−0.0002
Sample SD	0.0765	0.0639	0.0060	0.0685	0.0029	0.0001	0.0001

Differences are with-C minus without-C; negative favors C. Speed, velocity L2, v_x, and v_y are in m/s; yaw rate is in rad/s; BEV errors are feature-space values. SL1 denotes Smooth L1. Lower is better. Three training seeds (7, 8, 9), fixed epoch 3, 149 common chains from 30 scenes; mean ± sample SD is descriptive, not a confidence interval.

Shared persistence baseline. Persistence uses identical initial observations across the four models and three seeds. All seven paired C-inclusion differences are therefore exactly zero at H6 and over six steps. The four redundant zero-difference tables are omitted from typesetting; their complete values remain in the accompanying machine-readable results.

B.5 Semantic and interaction outputs: applicability and support

Tables 24 and 25 omit disabled C/I heads rather than treating their untrained numerical outputs as performance. The evaluator anchors node identities and ordered edge endpoints to the initial support. Relation types can change on that support; new nodes and edges are measured separately as coverage and birth statistics (Table 26). This protocol does not demonstrate open-world graph generation. E1 changes the training identity contract and must be distinguished from a same-target accuracy ablation.

In the seed-7 study, TF relation scores can be strong while AR changed-relation macro-F1 collapses. The main PCIE row has zero H6 AR changed-relation macro-F1 in this short H1-training protocol. Its C error also exceeds persistence at H6. These outcomes qualify the favorable common ego readouts and should not be hidden behind an aggregate graph or total loss.

Table 24: Applicable C/I readouts at H6, AR.

Model	C MAE	Node MAE	Edge MAE	Node F1	Edge F1	Type mF1	Chg. mF1
PCE	0.0997	—	—	—	—	—	—
PIE	—	0.8133	1.9974	0.8872	0.7867	0.7850	0.0000
PCIE	0.0976	0.8214	1.9882	0.8872	0.7867	0.7854	0.0000
Independent	0.1126	0.8677	2.0541	0.8872	0.7870	0.7857	0.0000
Monolithic	0.0960	0.8633	1.9761	0.8872	0.7865	0.7841	0.0003
E1 identity	0.0993	0.8286	2.0116	0.8872	0.7865	0.7850	0.0119
MILE-style	0.0944	0.8877	2.0722	0.8872	0.7865	0.7849	0.0003
OccWorld-style	0.1015	0.8441	1.9631	0.8872	0.7865	0.7852	0.0000
Persistence	0.0917	0.8557	1.6513	0.8872	0.7865	0.7852	0.0000

Continuous errors decrease; F1 scores increase and are fractions. Dashes denote disabled heads, never zero error. Graph scores cover the common initial support. E1 changes the training edge-identity contract; therefore its changed-relation score is not a pure same-target accuracy comparison with the legacy-identity rows.

Table 25: Applicable C/I readouts at H6, TF.

Model	C MAE	Node MAE	Edge MAE	Node F1	Edge F1	Type mF1	Chg. mF1
PCE	0.0978	—	—	—	—	—	—
PIE	—	0.1410	0.7805	0.9773	0.9571	0.9273	0.8055
PCIE	0.0968	0.1432	0.8135	0.9773	0.9570	0.9273	0.8055
Independent	0.1116	0.1504	0.8039	0.9773	0.9562	0.9273	0.8055
Monolithic	0.0959	0.1498	0.7937	0.9773	0.7864	0.9258	0.8052
E1 identity	0.0981	0.1452	0.7942	0.9773	0.9573	0.9274	0.8044
MILE-style	0.0948	0.1527	0.7878	0.8872	0.9573	0.9228	0.8052
OccWorld-style	0.1215	0.1472	0.7831	0.8872	0.9573	0.9228	0.8052
Persistence	0.0917	0.8557	1.6513	0.8872	0.7865	0.7852	0.0000

Continuous errors decrease; F1 scores increase and are fractions. Dashes denote disabled heads, never zero error. Graph scores cover the common initial support. E1 changes the training edge-identity contract; therefore its changed-relation score is not a pure same-target accuracy comparison with the legacy-identity rows.

Table 26: Coverage of the initially fixed interaction support.

Horizon	Node coverage	Edge coverage	New nodes/chain	New edges/chain
H1	0.9586	0.9275	1.2953	57.4295
H2	0.9231	0.8729	2.4228	100.5302
H3	0.8922	0.8230	3.3893	138.7450
H4	0.8638	0.7764	4.2617	173.2550
H5	0.8359	0.7290	5.0940	206.5436
H6	0.8056	0.6801	5.9866	240.6980

Coverage and births describe the held-out chains, not model-generated births. New participants and relations are reported separately and are not silently counted as tracked predictions. The same support is used by all nine evaluations.

B.6 Historical controls and sequence-training evidence

The original 12-epoch composition controls are recorded separately in Table 27. Their historical intent cache and full training-pair budget differ from the new four-epoch controls. Loss components are not interchangeable with observable MAEs, and totals or latent-token objectives are not used to rank different output sets. Table 28 compares E0-R and E1 at the same epoch 11. The largest changes are concentrated in edge prediction; node/BEV/C are nearly unchanged and ego loss is slightly worse. The 15.63% changed-edge figure sometimes associated with this comparison is an E1 within-run epoch-0-to-11 decrease, not an E0-R-to-E1 relative gain.

Table 27: Historical full-pair composition controls at the same final epoch 11.

State	BEV loss	Ego loss	C loss	Node loss	Edge loss	Graph loss
PE	0.051536	0.009452	—	—	—	—
PCE	0.051602	0.009064	0.069697	—	—	—
PIE	0.052144	0.010884	—	0.193935	1.478330	2.887253
PCIE	0.051912	0.010641	0.070371	0.193601	1.476625	2.882590

Original CSV loss components, not the observable MAEs in the H6 tables. All four ran 12 epochs; PE/PCE used two loader workers and PIE/PCIE eight. Historical C retains annotation-derived intent. No total or token loss is used to rank models with different output sets.

Table 28: Edge identity at the same epoch 11: E0-R versus E1.

Validation loss component	E0-R	E1	Relative change (%)
Total	0.645917	0.604131	-6.469300
Graph	2.884102	2.671392	-7.375238
Edge	1.481531	1.265175	-14.603586
Edge existence	0.517094	0.490892	-5.067237
Node	0.193663	0.193687	0.012750
BEV	0.051856	0.051920	0.122565
C	0.070102	0.070240	0.196471
Ego	0.010295	0.010571	2.674649

Negative relative changes are lower losses. E0-R uses (u, v, r) identity; E1 uses ordered endpoints (u, v). E0-R dev-best is epoch 9 and E1 dev-best is epoch 11; this table deliberately compares the same training endpoint. The often cited 15.63% changed-edge decrease is E1 epoch-0-to-11, not an E0-R-to-E1 effect; E0-R reports zero for that differently defined target.

Historical full214 contains 214 valid H5 chains from the 765-pair development subset. It is not the full official validation split. The E0 four-model and unified/branch-transition comparisons are shown in Tables 29 and 30. Unified PCIE has lower H5 ego amplification than branch-specific transitions, while several local C/I metrics favor the latter. The old structured/monolithic rollout used unmatched checkpoint budgets and is not used as an architecture advantage; the new common-endpoint comparison above supplies the matched-budget evidence.

Table 29: Historical H5 AR comparison on the same 214 chains.

Model	BEV MAE	C MAE	Ego MAE	Node MAE	Edge MAE	Edge F1	Chg. mF1
PCIE	0.2430	0.0770	0.1800	0.3205	1.5383	0.8232	0.0056
PIE	0.2458	—	0.1929	0.3220	1.5424	0.8216	0.0262
MILE-style	0.2457	0.0743	0.1834	0.3276	1.5493	0.8251	0.0056
OccWorld-style	0.2467	0.0731	0.0835	0.3179	1.5422	0.8211	0.0054
Independent	0.2439	0.0751	0.2087	0.3177	1.5273	0.8241	0.0056
Persistence	0.2459	0.0823	0.2324	0.7652	1.5283	0.8161	0.0000

These are the E0/paper-v0 family, not final E4. PCIE/PIE/Independent use epoch 11; internal MILE-/OccWorld-style controls use selected epoch 10 after 12-epoch training. Style models are project adaptations, not official reproductions. All modes receive the recorded action sequence. PIE C is not evaluated.

Table 30: Raw H5 AR/TF errors and their ratios in the historical controls.

Model	BEV AR	BEV TF	AR/TF	Ego AR	Ego TF	AR/TF
PCIE	0.24303	0.23777	1.02213	0.18003	0.06558	2.74537
PIE	0.24583	0.23870	1.02983	0.19293	0.07227	2.66938
MILE-style	0.24572	0.23995	1.02402	0.18343	0.06948	2.64010
OccWorld-style	0.24670	0.24284	1.01592	0.08349	0.08357	0.99904
Independent	0.24394	0.23710	1.02885	0.20872	0.06088	3.42818

Ratios divide observable MAEs, not weighted objectives. Independent uses branch-specific transition parameters over the same fused state and has 46.49M parameters versus 26.73M for unified PCIE. Reduced ego amplification supports this particular historical comparison, not uniformly better factor prediction.

Table 31 distinguishes identity correction, teacher-forced sequence training, predicted-state feedback and the changed-relation weight sweep. E2-TF, E2-AR and changed-loss variants are experimental branches rather than a sequential initialization chain. The .20 weight improves changed-relation F1 but degrades overall relation type F1 relative to .05. E4 is not uniformly better than every branch.

The historical E4 uses six epochs with H3 at epochs 0–1 and H5 at epochs 2–5, detached predicted-state feedback probability .50 to .75, seed 7 and feedback seed 1701. AdamW uses learning rate 5×10^{-5}, minimum 10^{-6}, batch 8 with eight-step accumulation, and changed-edge weight .05. Its original 14,837 pairs yield 4,291 H5 training sequences; the final reported checkpoint is epoch 5, selected by development AR loss. Its resolved configuration and evaluator output are recovered. The surviving launch-log inspection contains no initialization-checkpoint field, so the exact warm-start ancestor is not inferred from conflicting earlier recipe text. This provenance gap is not filled with a claimed serial E3-to-E4 ancestry. Table 32 reports all evaluated modes/horizons for that checkpoint.

Table 31: Historical H5 AR dynamics and changed-relation weighting, full214.

Stage	BEV MAE	C MAE	Ego MAE	Node MAE	Edge F1	Type mF1	Chg. mF1
E0	0.2430	0.0770	0.1800	0.3205	0.8232	0.8151	0.0056
E1	0.2433	0.0761	0.1949	0.3231	0.8282	0.8299	0.2781
E2-TF	0.2462	0.0803	0.1923	0.3344	0.8385	0.8325	0.2833
E2-AR	0.2429	0.0782	0.1601	0.3235	0.8583	0.8339	0.2985
E3, .05	0.2428	0.0773	0.1599	0.3209	0.8611	0.8455	0.4278
E3, .10	0.2434	0.0782	0.1605	0.3190	0.8617	0.8325	0.4639
E3, .20	0.2450	0.0794	0.1600	0.3222	0.8686	0.7771	0.5171
E4	0.2407	0.0767	0.1113	0.3306	0.8707	0.8460	0.4593

E2-TF/E2-AR/E3 are recipe branches, not a serial checkpoint chain. The .05/.10/.20 rows are distinct changed-relation weights; E2-AR provides the zero-weight reference. E4 uses .05 and a longer-horizon curriculum, so it is not a single-factor successor to the .20 row. Births remain outside tracked support. Errors decrease and F1 scores increase.

Table 32: Selected historical E4: every evaluated horizon and mode.

Mode	BEV MAE	C MAE	Ego MAE	Node MAE	Edge F1	Type mF1	Chg. mF1
H1 TF	0.23697	0.07620	0.07359	0.08053	0.96343	0.89797	0.52283
H1 AR	0.23697	0.07620	0.07359	0.08053	0.96343	0.89797	0.52283
H1 Persist.	0.17421	0.07043	0.05163	0.18209	0.96366	0.93114	0.00000
H3 TF	0.23786	0.07669	0.07215	0.07505	0.96192	0.89150	0.70782
H3 AR	0.23937	0.07514	0.09189	0.22623	0.91683	0.87562	0.47737
H3 Persist.	0.22639	0.08068	0.14267	0.50179	0.89379	0.85852	0.00000
H5 TF	0.23742	0.07803	0.07605	0.07066	0.95683	0.88567	0.77801
H5 AR	0.24072	0.07669	0.11132	0.33064	0.87071	0.84600	0.45934
H5 Persist.	0.24594	0.08226	0.23243	0.76521	0.81611	0.81146	0.00000

Checkpoint epoch 5, seed 7; 214 chains and recorded actions. This is the historical sequence-trained E4, separate from the new four-epoch H1 controls. Macro-F1 values are fractions.

B.7 Does the trained transition respond to its action input?

The final historical E4 is audited on all 765 one-step development pairs, using strict checkpoint loading and normal, zero, globally shuffled and persistence conditions. Shuffling is a cyclic shift across the complete evaluated ordering, not a within-batch operation. The original legacy-residual interaction transition is retained. Decoder tracing confirms that every model call receives the intended action; there are no missing or mismatched decoder interventions.

Table 33 reports factual target errors, and Table 34 separately reports differences between model predictions. The ego response is pronounced under zeroing; shuffling also changes the outputs. This establishes action dependence, not that every response has the correct semantic direction, nor that the forecast is accurate under an unobserved intervention. The comparisons retain persistence and explicitly mark C coverage differences. The old eight-sample pilot is not used as final-model evidence.

Table 33: Final historical E4 on all 765 factual one-step pairs.

Input	BEV MAE	C MAE	Ego MAE	Node MAE	Node XY	Edge MAE	Latent MSE
Normal	0.2363	0.0736	0.0775	0.2294	1.5548	1.3269	0.0203
Zero	0.2390	0.0780	0.7832	0.2795	2.0073	1.3313	0.0187
Shuffled	0.2363	0.0739	0.0877	0.2324	1.5812	1.3275	0.0204
Persistence	0.1773	0.0698	0.0553	0.3281	2.4650	1.2268	0.0030

The same recorded target scenes are retained for all conditions. Shuffling is a global cyclic shift, including the wraparound donor. Node XY is coordinate error, not trajectory ADE. C target coverage is .755714 for learned predictions and 1 for persistence; those C errors do not use identical matched coverage. Alternative-input factual error is sensitivity, not counterfactual ground-truth accuracy. Persistence is retained even where it is better.

Table 34: Prediction differences when the action input changes.

Pair	BEV MAE	C MAE	Ego MAE	Node MAE	Slot MAE
Normal–zero	0.026683	0.013540	0.761872	0.065720	0.110360
Normal–shuffle	0.002775	0.001263	0.080763	0.007371	0.012120

These compare two model outputs, not either output with a counterfactual target. All 765 actions changed under both zeroing and global shuffling; decoder action tracing recorded no missing or mismatched intervention.

Numerical provenance. The accompanying machine-readable table data records source-file/member SHA256 hashes, JSON pointers or CSV row pointers, unrounded cell values, and any explicit ratio operation. The complete source reports retain all available modes and horizons. Rebuilding these tables reformats archived results; it does not rerun model inference, add statistical tests, or retroactively establish unknown historical execution hashes.

C Planning Interfaces, Controls, and Complete Outcomes

C.1 From a state definition to a fixed intervention library

The archived factor-ablation records establish that PE/PCE and PIE/PCIE models had already been implemented before the first documented discussion of planning-metric correspondence. This ordering supports the account that the state definition motivated later planning design. It does not independently establish everything the authors knew at design time. The correspondence is also not one-to-one: collision risk can depend jointly on physical geometry, participant relations and ego motion, and traffic-rule compliance is not an isolated measurement of C.

The fixed transformations below implement the actual expansion, rather than an idealized factor intervention. Let $a_h = (\Delta x_h, \Delta y_h, \Delta\psi_h)$ be an anchor action at one of eight half-second steps. Braking-style variants multiply all three components, not only longitudinal velocity. Geometry variants leave Δx fixed and scale both lateral and yaw increments. No learned scene-wise gate decides which family to instantiate.

Table 35: Motion alternatives added to the learned K6 proposals. K10 adds the four middle transformations; K11 additionally includes CV. Names describe design intent, not guaranteed feasibility or safety.

Family	Implemented transformation
K6	Six action sequences proposed from composed current state and ego history
Hard deceleration	Anchor 0, all components scaled by $[1, .7, .4, .2, .1, 0, 0, 0]$
Soft deceleration	Anchor 5, all components scaled by $[1, .9, .8, .7, .6, .5, .4, .3]$
Reduced curvature	Anchor 0, $(\Delta x, .5\Delta y, .5\Delta\psi)$
Increased curvature	Anchor 5, $(\Delta x, 1.5\Delta y, 1.5\Delta\psi)$
Constant velocity	$(.5v, 0, 0)$ at each step, using the stored current speed

The original 186-token candidate audit reports oracle raw PDM of 80.9599 for K6, 90.1634 for K6+CV, 91.0703 for K10, and 94.6619 for K11. The four structured additions supply the winning candidate in 28 of these 186 cases under the recorded tie rule. In a separate historical 1,662-token validation subset excluding those 186 tokens, the corresponding K6/K10/K11 oracles are 77.44/89.15/92.41. These are candidate-coverage diagnostics, not selected scores or a claim of independent upstream design isolation. The later component-readout pool has a different oracle, 94.1455.

C.2 Independent training of the original raw-consequence readout

The controlled experiment retains the original NAVSIM V5 `two_stream_relative` encoder, scorer, optimizer and loss. The architecture has width 256, eight heads, two scorer layers, dropout .1, and world-stream weight 1. AdamW uses learning rate 10^{-4}, weight decay 10^{-4} and gradient clipping 1, with no scheduler or AMP. Each optimizer step contains one scene's eleven candidates. The soft ranking temperature is .25; the world-contrast coefficient and margin are .5 and .05. The state residual branch is disabled.

Both arms start from scratch. For each seed, all 130 initialized tensors are bitwise equal on the server, and actual training orders and protected input traces match. The trajectory-only arm zeros all five encoded future types at `WorldConsequenceEncoder.forward_from_raw` during training and evaluation. Actions and integrated candidate poses remain available. The world stream and its dropout/contrast computations remain in the original architecture; the masked contrast term is not assumed identically constant. Thus nominal architecture is held fixed, but this does not equalize active input-dependent representational capacity.

No new world or proposal inference is needed: the native caches contain raw predicted physical, semantic, graph, ego and unified-state sequences for each candidate. Runtime guards reject calls to world/proposal forward methods inside the readout training loop. These guards verify cache reuse, not absence of world-model information. All 605 tokens are checked against saved source

entries; add-on and CV poses match labels within absolute tolerance 10^{-4}. Original K6 labels lack generation-time pose hashes, so this audit cannot prove the immutability of those historical label inputs.

Table 36 reports all four epochs for seeds 1701–1705. Seeds 1704/1705 were added after inspecting the original three; all five are included without replacing the negative seed 1703 or changing the four-epoch budget. Primary results use epoch 3: aligned futures score 60.9315 ± 2.5878, versus 57.5777 ± 2.9270 for trajectory-only inputs (mean $\pm$ sample SD). Four paired differences are positive, with an overall mean of +3.3538 points. Table 37 preserves every pair. Development-best means are 62.6608 ± 2.0841 and 58.7018 ± 2.9775; the earliest-maximum rule uses the same development set and remains exploratory (Table 38). Original-three, appended-two and all-five descriptive differences are reported together in Table 39. An always-CV choice, evaluated from the same stored labels, gives 60.0796 PDM with a 34.95% zero-score rate. It requires no newly trained selector. The matched K11 oracle is 94.6619 and K10 oracle 91.0703 for every run.

Table 36: All forty epoch outcomes of the raw-consequence input study. PDM $\times 100$. Epoch 3 is primary; seeds 1704/1705 were appended after inspecting the original three. Every seed and epoch is retained.

Input	Seed	Epoch 0	Epoch 1	Epoch 2	Epoch 3
Aligned future	1701	60.690	56.313	58.591	60.706
Aligned future	1702	60.713	57.279	59.093	62.261
Aligned future	1703	56.722	60.410	60.720	60.841
Aligned future	1704	56.471	62.213	61.738	63.921
Aligned future	1705	60.572	62.037	65.575	56.929
Trajectory only	1701	55.577	54.399	52.257	55.227
Trajectory only	1702	50.523	53.125	57.409	57.470
Trajectory only	1703	51.073	54.822	56.768	61.672
Trajectory only	1704	52.421	62.078	60.809	59.066
Trajectory only	1705	53.603	56.713	55.455	54.454

Table 37: Fixed terminal raw-consequence controls, all five seeds. Zero rates use PDM $\leq 10^{-12}$. Paired differences are in PDM points.

Seed	Aligned	Trajectory	Difference	Aligned zero (%)	Trajectory zero (%)
1701	60.706	55.227	+5.480	31.183	37.634
1702	62.261	57.470	+4.791	30.108	34.409
1703	60.841	61.672	-0.831	31.183	22.581
1704	63.921	59.066	+4.855	29.032	31.183
1705	56.929	54.454	+2.474	36.559	38.710

Table 38: Exploratory development-best raw-consequence outcomes. Each run uses its earliest maximum over epochs 0–3 on the same development data; this does not replace fixed epoch 3. All values are PDM points.

Seed	Aligned epoch	Aligned	Trajectory epoch	Trajectory	Difference
1701	3	60.706	0	55.577	+5.130
1702	3	62.261	3	57.470	+4.791
1703	3	60.841	3	61.672	-0.831
1704	3	63.921	1	62.078	+1.843
1705	2	65.575	1	56.713	+8.862

Table 39: All retained Stage-C cohorts, descriptive paired differences (aligned minus trajectory-only). Mean ± sample SD in PDM points across the listed fitted seeds. The two appended seeds do not replace the original three.

Cohort	Seeds	Fixed epoch 3	Development best
Original three	1701–1703	$+3.147 \pm 3.462$	$+3.030 \pm 3.348$
Appended two	1704–1705	$+3.665 \pm 1.683$	$+5.353 \pm 4.964$
All five	1701–1705	$+3.354 \pm 2.604$	$+3.959 \pm 3.658$

The saved per-token outputs permit exact metric reaggregation: first-argmax selection, candidate counts, PDM, zero rate, oracle match, value match, regret, and CV selection rate. All 40 epoch outputs contain 186 tokens (7,440 rows), and their candidate labels are invariant across arms and epochs. The supplement includes the initialization audit record, sample order, trace, and full endpoint tables. Public result replay checks those saved records; without weight files it does not independently repeat the original bitwise tensor comparison.

C.3 Historical mechanism audit: dependence is not retrained utility

An earlier original NAVSIM K6/H8 readout was trained on 4,209 tokens and evaluated on 186. Its frozen-input diagnostic, shown at terminal epoch 11 in Table 40, distinguishes aligned consequences, zeroed world inputs, zeroed explicit trajectory inputs, and cyclically shifted worlds. These are inference-time changes to a fixed network, not independently trained selectors. Pairwise ranking uses the original composite cost with collision weight 4 and comfort weight .05; L2 regret uses a different, L2-based oracle. Neither is PDM.

Table 40: Original NAVSIM frozen-readout mechanism audit, fixed epoch 11. Zeroing explicit trajectory inputs does not remove trajectory information already reflected in candidate-conditioned future states.

Input condition	Pairwise (%) ↑	L2 regret (m) ↓	L2 oracle match (%) ↑
Aligned	69.928	.96754	33.871
World input zeroed	64.086	1.38825	20.968
Explicit trajectory input zeroed	65.842	1.02564	24.731
Cyclic world shift	55.914	1.55401	18.280

All twelve epochs are retained in the machine-readable results. Aligned regret is lower than shifted-world regret in 12/12 epochs, lower than world-zeroed regret in 11/12, and lower than explicit-trajectory-zeroed regret in 9/12. At epoch 8, the last of these is .98156 m versus .98813 m aligned. Epochs are correlated training checkpoints, not independent repetitions. This establishes sensitivity to correspondence without proving uniform fusion benefit; the independent PDM-trained control above addresses a different, stronger question.

C.4 Component-provider controls and selection gaps

The three-provider experiment uses the original V13 residual ranker. Predicted traffic, current-static traffic, and constant-velocity traffic generate each candidate's dynamic planning components, prior and fallback. The selected trajectories, motion history and candidate pool are shared. The component provenance differs between historical training caches and the later development caches, as recorded in Appendix A.

Table 41: Component readout: endpoint sensitivity and paired recording intervals. All values are PDM points. Intervals are marginal and conditional on the three fixed training instances.

Endpoint	Provider	Mean selected	Predicted minus provider	95% interval
Fixed final	Predicted	87.4557	—	—
	Static	84.6219	2.8339	[.7948, 5.0962]
	CV	84.4899	2.9659	[−.6924, 6.8110]
Development best	Predicted	88.7857	—	—
	Static	86.1498	2.6358	[.8171, 4.7063]
	CV	86.4873	2.2984	[−.7882, 5.4714]

The candidate oracle is 94.1455. The predicted-provider gap is 6.6898 at the fixed endpoint and 5.3598 after development selection. These are pool-specific ceilings; they must not be labeled ground-truth-future input experiments. The latter would replace the predictor's outputs inside the same selector, which this experiment does not do.

A separate historical V20 input diagnostic has fixed-final PDM 88.8737 with full features, 88.2683 with zeroed future features, and 87.5175 with mismatched features. Its development-best scores are 89.3365, 89.5675 and 89.2779: zeroing is slightly better at that endpoint. These mixed outcomes are retained, and do not support a universal statement that every future feature improves every trained readout. Reduced-data auxiliary summary selectors are indexed separately in the code archive and are not substituted for either original planning interface.

C.5 Equal-count pools: geometry, coverage, and selected utility

The primary 482-token evaluation contains sixteen recording groups; the secondary 186-token historical report contains 74. All three pools preserve the same K7 (six original candidates plus CV) and replace four positions. For generic lateral expansion, anchors 0 and 5 receive left/right smooth offsets with amplitude $\min(.75\,\mathrm{m}, .1 \times \text{arc length})$. Generic speed uses .8/1.2 time scaling on the same anchors. Exact float32 pose hashes establish eleven distinct candidates per primary token, but do not establish amplitude or effective behavioral diversity matching.

Six frozen selectors are reused: three seeds each for current-state and kinematic inputs, trained for 32 epochs on the guided pool and selected by development PDM. Current-state inputs repeat the current state. Kinematic inputs extrapolate observed participants, transform them to each candidate's ego frame, update geometry on fixed edges, and use an ego-motion proxy; physical and semantic features remain current. This is not a learned future world provider. Before pool transfer, all 668 guided examples and six models reproduce saved features, scores, choices and pose hashes exactly.

The guided additions have larger mean anchor-relative trajectory displacement than the lateral controls (2.435 vs. .426 m), and 2.243 m for speed controls. Corresponding endpoint distances are 5.710/.725/3.750 m, and mean absolute speed changes 1.225/.036/.944 m/s. Thus count matching alone does not isolate semantic design. A geometry-only fit on 3,323 training tokens in 128 recordings was stopped under its fixed checks: maximum and mean normalized quantile errors were .395736 and .120796, above limits .20 and .10, with joint-anchor checks also failing. No matched-pool PDM result was generated.

Table 42: Primary pool-transfer effects, guided minus comparator (PDM points). Each interval uses 10,000 paired recording bootstrap draws; all four gap-change intervals cross zero. There are only sixteen primary recording groups.

Input	Comparator	Selected difference	95% interval	Gap difference interval
Current	Lateral	3.721	[1.063,6.115]	[−2.654,2.099]
Current	Speed	3.459	[1.275,5.685]	[−1.893,2.600]
Kinematic	Lateral	4.277	[2.207,6.764]	[−2.726,.841]
Kinematic	Speed	3.628	[2.308,5.159]	[−1.521,.923]

The secondary cohort is weaker: selected PDM for guided/lateral/speed is 77.552/72.136/75.489 with current-state inputs and 77.121/73.789/74.752 with kinematic inputs. The selected differences against speed have intervals [−.645,4.991] and [−.630,5.645], respectively. The two secondary oracle difference intervals also cross zero. No multiplicity adjustment or upstream design isolation is inferred. These outcomes support candidate-pool utility under the tested transfer, not a proof of matched-amplitude semantic causality.

C.6 What changed in the planning readout

The conceptual pipeline retains four roles: proposing actions, predicting their consequences, evaluating candidates, and measuring the selected plan. The historical raw neural scorer and later value/component readouts occupy the evaluation role. Early V5–V7 implementations therefore continue the raw-consequence line; a version prefix alone does not mark a change of method. The later model-predictive evaluator replaces that learned scoring interface with explicit future-derived components, followed by utility refinements. The observed large oracle–selection gap motivated that transition. We do not use a new finite-budget ablation, or an invalid synthetic replay, to retroactively claim the earlier scorer was universally ineffective.

C.7 NAVHARD: official evaluation and candidate diversity

The NAVHARD/v2 evaluator extends the earlier NAVSIM/PDM setting (Dauner et al., 2023; Cao et al., 2025). Its two-stage combined score depends on linked branches and route weights, and is not the product of the marginal stage means. We retain the official combined output. The local experiments cover 5,912 successful scenarios with zero failures; three additional summary rows in the CSV are not scenarios. Synthetic inputs use their own available histories. An earlier erroneous replay of original-scene trajectories at synthetic tokens is excluded, rather than interpreted as a planner failure.

The later batch selectors load original/synthetic grouping and use synthetic starting positions and candidate endpoints to construct branch weights in joint search. This provides more information than independent per-scene sensor inference. Scores below are benchmark-aware development measurements, not blind leaderboard estimates or a certification of challenge compliance.

Table 43: Complete two-rule by two-pool NAVHARD comparison, combined EPDMS $\times 100$. The additional strict runs isolate candidate retention within each score rule. No uncertainty interval was computed.

Stage-1 rule	Full K11	Strict Top-3	Pruning change
Component rule (V16)	53.6617	49.1614	−4.5003
Learned unary rule (V21)	54.1712	48.9722	−5.1991
Rule change	+.5095	−.1892	—

Each strict run evaluates exactly nine candidate pairs for each of 225 linked groups, from 450 stage-1 support entries. Both strict conditions use the same input and evaluator configuration except their intended scoring rule and runtime bookkeeping. Historic full runs predate the new trace checks and do not have complete generation-time hashes; their saved inputs and outputs remain separately identified. The full-search winner is a predicted-utility winner under the saved tie rule, not a label oracle. Removing it in 95/225 and 84/225 groups accounts for the changed selections: all retained-winner groups reproduce the original downstream choice. This does not prove that all equally valued alternatives were eliminated.

The earlier V22-to-V23 comparison changed both candidate support and unary integration. Their combined EPDMS is 49.1614 and 54.1712, respectively; the new crossed comparison above separates those two design changes. An offline learned unary's Top-3 reachability does not ensure that those three candidates preserve joint-search compatibility. Neither this result nor the stage-specific component shifts isolate a change in the world predictor.

Table 44: Complete NAVHARD component decomposition for the fixed utility–support comparison. All values are on a 0–100 scale.

Metric	V16 / K11	V16 / Top-3	V21 / K11	V21 / Top-3
Stage 1				
NC	95.333	92.778	96.000	92.778
DAC	94.667	95.556	93.556	94.444
DDC	98.111	99.889	98.667	99.889
Traffic lights	99.778	99.556	99.778	99.556
Progress	69.471	73.455	68.643	73.361
TTC	95.333	92.889	96.667	93.111
Lane keeping	92.222	93.556	92.667	94.000
History comfort	94.444	93.556	95.111	94.222
Extended comfort	53.778	49.333	54.222	47.556
EPDMS	74.348	73.693	74.103	72.552
Stage 2				
NC	94.286	92.080	94.322	92.573
DAC	92.626	87.979	93.415	88.899
DDC	94.787	92.105	95.405	92.872
Traffic lights	99.114	99.086	99.154	99.086
Progress	63.038	65.640	62.251	64.722
TTC	92.285	89.738	92.346	90.360
Lane keeping	67.563	60.820	68.657	62.073
History comfort	96.797	96.065	96.821	96.150
Extended comfort	77.794	70.562	78.618	71.106
EPDMS	70.330	65.160	70.742	65.930
Combined EPDMS	53.662	49.161	54.171	48.972

V16 and V21 label the two Stage-1 utility rules; V21 with full K11 is the historical V23 variant. Top-3 means the new exact 3×3 allowed-pair enumeration; its Stage-1 summaries are not interchangeable with the earlier V22 finite-mask run. Stage-2 search remains full. Its component averages depend on the Stage-1 endpoint-induced branch weights. Combined EPDMS is not the product of the marginal stage means. These are local official-evaluator, benchmark-aware development results, not blind leaderboard scores.

Table 45 separates the reconstructed baseline from privileged candidate optimization and a NAVHARD-supervised full-set diagnostic. The joint candidate oracle uses future-dependent labels; its score is an opportunity bound for that historical pool. The supervised diagnostic includes its own training groups, so its higher score does not constitute generalization evidence.

Table 45: Historical NAVHARD candidate bounds and transductive diagnostic. EPDMS is shown on a 0–100 scale; the rows have different information privileges.

Setting	Stage 1	Stage 2	Combined
Reconstructed MPC baseline	76.703	53.594	41.414
NAVHARD-supervised V5 (full set)	77.609	73.895	58.548
Privileged joint K11 oracle	79.352	73.848	61.644

The oracle uses future-dependent official candidate utilities and joint branch search. It is a fixed-pool opportunity bound, not a planner given ground-truth future states. V5 was trained with NAVHARD candidate utilities and this full-set evaluation includes its training groups; it is transductive. The reconstructed baseline agrees with its saved reference to below 10^{-12}. Stage-1-only optimization reaches 84.256 on its own objective; that number is not the Stage-1 marginal of the joint oracle. No held-out V5 score is asserted here.

D PAIRED UNCERTAINTY FOR THE COMPLETED DYNAMICS AND PLANNING MODELS

Grouped uncertainty on completed models. We additionally report a post-hoc paired cluster bootstrap with 10,000 draws, resampling 30 nuScenes scenes or 74 NAVSIM recordings and retaining every sampled cluster's samples and multiplicities. Each draw uses the pooled sample mean, rather than an unweighted mean of cluster means. For H6, the four velocity-error contrasts in Table 46 have negative marginal intervals, conditional on the completed seed-7, four-epoch-prefix models. This does not estimate training-seed variation. The separate repetition in Appendix B.4 has mixed seed effects; its fitted-seed sample SDs and the intervals here describe different sources of variation.

For Stage-C, we first average each token's paired values equally across all five fitted seeds at epoch 3, then resample 74 recording groups. The aligned minus trajectory-only difference is 3.354 raw-PDM points, with a 95% interval of $[-0.357, 7.102]$; both comparisons with the existing always-CV label column also have intervals crossing zero. The positive mean and four positive seed differences therefore coexist with unresolved uncertainty on this development cohort. Seeds 1704/1705 were appended after the original results were observed; the five-seed analysis was specified before their results, not before the original study. No seed or endpoint is omitted based on outcome.

Table 46: Post-hoc paired cluster-bootstrap uncertainty on completed models. H6 reports velocity-vector error differences (m/s; lower is better) for the single seed-7 fixed-four-epoch comparison. Stage-C reports raw-PDM differences in percentage points (higher is better), averaging five fitted seeds per token at fixed epoch 3. All intervals are unadjusted marginal 95% intervals; they do not estimate training-seed uncertainty.

Comparison	Endpoint	Difference	95% interval
nuScenes: 149 chains, 30 scene clusters; m/s			
PCE − PE	H6	-0.1486	[-0.2890, -0.0345]
PCE − PE	Mean H1–H6	-0.0571	[-0.1189, -0.0026]
PCIE − PIE	H6	-0.1550	[-0.2459, -0.0570]
PCIE − PIE	Mean H1–H6	-0.0891	[-0.1419, -0.0249]
NAVSIM: 186 development tokens, 74 recording clusters; points			
Aligned − trajectory-only	Fixed e3	3.354	[-0.357, 7.102]
Aligned − always-CV	Fixed e3	0.852	[-6.450, 7.953]
Trajectory-only − always-CV	Fixed e3	-2.502	[-9.412, 4.055]

Endpoint and grouping sensitivity. Resampling the 124 scene identifiers gives a fixed-epoch aligned-minus-control interval of $[-1.184, 7.770]$ points. This sensitivity check does not assume adjacent scenes are independent. At the exploratory development-best endpoint, the mean difference is 3.959 points: its recording interval is $[0.239, 7.659]$, while its scene interval is $[-0.421, 8.145]$. Using development scores to choose epochs makes this selected-endpoint result exploratory; it does not replace the fixed-epoch primary comparison. Table 47 reports both groupings and all three comparisons at each endpoint.

Table 47: Stage-C grouping and endpoint sensitivity, all five fitted seeds. Paired raw-PDM differences and marginal 95% intervals, in points. Recording groups are the primary resampling unit; scene grouping and same-development-set endpoint selection are sensitivity analyses.

Endpoint	Comparison	Difference	Recording interval	Scene interval
Fixed epoch 3	Aligned − trajectory	+3.354	[-0.357, 7.102]	[-1.184, 7.770]
Fixed epoch 3	Aligned − CV	+0.852	[-6.450, 7.953]	[-6.266, 7.708]
Fixed epoch 3	Trajectory − CV	-2.502	[-9.412, 4.055]	[-9.529, 4.368]
Development best	Aligned − trajectory	+3.959	[0.239, 7.659]	[-0.421, 8.145]
Development best	Aligned − CV	+2.581	[-4.728, 9.596]	[-4.571, 9.399]
Development best	Trajectory − CV	-1.378	[-7.863, 4.704]	[-7.985, 5.068]

The resampling units address dependence among overlapping chains or samples from the same recording. They do not remove the documented train/development recording overlap, repeated development use, or model-selection history. The H6 intervals cover autoregressive H6 and equal-weight means over H1–H6; all seven common readouts use the same scene draws. No additional model training or inference is part of this analysis.

The supplement retains all 28 H6 marginal intervals and all 84 five-seed Stage-C intervals, including adverse outcomes, zero rates and per-seed contrasts. Stage-C JSON, per-token differences and bootstrap draws appear in `results/stagec_five_seed/`; the earlier three-seed analysis remains identified as historical in `results/paired_cluster_all41.csv`. Intervals are unadjusted marginal summaries conditional on the fitted models, not training-population uncertainty or blind-test guarantees.

E Additional Temporal and Readout Controls

E.1 Candidate-specific first-predicted-state persistence

This control asks whether the raw readout benefits from the later predicted states beyond information already present in each candidate's first prediction. For candidate k, its cached raw sequence $(\widehat{W}_1^k, \ldots, \widehat{W}_8^k)$ is replaced by $(\widehat{W}_1^k, \ldots, \widehat{W}_1^k)$ before the original consequence encoder. Candidate actions, complete trajectories, labels, architecture, optimizer and losses are unchanged. The first raw state retains its `trajectory_proposals`; this is neither a current-world nor a no-future-information control. Candidate-relative state variation is retained, unlike repeating one identical current-state tensor across all candidates, which would cancel its data-dependent content under candidate centering.

Three new fits use seeds 1701–1703, 419 training and 186 development tokens, four epochs and 1,676 optimizer calls. Each has the same 130 initial tensors and realized sample order as its original aligned counterpart. The raw-state repetition is applied in both training and evaluation. All 2,421 encoding calls per fit are finite and retain nonzero candidate-centered variation; no token is dropped for this diagnostic. This cached readout experiment creates no new proposals, world-model outputs or PDM labels. Only the consequence encoder and scorer are retrained, with the original world-contrast term still active.

The fixed epoch-3 comparison is adverse for the full sequence: its three differences from first-state persistence are $-0.4997, -2.5190, -8.3862$ PDM points, averaging -3.8016. The 74-recording paired interval is $[-8.9794, 1.1963]$. This does not establish a temporal-sequence gain and does not prove that longer prediction is generally harmful. The first-state arm's larger across-seed dispersion is retained in Table 48. All aligned and first-state development-best checkpoints happen to be epoch 3.

After observing these results, an additional comparison was requested against the original trajectory-only fits on the same three seeds. Its epoch-3 gains are $5.9792, 7.3100, 7.5553$ points: mean 6.9482, interval $[1.1087, 13.0876]$. Development-best gains average 6.8316 with interval $[1.4686, 12.7158]$. This is explicitly result-aware exploration; it neither replaces the negative primary comparison nor creates an independent test. Protected input traces, labels, sample orders and initial tensor hashes match across all three arms. Two archived server comparisons establish bitwise initialization equality through their common aligned model; lightweight public replay verifies the hash records, not absent checkpoint bytes.

Table 48: Same three seeds at fixed epoch 3. Raw PDM $\times 100$; original trajectory-only/full-sequence fits and the new first-predicted-state repetition. These rows do not replace the separate five-seed study.

Seed	Trajectory only	Full sequence	First-state repetition
1701	55.2270	60.7065	61.2061
1702	57.4695	62.2606	64.7796
1703	61.6717	60.8408	69.2270
Mean $\pm$ SD	58.1227 ± 3.2716	61.2693 ± 0.8611	65.0709 ± 4.0183

Table 49: Both endpoints and all paired seed differences, raw-PDM points. Full minus first-state is the temporal primary control. First-state minus trajectory-only was added after viewing the first-state results and is exploratory. Intervals condition on these three fitted pairs and resample 74 recordings.

Endpoint	Seed	Full − first-state	First-state − trajectory
Fixed e3	1701	-0.4997	+5.9792
Fixed e3	1702	-2.5190	+7.3100
Fixed e3	1703	-8.3862	+7.5553
Fixed e3	Mean [95% CI]	$-3.8016\,[-8.9794, +1.1963]$	$+6.9482\,[+1.1087, +13.0876]$
Dev-best	1701	-0.4997	+5.6294
Dev-best	1702	-2.5190	+7.3100
Dev-best	1703	-8.3862	+7.5553
Dev-best	Mean [95% CI]	$-3.8016\,[-8.9794, +1.1963]$	$+6.8316\,[+1.4686, +12.7158]$

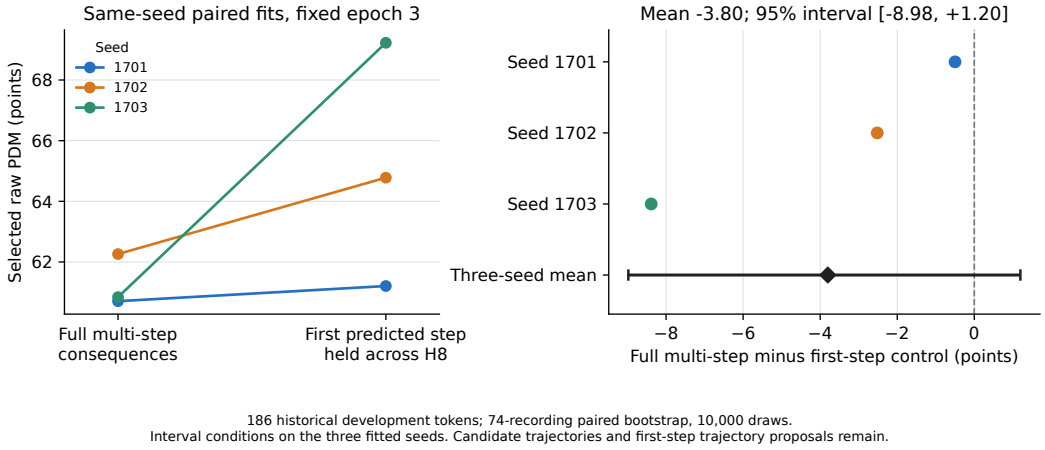

Figure 3: **Frozen temporal control on the original three readout seeds.** Full multi-step consequences versus repetition of each candidate's first predicted raw state, at fixed epoch 3. The interval resamples the 74 historical development recordings, conditional on the three fitted pairs. All seed differences are shown, including their adverse direction for the full sequence.

All intervals use 10,000 whole-recording draws with replacement, pooled token weighting and equal averaging of the three fitted seeds per token. Draws are shared across endpoints and metrics (random seed 2026092602). They do not estimate training-seed population uncertainty or remove development overlap. Zero means raw PDM $\leq 10^{-12}$. Complete epoch/endpoint rows, per-token values, and replay scripts are retained; the two arms' input change also changes the effective behavior of the unchanged contrast objective, preventing a pure causal-information attribution.

E.2 Last observation and frozen latent readout

The original persistence baseline copies the chain's initial state. We add a last-observation baseline which predicts each W_{t+h} by the true preceding W_{t+h-1}. It shares the preceding-state information available to TF; neither this baseline nor TF is open-loop AR performance. All nine seed-7 epoch-3 models are compared on exactly 149 chains from 30 scenes, with seven P/E readouts at each of H1–H6. Table 50 retains both the final and six-step mean contrasts. The last observation has six-step velocity error .204367 m/s and BEV MAE .167801; PCIE has .303804 and .279550. All nine models' average velocity errors are higher than this local baseline. This identifies short-budget conditional-prediction error in addition to long-horizon feedback error; it is not a measurement of information loss inside the composer.

For PCIE and monolithic PCIE, a further frozen-head diagnostic uses the first chain in each scene in original manifest order: 30 chains, six steps each, selected without errors or scores. The physical

head is applied directly to current slots with the native query-grid reshape and bilinear interpolation. Native future-slot decoding is replayed first to verify routing, and the TF outputs reproduce their archived targets and errors. The head was trained for future-state prediction, so applying it to current slots is an off-label readout, not a separately trained current-state reconstruction task.

Table 51 compares the same next-state targets: ordinary transition outputs have lower BEV error than direct current-slot readout, but both remain worse than last observation. Differences between these readouts cannot be decomposed into representation and transition error, because head supervision and latent distributions differ. The 30-chain means are compared only on their matched support. BEV scores concern encoder features, not occupancy accuracy; native speed, velocity and yaw-rate errors are not pose ADE/FDE. C/I last-observation readouts were not evaluated.

Table 50: TF error minus true last-observation error at seed 7, epoch 3: positive values favor the baseline. All 149 chains and seven readouts; full horizon tables accompany the source. Speed/velocity use m/s, yaw rate rad/s, BEV feature errors their native scale.

Model	Horizon	Speed	v_x	v_y	Yaw	Vel. L2	BEV MAE	BEV SL1
PE	H6	+0.0660	+0.0989	+0.0310	+0.0251	+0.1143	+0.1114	+0.0362
PE	Mean6	+0.0464	+0.0780	+0.0249	+0.0206	+0.0902	+0.1117	+0.0360
PCE	H6	+0.0983	+0.0987	+0.0310	+0.0265	+0.1150	+0.1113	+0.0362
PCE	Mean6	+0.0771	+0.0749	+0.0253	+0.0223	+0.0882	+0.1119	+0.0361
PIE	H6	+0.0963	+0.0958	+0.0302	+0.0300	+0.1097	+0.1122	+0.0366
PIE	Mean6	+0.0903	+0.0805	+0.0247	+0.0257	+0.0917	+0.1124	+0.0363
PCIE	H6	+0.1018	+0.1082	+0.0285	+0.0240	+0.1239	+0.1113	+0.0362
PCIE	Mean6	+0.0830	+0.0876	+0.0228	+0.0205	+0.0994	+0.1117	+0.0361
Independent	H6	+0.0975	+0.0977	+0.0306	+0.0277	+0.1113	+0.1108	+0.0360
Independent	Mean6	+0.0803	+0.0777	+0.0243	+0.0232	+0.0886	+0.1112	+0.0358
Monolithic	H6	+0.2289	+0.2403	+0.0363	+0.0381	+0.2534	+0.1133	+0.0370
Monolithic	Mean6	+0.1993	+0.2204	+0.0320	+0.0331	+0.2314	+0.1135	+0.0367
Ordered identity	H6	+0.1891	+0.1899	+0.0341	+0.0247	+0.1982	+0.1118	+0.0364
Ordered identity	Mean6	+0.1563	+0.1553	+0.0292	+0.0209	+0.1629	+0.1122	+0.0362
MILE-style	H6	+1.0317	+1.0623	+0.0886	+0.0303	+1.0788	+0.1269	+0.0428
MILE-style	Mean6	+1.0806	+1.1147	+0.0854	+0.0265	+1.1275	+0.1273	+0.0426
OccWorld-style	H6	+1.6779	+1.6941	+0.1188	+0.0744	+1.7010	+0.1302	+0.0446
OccWorld-style	Mean6	+1.7681	+1.7886	+0.1111	+0.0696	+1.7953	+0.1311	+0.0446

Table 51: Frozen physical-head diagnostic on the same 30 chains, six-step BEV MAE. Current-slot readout is off-label because the head was trained for next-state prediction. All full horizons and both BEV losses remain in the source.

Model	Last $P \to P^+$	TF $\to P^+$	Current $Z \to P$	Current $Z \to P^+$
PCIE	0.182763	0.283037	0.293151	0.292835
Monolithic	0.182763	0.284583	0.289422	0.289174

E.3 Does feedback explain the C-associated velocity tradeoff?

This saved-output diagnostic retains all 12 factor/seed groups and 32,184 chain–horizon–mode rows. Its contrasts and resampling rule were frozen before this new computation, but motivated by previously inspected results; it is post-hoc, not a preregistered confirmatory experiment. For each C pair, define $\Delta_h^m = e_h^m(\text{with } C) - e_h^m(\text{without } C)$ and $D_h = \Delta_h^{\mathrm{AR}} - \Delta_h^{\mathrm{TF}}$. The primary velocity diagnostic is the difference between the H4–H6 and H1–H3 averages of D_h for PCIE–PIE. It is -0.02828 m/s with interval $[-0.08135, 0.02951]$; seed values are $-0.11177, +0.07968, -0.05273$. There is no consistent late amplification in this contrast.

PCIE–PIE instead already has higher six-step TF velocity error in all three seeds: mean $+0.07384$ m/s, interval $[0.04443, 0.09929]$. Thus AR accumulation alone cannot explain this disadvantage. The direction remains observable specific: lateral-velocity error falls across all three fitted seeds in both modes at H6 and in the six-step mean, whereas longitudinal-velocity and speed errors rise on

average. The two BEV errors also fall across the three fits but their AR H6 and six-step scene intervals include zero. C inclusion changes both semantic inputs and their auxiliary supervision; these contrasts do not identify gradient conflict, redundancy or a causal error-propagation mechanism.

For each bootstrap draw, 30 scenes are sampled with replacement and all their chains retained, including multiplicity. The three fitted seeds are averaged equally per chain before pooled chain weighting. The same 10,000 draws (seed 2026092610) serve every metric and contrast. Table 52 reports the complete velocity panel; machine-readable files retain all 1,120 contrast/endpoint/metric rows, 3,360 seed differences and 840 absolute readouts, including AR/TF minus persistence. Intervals are marginal, unadjusted and conditional on these fitted models; descriptive horizon crossings do not select a new endpoint.

Table 52: Complete velocity-L2 feedback panel (m/s), three fitted seeds averaged per chain. Negative differences favor C; bracketed values are marginal scene-cluster intervals, not seed-population intervals. Both pairs and every specified endpoint are retained.

Endpoint	AR C difference	TF C difference	AR − TF difference
PCIE − PIE			
H1	$+0.0687\,[+0.0385, +0.0971]$	$+0.0687\,[+0.0385, +0.0971]$	$+0.0000\,[+0.0000, +0.0000]$
H2	$+0.0791\,[+0.0314, +0.1282]$	$+0.0666\,[+0.0363, +0.0938]$	$+0.0124\,[-0.0145, +0.0421]$
H3	$+0.0733\,[+0.0084, +0.1431]$	$+0.0744\,[+0.0433, +0.1024]$	$-0.0011\,[-0.0485, +0.0511]$
H4	$+0.0577\,[-0.0223, +0.1430]$	$+0.0764\,[+0.0449, +0.1032]$	$-0.0187\,[-0.0814, +0.0512]$
H5	$+0.0487\,[-0.0450, +0.1484]$	$+0.0776\,[+0.0453, +0.1044]$	$-0.0289\,[-0.1035, +0.0544]$
H6	$+0.0535\,[-0.0533, +0.1632]$	$+0.0793\,[+0.0486, +0.1053]$	$-0.0258\,[-0.1159, +0.0723]$
Mean H1–6	$+0.0635\,[-0.0038, +0.1346]$	$+0.0738\,[+0.0444, +0.0993]$	$-0.0103\,[-0.0595, +0.0444]$
Early H1–3	$+0.0737\,[+0.0280, +0.1207]$	$+0.0699\,[+0.0400, +0.0968]$	$+0.0038\,[-0.0207, +0.0310]$
Late H4–6	$+0.0533\,[-0.0389, +0.1512]$	$+0.0778\,[+0.0467, +0.1038]$	$-0.0245\,[-0.0993, +0.0593]$
Late − early	$-0.0204\,[-0.0763, +0.0368]$	$+0.0079\,[-0.0091, +0.0220]$	$-0.0283\,[-0.0814, +0.0295]$
PCE − PE			
H1	$+0.0027\,[-0.0190, +0.0266]$	$+0.0027\,[-0.0190, +0.0266]$	$+0.0000\,[+0.0000, +0.0000]$
H2	$-0.0076\,[-0.0570, +0.0442]$	$-0.0038\,[-0.0260, +0.0207]$	$-0.0038\,[-0.0336, +0.0275]$
H3	$-0.0072\,[-0.0792, +0.0695]$	$-0.0042\,[-0.0262, +0.0206]$	$-0.0030\,[-0.0564, +0.0546]$
H4	$-0.0094\,[-0.0971, +0.0871]$	$-0.0027\,[-0.0245, +0.0224]$	$-0.0067\,[-0.0773, +0.0714]$
H5	$-0.0284\,[-0.1292, +0.0854]$	$-0.0063\,[-0.0290, +0.0186]$	$-0.0221\,[-0.1037, +0.0730]$
H6	$-0.0610\,[-0.1700, +0.0634]$	$-0.0050\,[-0.0271, +0.0195]$	$-0.0560\,[-0.1496, +0.0514]$
Mean H1–6	$-0.0185\,[-0.0904, +0.0597]$	$-0.0032\,[-0.0246, +0.0205]$	$-0.0153\,[-0.0686, +0.0445]$
Early H1–3	$-0.0040\,[-0.0509, +0.0462]$	$-0.0018\,[-0.0234, +0.0220]$	$-0.0023\,[-0.0298, +0.0271]$
Late H4–6	$-0.0329\,[-0.1315, +0.0781]$	$-0.0047\,[-0.0267, +0.0200]$	$-0.0283\,[-0.1089, +0.0642]$
Late − early	$-0.0289\,[-0.0865, +0.0394]$	$-0.0029\,[-0.0113, +0.0063]$	$-0.0260\,[-0.0822, +0.0406]$

E.4 Finite-budget learning curves

The completed dynamics logs retain four epochs, 23 training batches per epoch, and the first four epochs of the original twelve-epoch cosine schedule. Table 53 reports all four factor models' mean one-step validation losses across seeds 7–9. PCIE's Ego loss falls by 66.45% from epoch 2 to 3, whereas BEV loss falls by .74%. These curves do not establish a common plateau. Although the absolute PCIE–PIE Ego gap shrinks from .082335 to .010980, its relative excess changes from 10.14% to 31.50%; this is not evidence of an imminent crossover. Semantic validation losses are not uniformly improving. Ego loss here is eight-dimensional Smooth-L1 with mixed units, not a velocity MAE. One-step loss trends cannot predict unmeasured long-horizon trends, and extra PE/PCE seed-7 epochs are kept separate from the common endpoint.

Table 53: Mean one-step validation losses across seeds 7–9 at each of four epochs. Ego is mixed-unit eight-dimensional Smooth-L1, not m/s. The twelve-epoch cosine schedule is unchanged; these are its first four epochs.

Model	Loss	e0	e1	e2	e3	e2 to e3
PE	BEV	0.087743	0.070200	0.068737	0.068257	−0.70%
PE	Ego	0.751192	0.318159	0.085645	0.032609	−61.93%
PCE	BEV	0.089121	0.070535	0.068649	0.068092	−0.81%
PCE	Ego	0.828734	0.377543	0.112074	0.032285	−71.19%
PIE	BEV	0.089834	0.071477	0.069243	0.068792	−0.65%
PIE	Ego	0.811698	0.365910	0.109684	0.034855	−68.22%
PCIE	BEV	0.089581	0.072935	0.069014	0.068502	−0.74%
PCIE	Ego	0.894033	0.419572	0.136603	0.045835	−66.45%

E.5 A PRECEDING-CHECKPOINT LONG-HORIZON CHECK

A targeted checkpoint comparison evaluates PCIE seed 7 at epochs 2 and 3 on exactly the same 149 chains. Target tensors, time alignment and persistence values match exactly; the native AR, TF and persistence protocols are retained. Table 54 gives both learned modes, both endpoints and all seven P/E readouts. Epoch-3 AR H6 velocity error is .617291 m/s versus 1.126015 at epoch 2 (−45.18%); the six-step mean falls from .941100 to .483352 (−48.64%). BEV MAE improves by only .57% at both endpoints.

This is direct evidence that this fitted model's long-horizon motion readout was still changing substantially in the last common training epoch. It does not establish convergence, predict a future ranking reversal, or remove the last-observation deficit under TF. The full native table also retains C/I results: AR six-step causal MAE changes from .099785 to .095226, node MAE from .532470 to .509444, and edge-attribute MAE from 1.522358 to 1.492383. Changed-relation macro-F1 remains zero; edge-type macro-F1 decreases slightly from .849371 to .849334. No claim of universal metric improvement follows. The evidence bundle retains all 147 common-readout rows and 1,092 native metric rows, including support and classification measures.

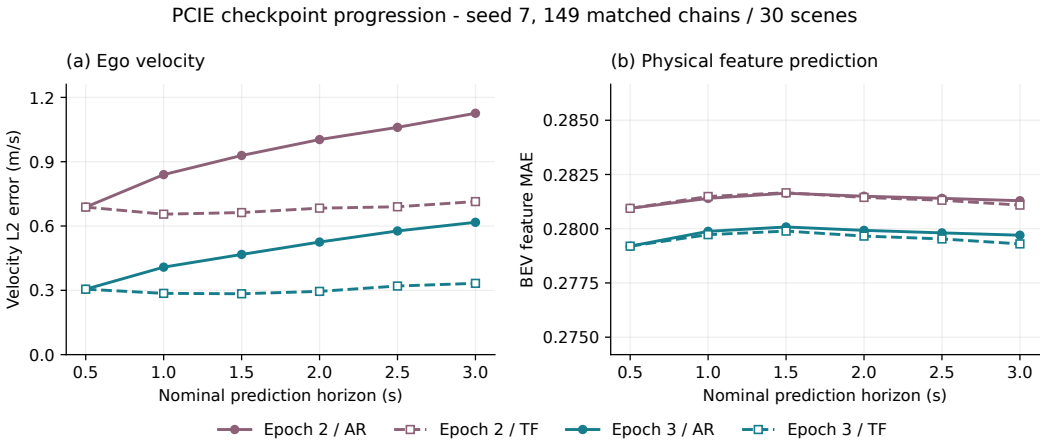

Figure 4: **The final common training epoch still improves motion.** PCIE seed 7 on the same 149 chains: checkpoints e2 and e3, native AR and TF. Persistence values are identical in the saved evaluations and are not plotted. The contrast concerns one fitted seed and two fixed checkpoints; it does not estimate converged capacity.

Table 54: PCIE seed 7 at checkpoints e2 and e3: same 149 chains and exact target/time/persistence alignment. Both modes and all seven P/E readouts are retained. Relative changes do not establish convergence or cross-model ranking at a longer budget.

Mode	Endpoint	Readout	e2	e3	Change
AR	H6	Speed	1.120109	0.583180	-47.94%
AR	H6	v_x	1.119851	0.600199	-46.40%
AR	H6	v_y	0.056038	0.053525	-4.48%
AR	H6	Yaw rate	0.042462	0.036854	-13.21%
AR	H6	Velocity L2	1.126015	0.617291	-45.18%
AR	H6	BEV MAE	0.281291	0.279700	-0.57%
AR	H6	BEV SL1	0.069416	0.068768	-0.93%
AR	Mean6	Speed	0.930386	0.454696	-51.13%
AR	Mean6	v_x	0.934517	0.467933	-49.93%
AR	Mean6	v_y	0.051777	0.048626	-6.08%
AR	Mean6	Yaw rate	0.039522	0.034161	-13.57%
AR	Mean6	Velocity L2	0.941100	0.483352	-48.64%
AR	Mean6	BEV MAE	0.281361	0.279764	-0.57%
AR	Mean6	BEV SL1	0.069414	0.068772	-0.92%
TF	H6	Speed	0.677994	0.303758	-55.20%
TF	H6	v_x	0.700604	0.309978	-55.76%
TF	H6	v_y	0.061233	0.050662	-17.26%
TF	H6	Yaw rate	0.041340	0.034351	-16.91%
TF	H6	Velocity L2	0.714065	0.332601	-53.42%
TF	H6	BEV MAE	0.281088	0.279300	-0.64%
TF	H6	BEV SL1	0.069297	0.068577	-1.04%
TF	Mean6	Speed	0.650004	0.280586	-56.83%
TF	Mean6	v_x	0.670872	0.285514	-57.44%
TF	Mean6	v_y	0.055408	0.046737	-15.65%
TF	Mean6	Yaw rate	0.039175	0.032494	-17.05%
TF	Mean6	Velocity L2	0.682242	0.303804	-55.47%
TF	Mean6	BEV MAE	0.281321	0.279550	-0.63%
TF	Mean6	BEV SL1	0.069369	0.068660	-1.02%