

Dynamic Sparsity: Challenging Common Sparsity Assumptions for Learning World Models in Robotic Reinforcement Learning Benchmarks

Muthukumar Pandaram^{*1}, Jakob Hollenstein^{*1}, David Drexel^{*1},
Samuele Tosatto^{1,2}, Antonio Rodríguez-Sánchez¹, Justus Piater^{1,2}

¹Department of Computer Science, University of Innsbruck, Austria

²Digital Science Center, University of Innsbruck, Austria

{muthukumar.pandaram,jakob.hollenstein,david.drexel}@uibk.ac.at

Abstract

The use of learned dynamics models, also known as *world models*, can improve the sample efficiency of reinforcement learning. Recent work suggests that the underlying *causal graphs* of such dynamics models are sparsely connected, with each of the future state variables depending only on a small subset of the current state variables, and that learning may therefore benefit from *sparsity priors*. Similarly, *temporal sparsity*, i.e. sparsely and abruptly changing local dynamics, has also been proposed as a useful inductive bias. In this work, we critically examine these assumptions by analyzing ground-truth dynamics from a set of robotic reinforcement learning environments in the MuJoCo Playground benchmark suite, aiming to determine whether the proposed notions of state and temporal sparsity actually tend to hold in typical reinforcement learning tasks. We study (i) whether the causal graphs of environment dynamics are sparse, (ii) whether such sparsity is state-dependent, and (iii) whether local system dynamics change sparsely. Our results indicate that global sparsity is rare, but instead the tasks show local, *state-dependent* sparsity in their dynamics and this sparsity exhibits distinct structures, appearing in temporally localized clusters (e.g., during contact events) and affecting specific subsets of state dimensions. These findings challenge common sparsity prior assumptions in dynamics learning, emphasizing the need for grounded inductive biases that reflect the state-dependent sparsity structure of real-world dynamics.

Supplementary Version — arxiv.org/abs/2511.08086

Code — github.com/jkbjh/dynamic-sparsity-paper-code

1 Introduction

Reinforcement Learning (RL) promises to enable robots to learn complex, adaptive behaviors autonomously. However, it is often sample-inefficient due to the large number of environment interactions required.

Model-Based Reinforcement Learning (MBRL) addresses this issue by learning models of the environment’s dynamics, also called *world models* (Sutton 1988; Schmidhuber 2015; Ha and Schmidhuber 2018). World models allow agents to simulate interactions and plan effectively with fewer real-world samples (Schmidhuber 1990a,b,c, 1991; Hafner et al.

2019a, 2025). These world models learn *dense dynamics* where future state values are predicted based on the whole set of the current state values, leading to learning spurious correlations between the states which leads to poor generalization and prediction accuracy (Wang et al. 2021). To address this issue, recent research has proposed incorporating sparsity as an inductive bias into learning of dynamics models for Model-Based Reinforcement Learning (MBRL), framing them as *causal models* (Wang et al. 2022, 2024; Lei, Schölkopf, and Posner 2024). These causal models are often assumed to be *sparsely connected*, meaning that each future state variable depends only on a limited subset of the current state variables and actions (Wang et al. 2022, 2024; Hwang et al. 2024; Lange and Kording 2025). Additionally, some studies advocate for modeling *temporal sparsity*, where dynamics change abruptly due to discrete latent transitions in the environment (Gumbsch, Butz, and Martius 2021; Jain et al. 2021; Orujlu et al. 2025).

Despite these promising directions, such assumptions are often validated only in controlled or synthetic environments, e.g. manipulation tasks with unmodifiable state variables and unmovable objects (Wang et al. 2022). It remains uncertain whether these sparsity assumptions extend from custom problems with known causal graphs to complex robotic environments featuring contacts and dynamic interactions.

In this paper, we explore the validity of these sparsity-based priors using ground-truth dynamics from the *MuJoCo Playground* (Zakka et al. 2025a), a suite of physics-based robotic benchmarks widely used in RL research. We focus our analysis on the *sparsity* of the transition functions’ *Jacobians*, which capture how future states linearly change with respect to current states and actions. This analysis enables the characterization of sparsity in the underlying causal structure, as the absence of causal influence corresponds to zero entries in the Jacobian. Specifically, we address four key questions:

- (Q1) Do the true dynamics exhibit sparsity in their causal structure?
- (Q2) Is such sparsity dependent on the current state?
- (Q3) Do the dynamics undergo temporally sparse transitions?
- (Q4) Does naive Multi-Layer Perceptron (MLP) training recover the ground-truth dynamics’ sparsity?

^{*}These authors contributed equally.

2 Related Work

This section reviews prior research most relevant to our work, beginning with general developments in continuous control environments and reinforcement learning, then narrowing our focus to methods concerned with modeling and learning environment dynamics with sparsity priors.

Continuous Control Environments: Continuous control tasks, such as those modeled by MuJoCo (Todorov, Erez, and Tassa 2012), DeepMind Control Suite (Tassa et al. 2018), Brax (Freeman et al. 2021) and MuJoCo Playground (Zakka et al. 2025a), have become standard benchmarks for evaluating reinforcement learning (RL) algorithms due to challenging high-dimensional, continuous state and action spaces, requiring agents to learn smooth control policies for locomotion or manipulation. In this paper, we focus on the continuous control tasks shown in Figure 1 from MuJoCo Playground.

Continuous Control through RL: Model-free approaches, such as DDPG (Lillicrap et al. 2016), PPO (Schulman et al. 2017), and SAC (Haarnoja et al. 2018), have achieved impressive results in continuous control. These methods typically rely on large amounts of interaction data, which limits their sample efficiency. In contrast, model-based RL approaches seek to improve sample efficiency by learning a predictive model of the environment’s dynamics. Notable methods include PETS (Chua et al. 2018), PlaNet (Hafner et al. 2019b), and Dreamer (Hafner et al. 2019a, 2021, 2025), which learn latent dynamics models to plan actions or improve policy learning. Efforts have also explored hybrid model-based/model-free strategies, where learned models are used to generate synthetic rollouts or augment training data (Janner et al. 2019).

World Models and Sparse Dynamics Learning: A growing body of work (Pitis, Creager, and Garg 2020; Wang et al. 2021, 2022, 2024; Hwang et al. 2024; Zhao et al. 2025) focuses on learning structured models of the world by combining world model learning (Ha and Schmidhuber 2018) with causal representation learning (Schölkopf et al. 2021). The main motivation for these methods is that incorporating causal structure into learning dynamics models can increase generalizability, particularly under distribution shifts or changes in controllable factors (Wang et al. 2021). Recent papers (Wang et al. 2021, 2022, 2024) propose inducing causal structure by incorporating sparsity into dynamics model learning. These papers argue that dense models, which use all current states and actions to predict future states, are prone to capturing spurious correlations between unrelated features which reduces prediction accuracy and hinders generalization. They propose learning a global context *independent* causal graph and assume that the local dependencies do not change over time. In contrast, another line of work focuses on modeling fine-grained local *context-specific independence* i.e., state-dependent sparsity within causal graphs by learning local causal models based on individual contexts. (Pitis, Creager, and Garg 2020; Pitis et al. 2022; Chitnis et al. 2021; Hwang et al. 2024; Zhao et al. 2025). These papers learn the local causal graphs for different purposes, such as data aug-




Figure 1: Benchmark Environments from the DeepMind Control Suite (Tassa et al. 2018), implemented in MuJoCo Playground (Zakka et al. 2025a): (top) BallInCup, CartpoleBalance, CheetahRun, ReacherHard, (bottom) FingerSpin, FingerTurnEasy, WalkerRun, SwimmerSwimmer6

mentation for RL agents (Pitis, Creager, and Garg 2020; Pitis et al. 2022) or exploration (Wang et al. 2023). These local causal models are learned in various ways such as using attention scores (Pitis, Creager, and Garg 2020), examining the Jacobians of the learned dynamics model (Wang et al. 2023; Zhao et al. 2025), vector quantization of local subgraphs (Hwang et al. 2024; Zhao et al. 2025), or using hard-attention in a Transformer world model with sparsity regularization (Lei, Schölkopf, and Posner 2024). Similarly, Orujlu et al. (2025) use RL agents to dynamically construct sparse, time-varying causal graphs, instead of the soft, dense connections typical of Transformers. Gumbsch, Butz, and Martius (2021) and Jain et al. (2021) model sparsity in the temporal evolution of latent states through the use of L0 regularization and variational sparse gating mechanisms respectively.

All the papers discussed in this section so far learn dynamics models predominantly in controlled or synthetic environments with known causal graphs, under the assumption that sparsity provides a useful inductive bias for MBRL in these environments. However, it was unclear if this assumption actually holds for the true dynamics of *common* continuous-control reinforcement learning environments as well. In this paper, we systematically investigate this assumption.

3 Background and Notation

This paper focuses on robotic RL benchmarks and in particular DM Control Suite environments (Tassa et al. 2018), as implemented in the MuJoCo Playground (Zakka et al. 2025b), for the differentiable MJX simulator, a version of MuJoCo (Todorov, Erez, and Tassa 2012) written in JAX (Bradbury et al. 2018). The considered systems are vector-valued, discrete-time and time-invariant, and their evolution is fully captured by $s_{t+1} = \text{step}(s_t, a_t)$ where $s_t, s_{t+1} \in \mathbb{R}^{d_s}$ are d_s -dimensional vectors representing the ground truth states, $a_t \in \mathbb{R}^{d_a}$ is a d_a -dimensional vector representing the control signal (a.k.a. *action*), and $\text{step} : \mathbb{R}^{d_s} \times \mathbb{R}^{d_a} \rightarrow \mathbb{R}^{d_s}$ is the simulator’s step function¹ that calculates the next state. Figure 2 illustrates how the states change with the actions. The observations (e.g., pixel-space) depend non-invertibly on the states. Performing a similar analysis on observations requires the ground-truth Jacobians between successive observations,

¹Up to the indeterminism by modern GPU implementations.

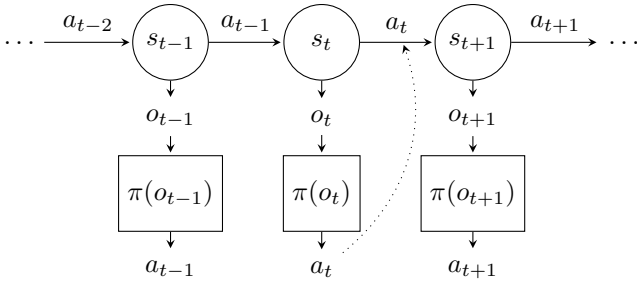


Figure 2: The simulator’s ground-truth state s_t is advanced by the simulator’s $\text{step}(s_t, a_t)$ function to produce the next state s_{t+1} . Each of these states s_t non-invertibly produces a corresponding observation o_t that is used by the trained agent to choose an action. In this work we analyze the dynamics of the ground-truth next state s_{t+1} with respect to the current state s_t .

which do not exist generally. Thus, we restrict our analysis to state space only. To interact with the environment, actions are sampled from the stochastic policy of the trained agent, i.e., $a_t \sim \pi(\cdot|o_t)$. The observations o_t capture some relevant information about the ground-truth state, but the mapping $o_t = f_o(s_t)$ is, in general, not invertible. Note that all vectors in this manuscript are column vectors, if not stated otherwise.

Sparse Causal Graphs The dynamics model, i.e., the step function, can be viewed as a vector of scalar-valued component functions, mapping the vector valued inputs to the scalar output of each separate state dimension i , i.e., $(s_t, a_t) \mapsto s_{t+1}^{(i)}$. These scalar-valued output functions may depend only on a subset of the input variables (s_t^i, a_t^i) where $s_t^i \subseteq s_t = \{s_t^{(1)}, \dots, s_t^{(d_s)}\}$. If the output for state variable s_{t+1}^j does not depend on s_t^i , then no connecting edge exists in the causal graph representing the step function. *Sparse* causal graphs (i.e., graphs with few edges) are in principle simpler to learn than dense ones, since the graph contains less information and the output can be inferred by using only a subset of the inputs. However, the knowledge of which edges are absent is generally unavailable, making it hard to harness this sparsity to simplify the learned model.

Differentiable Dynamics and Jacobians We consider differentiable step functions, and we define the first-order derivatives (a.k.a. *Jacobians*) w.r.t. states and actions as

$$J_s = \frac{\delta}{\delta s_t} \text{step}(s, a) = \begin{bmatrix} \frac{\delta \text{step}_1}{\delta s_t^{(1)}} & \cdots & \frac{\delta \text{step}_1}{\delta s_t^{(d_s)}} \\ \vdots & \ddots & \vdots \\ \frac{\delta \text{step}_{d_s}}{\delta s_t^{(1)}} & \cdots & \frac{\delta \text{step}_{d_s}}{\delta s_t^{(d_s)}} \end{bmatrix}$$

$$J_a = \frac{\delta}{\delta a_t} \text{step}(s, a) = \begin{bmatrix} \frac{\delta \text{step}_1}{\delta a_t^{(1)}} & \cdots & \frac{\delta \text{step}_1}{\delta a_t^{(d_a)}} \\ \vdots & \ddots & \vdots \\ \frac{\delta \text{step}_{d_s}}{\delta a_t^{(1)}} & \cdots & \frac{\delta \text{step}_{d_s}}{\delta a_t^{(d_a)}} \end{bmatrix}$$

respectively. In other words, the Jacobians above capture the local variability of the future state s_{t+1} given infinitesimal

variations of the current state s_t and action a_t . The goal of this paper is to study the system’s *Jacobians*, which give useful information about the local interactions between variables, in order to assess the sparsity of the underlying causal graph. To this end, we collect and analyze a dataset \mathcal{D} consisting of states, actions, next states and corresponding state and action Jacobians. To make sure that we cover relevant parts of the state space, we collect the dataset by using an expert policy trained via reinforcement learning (Section A1), and using MJX, we auto-differentiate $\text{step}(s, a)$ to obtain the state and action Jacobians J_s and J_a .

4 Using Jacobians to Assess Sparsity

To see how the Jacobians of the environment dynamics relate to sparsity in the causal graph of state variables, consider a differentiable function

$$f : \mathbb{R}^m \rightarrow \mathbb{R}^n$$

$$x = (x_1, \dots, x_m) \mapsto f(x) = (f_1(x), \dots, f_n(x)),$$

where x_i denotes the i -th input variable and f_j the component function mapping x to the j -th output variable y_j . Furthermore, suppose the existence of a directed, bipartite graph $\mathcal{G} = (V, E)$ with vertex set

$$V = \{x_1, \dots, x_m, y_1, \dots, y_n\}$$

encoding direct causal relationships from inputs to outputs. An edge $(x_i \rightarrow y_j) \in E$ indicates that the output y_j directly depends on the input x_i . As a consequence, for each $j \in \{1, \dots, n\}$, f_j is a function only of those inputs x_i for which $(x_i \rightarrow y_j) \in E$. Hence, if there is no such edge, the partial derivative of f_j with respect to x_i must be zero, i.e.,

$$(x_i \not\rightarrow y_j) \implies \frac{\partial f_j}{\partial x_i} = 0.$$

A zero partial derivative expresses that f_j is invariant to infinitesimal changes in x_i when x_i is not a direct cause of y_j in the causal graph \mathcal{G} .

Consequently, the absence of edges between certain inputs and outputs in the causal graph \mathcal{G} directly translates into sparsity in the Jacobian matrix:

$$J_f(x) := \left(\frac{\partial f_j}{\partial x_i}(x) \right)_{\substack{1 \leq j \leq n \\ 1 \leq i \leq m}} \in \mathbb{R}^{n \times m}.$$

More precisely, if an edge $(x_i \rightarrow y_j)$ is missing in \mathcal{G} , the corresponding element of the Jacobian must be zero. In turn, the number of zero elements present in $J_f(x)$ provides an upper bound on the sparsity of the causal graph. Each zero element in the Jacobian is a necessary condition for the absence of direct causal influence but not a sufficient one, due to the possibility of higher-order derivatives.

If we consider not only the derivative at a single point, but across the entire domain of f we can make a stronger statement. A Jacobian element that is zero everywhere, i.e.,

$$\frac{\partial f_j}{\partial x_i}(\mathbf{x}) = 0 \quad \forall \mathbf{x} \in \mathbb{R}^m,$$

implies that all higher-order derivatives with respect to x_i are zero as well. A global zero element is thus both necessary and sufficient for the absence of a causal edge $(x_i \rightarrow y_j)$.

Environment	Dimension		Jacobian Zero Elements	
	State	Action	State (%)	Action
BallInCup	8	2	0 (0.00)	0
CartpoleBalance	4	1	0 (0.00)	0
CheetahRun	18	6	17 (5.25)	0
ReacherHard	4	2	3 (18.75)	0
FingerSpin	6	2	0 (0.00)	0
FingerTurnEasy	6	2	0 (0.00)	0
WalkerRun	18	6	17 (5.25)	0
SwimmerSwimmer6	16	5	30 (11.72)	0

Table 1: This table shows the counts of elements in the state and action Jacobian that remain constantly zero across multiple rollouts, per environment. As this gives an upper bound for the sparsity in the underlying global causal graph, it indicates that global sparsity is mostly absent from these environments.

5 Experiments

For our experimental analysis, we collect trajectories using expert reinforcement-learning policies trained with PPO (Schulman et al. 2017). To encourage exploration, we inject colored noise following the method of (Hollenstein, Martius, and Piater 2024). Additional training details are provided in Section A1. The state and action Jacobians of the ground-truth transition function $\text{step}(s, a)$ are obtained via automatic differentiation. We define the *sparsity value* of the Jacobian matrix as the number of zeros in the Jacobian matrix divided by the total number of elements in the Jacobian matrix.

5.1 (Q1) Do the True Dynamics Exhibit Sparsity in Their Causal Structure?

This section investigates whether common Reinforcement Learning environments exhibit globally consistent sparse dynamics. We examined the Jacobians of step with respect to both s and a for sparsity, i.e., the presence of zero elements that persist across all collected samples. As discussed in Section 4, the condition

$$\frac{\partial}{\partial s^{(i)}} \text{step}^{(j)}(s, a) = 0 \quad \forall s, a$$

is both necessary and sufficient for the (j, i) element of the state Jacobian to be zero.

Since we cannot exhaustively evaluate this condition, we only evaluate a necessary condition for sparsity, and provide an upper bound for the true sparsity. The results for this experiment are listed in Table 1. We count the zero elements in both the state ($\frac{\delta}{\delta s}$) and action ($\frac{\delta}{\delta a}$) Jacobians. To account for floating point precision, we use a threshold of $|x| < 10^{-12}$ to determine whether an element x is zero. Our experiments show that there are indeed environments that exhibit globally zero elements in the Jacobians for all samples tested, but that this is only the case for very few elements (5.25% – 18.75%) and only in a handful of environments (CheetahRun, ReacherHard, WalkerRun, SwimmerSwimmer6). This indicates that

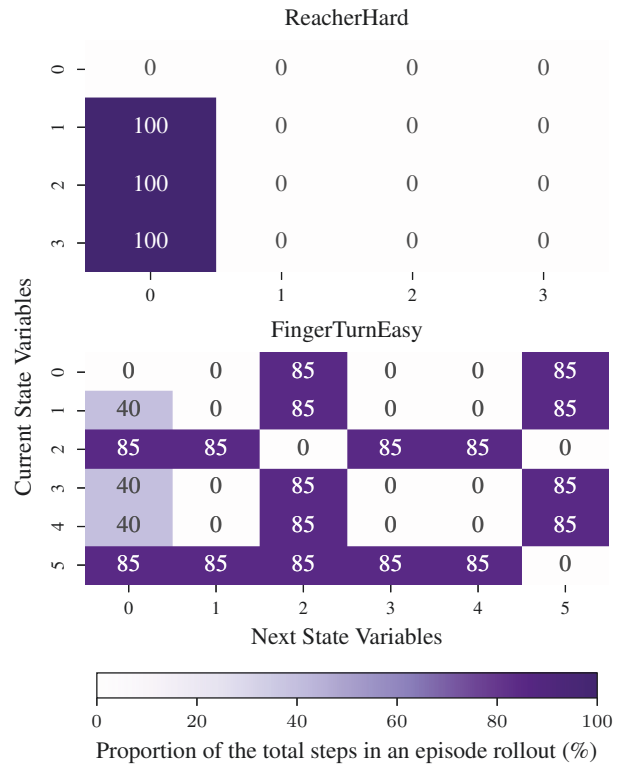


Figure 3: For the two environments ReacherHard and FingerTurnEasy, the heatmap illustrates the proportion of time each element of the Jacobians $J_s = \frac{\delta}{\delta s} \text{step}(s, a)$ and $J_a = \frac{\delta}{\delta a} \text{step}(s, a)$ remains zero (indicating the independence of the variables) during an episode rollout, expressed as a percentage of the total episode duration averaged across rollouts and seeds. The heatmap values are rounded to the nearest integer. Most Jacobian elements remain non-zero throughout the episode, a small number stay zero for the entire duration and the remaining elements are zero for only a fraction of the timesteps. Similar heatmaps for the remaining environments considered for analysis are shown in Section A3.1.

globally consistent sparsity is rare and is thus unlikely to be a generally important inductive bias. Instead of requiring a Jacobian element to be zero for all state and action samples, we can investigate the relaxed problem of looking at the percentage of samples where the element is zero. This is illustrated for the two environments ReacherHard and FingerTurnEasy in Figure 3. The results clearly show the globally zero elements in the ReacherHard environment, but further highlight that specific elements can be zero for a certain percentage of the samples, as for FingerTurnEasy. This hints at a different, potentially more widely applicable sparsity assumption: state-dependent sparsity, which we investigate next.

5.2 (Q2) Is Such Sparsity Dependent on the Current State?

In the previous section, we examined globally zero elements in the Jacobians. We found that some environments do have

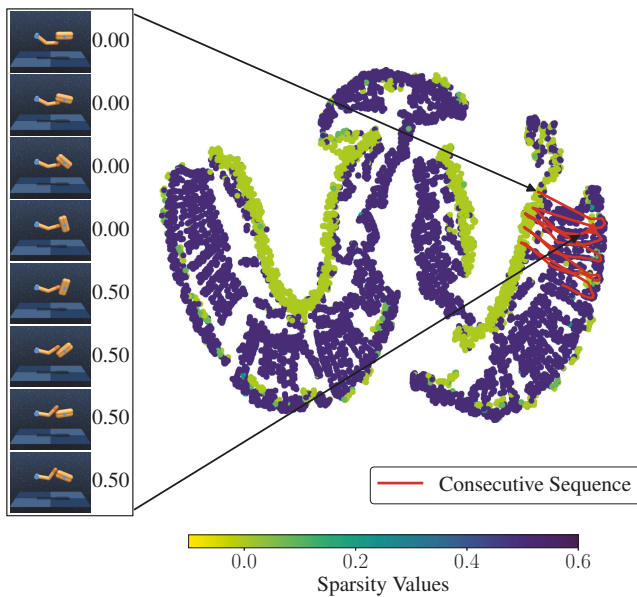


Figure 4: A 2-dimensional t-SNE embedding of state, action and next state tuples with a perplexity value of 50 colored by the *combined* sparsity values of state and action Jacobians across 10 episodes of FingerTurnEasy. Sparsity in the Jacobians is often related to contacts: When the Finger is not moving the object, we observe higher sparsity compared to when the Finger pushes the object in the process. The sparsity values are given near the images of the states.

such zero elements, but most do not. Even when they do, there are only a few of these globally zero elements. However, relaxing the requirement of global sparsity for the Jacobian elements, we can investigate partial sparsity in the Jacobian, i.e., when the Jacobian elements are zero only for fraction of the samples.

Sparsity (i.e., zero elements) in the Jacobian for a fraction of the sampled states is a necessary condition for state-dependent sparsity, which, as indicated by Figure 3, is present in at least some of the environments. In this section, we turn to looking at state-dependent sparsity in more detail. Figure 4 illustrates state dependent sparsity for the FingerSpin environment, through a t-SNE embedding of the state, action and next state (s, a, s') tuple and color coding the combined state and action Jacobian’s sparsity at each point. The figure illustrates a sequence of frames, the corresponding transitions through the t-SNE embedded state-space, and the sparsity values. As can be observed in the first frames, the sparsity is lower when the finger interacts with the object. This sequence of steps with low sparsity followed by a few steps of high sparsity repeats as the object briefly comes into contact with the finger (=low sparsity) and then spins freely (=high sparsity). Sparsity could be induced in a learned representation, e.g. through regularization (similar to Lei, Schölkopf, and Posner (2024)). However, this method requires tuning the amount of sparsity that is induced. Figure 5 illustrates histograms of the observed sparsity across all samples for all tested environments. The histograms indicate that the

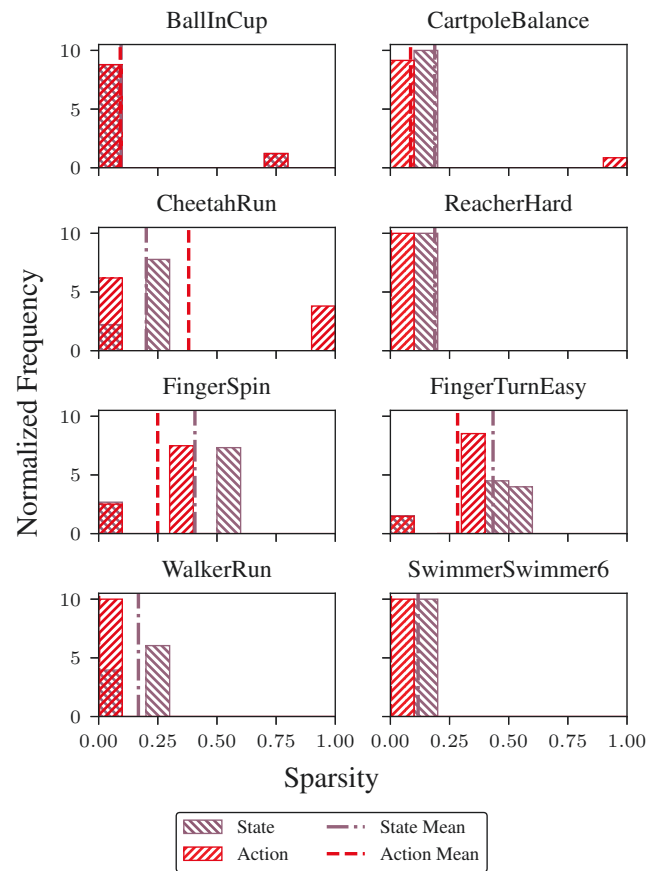


Figure 5: Histograms showing the distribution of state and action Jacobians’ sparsity values in the whole dataset with a bin width of 0.1. The sparsity values are mostly concentrated in a small number of bins, indicating a repetition of similar sparsity patterns in the Jacobians over the trajectories.

sparsity only assumes certain values, presumably based on whether the agent touches the ground (CheetahRun, WalkerRun), or interacts with the object (BallInCup, FingerSpin, FingerTurnEasy). The figure also illustrates the mean sparsity for each environment. Presumably, if sparsity regularization is used, it should be tuned to induce a similar amount of sparsity on average. While Figure 5 indicates that the sparsity only assumes specific values, it is unclear how the sparsity changes temporally.

5.3 (Q3) Do the Dynamics Undergo Temporally Sparse Transitions?

The previous section established that the system assumes specific sparsity values, likely due to the structure of the interacting objects and the robot, such as objects coming into contact or airborne phases of the gait cycle. If sparsity relates to the gait cycle or contact events, its behavior is expected to be temporally consistent—remaining stable over multiple steps before switching—which can be observed in Figure 4. The system’s dynamics can be described by differential equations where changes in contact correspond to transitions between

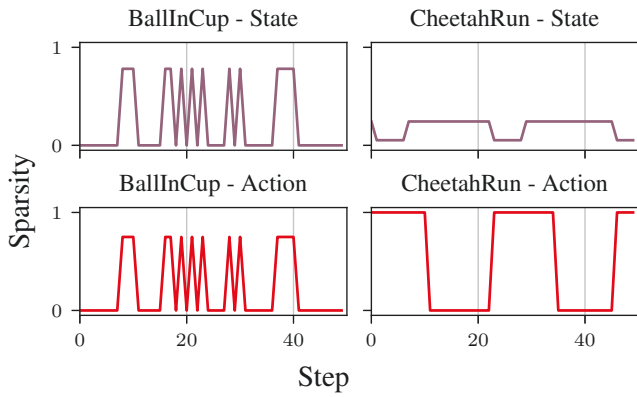


Figure 6: An illustrative example showing how sparsity values evolve over time in the BallInCup and CheetahRun environments. In BallInCup, sparsity pattern varies with the tautness of the string connecting the ball and the cup, while in CheetahRun, it reflects changes in the cheetah’s gait. Similar plots for the remaining environments considered for analysis are shown in Section A3.2.

different underlying differential equations. Therefore, incorporating inductive biases favoring modeling of temporally sparse switches between underlying dynamics may improve the learned performance (Gumbsch, Butz, and Martius 2021; Jain et al. 2021).

Gumbsch, Butz, and Martius (2021) have pursued a similar approach, employing a temporally sparsely changing recurrent network to model both agent behavior in partially observable environments and environment dynamics. In their method, a hyper-parameter penalizes the changes and thus implicitly regulates how often the system switches between states, or inversely, stays in the same state. In this section, we examine whether sparsity changes occur in a temporally sparse manner and analyze the distribution of durations over which these sparsity states persist. Figure 6 depicts the evolution of the percentage of zero elements over time during a single rollout with the expert agent, using the stochastic Gaussian policy. In the BallInCup environment, a ball tethered to a controllable cup by a string must be caught by moving the cup. Sparsity is high when the string is loose and low when the ball is inside the cup or the string is taut, indicating temporally sparse changes. Figure 7 presents the distribution of durations, measured by consecutive timesteps, during which Jacobian elements remain in the close to zero or non-zero state. The durations of elements that remain either zero or non-zero for the maximum length of the episode are not shown. Many of the elements in these Jacobians remain in the same state for durations of multiple steps, hinting at underlying structure that sparsely changes in time. The mean duration of change likely poses as a useful default hyperparameter for methods attempting to regularize for temporally sparse switching.

5.4 (Q4) Does Naive MLP Training Recover the Ground-truth Dynamics’ Sparsity?

The previous sections demonstrated that even though global sparsity is rare but sparsity does occur in the ground-truth

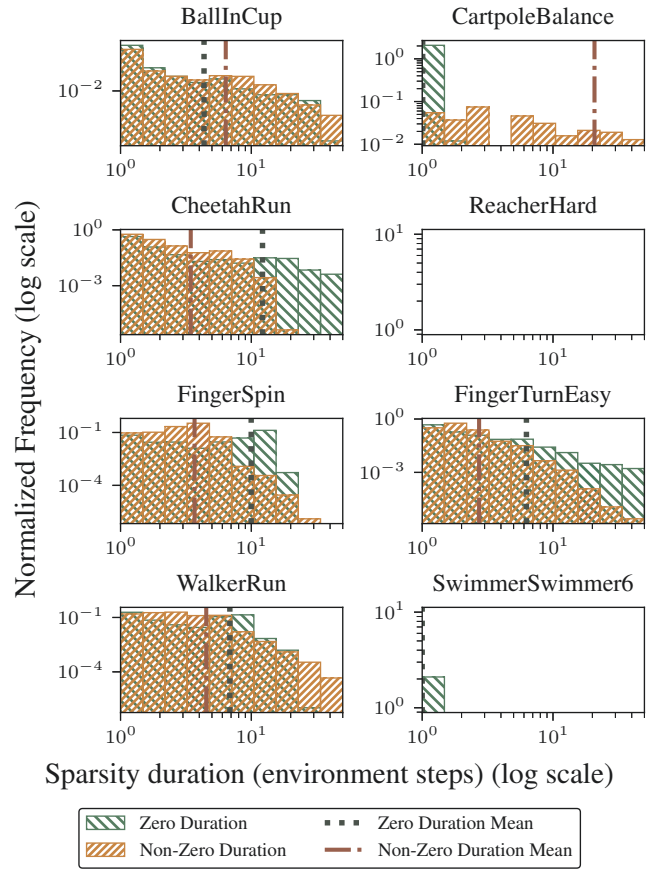


Figure 7: Histograms display the duration distributions of Jacobian elements being zero/non-zero across environments during episode rollouts. The elements mostly remain in the same state over multiple timesteps, indicating temporal sparsity. The Jacobian elements in ReacherHard do not switch between zero/non-zero.

dynamics, though only to a limited, state-dependent extent.

A natural follow-up question is whether a multi-layer perceptron (MLP) trained to approximate $\text{step}(s, a)$ can capture this sparsity. To investigate this question, we tested datasets collected from our benchmark environments. Specifically, we trained a two-layer neural network with ELU activations and 512 units per layer (Section A2), using mean squared error (MSE) loss to *predict the next state*. Datasets were normalized before training (Section A2.2). We consider a Jacobian entry of the MLP with an absolute value below a threshold of $|x| < 10^{-6}$ to be effectively zero.

In addition to the baseline MSE loss, we tested the separate inclusion of ground-truth Jacobian losses (MSE and mean absolute error (MAE)) and an L1 penalty applied exclusively to the predicted Jacobians. We also introduced a sparsity-aware error (SAE) loss, which applies an MAE loss *only* to the Jacobian elements that are expected to be zero based on the ground-truth Jacobian. The results, summarized in Figure 8, revealed a slight trend: using an MAE loss for the Jacobian appeared to improve next-state prediction performance, while




Figure 8: Simple MLP architectures are insufficient to capture ground-truth sparsity. The box plots show the quartiles and the median. (Left) Test-set next-state prediction error aggregated across errors normalized per-environment: adding ground-truth Jacobian loss terms (MSE(J), MAE(J), SAE(J)) barely affected prediction accuracy (only MSE(J) caused a slight reduction), while regularization (L1) reduced performance. (Right) Test set average MLP Jacobian sparsity value compared to ground-truth. Training increases the MLP Jacobian sparsity value over the untrained (Initial) value—even with *ground-truth losses* the sparsity value remains far below the ground-truth (Target).

MSE loss slightly reduced it. However, these differences were not significant. In contrast, the induced sparsity increased significantly when ground-truth Jacobian losses were included, but still remained multiple orders of magnitude below the target sparsity. Notably, the SAE loss increased sparsity without degrading prediction performance, whereas applying L1 regularization to the Jacobians increased sparsity but negatively impacted state prediction accuracy. A reduced sparsity value indicates more entangled predictions compared to the ground truth, and consequently higher prediction errors from spuriously learned correlations. While the naive MLP recovers some sparsity in its predictions, it is insufficient to fully capture the ground-truth sparsity, indicating a need for improved world model architectures.

6 Discussion

Main Findings Our study provides insights into the role of sparsity in learning the dynamics of classical reinforcement learning (RL) environments and its implications for model-based reinforcement learning (MBRL). First, we observed that globally sparse causal structures, as indicated by consistently zero Jacobian elements, are rare across most environments: While some environments, such as ReacherHard and CheetahRun, exhibited limited global sparsity, the majority showed dense interactions between state and action variables. This suggests that enforcing strong global sparsity priors in learned dynamics models may not be universally beneficial. Interestingly, we found evidence of state-dependent sparsity i.e., the causal structure of the dynamics changes based on the current state. For example, in the FingerSpin environment,

sparsity was higher when the finger was not in contact with the object and lower during interactions.

In our learning experiments, we observed that while a two-layer MLP trained to approximate step(s, a) could recover some sparsity in its predictions, the induced sparsity remained far below the target sparsity. Using MAE-based Jacobian losses slightly improved next-state prediction performance compared to MSE-based losses, though the differences were not statistically significant. Including ground-truth Jacobian losses increased induced sparsity but did not fully capture the sparsity present in the ground-truth dynamics, leaving the network’s predictions more entangled. This likely leads to reduced generalization and more compounding errors during multistep prediction.

Limitations While this study provides valuable insights, there are limitations that leave room for further improvement. Due to computational budget reasons, the experiments were conducted with a limited set of agents and environments. Increasing the set of agents and environments would further generalize the findings. Similarly, the architectural space is vast and architectural exploration was thus limited to the most intuitive choice, i.e., equal-width MLPs. While expert trajectories are arguably the most relevant data for modeling the environment and achieving good task performance—the collected data were limited to stochastic-policy rollouts of trained agents—more diverse data could further validate the results.

7 Conclusion

While sparsity priors hold promise for improving sample efficiency and generalization, their effectiveness depends on alignment with the true structure of the environment. Overly strong or misaligned priors could hinder learning by blocking important interactions. Our findings indicate that sparsity is indeed present in many reinforcement learning environments, but requires modeling in a state-dependent way. Additionally, we observed that changes in sparsity often occur in a temporally sparse manner, with periods of stable sparsity interspersed with abrupt transitions such as contact events or phase changes in locomotion. Our results also indicate that naive MLP implementations are insufficient to fully capture and exploit these sparsity structures. These findings highlight the need for further development of dynamics model architectures that can explicitly model sparsity and thus dynamically adapt their causal structure based on state and time, thereby improving generalization and interpretability.

Acknowledgements

We thank Christof Beck, Ananth Rachakonda, and Henri Geiß for their helpful comments on earlier drafts.

This research was partially funded by the Austrian Science Fund (FWF): I 5755-N (ELSA), and by the Autonomous Province of Bolzano-Bozen - South Tyrol under Funding Agreement 10/2024, Abstractron.



References

- Bradbury, J.; Frostig, R.; Hawkins, P.; Johnson, M. J.; Leary, C.; Maclaurin, D.; Necula, G.; Paszke, A.; VanderPlas, J.; Wanderman-Milne, S.; and Zhang, Q. 2018. JAX: Composable Transformations of Python+NumPy Programs.
- Chitnis, R.; Silver, T.; Kim, B.; Kaelbling, L.; and Lozano-Perez, T. 2021. CAMPs: Learning Context-Specific Abstractions for Efficient Planning in Factored MDPs. In *Proceedings of the 2020 Conference on Robot Learning*, 64–79. PMLR.
- Chua, K.; Calandra, R.; McAllister, R.; and Levine, S. 2018. Deep Reinforcement Learning in a Handful of Trials Using Probabilistic Dynamics Models. In *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc.
- Freeman, C. D.; Frey, E.; Raichuk, A.; Girgin, S.; Mordatch, I.; and Bachem, O. 2021. Brax - a Differentiable Physics Engine for Large Scale Rigid Body Simulation. In *Thirty-Fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.
- Gumbsch, C.; Butz, M. V.; and Martius, G. 2021. Sparsely Changing Latent States for Prediction and Planning in Partially Observable Domains. In *Advances in Neural Information Processing Systems*, volume 34, 17518–17531. Curran Associates, Inc.
- Ha, D.; and Schmidhuber, J. 2018. Recurrent World Models Facilitate Policy Evolution. In *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc.
- Haarnoja, T.; Zhou, A.; Abbeel, P.; and Levine, S. 2018. Soft Actor-Critic: Off-Policy Maximum Entropy Deep Reinforcement Learning with a Stochastic Actor. In *Proceedings of the 35th International Conference on Machine Learning*, 1861–1870. PMLR.
- Hafner, D.; Lillicrap, T.; Ba, J.; and Norouzi, M. 2019a. Dream to Control: Learning Behaviors by Latent Imagination. In *International Conference on Learning Representations*.
- Hafner, D.; Lillicrap, T.; Fischer, I.; Villegas, R.; Ha, D.; Lee, H.; and Davidson, J. 2019b. 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*, 2555–2565. PMLR.
- Hafner, D.; Lillicrap, T. P.; Norouzi, M.; and Ba, J. 2021. Mastering Atari with Discrete World Models. In *International Conference on Learning Representations*.
- Hafner, D.; Pasukonis, J.; Ba, J.; and Lillicrap, T. 2025. Mastering Diverse Control Tasks through World Models. *Nature*, 640(8059): 647–653.
- Hollenstein, J.; Martius, G.; and Piater, J. 2024. Colored Noise in PPO: Improved Exploration and Performance through Correlated Action Sampling. *Proceedings of the AAAI Conference on Artificial Intelligence*, 38(11): 12466–12472.
- Hwang, I.; Kwak, Y.; Choi, S.; Zhang, B.-T.; and Lee, S. 2024. Fine-Grained Causal Dynamics Learning with Quantization for Improving Robustness in Reinforcement Learning. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, 20842–20870. PMLR.
- Jain, A. K.; Sujit, S. K.; Joshi, S.; Michalski, V.; Hafner, D.; and Kahou, S. E. 2021. Learning Robust Dynamics through Variational Sparse Gating. In *Deep RL Workshop NeurIPS 2021*.
- Janner, M.; Fu, J.; Zhang, M.; and Levine, S. 2019. When to Trust Your Model: Model-based Policy Optimization. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Lange, R. D.; and Kording, K. P. 2025. Causality in the Human Niche: Lessons for Machine Learning. arXiv:2506.13803v1.
- Lei, A.; Schölkopf, B.; and Posner, I. 2024. SPAR-TAN: A Sparse Transformer Learning Local Causation. arXiv:2411.06890v2.
- Lillicrap, T. P.; Hunt, J. J.; Pritzel, A.; Heess, N.; Erez, T.; Tassa, Y.; Silver, D.; and Wierstra, D. 2016. Continuous Control with Deep Reinforcement Learning. In *International Conference on Learning Representations*.
- Orujlu, T.; Gumbsch, C.; Butz, M. V.; and Wu, C. M. 2025. Reframing Attention as a Reinforcement Learning Problem for Causal Discovery. In *Presented at the Causal Reinforcement Learning Workshop*.
- Pitis, S.; Creager, E.; and Garg, A. 2020. Counterfactual Data Augmentation Using Locally Factored Dynamics. In *Advances in Neural Information Processing Systems*, 3976–3990.
- Pitis, S.; Creager, E.; Mandlkar, A.; and Garg, A. 2022. MoCoDA: Model-based Counterfactual Data Augmentation. In *Advances in Neural Information Processing Systems*, volume 35, 18143–18156. Curran Associates, Inc.
- Schmidhuber, J. 1990a. Making the World Differentiable: On Using Self Supervised Fully Recurrent Neural Networks for Dynamic Reinforcement Learning and Planning in Non-Stationary Environments. *Forschungsberichte, TU Munich, FKI 126 90*: 1–26.
- Schmidhuber, J. 1990b. An On-Line Algorithm for Dynamic Reinforcement Learning and Planning in Reactive Environments. In *1990 IJCNN International Joint Conference on Neural Networks*, 253–258 vol.2.
- Schmidhuber, J. 1990c. Reinforcement Learning in Markovian and Non-Markovian Environments. In *Advances in Neural Information Processing Systems*, volume 3. Morgan-Kaufmann.
- Schmidhuber, J. 1991. Learning Algorithms for Networks with Internal and External Feedback. In *Connectionist Models*, 52–61. Morgan Kaufmann. ISBN 978-1-4832-1448-1.
- Schmidhuber, J. 2015. On Learning to Think: Algorithmic Information Theory for Novel Combinations of Reinforcement Learning Controllers and Recurrent Neural World Models. arXiv:1511.09249v1.
- Schölkopf, B.; Locatello, F.; Bauer, S.; Ke, N. R.; Kalchbrenner, N.; Goyal, A.; and Bengio, Y. 2021. Toward Causal

Representation Learning. *Proceedings of the IEEE*, 109(5): 612–634.

Schulman, J.; Wolski, F.; Dhariwal, P.; Radford, A.; and Klimov, O. 2017. Proximal Policy Optimization Algorithms. arXiv:1707.06347v2.

Sutton, R. S. 1988. Learning to Predict by the Methods of Temporal Differences. *Machine Learning*, 3(1): 9–44.

Tassa, Y.; Doron, Y.; Muldal, A.; Erez, T.; Li, Y.; Casas, D. d. L.; Budden, D.; Abdolmaleki, A.; Merel, J.; Lefrancq, A.; Lillicrap, T.; and Riedmiller, M. 2018. DeepMind Control Suite. arXiv:1801.00690v1.

Todorov, E.; Erez, T.; and Tassa, Y. 2012. MuJoCo: A Physics Engine for Model-Based Control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 5026–5033.

Wang, Z.; Hu, J.; Stone, P.; and Martín-Martín, R. 2023. ELDEN: Exploration via Local Dependencies. In *Advances in Neural Information Processing Systems*, volume 36, 15456–15474. Curran Associates, Inc.

Wang, Z.; Wang, C.; Xiao, X.; Zhu, Y.; and Stone, P. 2024. Building Minimal and Reusable Causal State Abstractions for Reinforcement Learning. *Proceedings of the AAAI Conference on Artificial Intelligence*, 38(14): 15778–15786.

Wang, Z.; Xiao, X.; Xu, Z.; Zhu, Y.; and Stone, P. 2022. Causal Dynamics Learning for Task-Independent State Abstraction. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, 23151–23180. PMLR.

Wang, Z.; Xiao, X.; Zhu, Y.; and Stone, P. 2021. Task-Independent Causal State Abstraction. In *NeurIPS 2021 Workshop on Robot Learning: Self-supervised and Lifelong Learning*.

Zakka, K.; Tabanpour, B.; Liao, Q.; Haiderbhai, M.; Holt, S.; Luo, J. Y.; Allshire, A.; Frey, E.; Sreenath, K.; Kahrs, L. A.; Sferrazza, C.; Tassa, Y.; and Abbeel, P. 2025a. MuJoCo Playground. arXiv:2502.08844v1.

Zakka, K.; Tabanpour, B.; Liao, Q.; Haiderbhai, M.; Holt, S.; Luo, J. Y.; Allshire, A.; Frey, E.; Sreenath, K.; Kahrs, L. A.; Sferrazza, C.; Tassa, Y.; and Abbeel, P. 2025b. MuJoCo Playground: An Open-Source Framework for GPU-accelerated Robot Learning and Sim-to-Real Transfer. https://github.com/google-deepmind/mujoco_playground.

Zhao, Z.; Li, H.; Zhang, H.; Wang, J.; Faccio, F.; Schmidhuber, J.; and Yang, M. 2025. Curious Causality-Seeking Agents Learn Meta Causal World. In *The Thirty-Ninth Annual Conference on Neural Information Processing Systems*.