

# RNM-TD3: N:M Semi-structured Sparse Reinforcement Learning From Scratch

Isam Vrce<sup>1</sup><sup>a</sup>, Andreas Kassler<sup>1</sup><sup>b</sup> and Gökçe Aydos<sup>2</sup><sup>c</sup>

<sup>1</sup>*Deggendorf Institute of Technology, Dieter-Görлиз-Platz 1, Deggendorf, Germany*

<sup>2</sup>*Department of Engineering Technology, Technical University of Denmark, Ballerup, Denmark*

{isam.vrce, andreas.kassler}@th-deg.de, gokay@dtu.dk

**Keywords:** N:M sparsity, TD3, Reinforcement Learning, acceleration

**Abstract:** Sparsity is a well-studied technique for compressing deep neural networks (DNNs) without compromising performance. In deep reinforcement learning (DRL), neural networks with up to 5% of their original weights can still be trained with minimal performance loss compared to their dense counterparts. However, most existing methods rely on unstructured fine-grained sparsity, which limits hardware acceleration opportunities due to irregular computation patterns. Structured coarse-grained sparsity enables hardware acceleration, yet typically degrades performance and increases pruning complexity. In this work, we present, to the best of our knowledge, the first study on *N:M* structured sparsity in RL, which balances compression, performance, and hardware efficiency. Our framework enforces row-wise *N:M* sparsity throughout training for all networks in off-policy RL (TD3), maintaining compatibility with accelerators that support *N:M* sparse matrix operations. Experiments on continuous-control benchmarks show that RNM-TD3, our *N:M* sparse agent, outperforms its dense counterpart at 50%-75% sparsity (e.g., 2:4 and 1:4), achieving up to a 14% increase in performance at 2:4 sparsity on the Ant environment. RNM-TD3 remains competitive even at 87.5% sparsity (1:8), while enabling potential training speedups.

## 1 INTRODUCTION

Modern deep learning (DL) largely relies on over-parameterized networks, whose performance scales with data, model size, and compute (Hestness et al., 2017; Kaplan et al., 2020). Yet, practical deployment demands low latency, low energy consumption, and small memory footprints. Model compression techniques, such as pruning/quantization (Han et al., 2016; Gale et al., 2019), distillation (Hinton et al., 2015), and sparse training (Frankle and Carbin, 2018; Evci et al., 2020), address these challenges by reducing model size while preserving accuracy. In this work, we focus on sparse training in DRL.

A neural network is considered sparse when a large fraction of its parameters are zero. Based on the distribution of zero-valued parameters, sparsity can be classified into three categories: **Unstructured sparsity** zeroes individual weights distributed across the network without any structure, often achieving

high compression with minimal accuracy loss (Zhu and Gupta, 2017; Molchanov et al., 2017; Frantar and Alistarh, 2023). However, the irregular pattern limits hardware speedups. **Structured sparsity** removes entire channels, filters, blocks. This process shrinks tensors, which reduces the number of multiply and accumulate (MAC) operations. However, coarser granularity often yields lower accuracy than unstructured sparse networks at the same sparsity level (Cheng et al., 2024), with some mitigation via dependency-aware pruning at an additional computational cost (Fang et al., 2023). **Semi-structured *N:M* sparsity** enforces a constraint whereby each contiguous group of *M* weights contains no more than *N* nonzero weights (Zhou et al., 2021). When supported in hardware (e.g., 2:4 on NVIDIA Ampere Tensor Cores), this regularity enables accelerated matrix multiplications during the forward pass (Mishra et al., 2021). However, accelerating backpropagation requires an *N:M* sparsity pattern along the columns of the original weight matrix *W* (Zhang et al., 2023; Yang et al., 2025). Beyond the 2:4 sparsity pattern, finer sparsity patterns such as 1:3 and 1:4 have also been explored (Taka et al., 2025), enabling greater flexibility.

---

<sup>a</sup> <https://orcid.org/0009-0000-1722-2812>

<sup>b</sup> <https://orcid.org/0000-0002-9446-8143>

<sup>c</sup> <https://orcid.org/0000-0002-4183-9086>

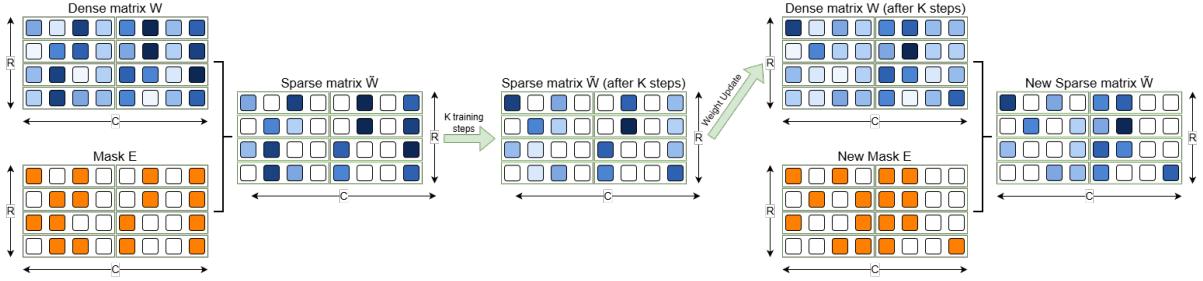


Figure 1: Illustration of a single cycle of the RNM-TD3 algorithm, demonstrated with an  $N : M$  sparsity pattern where  $N = 2$  and  $M = 4$ . The weight matrices depicted have dimensions  $R \times C$  (rows  $\times$  columns). (Left) A dense matrix  $W$  and an  $N : M$  sparse mask  $E$  are combined to create the sparse matrix  $\tilde{W}$ . This sparse matrix is trained for  $K$  steps, during which only the active weights update the underlying dense matrix. (Right) After  $K$  training steps, the weights in the original matrix  $W$  change in magnitude (only at the positions of active weights), resulting in a different mask  $E$ . The updated matrix  $W$  and the new mask  $E$  create a new sparse matrix  $\tilde{W}$ . This cycle repeats until the end of training.

In contrast to supervised learning, DRL involves training on non-stationary data induced by a continually evolving policy, using bootstrapping with delayed rewards and moving targets (Sutton and Barto, 2018; Mnih et al., 2015). These characteristics make DRL training inherently unstable, and pruning can further exacerbate this instability (Livne and Cohen, 2020). Existing sparse DRL approaches either prune pre-trained dense models (Yu et al., 2019) or employ topology-evolution strategies that involve neuron pruning and regrowth (Sokar et al., 2021; Tan et al., 2022). While effective for compression, both strategies limit training-time acceleration, as post-training pruning does not exploit sparse matrices during training.

In this work, we introduce an end-to-end  $N : M$  structured sparse DRL framework that maintains a hardware-aware  $N : M$  sparsity pattern throughout training for both actor and critic networks of the state-of-the-art DRL algorithm Twin Delayed Deep Deterministic Policy Gradient (TD3). We demonstrate that the period of mask updates correlates with Sparse Architecture Divergence (SAD) (Zhou et al., 2021), providing new insights into the dynamics of sparsity during DRL training. Results indicate that training an  $N : M$  sparse agent can match the performance of dense and unstructured baselines. We make the following contributions:

- **Framework for  $N : M$ -sparse DRL training.** We maintain strict  $N : M$  constraints throughout training so that both actor and critic are compatible with hardware-aware sparse matrix-matrix multiplications. (e.g., 2:4 on NVIDIA Ampere Tensor Cores) (Mishra et al., 2021).
- **SAD mask dynamics study.** We relate the period of mask update to the SAD between successive masks. We find that DRL favors less frequent

mask updates compared to supervised learning settings. Moreover, we show that maintaining a stable, non-zero SAD, which is achieved naturally through training, is correlated with improved cumulative reward.

- **Empirical validation.** We demonstrate that on MuJoCo continuous-control benchmarks, our RNM-TD3 agent outperforms dense baselines on average return and rivals unstructured sparse baselines, while offering theoretical acceleration on hardware that supports  $N : M$  sparse computation.

The rest of the paper is structured as follows. Section 2 reviews related work. Section 3 details our methodology and introduces our sparse training strategy. We evaluate our approach in Section 4 and conclude the paper in Section 5.

## 2 RELATED WORK

Early work on compressing neural networks in RL focused on knowledge distillation. A trained agent served as the source model, and its behavior was distilled into a compact student network, resulting in a model size reduction of up to 15 times with no performance loss (Rusu et al., 2015). Building on this, (Schmitt et al., 2018) introduced an unconstrained teacher-student paradigm, reporting a 42% performance gain while requiring  $\sim 10 \times$  fewer training steps for the student.

The first attempts to induce sparsity in DRL focused on post-training pruning: (i) train a dense model, (ii) remove small-magnitude weights, (iii) fine-tune the sparse model. This strategy closely mirrors iterative magnitude pruning techniques developed in deep supervised learning (Han et al., 2015),

and was later adapted to DRL (Livne and Cohen, 2020).

Works inspired by the lottery ticket hypothesis (LTH) (Frankle and Carbin, 2018) applied in RL show that sparse sub-networks, when trained from scratch, can match or even surpass their dense counterparts (Yu et al., 2019; Vischer et al., 2021). In offline RL, single-shot pruning methods (i.e., sparsity is defined at initialization and remains fixed throughout training) such as SNIP (Lee et al., 2018) and GraSP (Wang et al., 2020) remain competitive even at  $\sim 95\%$  sparsity (Arnab et al., 2021). More recently, (Ma et al., 2025) adopted a pre-training pruning scheme based on one-shot random pruning, showing that highly sparse DRL networks can outperform dense baselines.

Dynamic Sparse Training (DST) trains sparse models from scratch by dynamically adapting the network topology, pruning low-importance weights and regrowing new connections during training. While static sparse training can be surprisingly effective in RL (Graesser et al., 2022), DST methods (Mocanu et al., 2018; Evci et al., 2020) aim to optimize performance by evolving the structure alongside the weights. RLx2 (Tan et al., 2022) achieves a compression of  $7.5\times$  to  $20\times$  with  $< 3\%$  performance loss compared to the original dense counterpart.

Following a similar philosophy of dynamic topology adaptation, (Sokar et al., 2021) learn binary masks to select sparse sub-networks, optimizing the topology via periodic updates. Extending the concept of learning sub-networks within a dense network, (Arnab et al., 2024) discover binary task-specific masks, enabling multiple specialized pathways within a single model.

Most existing sparse DRL approaches focus primarily on two objectives: model compression and optimizing policy performance under sparsity constraints. However, these methods typically rely on unstructured sparsity, which suffers from irregular memory access patterns and fails to translate theoretical reductions in floating-point operations (FLOPs) into real-world speedups on standard hardware (Wen et al., 2016). In contrast, we aim to optimize on three fronts simultaneously: model compression, policy quality, and training acceleration. Our method enforces hardware-aware  $N:M$  sparsity from the beginning of the training with no dense initialization. This ensures that the computational efficiency of the  $N:M$  pattern is utilized throughout the entire training, enabling acceleration from the very first step.

### 3 METHODOLOGY

In this section, we present RNM-TD3, an algorithm designed to train a row-wise  $N:M$  structured sparse agent from scratch. We build on TD3, a widely used off-policy actor-critic algorithm. Using the deterministic TD3 algorithm as our backbone enables us to isolate the impact of  $N:M$  structured sparsity on performance without the additional factors introduced by stochastic policy methods. Figure 1 provides a high-level description of the key idea, how the row-wise  $N:M$  structured sparsity is maintained during training. We first introduce the problem formulation and notation, describe our approach in more detail, highlight the key differences between sparse training in DRL and supervised learning and our method to stabilize training through soft-reset.

#### 3.1 Problem Formulation & Notation

We frame the RL problem as a standard Markov decision process (MDP)  $\mathcal{M} = (\mathcal{S}, \mathcal{A}, \mathcal{P}, r, \gamma)$ . Here,  $\mathcal{S}$  denotes the state space,  $\mathcal{A}$  is the continuous action space,  $\mathcal{P}$  represents the transition probability distribution,  $r$  is the reward function, and  $\gamma \in [0, 1)$  is the discount factor.

The TD3 agent consists of a deterministic actor  $\mu_\theta(s)$  and twin critics  $Q_{\phi_1}, Q_{\phi_2}$ , parameterized by  $\theta$  and  $\phi_1, \phi_2$ , respectively. To enable hardware acceleration, we enforce a row-wise  $N:M$  sparsity pattern, meaning that for each row of a weight matrix and each group of  $M$  consecutive weights in that row, at most  $N$  weights are non-zero.

Let  $W^{(\ell)} \in \mathbb{R}^{d_{\text{out}} \times d_{\text{in}}}$  denote the weight matrix of a fully connected layer  $\ell$ , where  $d_{\text{in}}$  and  $d_{\text{out}}$  represent the number of input and output neurons, respectively. Our objective is to maximize the expected return while ensuring that  $W^{(\ell)}$  satisfies the row-wise  $N:M$  sparsity pattern:

$$\forall i \in \{1, \dots, d_{\text{out}}\}, \forall j \in \{1, \dots, \lceil d_{\text{in}}/M \rceil\} : \|W_{i, (j-1)M+1 : \min\{jM, d_{\text{in}}\}}^{(\ell)}\|_0 \leq N. \quad (1)$$

where  $\|\cdot\|_0$  denotes the  $L_0$  norm, which counts the number of non-zero elements.

Intuitively, Eq. (1) partitions each row of  $W^{(\ell)}$  into contiguous blocks of size  $M$ . Within each block, at most  $N$  entries are non-zero. Sparsity is enforced via binary masks  $E^{(\ell)} \in \{0, 1\}^{d_{\text{out}} \times d_{\text{in}}}$ , yielding effective weights  $\tilde{W}^{(\ell)} = W^{(\ell)} \odot E^{(\ell)}$ . The projection operator  $\mathcal{P}_{N:M}$  used to construct the mask is defined as:

$$\mathcal{P}_{N:M}(W)_{i,j} = \begin{cases} 1, & \text{if } |W_{i,j}| \text{ is in the top-}N \\ & \text{of block } (i, \lfloor (j-1)/M \rfloor) \\ 0, & \text{otherwise.} \end{cases} \quad (2)$$

Masks are initialized to satisfy Eq. (1) and are applied to all six TD3 networks (i.e., the online actor network, the twin critic networks, and their corresponding target networks). Online masks are updated periodically during training, while remaining fixed between updates. Target networks are updated by Polyak averaging (Lillicrap et al., 2015):

$$\bar{\Psi} \leftarrow \tau \Psi + (1 - \tau) \bar{\Psi}, \quad \Psi \in \{\theta, \phi_1, \phi_2\}. \quad (3)$$

where  $\tau$  is a small constant (typically  $\tau = 0.005$ ) that controls the interpolation rate. However, during the online mask update, the online mask is immediately copied to the corresponding targets to maintain stability.

### 3.2 Training under $N:M$ Sparsity

Training starts with  $N:M$  sparse neural networks. The weight matrix  $W$  is initialized using a Kaiming uniform distribution, adjusted to account for the reduced effective fan-in induced by the  $N:M$  sparsity constraint. Immediately after initialization and prior to the first training step, we apply the projection operator from Eq. (2):

$$E = \mathcal{P}_{N:M}(W), \quad \tilde{W} = W \odot E. \quad (4)$$

The resulting sparse matrix  $\tilde{W}$  is now  $N:M$  sparse. To update the weights in the dense weights  $W$ , we propagate the gradients from the loss  $\mathcal{L}$  only to the active weights, while the gradients for inactive weights are zeroed out:

$$\frac{\partial \mathcal{L}}{\partial W} = \frac{\partial \mathcal{L}}{\partial \tilde{W}} \odot E. \quad (5)$$

After  $K$  training steps, if the magnitude of an active weight becomes smaller than that of any inactive weight within the same group of  $M$  weights, the inactive weight with the largest magnitude is reactivated at the next mask update, thereby altering the network topology.

However, this approximation can induce unstable mask updates, often referred to as "mask flickering". To quantify topological evolution, we measure the Sparse Architecture Divergence (SAD) (Zhou et al., 2021), defined as the Hamming distance between consecutive masks:

$$\text{SAD}_{t,t+k} = \|E_{t+k} - E_t\|_1. \quad (6)$$

High SAD indicates significant topological changes, while low SAD suggests convergence to a fixed sparsity pattern. While supervised learning typically benefits from minimizing SAD to ensure convergence (Zhou et al., 2021; Yang et al., 2025), our analysis suggests that DRL requires a different regime due to the non-stationary data distribution.

### 3.3 Differences Between $N:M$ Sparse Training in DRL and Supervised Learning

In supervised learning,  $N:M$  sparse training typically updates masks at every optimization step. To reduce SAD between successive mask updates and improve final accuracy, supervised learning techniques introduced a weight factor  $\lambda_W$  that progressively drives inactive weights toward zero at each step: Sparse Refined Straight-Through Estimator (SR-STE) (Zhou et al., 2021).

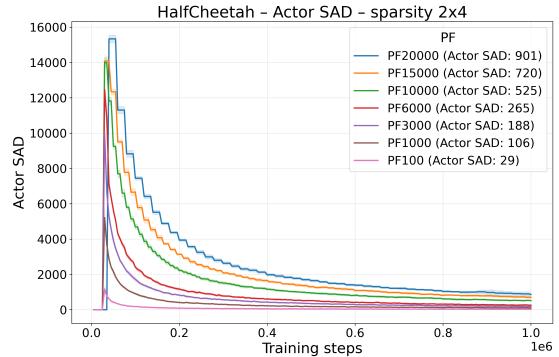


Figure 2: SAD for 7 different mask update periods for the actor network in the Half Cheetah environment. SAD decays approximately exponentially regardless of the mask update period.

In DRL, however, we find that superior performance is achieved with much less frequent mask updates, occurring every  $K \gg 1$  environment steps. A controlled ablation study over  $K$  (Section 4.2) reveals an optimal mask update period of around  $K = 4000$  for HalfCheetah environment with 1:4 sparsity. The optimal "K" depends on both the environment and the sparsity level. Values around  $K = 5000$  environment steps work well across tasks and sparsity patterns. SAD decays approximately exponentially toward a stable non-zero value (see Figure 2), in contrast to the concave decay trajectories observed in supervised learning (see Figure 4a in (Zhou et al., 2021)). This suggests that explicit SAD reduction mechanisms, such as those in SR-STE, are unnecessary or even detrimental in the DRL setting. We ex-

amined two natural variants of applying the SR-STE shrinkage factor  $\lambda_W$  within DRL:

1. **Per-step application.** Applying  $\lambda_W$  at every training step causes inactive weights that are masked by chance at initialization to be repeatedly shrunk for the next  $K$  steps (i.e., until the next mask update). This acts as an one-shot pruning of potentially important connections, leading to the premature suppression of useful weights and degraded performance.
2. **Per-mask-update application.** Applying  $\lambda_W$  only once when the mask is updated avoids per-step shrinkage. However, in momentum-based optimizers (e.g., Adam (Kingma and Ba, 2014)), this factor still contributes across many subsequent updates. Since the gradient for inactive weights is zero,  $\lambda_W$  becomes the only factor affecting their magnitude, progressively driving them toward zero between mask updates. This reduces the likelihood of weight reactivation and, consequently, limits topological flexibility.

Our empirical results suggest that in DRL, the goal is not to drive SAD to zero. Instead, a stable, non-zero SAD emerges naturally when masks are updated at the appropriate frequency. Configurations with larger SAD exhibit greater adaptability to changing policy dynamics, whereas frequent mask updates result in low SAD, which reduces topological flexibility and harms final performance.

### 3.4 Stabilizing $N:M$ Training Through Soft Reset

Training highly sparse  $N:M$  networks in inherently unstable environments, such as *Humanoid*, can lead to training instabilities. The combination of an unstable environment, high sparsity ratios (e.g., 1:8), and dynamic sparse training from initialization often results in a network topology that lacks the capacity to model complex locomotion dynamics during the early stages of training. There is a high probability of severing the information flow from key sensory inputs (e.g., knee angle sensors) to critical actuators (e.g., leg motors) by masking important weights. Consequently, the agent attempts to maintain balance by activating non-critical actuators, such as those in the arms, which fails to produce meaningful locomotion. The agent falls shortly after the start of the episode, generating a series of short episodes that prevent the accumulation of meaningful learning experiences.

After a mask change, the critic suggests target values for the actor, which the actor attempts to match by minimizing the error. However, due to high sparsity,

the few remaining active connections must compensate for the missing connections. This is especially problematic at the beginning of training when weights are randomly initialized. To achieve the output shifts required by the critic using such limited connectivity, the actor aggressively increases the magnitude of the active weights. Given the tanh nonlinearity at the output, these high-magnitude pre-activations cause the neurons to saturate (locking at  $\pm 1$ ). As a result, the agent repeatedly selects identical actions, reducing the SAD to zero.

Empirically, with a mask-update period of  $K = 15,000$ , early performance collapse occurred in approximately 80% of runs in the *Humanoid* task. To address this, we introduce a simple *soft reset* mechanism. Whenever SAD falls below a critical threshold (e.g.,  $SAD < 10$ ) during a mask update, we resample a new random  $N:M$  mask. This forces topological exploration. The soft reset completely eliminates early collapse. After recovery, training proceeds normally, and SAD maintains a healthy non-zero value throughout the remainder of the training.

## 4 EXPERIMENTAL EVALUATION

We evaluate our proposed sparse TD3 variant, **RNM-TD3**, on continuous-control tasks from the Gymnasium (Towers et al., 2024) MuJoCo benchmark suite (Todorov et al., 2012). RNM-TD3 introduces a general  $N:M$  structured sparsity framework built on top of the TD3 algorithm, agnostic to both network architecture and specific hyperparameters. Our experimental study addresses two key questions: (i) how different  $N:M$  sparsity patterns impact performance across diverse environments relative to dense and unstructured baselines, and (ii) how the mask-update period influences both the return and the dynamics of the sparse topology.

We compare our method against four baselines:

- **Dense**, TD3 with dense neural networks.
- **SSN N:M**, a static sparse network with a fixed  $N:M$  mask sampled at initialization and never updated.
- **DS-TD3** (Sokar et al., 2021): A dynamic sparse training method utilizing random regrowth.
- **RLx2** (Tan et al., 2022): A method that prunes weights with the smallest magnitude and regrows connections based on gradient magnitude.

To ensure a fair comparison, all methods use identical architectures and hyperparameters. The actor and both critic networks consist of two hidden layers

Table 1: Comparison of RNM-TD3 with baselines. Results are reported as mean  $\pm$  standard deviation, normalized relative to the dense baseline (i.e., scaled such that the dense mean and standard deviation are 1.0). Raw dense baseline scores are provided for reference: Ant-v5 ( $4453 \pm 1320$ ), HalfCheetah-v5 ( $10211 \pm 1110$ ), Humanoid-v5 ( $4868 \pm 659$ ), and Walker2d-v5 ( $3767 \pm 833$ ). TS denotes Target Sparsity (%). RNM-TD3 and SNM has the same realized sparsity which slightly deviates from TS because the last layer remains dense (note that last layer is dense in all methods). DS-TD3 method uses the same layer sparsity parameter for actor and critic networks, so realized sparsity will always differ from target sparsity.

Environment	TS(%)	RNM-TD3(best) Act.(%) Crit.(%)		SSN-N:M (best) Act.(%) Crit.(%)		DS-TD3 Act.(%) Crit.(%)		RLx2 Act./Crit.(%)
Ant-v5	50	$1.14 \pm 0.05$ (2:4)		$1.05 \pm 0.21$ (4:8)		$0.91 \pm 0.18$		$0.93 \pm 0.16$
		48.92	49.86	48.92	49.86	48.24	49.94	50
	75	$1.07 \pm 0.15$ (1:4)		$0.93 \pm 0.28$ (1:4)		$0.78 \pm 0.16$		$1.01 \pm 0.12$
		73.37	74.79	73.37	74.79	72.84	74.64	75
	87.5	$0.86 \pm 0.16$ (1:8)		$0.75 \pm 0.18$ (1:8)		$0.83 \pm 0.19$		$0.82 \pm 0.19$
		85.60	87.26	85.60	87.26	85.34	87.19	87.5
H.Cheetah-v5	50	$1.01 \pm 0.08$ (4:8)		$0.96 \pm 0.08$ (2:4)		$1.05 \pm 0.09$		$0.95 \pm 0.10$
		48.92	49.82	48.92	49.82	47.77	49.65	50
	75	$0.90 \pm 0.05$ (1:4)		$0.82 \pm 0.07$ (1:4)		$0.95 \pm 0.17$		$0.92 \pm 0.10$
		73.39	74.73	73.39	74.73	73.00	74.84	75
	87.5	$0.72 \pm 0.06$ (1:8)		$0.68 \pm 0.05$ (1:8)		$0.81 \pm 0.08$		$0.80 \pm 0.05$
		85.62	87.18	85.62	87.18	85.23	87.05	87.5
Humanoid-v5	50	$1.05 \pm 0.03$ (2:4)		$1.03 \pm 0.05$ (2:4)		$1.04 \pm 0.02$		$0.99 \pm 0.03$
		48.63	49.92	48.63	49.92	48.16	50.01	50
	75	$1.04 \pm 0.03$ (1:4)		$1.01 \pm 0.12$ (1:4)		$0.99 \pm 0.06$		$0.93 \pm 0.07$
		72.95	74.88	72.95	74.88	72.52	74.73	75
	87.5	$0.28 \pm 0.20$ (1:8) *		$0.82 \pm 0.35$ (1:8)		$1.02 \pm 0.03$		$0.64 \pm 0.26$
		85.10	87.36	85.10	87.36	84.89	87.28	87.5
Walker2d-v5	50	$1.04 \pm 0.13$ (4:8)		$0.94 \pm 0.22$ (4:8)		$1.18 \pm 0.13$		$0.97 \pm 0.15$
		48.92	49.82	48.92	49.82	47.77	49.65	50
	75	$0.96 \pm 0.09$ (1:4)		$0.89 \pm 0.12$ (1:4)		$0.99 \pm 0.26$		$0.82 \pm 0.24$
		73.39	74.73	73.39	74.73	73.00	74.84	75
	87.5	$0.58 \pm 0.29$ (1:8)		$0.53 \pm 0.26$ (2:16)		$0.86 \pm 0.30$		$0.83 \pm 0.20$
		85.62	87.18	85.62	87.18	85.23	87.05	87.5

with 256 neurons. All experiments run for 1M environment steps using a standard static replay buffer and 1-step returns. Except for their specific sparsity mechanisms, all algorithms share identical TD3 settings (e.g., learning rate, batch size). Performance is reported as the average return over 11 random seeds for each environment over 1M steps.

#### 4.1 Comparative evaluation on continuous RL tasks

As shown in Table 1, RNM-TD3 exhibits superior performance across nearly all evaluated setups compared to static  $N:M$  sparse training (SSN), with the exception of Humanoid-v5 at 1:8 sparsity ("\*\*" in the table). Our method consistently achieves higher average returns with lower variance, confirming that dynamic mask updates significantly improve upon static initialization. Notably, at 2:4 sparsity, RNM-TD3 outperforms the dense baseline in all environments and

remains competitive at 1:4 sparsity.

In most configurations, our structured sparse approach yields performance gains over unstructured dynamic methods (DS-TD3 and RLx2), achieving up to a 25% improvement in the complex Ant environment at 50% sparsity. While unstructured methods occasionally perform better in lower-dimensional environments or at extreme sparsity levels, they lack the hardware-acceleration potential inherent to our  $N:M$  structured approach.

We note that the original RLx2 implementation (Tan et al., 2022) reports higher scores by utilizing a dynamic replay buffer and 3-step returns after 300K steps. However, to isolate the effect of the sparsity mechanism itself, we evaluate all methods under a unified, standard TD3 configuration.

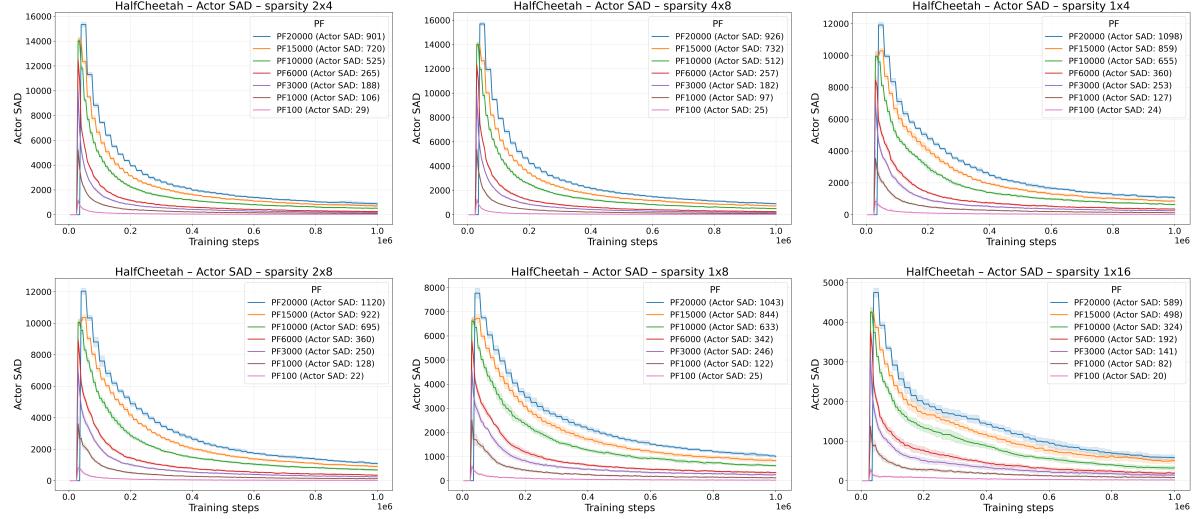


Figure 3: SAD for the HalfCheetah actor network for different sparsity patterns and mask update periods. Each curve is averaged over 11 seeds. Overall 462 runs (7 update periods  $\times$  6 sparsity patterns  $\times$  11 seeds) confirm a clear correlation between the mask update period and SAD: more frequent mask updates (smaller update periods) lead to lower SAD, while less frequent updates (larger update periods) result in higher SAD.

## Ablation Study: Impact of Initialization on N:M Sparse Training

To validate the effectiveness of the proposed initialization strategy, we compare our method, which adjusts the Kaiming distribution to account for the effective fan-in of  $N:M$  sparsity patterns, against initialization sampling from the standard Kaiming distribution used for dense networks. Adjusting for effective fan-in results in weight values that are scaled by a factor of  $\sqrt{M/N}$  at initialization. The results, normalized to the dense baseline, are presented in Table 2.

The most significant observation across environments and sparsity levels is the reduction in variance when using the  $N:M$  adjusted initialization. However, under high sparsity regimes and in unstable environments (e.g., *Humanoid*),  $\sqrt{M/N}$  scaling can produce initial weights with excessively large magnitudes, causing the actor’s outputs to saturate shortly after training begins. This saturation leads to the vanishing gradients problem. When neurons operate in the saturated regime of the tanh nonlinearity, where the derivative approaches zero, gradients diminish significantly during backpropagation, effectively preventing weight updates in the earlier layers.

Table 2: Performance comparison between Kaiming initialization adjusted for N:M fan-in (N:M Adj.) against standard Kaiming initialization (Standard K.). All results are reported as normalized mean scores  $\pm$  normalized standard deviation relative to the dense baselines. A mean  $> 1.0$  indicates performance superior to the dense model.

Env.	Spar.	N:M Adj.	Standard K.
Ant	(2:4)	$1.14 \pm 0.05$	$1.09 \pm 0.45$
	(1:4)	$1.07 \pm 0.15$	$0.94 \pm 0.39$
	(1:8)	$0.86 \pm 0.16$	$0.75 \pm 0.44$
HC	(2:4)	$1.01 \pm 0.08$	$0.96 \pm 1.01$
	(1:4)	$0.90 \pm 0.05$	$0.88 \pm 0.50$
	(1:8)	$0.72 \pm 0.06$	$0.71 \pm 0.61$
Hum.	(2:4)	$1.05 \pm 0.03$	$1.04 \pm 0.37$
	(1:4)	$1.04 \pm 0.03$	$1.03 \pm 0.15$
	(1:8)	$0.28 \pm 0.20$	$0.90 \pm 1.10$

## 4.2 Ablation study on mask update period

We investigate the relationship between the mask update period, SAD, and agent performance. We vary the mask update period from 10 to 20,000 environment steps in the HalfCheetah environment using a 1:4 sparsity pattern. Results are averaged over 21 random seeds for 1M steps.

Table 3 reveals a clear trend: as the mask update period increases, the average return improves, peaking at 4,000 environment steps before degrading. Crucially, the worst-performing configurations, which correspond to very frequent updates, are asso-

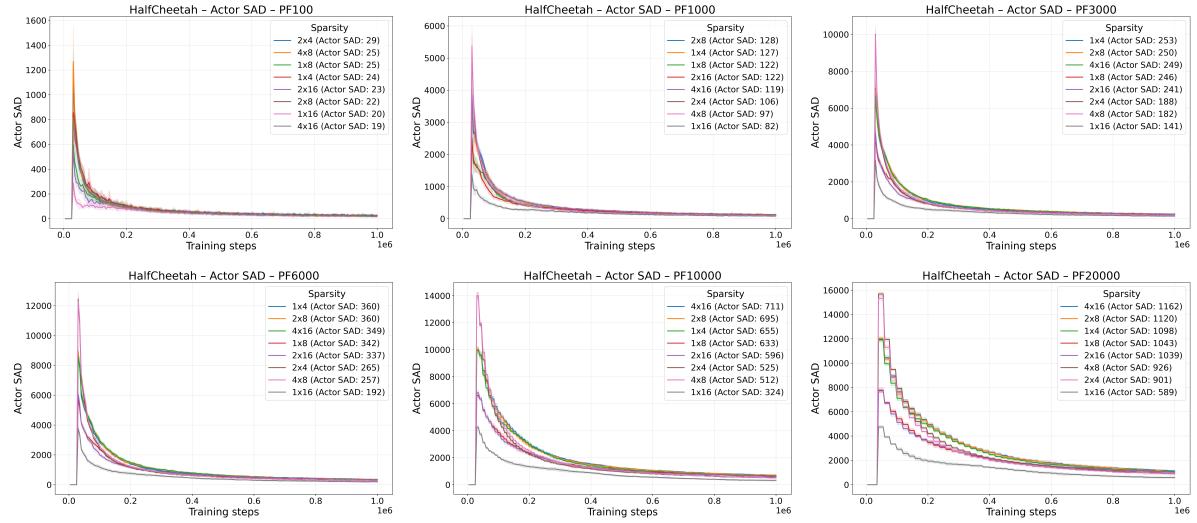


Figure 4: SAD for the HalfCheetah actor network across different N:M sparsity patterns and mask update periods. Each curve is averaged over 11 seeds. Across 462 runs, the sparsity level has only a minor effect on SAD for a fixed mask-update period

Table 3: Impact of the mask update period on HalfCheetah performance and SAD with 1:4 sparsity.

MA period	Avg. $\pm$ Std.	Act. SAD	Crit. SAD
10	$7895 \pm 816$	6	6
50	$8251 \pm 570$	17	21
100	$8298 \pm 895$	26	36
1000	$8057 \pm 1120$	128	185
2000	$8286 \pm 1139$	216	307
4000	$9036 \pm 534$	346	495
6000	$8861 \pm 805$	477	679
8000	$8717 \pm 753$	566	809
10000	$8453 \pm 782$	654	935
15000	$8568 \pm 817$	867	1243
20000	$8570 \pm 634$	1071	1469

ciated with very low SAD values. The intuition behind the observed performance trend is as follows:

- **Too small update period:** the network architecture changes too frequently, and the mask is updated before the underlying weights can sufficiently adapt.
- **Too large update period:** the architecture adapts too slowly, and each mask update produces disruptive jumps in average return.
- **Intermediate update period:** strikes a balance between stability and adaptability, allowing the sparse topology to remain stable long enough to learn while still updating frequently enough to remain responsive, resulting in peak performance.

An additional pattern can be observed: SAD appears to be correlated with the mask update period. To test this hypothesis, we conducted experiments in

all four environments across eight different sparsity patterns. Figure 3 shows the SAD during the training for the actor network in the HalfCheetah environment. We observe that low SAD correlates with high frequency mask update periods. The same qualitative behavior is observed across all environments and for both actor and critic networks. Furthermore, Figure 4 indicates that, for a fixed mask update period, the sparsity level has only a minor effect on SAD. This implies that the choice of mask update period has more influence on SAD than the sparsity level.

**Computational Efficiency Analysis.** Although our implementation does not utilize bidirectional N:M sparse masks (i.e., row-wise and column-wise N:M sparsity pattern), the computational cost is dominated by forward passes. In our setup (one gradient update per environment step and a policy delay of two), each environment step induces 7 forward passes and 2.5 backward passes across the actor, critics, and target networks. Consequently, even if hardware acceleration is applied only to the forward pass, the end-to-end training speedup remains substantial, but lower than the theoretical  $M/N$  factor.

## 5 CONCLUSIONS & FUTURE WORK

In this work, we introduce RNM-TD3, a semi-structured sparse DRL algorithm that maintains strict N:M sparsity throughout training. By enforcing a hardware-aware sparsity pattern across the actor and

critic networks, our method is aligned with the sparse matrix-matrix multiplication accelerators. Given that the computational cost of TD3 is dominated by forward passes, our approach offers practical end-to-end training speedups, addressing a key limitation of unstructured sparse methods, which often fail to translate theoretical FLOP reductions into real-world acceleration.

Our empirical analysis highlights the critical role of mask update dynamics. We found that a moderate, non-zero Sparse Architecture Divergence (SAD) correlates strongly with improved performance. We hypothesize that maintaining a higher SAD helps preserve topological flexibility, which may be beneficial in long-term training regimes where plasticity typically diminishes. Furthermore, in environments with changing dynamics, agents with larger SAD are expected to exhibit greater adaptability.

Finally, we observe that in high sparsity regimes, weight initialization plays a crucial role, making the early training phase particularly fragile. To address this, future work should investigate novel initialization schemes specifically tailored for high-sparsity regimes in complex environments and adaptive mask update schedules.

## ACKNOWLEDGEMENTS

Parts of this work were supported by the project EFBL, funded by the Bavarian State Ministry for Science and the Arts, and by the High-Tech Agenda of Bavaria.

## REFERENCES

Arnob, S. Y., Ohib, R., Plis, S., and Precup, D. (2021). Single-shot pruning for offline reinforcement learning. *arXiv preprint arXiv:2112.15579*.

Arnob, S. Y., Ohib, R., Plis, S., Zhang, A., Sordoni, A., and Precup, D. (2024). Efficient reinforcement learning by discovering neural pathways. *Advances in Neural Information Processing Systems*, 37:18660–18694.

Cheng, H., Zhang, M., and Shi, J. Q. (2024). A survey on deep neural network pruning: Taxonomy, comparison, analysis, and recommendations. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.

Evci, U., Gale, T., Menick, J., Castro, P. S., and Elsen, E. (2020). Rigging the lottery: Making all tickets winners. In *International conference on machine learning*, pages 2943–2952. PMLR.

Fang, G., Ma, X., Song, M., Mi, M. B., and Wang, X. (2023). Depgraph: Towards any structural pruning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 16091–16101.

Frankle, J. and Carbin, M. (2018). The lottery ticket hypothesis: Finding sparse, trainable neural networks. *arXiv preprint arXiv:1803.03635*.

Frantar, E. and Alistarh, D. (2023). Sparsegpt: Massive language models can be accurately pruned in one-shot. In *International conference on machine learning*, pages 10323–10337. PMLR.

Gale, T., Elsen, E., and Hooker, S. (2019). The state of sparsity in deep neural networks. *arXiv preprint arXiv:1902.09574*.

Graesser, L., Evci, U., Elsen, E., and Castro, P. S. (2022). The state of sparse training in deep reinforcement learning. In *International Conference on Machine Learning*, pages 7766–7792. PMLR.

Han, S., Mao, H., and Dally, W. J. (2016). Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding. In *International Conference on Learning Representations (ICLR)*.

Han, S., Pool, J., Tran, J., and Dally, W. (2015). Learning both weights and connections for efficient neural network. *Advances in neural information processing systems*, 28.

Hestness, J., Narang, S., Ardalani, N., Diamos, G., Jun, H., Kianinejad, H., Patwary, M. M. A., Yang, Y., and Zhou, Y. (2017). Deep learning scaling is predictable, empirically. *arXiv preprint arXiv:1712.00409*.

Hinton, G., Vinyals, O., and Dean, J. (2015). Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*.

Kaplan, J., McCandlish, S., Henighan, T., Brown, T. B., Chess, B., Child, R., Gray, S., Radford, A., Wu, J., and Amodei, D. (2020). Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361*.

Kingma, D. P. and Ba, J. (2014). Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980.

Lee, N., Ajanthan, T., and Torr, P. H. (2018). Snip: Single-shot network pruning based on connection sensitivity. *arXiv preprint arXiv:1810.02340*.

Lillicrap, T. P., Hunt, J. J., Pritzel, A., Heess, N., Erez, T., Tassa, Y., Silver, D., and Wierstra, D. (2015). Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971*.

Livne, D. and Cohen, K. (2020). Pops: Policy pruning and shrinking for deep reinforcement learning. *IEEE Journal of Selected Topics in Signal Processing*, 14(4):789–801.

Ma, G., Li, L., Wang, Z., Shen, L., Bacon, P.-L., and Tao, D. (2025). Network sparsity unlocks the scaling potential of deep reinforcement learning. *arXiv preprint arXiv:2506.17204*.

Mishra, A. K., Latorre, J. A., Pool, J., Stosic, D., Stosic, D., Venkatesh, G., Yu, C., and Micikevicius, P. (2021). ngl sparse deep neural networks. *CoRR*, abs/2104.08378. uses 2:4 sparsity pattern, metadata to encode the de-multiplexer control line.page 5.

Mnih, V., Kavukcuoglu, K., Silver, D., et al. (2015). Human-level control through deep reinforcement learning. *Nature*, 518:529–533.

Mocanu, D. C., Mocanu, E., Stone, P., Nguyen, P. H., Gibescu, M., and Liotta, A. (2018). Scalable training of artificial neural networks with adaptive sparse connectivity inspired by network science. *Nature communications*, 9(1):2383.

Molchanov, D., Ashukha, A., and Vetrov, D. (2017). Variational dropout sparsifies deep neural networks. In *International conference on machine learning*, pages 2498–2507. PMLR.

Rusu, A. A., Colmenarejo, S. G., Gulcehre, C., Desjardins, G., Kirkpatrick, J., Pascanu, R., Mnih, V., Kavukcuoglu, K., and Hadsell, R. (2015). Policy distillation. *arXiv preprint arXiv:1511.06295*.

Schmitt, S., Hudson, J. J., Zidek, A., Osindero, S., Doersch, C., Czarnecki, W. M., Leibo, J. Z., Kuttler, H., Zisserman, A., Simonyan, K., et al. (2018). Kick-starting deep reinforcement learning. *arXiv preprint arXiv:1803.03835*.

Sokar, G., Mocanu, E., Mocanu, D. C., Pechenizkiy, M., and Stone, P. (2021). Dynamic sparse training for deep reinforcement learning. *arXiv preprint arXiv:2106.04217*.

Sutton, R. S. and Barto, A. G. (2018). *Reinforcement Learning: An Introduction*. A Bradford Book, Cambridge, MA, USA.

Taka, E., Huang, N.-C., Chang, C.-C., Wu, K.-C., Arora, A., and Marculescu, D. (2025). Systolic sparse tensor slices: Fpga building blocks for sparse and dense ai acceleration. In *Proceedings of the 2025 ACM/SIGDA International Symposium on Field Programmable Gate Arrays*, pages 159–171.

Tan, Y., Hu, P., Pan, L., Huang, J., and Huang, L. (2022). Rlx2: Training a sparse deep reinforcement learning model from scratch. *arXiv preprint arXiv:2205.15043*.

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*, pages 5026–5033.

Towers, M., Kwiatkowski, A., Terry, J., Balis, J. U., De Cola, G., Deleu, T., Goulão, M., Kallinteris, A., Krimmel, M., KG, A., et al. (2024). Gymnasium: A standard interface for reinforcement learning environments. *arXiv preprint arXiv:2407.17032*.

Vischer, M. A., Lange, R. T., and Sprikeler, H. (2021). On lottery tickets and minimal task representations in deep reinforcement learning. *arXiv preprint arXiv:2105.01648*.

Wang, C., Zhang, G., and Grosse, R. (2020). Picking winning tickets before training by preserving gradient flow. *arXiv preprint arXiv:2002.07376*.

Wen, W., Wu, C., Wang, Y., Chen, Y., and Li, H. (2016). Learning structured sparsity in deep neural networks. *Advances in neural information processing systems*, 29.

Yang, C., Nie, W., Zhang, Y., Wu, Y., Zheng, X., Jiang, G., and Ji, R. (2025). Bame: Block-aware mask evolution for efficient n:m sparse training. In *International conference on machine learning*.

Yu, H., Edunov, S., Tian, Y., and Morcos, A. S. (2019). Playing the lottery with rewards and multiple languages: lottery tickets in rl and nlp. *arXiv preprint arXiv:1906.02768*.

Zhang, Y., Luo, Y., Lin, M., Zhong, Y., Xie, J., Chao, F., and Ji, R. (2023). Bi-directional masks for efficient n: M sparse training. In *International conference on machine learning*, pages 41488–41497. PMLR.

Zhou, A., Ma, Y., Zhu, J., Liu, J., Zhang, Z., Yuan, K., Sun, W., and Li, H. (2021). Learning N: M fine-grained structured sparse neural networks from scratch. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.

Zhu, M. and Gupta, S. (2017). To prune, or not to prune: exploring the efficacy of pruning for model compression. *arXiv preprint arXiv:1710.01878*.