Towards Implementing Truly Sparse Connections in Deep RL Agents

Bram Grooten[1], Ghada Sokar[1], Elena Mocanu[2], Shibhansh Dohare[3], Matthew Taylor[3], Mykola Pechenizkiy[1] and Decebal Constantin Mocanu[1,2]

[1]Department of Mathematics and Computer Science, Eindhoven University of Technology, the Netherlands
[2]Faculty of Electrical Engineering, Mathematics and Computer Science, University of Twente, the Netherlands
[3]Department of Computing Science, University of Alberta, Canada

Context

There is considerable room for improvement in the efficiency of reinforcement learning (RL) algorithms, both during training and inference [2]. Sparse neural networks have the potential to decrease the running time substantially. Dynamic sparse training (DST), which drops and regrows a proportion of connections regularly during training, is proven to be an effective method to train highly sparse networks in supervised learning [3]. Still, it has been explored just in very few recent works in deep RL.

- Sokar et al. [4] proposed DS-TD3 and DS-SAC, by combining the sparse evolutionary training (SET) procedure, a DST method, with the sample efficient algorithms TD3 and SAC.
- DS-TD3 and DS-SAC achieved 50% sparsity in five continuous control tasks, while outperforming dense training in 7 out of 10 cases.

Contributions

1. We develop an algorithmic framework with the aim of opening the path for truly sparse training implementations (without binary masks) in deep RL. Accordingly, we enhance the previous work [4] when using the Adam optimizer, by masking not only the weight and current gradient of non-existing connections, but also the running averages of the first and second moment of the gradients. We name our dynamic sparse training algorithms TowS-SAC and TowS-TD3, working ***towards truly sparse*** implementations.
2. We explore the impact of LeakyReLU on sparse training for deep RL, as the sparse activations of ReLU might be detrimental when training neural networks that are already sparse.

Dynamic sparse SAC

In Humanoid-v3, considered the most difficult environment of the MuJoCo Gym suite, we reach 90% sparsity while outperforming dense TD3 and SAC.

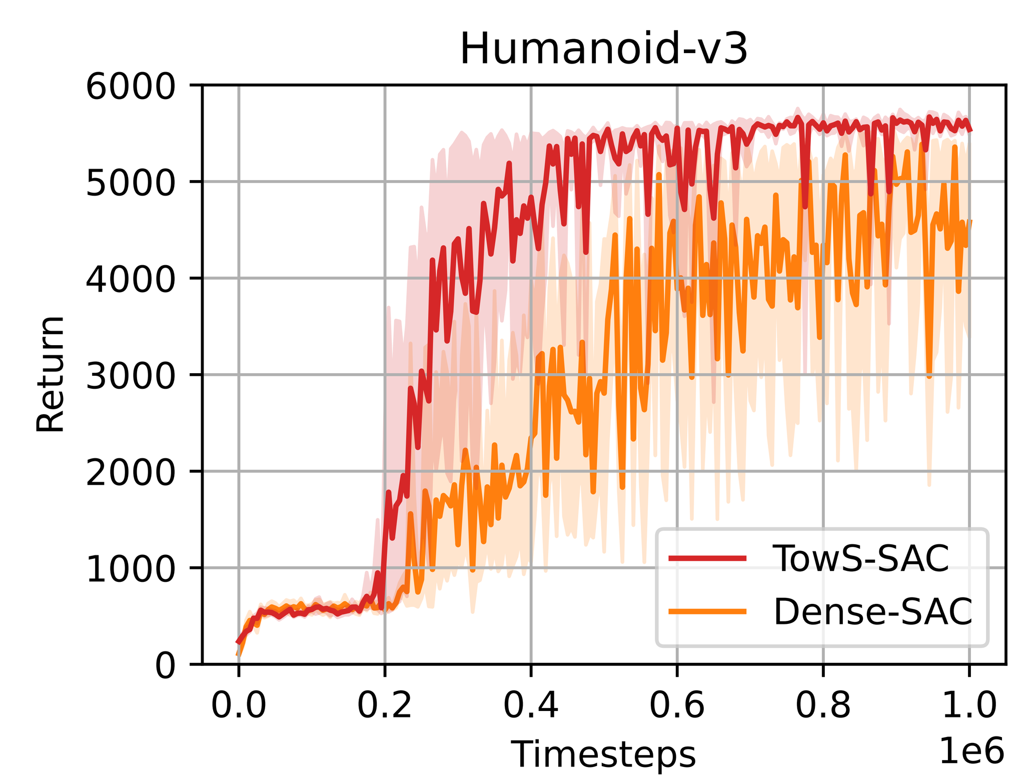

Figure 1: TowS-SAC (90% sparsity) outperforms its dense counterpart on the difficult Humanoid environment.

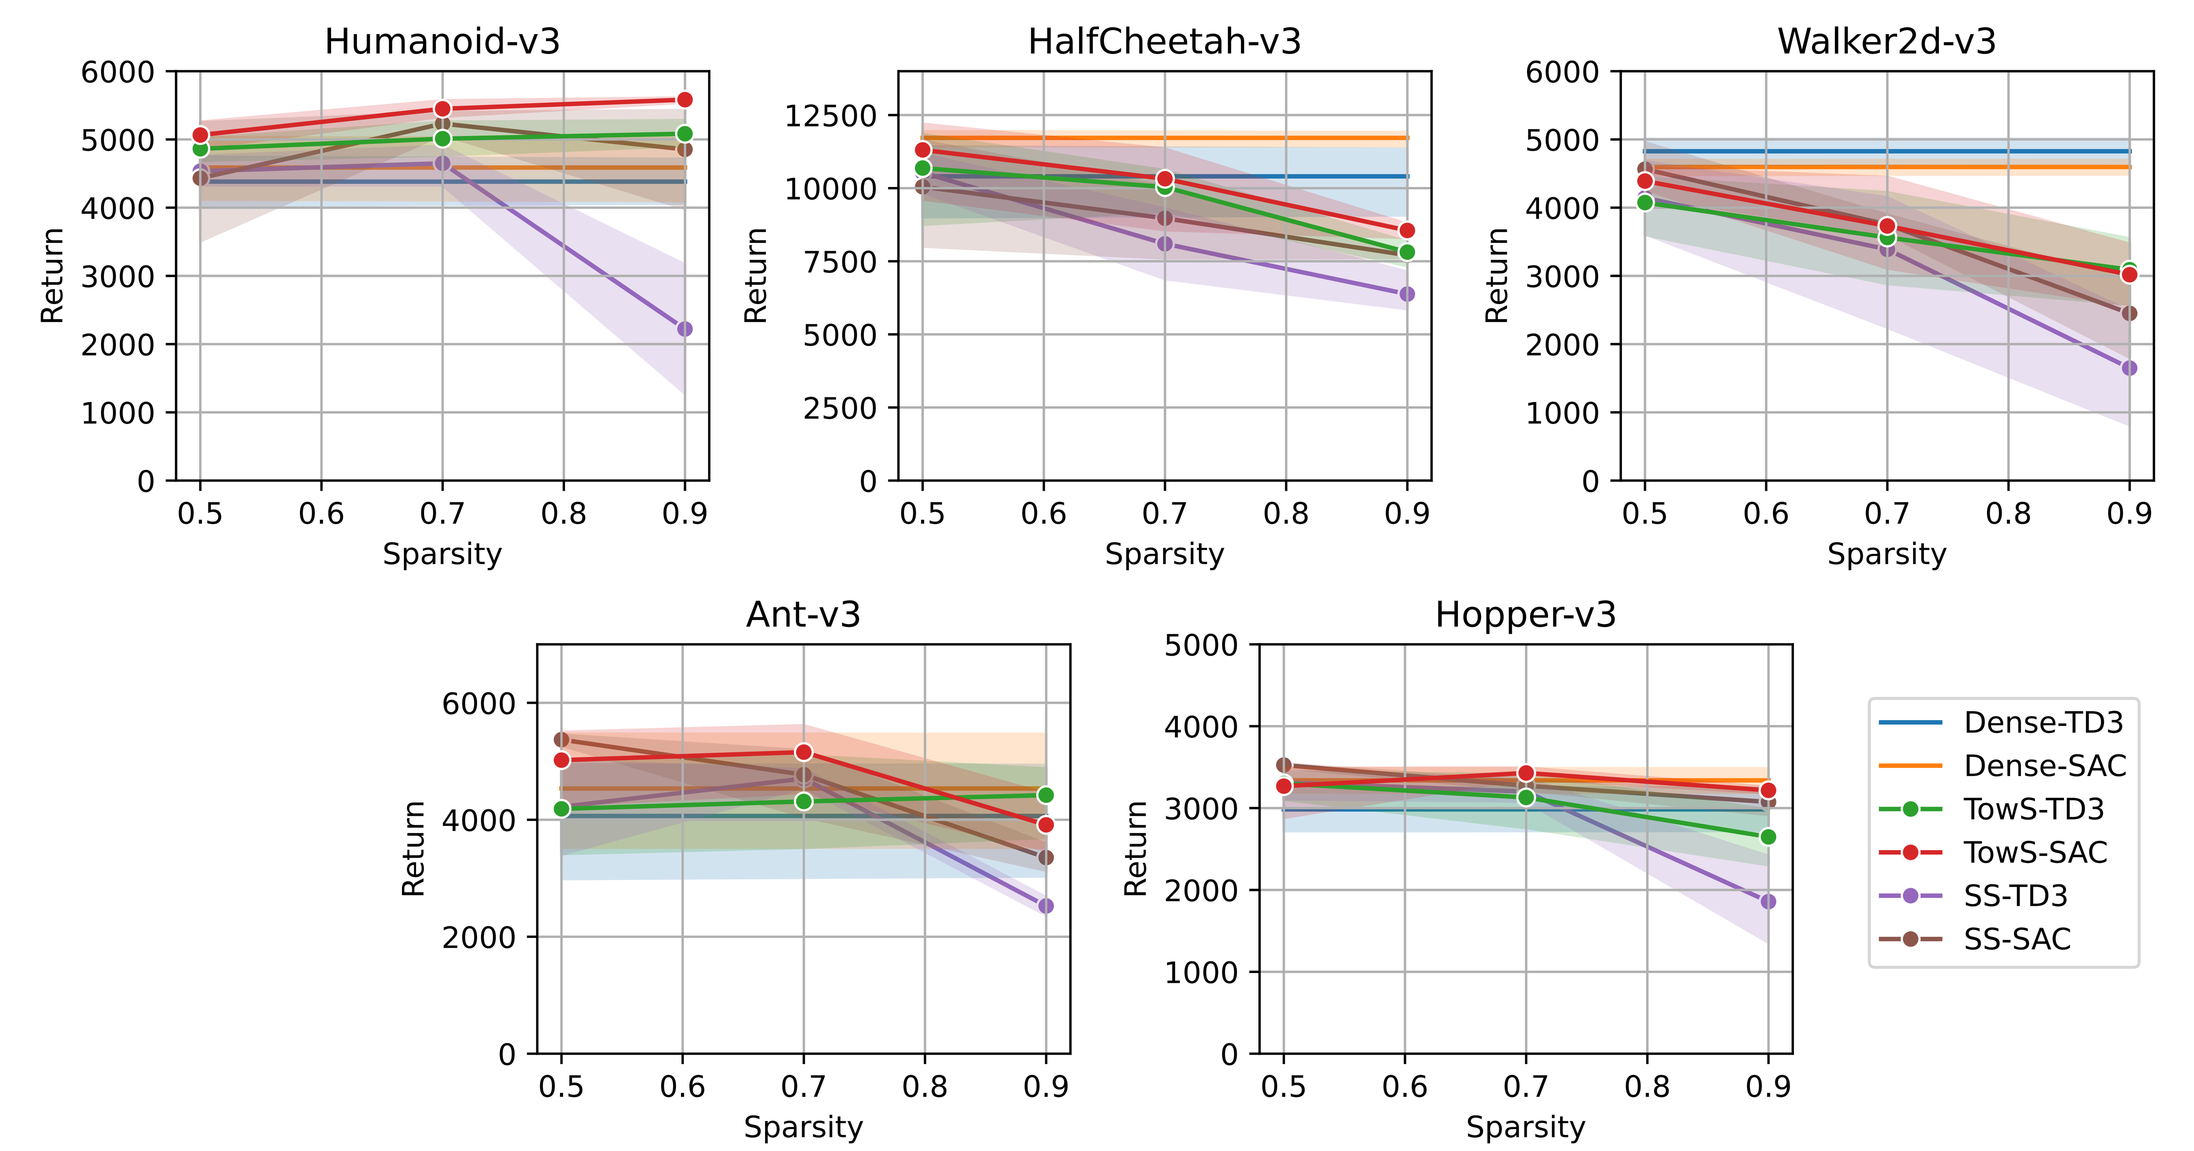

Figure 2: Towards truly dynamic sparse (TowS) and static sparse (SS) methods vs their dense counterparts (horizontal lines). On Humanoid, Ant, Hopper we reach higher sparsity levels than current state of the art, while increasing performance.

Main Result

We increase the sparsity level of the current state-of-the-art in online reinforcement learning by a large margin, while proposing a sparse training framework suitable for truly sparse implementations [1].

Policy	Sparsity	Return (↑)					FLOPs (↓)	Params. (↓)
		Humanoid	HalfCheetah	Walker2d	Ant	Hopper		
TowS-SAC	90%	5582.7	8553.8	3017.3	3913.2	3216.4	0.1×	35K
TowS-TD3	90%	5083.2	7813.7	3091.3	4421.2	2648.0	0.1×	34K
SS-SAC	90%	4854.2	7706.4	2450.3	3353.6	3075.0	0.1×	35K
SS-TD3	90%	2221.2	6378.3	1648.4	2526.0	1857.6	0.1×	34K
TowS-SAC	70%	5449.0	10321.7	3731.9	5156.6	3428.9	0.3×	102K
TowS-TD3	70%	5011.0	10037.0	3562.8	4313.3	3129.3	0.3×	101K
SS-SAC	70%	5231.5	8971.4	3753.7	4770.5	3272.7	0.3×	102K
SS-TD3	70%	4650.8	8098.2	3389.0	4704.7	3200.4	0.3×	101K
Dense-SAC	-	4587.8	11719.6	4596.2	4535.3	3338.5	1.01×	338K
Dense-TD3	-	4381.6	10404.2	4826.6	4064.9	2984.2	1×	334K

LeakyReLU vs. ReLU

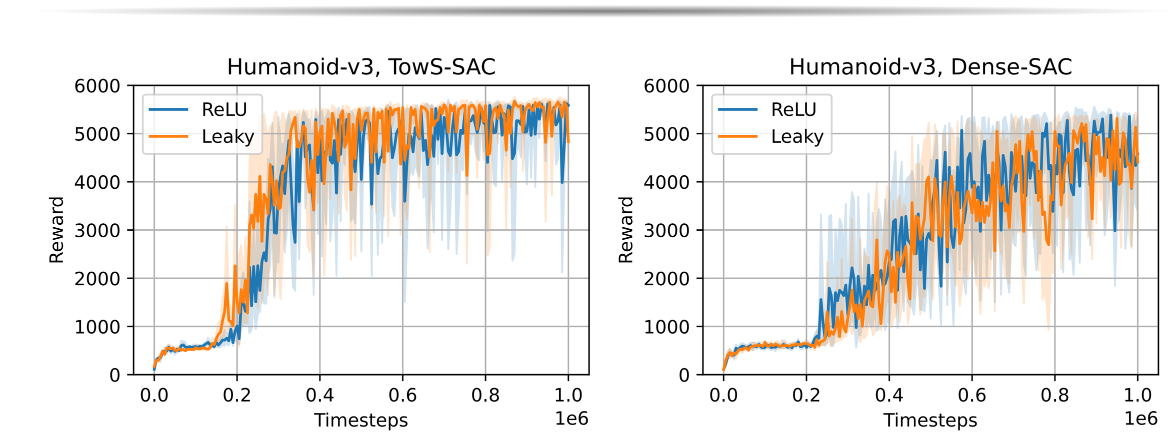

References

[1] Selima Curci et al. Truly Sparse Neural Networks at Scale. *arXiv:2102.01732*, 2021.

[2] Jonas Degrave et al. Magnetic control of tokamak plasmas through deep reinforcement learning. *Nature*, 2022.

[3] Decebal Constantin Mocanu et al. Scalable Training of Artificial Neural Networks with Adaptive Sparse Connectivity inspired by Network Science. *arXiv:1707.04780, Nature Communications*, 2018.

[4] Ghada Sokar et al. Dynamic Sparse Training for Deep Reinforcement Learning. *arXiv:2106.04217, IJCAI*, 2022.

Contact Information

- Web: www.bramgrooten.nl
- Email: b.j.grooten@tue.nl

TU/e UNIVERSITY OF ALBERTA UNIVERSITY OF TWENTE.