

REPRODUCIBILITY STATEMENT

- We utilized open-source implementations of the selected DRL algorithms from public repositories (not including our technique):
 - AMIGO was from here: <https://github.com/facebookresearch/adversarially-motivated-intrinsic-goals>
 - BT is here: https://github.com/andreneco/minigrid_bt
 - DQN, A2C, TRPO, and PPO were established on Stable Baselines3 (SB3)'s baselines repository (Raffin et al., 2021): <https://stable-baselines3.readthedocs.io/>
 - DreamerV3 was from here: <https://github.com/qxcv/dreamerv3>
 - All the other were from here: <https://github.com/sparisi/cbet>
- We used the the Minigrid package for the environments in our comparison, which is available here: <https://github.com/Farama-Foundation/Minigrid>
- For NACE we provide a stand-alone zip archive for reviewers to reproduce our results, which is runnable on a regular computer with Python interpreter. It includes a README.txt in the NACE folder, as well as scripts to generate the tables and the plots present in the paper.

REFERENCES

Antonin Raffin, Ashley Hill, Adam Gleave, Anssi Kanervisto, Maximilian Ernestus, and Noah Dormann. Stable-baselines3: Reliable reinforcement learning implementations. *Journal of Machine Learning Research*, 22(268):1–8, 2021. URL <http://jmlr.org/papers/v22/20-1364.html>.