Michał Bortkiewicz [1], Władysław Pałucki [2], Vivek Myers [3], Tadeusz Dziarmaga [4], Tomasz Arczewski [4], Łukasz Kuciński [2,5,6], Benjamin Eysenbach [7]
michalbortkiewicz8@gmail.com, wladek.palucki@gmail.com
[1] Warsaw University of Technology [2] University of Warsaw [3] UC Berkeley [4] Jagiellonian University
[5] Polish Academy of Sciences [6] Ideas NCBR [7] Princeton University

# Accelerating Goal-Conditioned RL Algorithms and Research

ICLR

Paper and Code

Reach out!

# JaxGCRL, or: How to run a meaningful RL research on a single GPU

## Fast Training

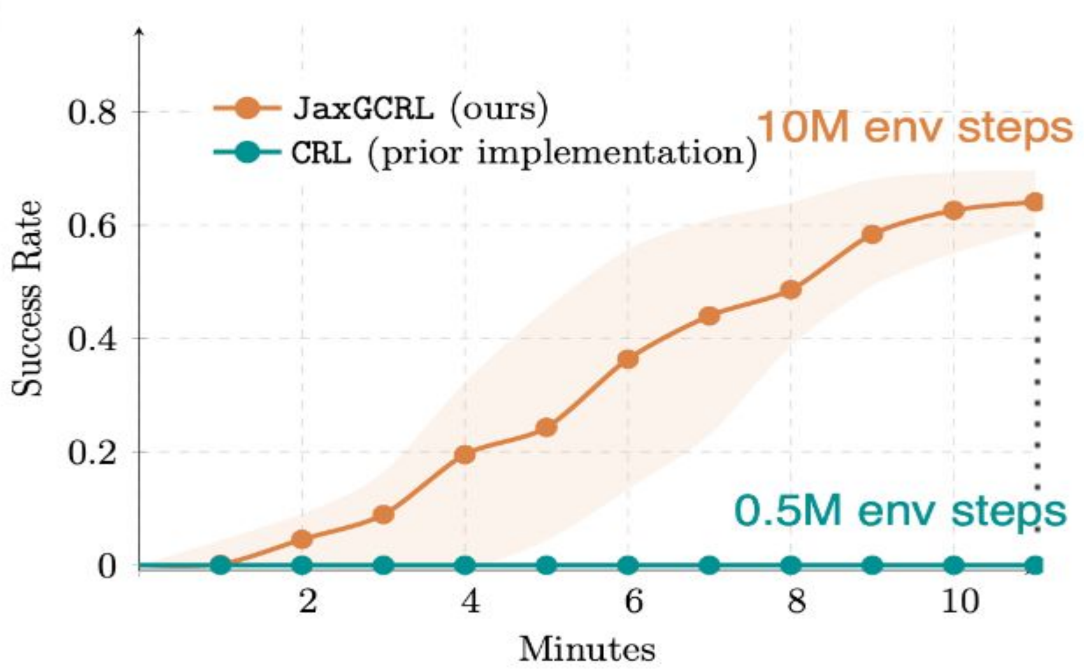


> 20 times faster training

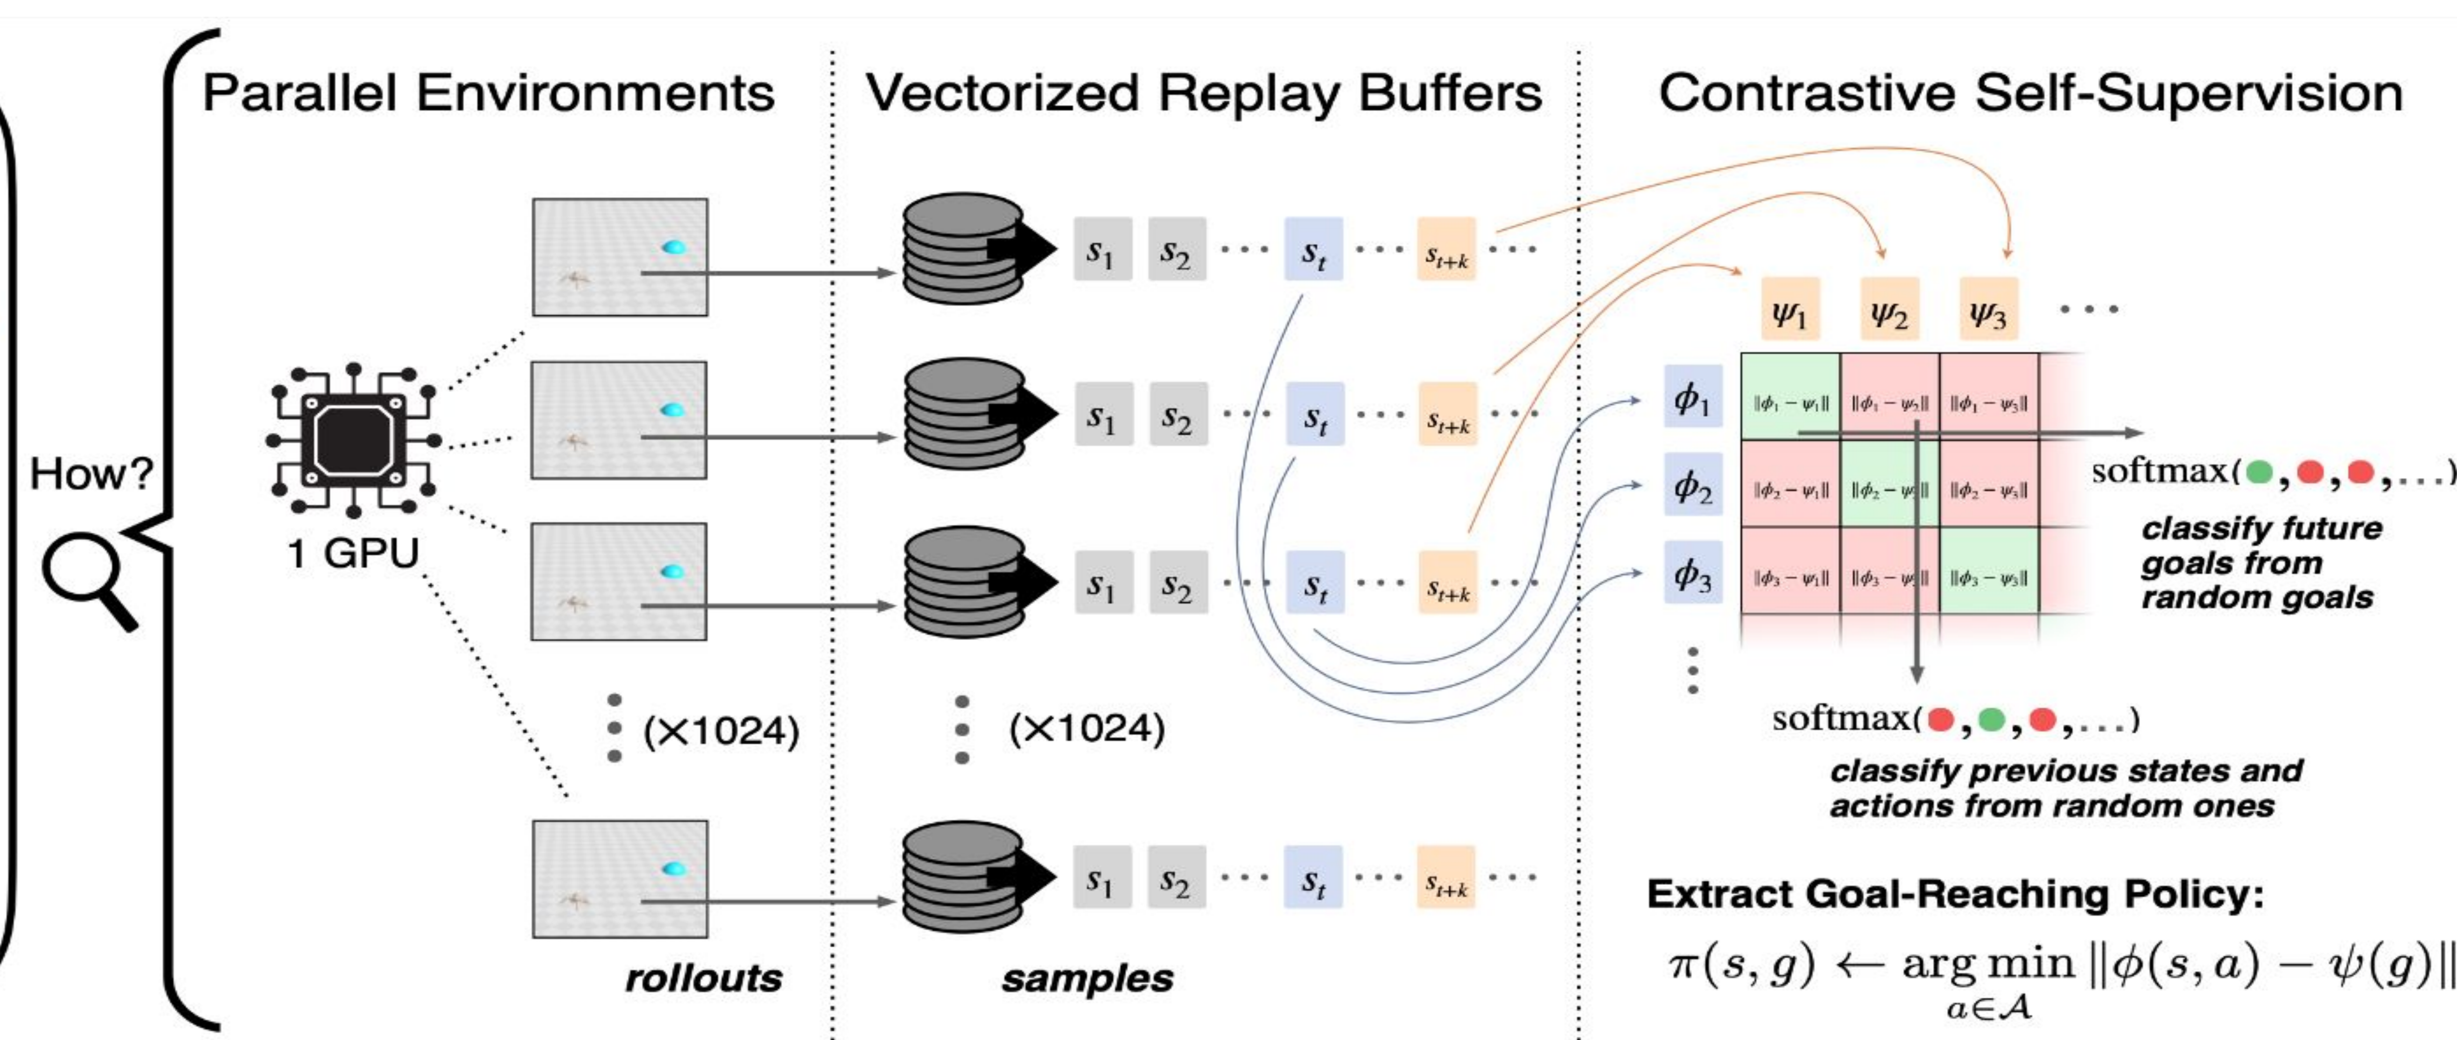


## Contributions

**Speed**: **JaxGCRL** runs 10M-step experiments in minutes and billion-step ones in hours on a single GPU.

**Codebase:** a blazingly fast JIT-compiled training pipeline for GCRL experiments.

**Benchmark:** a suite of >15 GPU-accelerated state-based environments that help to accurately assess GCRL algorithm capabilities and limitations.

**Extensive empirical analysis** of Contrastive RL design choices.

## Results

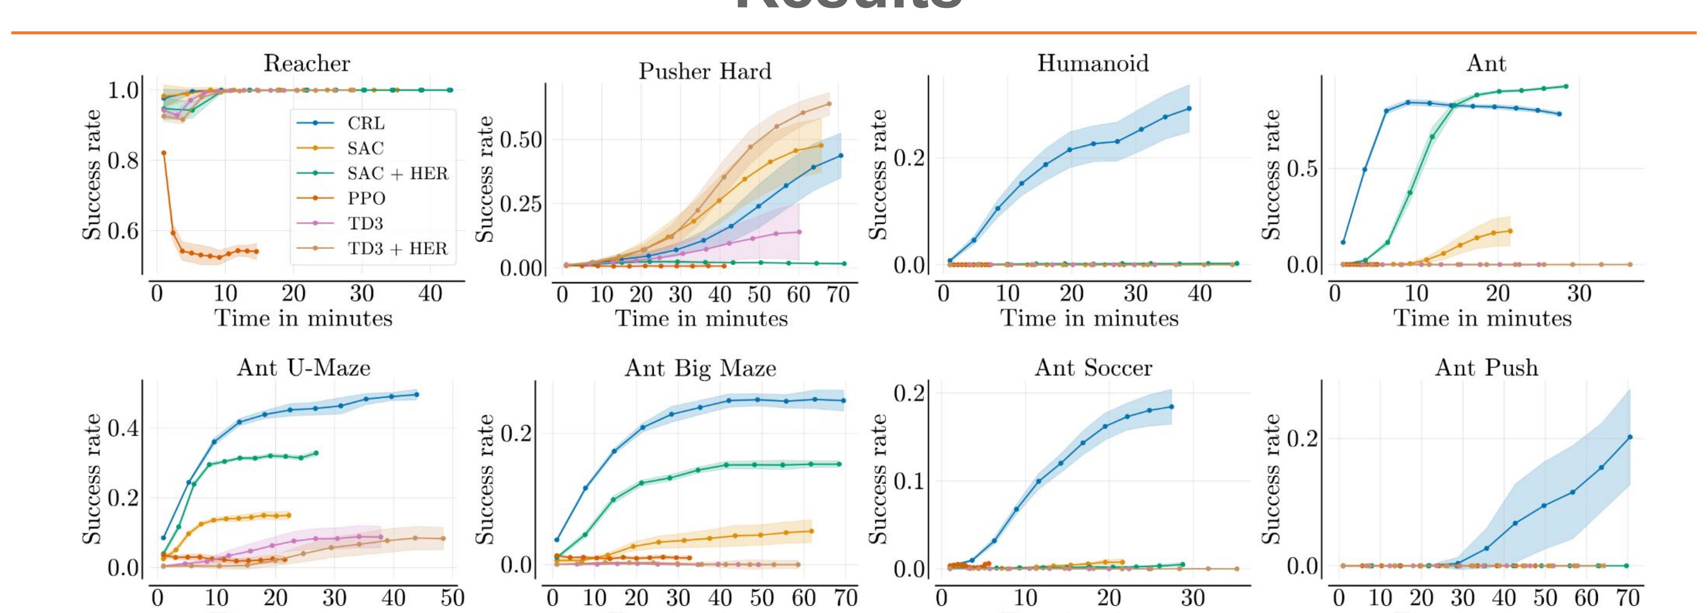


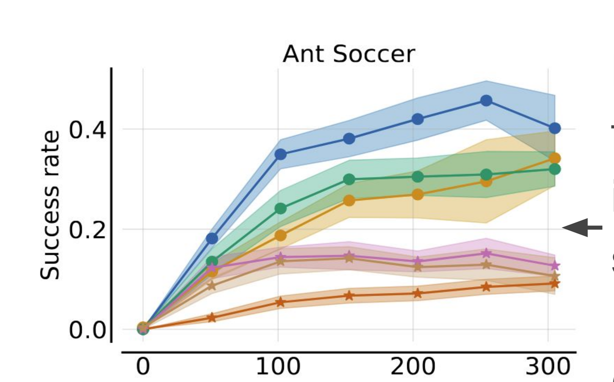


Different combinations of energy and contrastive functions affect learning in distinct ways: some improve success rates, while others increase the time spent near goals.

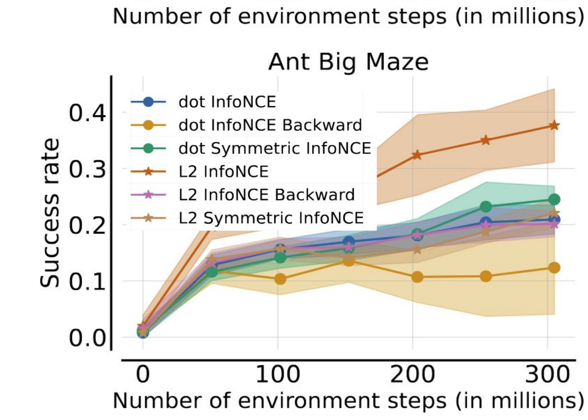


CRL thrives with large architectures and LayerNorm.

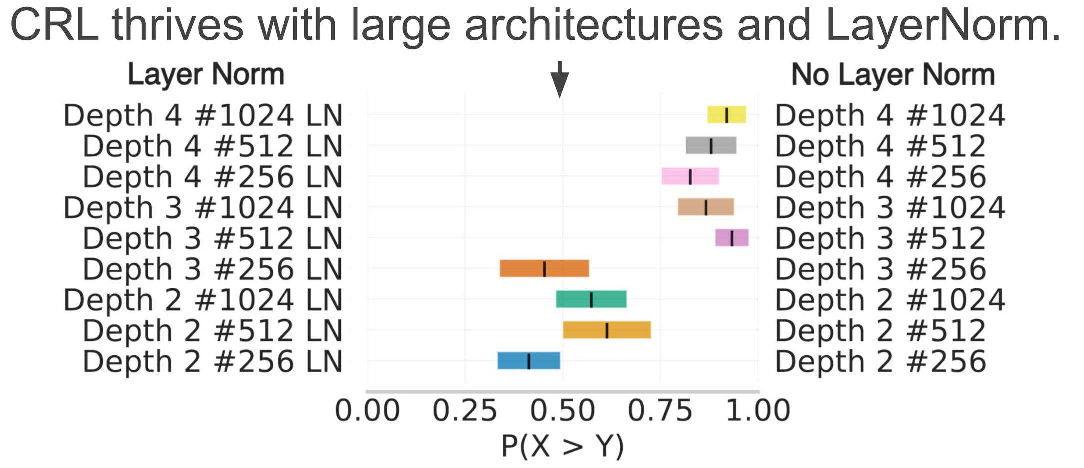


CRL is the only method that consistently learns across all sparse-reward environments, even when using a low update-to-data ratio. However, with relatively shallow architectures, performance tends to plateau early in data-rich settings. This suggests that model size might be a critical bottleneck for achieving high-performing agents.

```
> pip install jaxgcrl
```

## Environments and Features

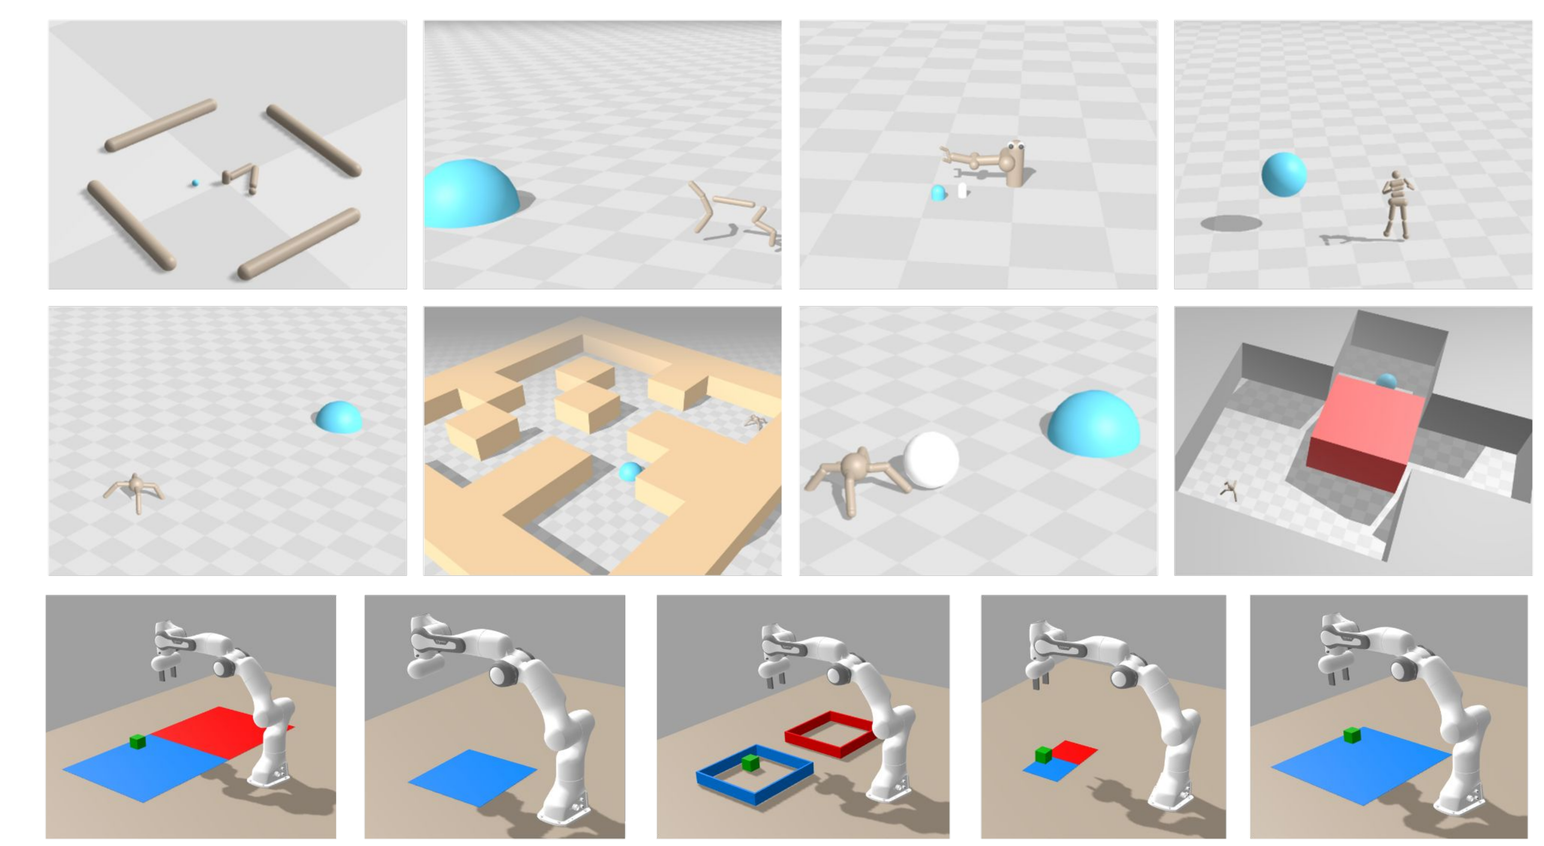

✅ 15 environments
✅ Easy installation and extensive documentation
✅ 5 baselines
✅ wandb integration and rendering

## Contributing and open challenges

We invite you to collaborate on number of exciting challenges:

- Integration of discrete environments and new arm manipulation tasks.
- Inverse RL and preference learning pipelines based on **JaxGCRL**.
- Autonomous goal sampling and exploration strategies.

Research conducted so far with the help of **JaxGCRL**:

- V. Myers and C. Ji et al. **(ICLR 2025)** Horizon Generalization in Reinforcement Learning
- Y. Jiang and Q. Liu et al. **(ICLR 2025)** Episodic Novelty Through Temporal Distance
- K. Wang et al. **(preprint)** 1000 Layer Networks for Self-Supervised RL: Scaling Depth Can Enable New Goal-Reaching Capabilities

And more!