

# rewardGym

## A framework for streamlining experiments in cognitive neuroscience

Simon R. Steinkamp<sup>1</sup>, David Meder<sup>1</sup>, & Oliver J. Hulme<sup>1,2,3</sup>

Mail-to: [simons@drcmr.dk](mailto:simons@drcmr.dk)

<sup>1</sup>Danish Research Centre for Magnetic Resonance, Copenhagen University Hospital - Amager and Hvidovre, Copenhagen, Denmark; <sup>2</sup>London Mathematical Laboratory, London, United Kingdom; <sup>3</sup>Department of Psychology, University of Copenhagen, Copenhagen, Denmark

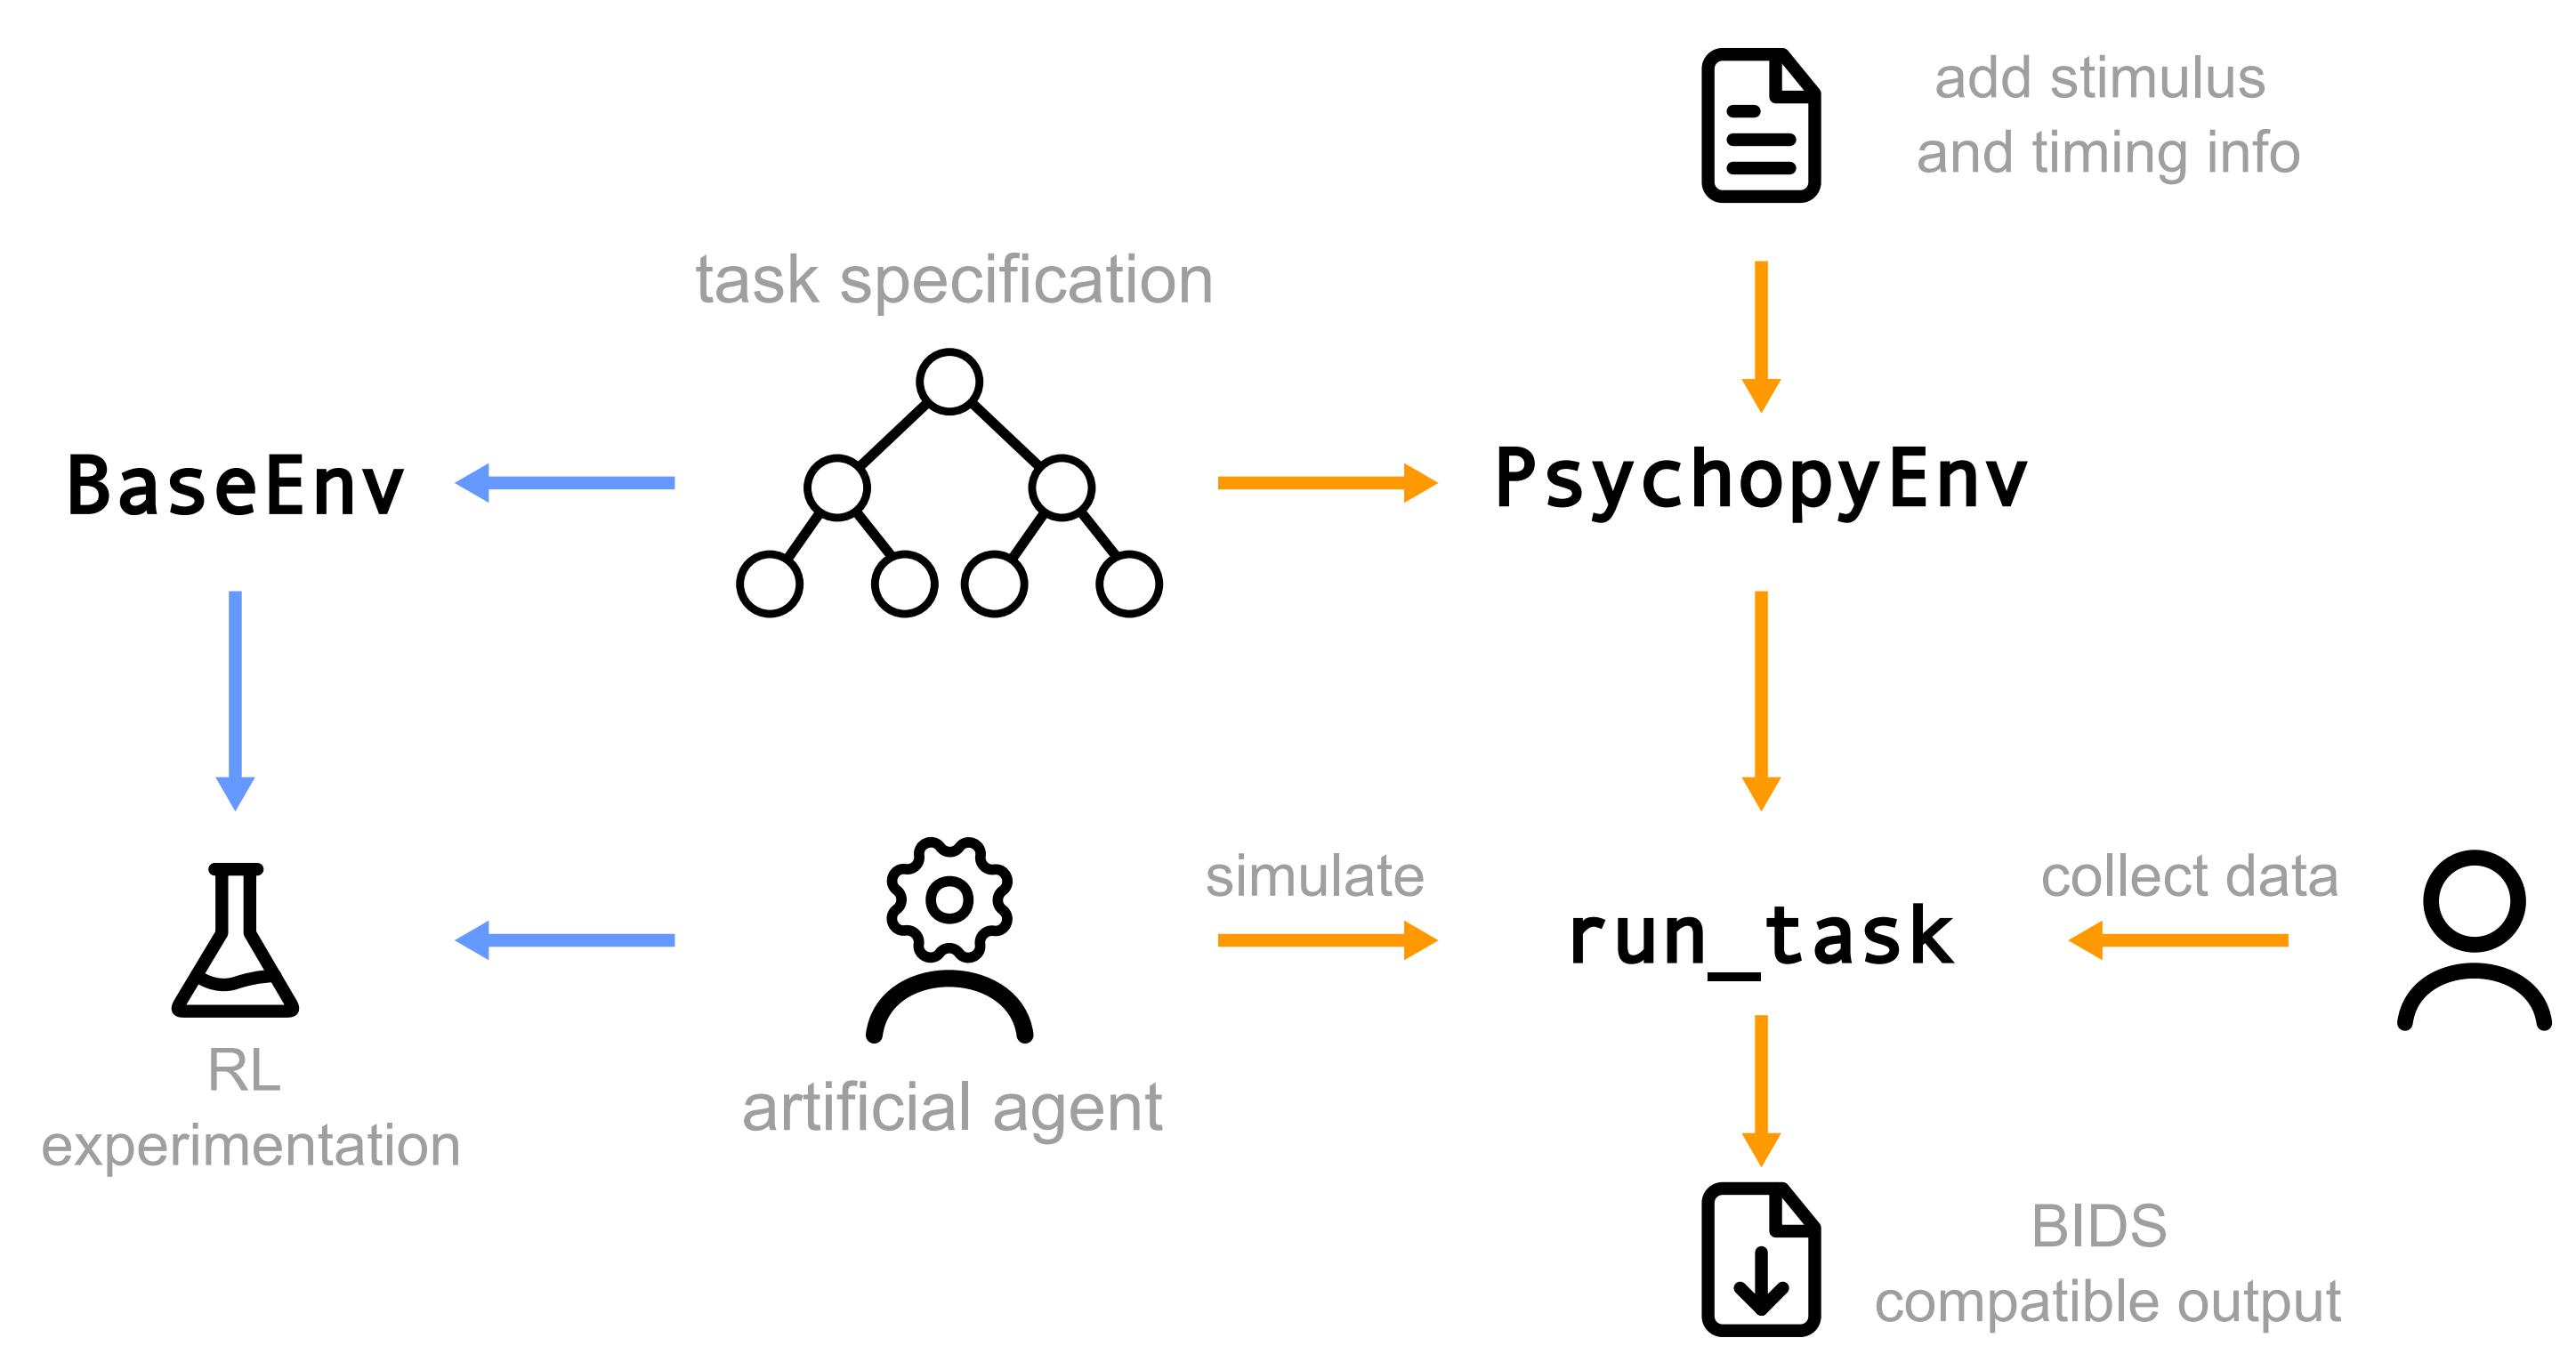


### Scope

rewardGym is being developed as part of the rewardMap project, testing multiple reward paradigms in the same subjects to reveal common signatures of reward processing. It is aimed primarily at cognitive neuroscientists and psychologists, who study learning and decision-making in human participants.

- ▶ Streamline cognitive experiments by using a common graphical language for each trial.
- ▶ Use the same logic for artificial and biological agents to interact with the experiment.
- ▶ Build on standards from reinforcement learning using the gymnasium API (Towers et al., 2024) and is further inspired by neuro-nav (Juliani et al., 2022).
- ▶ Common API allows for reuse of agents for both simulation and inference across many tasks.
- ▶ We address experimentation needs: high control over the environment, running with standalone PsychoPy (Peirce et al., 2019).

### Overview



**Overview of the framework.** At its core, is the task specification using a graph structure. Using this structure, the user can do classical reinforcement learning experiments (left) using the BaseEnv class. By augmenting the basic graph with stimulus information, the PsychopyEnv can collect data from human participants and deploy artificial agents to simulate data. This can also be done using the convenience function run\_task, which will store simulated and real data in BIDS format.

#### Task Abstraction

two-step task

risk-sensitive task

#### Experiment Control

Graph representation: two-step task

reward feedback

first decision stage

second decision stage

- ▶ stimulus timing
- ▶ randomization
- ▶ staircasing
- ▶ ...

#### Reinforcement Learning Experiments

```
env = gym.make('two-step')
n_episodes = 100
agent = RLAgent(action_space=env.action_space, state_space=env.observation_space, learning_rate=0.5, temperature=4.0, )

for t in range(n_episodes):
    obs, info = env.reset(agent_location=0)
    done = False
    while not done:
        action = agent.get_action(obs)
        next_obs, reward, terminated, _, info = env.step(action)
        agent.update(action=action, obs=obs, next_obs=next_obs, reward=reward, terminated=terminated)
    done = terminated
```

#### Hybrid: Simulate Realistic Files

- ▶ Prepare visualization
- ▶ behavioral analyses
- ▶ check experiment
- ▶ ...

#### Data Collection with PsychoPy

Build for standalone - no additional installation needed

**rewardGym.** The common task abstraction, allows integrating implemented tasks into classic reinforcement learning (RL) workflows, for experimentation and benchmarking.

**Features:** Adding experiment details (stimuli, timings, randomization), to simulate realistic data using agents, allowing to implement analysis pipelines and debugging before data collection. Data collection supports fMRI. Tasks implemented: MID, HCP, risk-sensitive, two-step, Go-NoGo, Posner. Common task set-up allows for rapid prototyping and automatic tests.

**Limitations:** Early development. Toolbox does not prioritize reinforcement-learning studies (deep RL, etc.). Part of a larger project: hard coded intricacies, but working on modularization.

#### Behavioral Simulation

Risk-sensitive: Taking safe options

agent

Two-step: model-free

Two-step: model-based

#### Model & Parameter Recovery

Risk-sensitive Model recovery

Recovered Model	original	valence
risk-ambivalent	1	0.033
risk-neutral	0	0.97
risk-seeking	0	0.1
random	0	1

Generative model

Risk-sensitive learning rates

Two-step Model recovery

Recovered Model	original	valence
model-free	1	0
model-based	0	1
random	0	1

Two-step: weighting parameter

**Study workflow.** Using a hybrid, valence-based agent (Niv et al., 2012; Gläscher et al., 2010) to simulate data and perform model and parameter recovery on the two-step (Daw et al., 2011) and a risk-sensitive decision making task (Rosenbaum et al., 2022). The agent's code can be reused for simulation and inference - the simulation study is available on binder. Code developed for the simulation is then applied to real data of a pilot study ( $n = 15$ ).

<https://mybinder.org/v2/gh/rewardMap/exampleWorkflow/binder?urlpath=%2Fdoc%2Ftree%2FCCNRLDMSims.ipynb>

### rewardGym on Github



<https://github.com/rewardMap/rewardGym>

### Future

- ▶ Better and more documentation and examples.
- ▶ Alternative data collection interfaces: pygame.
- ▶ rewardCoach - inference package for rewardGym
- ▶ Hopefully larger adaption of standardized modeling in psychology / cognitive science

Looking forward to discussions, ideas, and contributions!

### Example on Binder



### Acknowledgements

SRS was funded by a Lundbeck Foundation postdoctoral grant (ref R402-2022-1411). DM was funded by a Sapere Aude DFF-starting grant from the Independent Research Fund Denmark (Grant No. 1052-00054B). OJH was funded by a Novo Nordisk Foundation Exploratory Interdisciplinary Synergy Grant (ref NNF200C0064869).

### References

Mark Towers, Ariel Kwiatkowski, Jordan Terry, John U. Balis, Gianluca De Cola, Tristan Deleu, Manuel Goulão, Andreas Kallinteris, Markus Krimmel, Arjun KG, Rodrigo Perez-Vicente, Andrea Pierré, Sander Schulhoff, Jun Jet Tai, Hannah Tan, and Omar G. Younis. Gymnasium: A Standard Interface for Reinforcement Learning Environments, November 2024.

Arthur Juliani, Samuel Barnett, Brandon Davis, Margaret Sereno, and Ida Momennejad. NeuroNav: A Library for Neurally-Plausible Reinforcement Learning, June 2022.

Jonathan Peirce, Jeremy R. Gray, Sol Simpson, Michael MacAskill, Richard Höchenberger, Hiroyuki Sogo, Erik Kastman, and Jonas Kristoffer Lindeløv. PsychoPy2: Experiments in behavior made easy. *Behavior Research Methods*, 51(1):195–203, February 2019. ISSN 1554-3528. doi: 10.3758/s13428-018-01193-y.

Yael Niv, Jeffrey A. Edlund, Peter Dayan, and John P. O'Doherty. Neural Prediction Errors Reveal a Risk-Sensitive Reinforcement-Learning Process in the Human Brain. *Journal of Neuroscience*, 32(2):551–562, January 2012. ISSN 0270-6474, 1529-2401. doi: 10.1523/JNEUROSCI.5498-10.2012.

Jan Gläscher, Nathaniel Daw, Peter Dayan, and John P. O'Doherty. States versus Rewards: Dissociable Neural Prediction Error Signals Underlying Model-Based and Model-Free Reinforcement Learning. *Neuron*, 66(4):585–595, May 2010. ISSN 0896-6273. doi: 10.1016/j.neuron.2010.04.016.

Nathaniel D. Daw, Samuel J. Gershman, Ben Seymour, Peter Dayan, and Raymond J. Dolan. Model-Based Influences on Humans' Choices and Striatal Prediction Errors. *Neuron*, 69(6):1204–1215, March 2011. ISSN 0896-6273. doi: 10.1016/j.neuron.2011.02.027.

Gail M. Rosenbaum, Hannah L. Grassie, and Catherine A. Hartley. Valence biases in reinforcement learning shift across adolescence and modulate subsequent memory. *eLife*, 11:e64620, January 2022. ISSN 2050-084X. doi: 10.7554/eLife.64620.