

The Cell Must Go On: Agar.io for Continual Reinforcement Learning

Mohamed A. Mohamed, Kateryna Nekhomiaz, Vedant Vyas,
Marcos M. José, Andrew Patterson, Marlos C. Machado

Keywords: Continual Reinforcement Learning, Evaluation Platform, Benchmark, Simulator, Non-Stationarity, Partial Observability, Game, Agar.io, AgarCL.

Summary

Continual reinforcement learning (RL) concerns agents that are expected to learn continually, rather than converge to a policy that is then fixed for evaluation. This setting is well-suited to environments that the agent perceives as *changing* over time, rendering any static policy ineffective. In continual RL, researchers often simulate such changes either by modifying episodic environments to incorporate task shifts during interaction or by designing simulators that explicitly model continual dynamics. However, transforming episodic problems into continual ones primarily captures scenarios involving abrupt changes in the data stream and still relies on episodic structure. Meanwhile, the few simulators explicitly designed for empirical continual RL research are often limited in scope or complexity. In this paper, we introduce *AgarCL*, a research platform for continual RL that enables agents to progress toward increasingly sophisticated behaviour. *AgarCL* is based on the game *Agar.io*, a non-episodic, high-dimensional problem with stochastic, ever-evolving dynamics, continuous actions, and partial observability. The contribution(s) box below describes our main empirical contributions in more detail.

Contribution(s)

1. We introduce *AgarCL*, a non-episodic, high-dimensional reinforcement learning (RL) environment based on the game *Agar.io* with non-stationary dynamics and a hybrid action space.
Context: *GoBIGGER* (Zhang et al., 2023) is the only existing *Agar.io*-based platform, but it was designed for *episodic* multi-agent RL. Earlier work explored *Agar*-style games in much simpler settings (Ansó et al., 2019; Wiehe et al., 2018). *AgarCL* is much faster than alternative platforms, has different observation streams, and is well-suited to continual RL.
2. We introduce a suite of diagnostic mini-games in *AgarCL* that isolate specific challenges.
Context: Mini-games allow us to focus on specific challenges posed by *AgarCL* and to benchmark agents with respect to the particular skill a mini-game requires. They enable faster scientific iteration and may be viewed as intermediate steps toward an effective agent.
3. In one setting where agents are able to learn to maximize their return, we show that the performance of a deployed fixed policy degrades over time in *AgarCL*.
Context: The policies learned by PPO, our best-performing agent, that were checkpointed at 32M and 48M steps, collapse over time when deployed in *AgarCL*. We see such a collapse in more challenging settings with more non-stationarity, not in the easier settings.
4. We evaluate DQN, PPO, and SAC on the full-game setting and across all the mini-games. We also evaluated PPO with GRUs, Continual Backprop, ReDo, and Shrink & Perturb, and we observed that they were also unable to achieve strong performance in *AgarCL*.
Context: We considered multiple settings, and the different combinations were evaluated in specific settings. We generally report the performance over 10 independent runs after hyperparameter tuning (when appropriate). These results are intended to characterize the difficulty of the environment rather than to establish state-of-the-art algorithmic performance. Results for the continual learning baselines were obtained with standard hyperparameters.

The Cell Must Go On: Agar.io for Continual Reinforcement Learning

Mohamed A. Mohamed^{1,2}, Kateryna Nekhomiaz³, Vedant Vyas¹,
 Marcos M. José^{1,2}, Andrew Patterson⁴, Marlos C. Machado^{1,2,5}

¹University of Alberta
³NVIDIA

²Alberta Machine Intelligence Institute (Amii)
⁴RLCore ⁵Canada CIFAR AI Chair

Abstract

Continual reinforcement learning (RL) concerns agents that are expected to learn continually, rather than converge to a policy that is then fixed for evaluation. This setting is well-suited to environments that the agent perceives as *changing* over time, rendering any static policy ineffective. In continual RL, researchers often simulate such changes either by modifying episodic environments to incorporate task shifts during interaction or by designing simulators that explicitly model continual dynamics. However, transforming episodic problems into continual ones primarily captures scenarios involving abrupt changes in the data stream and still relies on episodic structure. Meanwhile, the few simulators explicitly designed for empirical continual RL research are often limited in scope or complexity. In this paper, we introduce *AgarCL*, a research platform for continual RL that enables agents to progress toward increasingly sophisticated behaviour. *AgarCL* is based on the game *Agar.io*, a non-episodic, high-dimensional problem with stochastic, ever-evolving dynamics, continuous actions, and partial observability. We provide benchmark results for DQN, PPO, and SAC on the primary continual RL challenge, as well as across a suite of smaller tasks within *AgarCL*. These smaller tasks isolate aspects of the full environment and allow us to characterize the distinct challenges posed by different components of the game. We further evaluate three continual learning methods—Shrink and Perturb, ReDo, and Continual Backpropagation—and observe they are also unable to achieve strong performance, suggesting that the challenges posed by *AgarCL* extend beyond the stability–plasticity dilemma.

1 Introduction

Continual reinforcement learning (RL) is the setting in which the agent is expected to learn continually, rather than converge to a policy that is subsequently fixed for evaluation or deployment. It can be viewed either as a problem formulation (Khetarpal et al., 2022; Abel et al., 2023; Kumar et al., 2025) or as a solution method for problems perceived by the agent as non-stationary. Such scenarios are often motivated by the big world hypothesis, which posits that the world is “bigger” than the agent (Javed & Sutton, 2024; Lewandowski et al., 2025). In this regime, continual adaptation is more effective than committing to a static policy (Sutton et al., 2007; Janjua et al., 2024).

Much of the progress in RL has been driven by empirical advances, with experimental platforms often shaping algorithmic innovation and research directions. Consequently, evaluation platforms play a central role in determining what the field studies and how progress is measured (e.g., Todorov et al., 2012; Bellemare et al., 2013; Beattie et al., 2016). In continual RL, most evaluation platforms are adaptations of traditional RL benchmarks, modified to introduce non-stationarity and thereby reflect the idea that the world is bigger than the agent. This is typically done by artificially switching tasks during interaction (Powers et al., 2022; Abbas et al., 2023; Anand & Precup, 2023; Tomilin




Figure 1: **Agent-environment interface and main entities in AgarCL.** The agent receives either pixel-based or symbolic observations. The pixel-based input consists of four channels representing pellets, bots, viruses, and the agent itself, while the symbolic alternative provides pre-processed features such as distances to nearby entities. The reward is defined as the change in the agent’s mass between consecutive time steps. The action space is hybrid: at each step, the agent selects a continuous $\langle x, y \rangle$ coordinate (analogous to mouse control) and chooses whether to move, split, or eject mass. Section 3 provides further details on the entities, dynamics, and action definitions.

et al., 2023), or, less commonly, by designing new environments specifically for continual RL that abandon episodic structure altogether (Platanios et al., 2020).

However, both approaches have limitations. Reusing existing environments—particularly complex simulators—is appealing, yet introducing non-stationarity via hand-crafted task switches that are disconnected from the agent’s behaviour primarily models exogenous, abrupt regime shifts. Although such shifts occur in practice, many real-world environments evolve gradually and continuously through interaction. This endogenous, interaction-driven non-stationarity remains comparatively underexplored in continual RL benchmarks. Designing environments specifically for continual RL is promising, but existing instantiations often remain limited in complexity or scope. Our work takes a complementary direction: we retain the non-episodic structure of purpose-built environments while introducing rich, evolving stochastic dynamics, continuous control, partial observability, resource-driven competition, and progressively increasing behavioural demands within a spatially structured world.

Concretely, we introduce a new evaluation platform based on the game Agar.io (Figure 1), which we call AgarCL. The platform was designed explicitly to support continual RL research. In this environment, the agent controls circular cells within a bounded, Petri dish-like arena and must accumulate mass in an ever-changing ecosystem of competing entities. Observations are high-dimensional and pixel-based,¹ and the action space combines continuous navigation—parameterized as $\langle x, y \rangle$ movement—with discrete affordances such as splitting or ejecting mass. The reward is defined as the change in the agent’s mass between consecutive time steps, so effective behaviour corresponds to sustained growth over an indefinite horizon.

The dynamics of Agar.io make it particularly well suited to studying continual RL. The agent operates in an ecosystem populated by other agents,² competing for mass in a bounded arena. Agents

¹Alternative symbolic observations, providing backward compatibility with related platforms (Zhang et al., 2023), are also supported; see Section 3.

²In AgarCL, the behaviour of other agents is governed by hand-coded policies. Extending the platform to support learning across multiple agents is a promising direction for future work.

grow by collecting stationary food pellets or absorbing smaller cells, and to do so, they have to simultaneously avoid larger threats. At every time step, mass decays at a rate that increases with size, creating sustained pressure to act. Crucially, beyond its potential never-ending nature and the complexity induced by other agents’ behaviours, the environment’s dynamics change as a function of the agent’s mass. Larger agents move more slowly, and the observation space progressively zooms out to keep the agent’s body visible, altering the scale and distribution of perceptual input. As a result, both the consequences of actions and the observation stream evolve continuously with the agent’s state. Although abrupt events also occur—such as absorbing another cell or splitting due to strategic choice or viral interaction—the dominant form of non-stationarity is smooth and interaction-driven. Figure 2 and a YouTube video³ depict some of the environment dynamics supported by AgarCL.

In this paper, we introduce AgarCL as both an evaluation platform and a challenge problem for continual RL. We evaluate three widely used RL algorithms—DQN (Mnih et al., 2015), PPO (Schulman et al., 2017), and SAC (Haarnoja et al., 2018)—as well as three approaches explicitly designed for continual learning: Shrink and Perturb (Ash & Adams, 2020), ReDo (Sokar et al., 2023), and Continual Backpropagation (Dohare et al., 2024). Despite their strong performance in many established benchmarks, none of these methods achieves sustained competence in the full game. To better understand these failures, we introduce a suite of mini-games that isolate specific challenges embedded in the full environment, including non-stationarity, exploration, and long-horizon credit assignment. These controlled experiments reveal that the difficulties in AgarCL cannot be attributed solely to stability–plasticity trade-offs. Finally, our empirical study exposes persistent challenges in continual RL research, particularly regarding evaluation methodology and hyperparameter sensitivity, suggesting that progress in the field depends not only on new algorithms but also on more robust benchmarking practices.

2 Background

We adopt the standard RL formalism to describe the sequential decision-making problem posed by AgarCL. Interactions occur at discrete time steps. The agent begins in a state $S_0 \sim \mu$, where μ is a start-state distribution over $\mathcal{S}$, but only observes $\omega_0 \in \Omega$, generated by an observation function $\Phi : \mathcal{S} \rightarrow \Omega$. At each time step t , the agent selects an action $A_t \in \mathcal{A}$ according to a (possibly stochastic) policy π that depends on its observation history, and receives a reward $R_{t+1} \in \mathbb{R}$ and a new observation ω_{t+1} . The agent’s goal is to maximize some variant of the expected return, G_t , that is, of the expected total sum of rewards it receives. The most common variant is the *discounted* return, G_t^γ , in which later rewards are discounted by $0 \leq \gamma < 1$, such that $G_t^\gamma \doteq \sum_{k=0}^{\infty} \gamma^k R_{t+k+1}$. Alternatively, one may consider the *average-reward* setting: $G_t^r \doteq \sum_{k=0}^{\infty} (R_{t+k+1} - r(\pi))$, where $r(\pi) \doteq \lim_{h \rightarrow \infty} \frac{1}{h} \sum_{t=1}^h \mathbb{E}[R_t | \omega_0, A_{0:t-1} \sim \pi]$, with ω_0 denoting the initial observation.

The environment evolves according to a transition function $p : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{P}(\mathcal{S})$, where $\mathcal{P}(\mathcal{S})$ denotes the set of probability measures over the state space. In MDPs (Puterman, 2014), observations coincide with states. However, AgarCL is partially observable and more appropriately modelled as a POMDP (Kaelbling et al., 1998), with complex dynamics hidden from the agent. This partial observability can induce apparent non-stationarity from the agent’s perspective, even when the underlying transition dynamics are stationary. As discussed above, we adopt the view that continual RL is beneficial in such settings because continual adaptation may outperform any fixed policy (see Sutton et al., 2007; Janjua et al., 2024). This perspective allows us to remain agnostic to restrictive assumptions about agent capacity when invoking the “big world” hypothesis (Javed & Sutton, 2024).

In this work, we evaluate performance using the undiscounted return, while training algorithms derived from the discounted formulation. Although the average-reward objective is arguably more natural for continuing tasks, we leverage the substantially more mature deep RL literature developed for the discounted setting. This choice mirrors common practice. For example, in Atari 2600 games, agents are evaluated on accumulated scores while trained to optimize discounted return (Bellemare

³<https://www.youtube.com/watch?v=CGpvzHIqFLA>.




Figure 2: **Environment dynamics and actions.** ① The agent (black) consumes smaller cells to gain mass. ② The SPLIT action divides each controlled cell in half and propels one of the resulting cells in a chosen $\langle x, y \rangle$ direction, enabling slower, larger cells to reach faster opponents. Cell speed decreases as mass increases. ③ Separated cells can later merge when brought sufficiently close together. ④ Depending on its mass, a collision with a virus either fragments the agent into multiple cells or ⑤ allows the agent to absorb the virus. ⑥ The agent may also EJECT mass in a chosen direction. Ejected mass can be consumed by any cell and, if sufficient mass is fed into a virus, causes it to spawn a new virus while propelling the original, enabling smaller agents to attack larger ones.

et al., 2013; Machado et al., 2018). Moreover, for every problem, there exists a critical discount factor such that any solution using a higher value will maximize the average reward (Blackwell, 1962).

3 AgarCL: Agar.io for Continual Reinforcement Learning

We begin by discussing our main contribution: Agar.io for Continual RL (AgarCL), a research environment for sequential decision-making in which the dynamics evolve continuously as a function of the agent’s state and its interactions with other agents. Unlike benchmarks that rely on explicit task switches, AgarCL induces endogenous, state-dependent changes in both control and observation, while other agents in the environment further shape the experienced dynamics. Together, these factors result in a setting where the effective environment shifts over time without externally imposed task boundaries. This aligns with the “big world” hypothesis (Javed & Sutton, 2024; Lewandowski et al., 2025), as the agent operates in a world whose complexity is not exhaustively captured within a fixed training horizon and where ongoing adaptation is expected to remain beneficial.

3.1 The Game: Agar.io

AgarCL is based on the multiplayer online game Agar.io, in which each player controls one or more circular cells in a shared two-dimensional arena. The setting resembles a Petri dish populated by interacting cells, food pellets, and viruses. Players increase their mass by absorbing smaller entities, including stationary pellets and other players’ cells, while avoiding larger opponents. Viruses introduce additional strategic structure: depending on a cell’s size, colliding with a virus can either fragment it into multiple smaller cells or provide temporary protection against larger threats (see Figure 1). Players may also voluntarily split their cells for tactical purposes, allowing simultaneous control of multiple bodies. Although the underlying rules are simple, complex dynamics emerge from these local interactions. Figure 2 illustrates several of these behaviours.

The arena contains three types of entities: *pellets*, *cells*, and *viruses*. Pellets are randomly scattered, static, have fixed size, and grant one unit of mass when consumed. Cells (including both the learning agent and bots) can consume pellets, viruses, and smaller cells; when a cell absorbs another cell, it gains its entire mass. Viruses have a fixed mass of 100, and depending on the size of the colliding cell, a virus may either be absorbed or cause the cell to fragment into multiple smaller cells.




Figure 3: **The first 5,000 time steps of an expert trajectory showing game progression.** The agent (pink) steadily gains mass: ① starting small, it eventually passes over a virus; ② splits into smaller cells; ③ grows sufficiently large to consume viruses and other agents; ④ splits to attack, increasing its speed. As mass increases, the field of view expands to keep all controlled cells visible. ⑤ The agent ultimately surpasses all opponents in mass and ⑥ can see the entire arena.

Players begin as a small cell with mass 25. If a player’s cell is consumed by a larger cell, it is eliminated and subsequently respawns with the same initial mass. Although this repeated elimination might suggest an episodic structure, *AgarCL* is fundamentally a continual task: successful agents could survive forever, and, crucially, the environment does not reset upon elimination. The cell that absorbed the player retains its newly acquired mass, so the consequences of actions taken in one “life” persist into subsequent ones. Our current implementation focuses on the single-agent setting, where other players, called *bots*, follow heuristic-based policies. This design isolates the continual adaptation problem while maintaining rich interaction dynamics. For researchers interested in fully multi-agent variants of the game, we refer to the *GOBIGGER* platform (Zhang et al., 2023).

Several mechanisms contribute to the game’s evolving dynamics. New pellets are generated every 600 ticks (10 seconds) whenever fewer than 500 pellets are present in the arena, and viruses are generated at each time step if their total number falls below 10. At the same time, each cell loses 0.2% of its mass per second, thereby imposing a constant pressure to grow. As mass increases, mobility decreases: a cell’s speed v is determined by $v = 100/\text{mass}^{0.439}$, resulting in a slower movement for larger cells. Consequently, smaller players must evade larger opponents, avoid being trapped near boundaries, or strategically exploit viruses for protection. Players can feed viruses mass; once sufficiently large—typically after seven feedings—a virus splits and can be directed toward nearby cells, potentially fragmenting larger opponents. A player’s field of view also expands with its mass to ensure that all controlled cells remain visible.⁴ Figure 3 illustrates a typical progression of play.

These dynamics create persistent trade-offs between short-term gains and long-term survival. For example, a common tactic is to split a large cell to increase the likelihood of consuming smaller opponents that would otherwise remain out of reach due to reduced mobility. However, this aggressive move introduces significant risks as the resulting smaller cells can be consumed by larger opponents. Such decisions require balancing offensive opportunities against defensive stability, with consequences that unfold over extended time scales. Because no single strategy is uniformly effective across mass regimes or local configurations of opponents, agents must continually adjust their behaviour in response to both their own evolving state and the surrounding ecosystem.

⁴The original game contains additional rules. For example, a player cannot split beyond 14 cells, and consuming multiple viruses in succession activates a penalty that increases mass decay; this penalty persists across respawns. We include these rules to faithfully replicate the original game mechanics and minimize experimenter-induced bias, as *Agar.io* was not designed for AI research, and we wanted to preserve that as much as possible.

3.2 Agent-Environment Interface

Having specified the underlying dynamics, we now formalize the interaction protocol between the agent and the environment. We adopt the standard RL formulation in which interaction proceeds through discrete time steps via actions, observations, and rewards. We detail the action space, observation space, and reward structure below and summarize the key entities in Figure 1.

Reward Function. To keep the rewards bounded, we define the reward at time step t as the change in the agent’s mass between consecutive time steps:

$$R_t = m_t - m_{t-1}. \quad (1)$$

Importantly, the problem is continuing (non-episodic), so the agent’s death does not terminate the game. Upon respawning, the agent receives a reward equal to its death mass less its initial mass. This choice discourages death as a cost-free reset when the agent is bigger than its initial mass. Additionally, each cell loses 0.2% of its mass per second, which we do not represent in Equation 1.

Observation Space. How the agent perceives the environment critically determines the difficulty and structure of the learning problem. We primarily consider a pixel-based, top-down rendering of the environment, requiring the agent to operate over a high-dimensional, partially observable, and non-stationary visual stream. Figures throughout the paper illustrate examples of the agent’s observations, with Figure 3 highlighting the evolution of the observation stream.

Variants of Agar.io have previously been explored in RL research (Wiehe et al., 2018; Ansó et al., 2019; Zhang et al., 2023); however, those works rely on structured grid-like or symbolic state representations. To the best of our knowledge, a high-dimensional pixel-based formulation has not been studied. While our experiments focus on the pixel-based setting, AgarCL also supports more symbolic observation variants for controlled comparisons. We describe both below.

Pixel-Based Observation: The observation at time step t is represented as a tensor $\mathbf{O}_t \in \mathbb{R}^{N \times N \times 4}$, where $N \times N$ denotes the spatial resolution (default $N = 128$). The four channels encode pellets, viruses, enemy cells, and the agent (including gridlines), respectively. Figure 1 provides an example rendering (gridlines are faint and may not be visible).

Grid-Like Observation: We additionally adapt the observation design from GoBigger (Zhang et al., 2023) to the single-agent setting. This representation consists of a global state and a player state. The global state includes coarse environmental information such as map size and elapsed time steps. The player state captures agent-specific information, including its field of view, visible entities within that region, current score, and available actions. The overlap field is critical because it encodes nearby pellets, viruses, and cells together with their positions, velocities, and associated attributes.

Action Space. The action space is hybrid. At each time step, the agent selects both a continuous control vector and a discrete action. The continuous component corresponds to a cursor location $\langle x, y \rangle \in [-1, 1]^2$, which determines the direction of movement for all controlled cells. This design mimics human play, in which movement is specified by pointing to a location in the arena.

Simultaneously, the agent might choose to select one discrete action from $\{\text{SPLIT}, \text{EJECT}\}$. Not selecting any actions applies only the continuous cursor control. The SPLIT action divides each eligible cell into two equal-mass cells, provided the resulting cells have mass at least 25 (otherwise the action has no effect). One of the resulting cells is propelled toward the cursor direction with substantial momentum. After splitting, multiple cells must be controlled jointly through the same cursor input. The EJECT action expels a small mass unit from each cell toward the cursor direction. Ejected mass can be consumed by cells or viruses. All selected actions are executed for four consecutive environment steps before a new action is sampled. The observation returned to the agent corresponds to the state after these four steps (i.e., a frame skip of 4).

Finally, the environment includes stochasticity: Gaussian noise $\mathcal{N}(0, 1)$ is added to the continuous control signal before execution, preventing perfectly deterministic actuation.

Simulation Speed. We report performance in terms of the interquartile mean (IQM) over ten independent trials. In our experiments, a random Python agent processes 2,016 frames per second with a frame skip of 1. With the default frame skip of 4, the agent receives 1,163 observations per second, corresponding to 4,652 environment frames and 1,163 agent decisions per second. This throughput is substantially higher than prior Agar.io-based benchmarks such as GOBIGGER. The performance gains result from several implementation choices, including EGL-based rendering, optimized collision detection, efficient observation generation, and careful memory management.

Technical Details. Appendix C provides implementation details regarding the release of AgarCL, including its software architecture, agent interface, configuration options, and performance.

4 Experiments and Results

To evaluate the learning capabilities of agents in AgarCL, we begin with standard deep reinforcement learning baselines: DQN (Mnih et al., 2015), PPO (Schulman et al., 2017), and SAC (Haarnoja et al., 2018).⁵ These algorithms span value-based and policy-gradient methods, and together represent widely used approaches for high-dimensional control. By starting with these methods, we assess whether existing deep RL techniques are sufficient to cope with the non-stationary, partially observable, and multi-agent dynamics induced by AgarCL.

All agents process $128 \times 128 \times 4$ pixel observations through a shared convolutional encoder. The encoder consists of three conv. layers with kernel sizes 8×8 , 4×4 , and 3×3 , and strides 4, 2, and 1, respectively. Each layer is followed by layer normalization (Ba et al., 2016) and a ReLU activation. The resulting $32 \times 12 \times 12$ feature map is flattened and passed through a fully connected layer to obtain an embedding. This embedding is then used to predict both discrete and continuous actions, enabling hybrid control. Appendix F details the algorithm-specific adaptations required for this setting.

4.1 Benchmark Results for Continual RL in AgarCL

The full game constitutes the reference and default configuration of AgarCL. We intentionally evaluate agents in this setting first, as it combines all the key challenges posed by AgarCL. The arena size is 350×350 and contains 10 viruses, 8 bots, and 500 pellets that regenerate every 600 ticks. We train DQN, PPO, and SAC for 160 million environment frames, averaging results over 10 independent runs. Training is computationally demanding: running SAC in this setting, for example, requires more than seven days on a single machine.

The results highlight the difficulty of AgarCL for standard deep RL methods. Across all algorithms, we observe a consistent failure to learn an effective policy in the full-game configuration. Figure 4 shows the learning curves, and exact numerical results are provided in Table 9 (Appendix I). These outcomes indicate that traditional approaches struggle to address the challenges surfaced by AgarCL, which are central to continual RL. We analyze these challenges in more detail in the following sections.

Hyperparameter tuning is a central yet often overlooked challenge in continual RL. In continual settings, standard tuning strategies, such as optimizing performance over a fixed horizon, are inherently problematic (Mesbahi et al., 2024). If an agent is expected to operate indefinitely, what horizon should guide tuning? Selecting a shorter interval risks overfitting to that duration, contradicting continual learning. This issue




Figure 4: **Performance of established deep RL algorithms in the full game of AgarCL.** Each curve is the average over 10 seeds, with a moving average computed over a window of 100 steps.

⁵Our implementations are available in <https://github.com/machado-research/AgarCL-benchmark>.

is not merely theoretical: e.g., hyperparameters that perform well on 400k frames in Atari 2600 games often differ substantially from those used when training for 200M frames. Moreover, RL performance is frequently highly sensitive to small hyperparameter changes (see Appendix H.2.1).

Because tuning specifically for the 160M-frame horizon of the full game would implicitly bias the evaluation toward that timescale, we avoid horizon-specific tuning. At the same time, hyperparameters must be chosen. We therefore reuse the best hyperparameters identified for AgarCL’s continual MINI-GAME 4, the setting most similar to the full game but considerably shorter. This pragmatic choice reflects a broader open question in continual RL: how should hyperparameters be selected when no natural training horizon exists?

4.2 AgarCL as a Continual Reinforcement Learning Testbed

We have not yet provided direct evidence that AgarCL meaningfully captures the challenges of continual RL. We do so here by examining whether continual learning confers an advantage. Specifically, we compare agents whose policies are fixed after training with agents that learn continually.

Demonstrating this effect in the full-game setting is difficult, as the algorithms evaluated in Section 4.1 did not learn a reasonable policy. Thus, we consider an easier configuration in which pellets and viruses regenerate every 120 ticks (instead of 600), and pellet density is increased to 1024 (instead of 500). We focus on PPO here because it consistently outperformed the other baselines across multiple configurations, as discussed in the next sections. Performance across different pellet densities is reported in Appendix J.

Having identified a setting in which PPO learns a reasonable policy, we evaluate the robustness of that policy when learning is halted. Specifically, we freeze policies trained for 32 million and 48 million steps and evaluate them in the configuration with 1024 pellets. As shown in Figure 5, although the frozen agents initially perform competitively, their performance eventually collapses. This behavior indicates that policies learned at a fixed point in time fail to remain effective as the environment evolves, supporting the use of AgarCL as a testbed for continual adaptation.

It is also informative to consider a configuration where this collapse does not occur. When reducing the number of bots from eight to four, PPO learns a relatively stable policy that collects pellets and reliably avoids opponents (see Appendix K). In this simpler setting, freezing the policy does not lead to the same degradation in performance. This contrast highlights the role of environmental complexity in inducing continual adaptation and underscores how specific design elements of AgarCL contribute to its suitability as a continual RL evaluation framework.

4.3 Existing Continual RL algorithms in AgarCL

While the primary goal of this paper is not to benchmark continual RL methods in AgarCL, it is natural to ask whether existing approaches transfer to this setting.

Many established continual learning methods rely on explicit task boundaries, which are not available in AgarCL. Representative examples include EWC (Kirkpatrick et al., 2017), MAS (Aljundi et al., 2018), and LwF (Li & Hoiem, 2016). EWC and MAS incorporate task identity into regularization constraints on parameter updates, while LwF uses task structure to define distillation losses that preserve performance on earlier tasks.




Figure 5: **Performance of fixed-policy agents initialized from checkpoints at 32M and 48M steps.** We report the moving average over 10 random seeds with a window size of 1000 steps.

We therefore focus on task-agnostic methods that can be applied to AgarCL without structural modifications. These approaches primarily aim to preserve plasticity during learning. Concretely, we augment PPO with Shrink and Perturb (Ash & Adams, 2020), ReDo (Sokar et al., 2023), and Continual Backpropagation (Dohare et al., 2024), and evaluate them in the easier setting introduced in Section 4.2.⁶ As shown in Figure 6, the performance of the baseline, PPO, starts to drop before 100 million frames, while we do not observe this behaviour with PPO with techniques designed to mitigate loss of plasticity. Hyperparameters and numerical results are reported in Appendix K.

We did not perform any additional investigation to determine the extent to which PPO’s performance degradation is attributable to loss of plasticity. Nevertheless, the observation that equipping PPO with techniques designed to mitigate loss of plasticity tend to achieve higher average performance, albeit with substantial variance, suggests that loss of plasticity may contribute to the observed performance degradation. More importantly, we observe that none of the methods achieve strong performance, and we saw little evidence of learning in the more challenging configuration with 500 pellets and 600-tick regeneration. These findings suggest that while stability–plasticity may be one contributing factor in AgarCL, the primary bottlenecks likely lie elsewhere, including exploration and long-horizon credit assignment.




Figure 6: **Performance of different continual learning approaches.** We report the moving average over 10 runs with a window size of 10,000 steps.

4.4 Validating AgarCL through Mini-Games

As shown above, AgarCL presents a highly challenging learning problem. Beyond requiring sustained online adaptation, it exposes additional difficulties central to continual RL, including exploration without resets, long-horizon credit assignment, and stable representation learning under a continuously evolving observation stream.

To better understand these challenges, we evaluate agents on a set of mini-games designed within Agar.io. Each mini-game isolates specific aspects of the full environment, such as non-stationarity and its non-episodic dynamics, allowing us to analyze their impact in a more controlled manner.

Importantly, we tuned each algorithm’s hyperparameters in every mini-game. We did so over 3 random seeds, selecting the configuration that maximized the mean return across the final 100 episodes of a 20-million-frame run. To avoid maximization bias, we evaluated this configuration with 10 new trials. Details on the hyperparameter ranges and selection criteria are provided in Appendix G.

Non-Stationarity through Mass Variation. We begin with a set of progressively challenging pellet-collection mini-games. These environments exclude bots and viruses, isolating the effect of mass dynamics. For simplicity, this first group is *episodic*, with fixed episode lengths of 500 (first three mini-games) and 3000 (last three mini-games) time steps, and a single start state.

These mini-games were designed to isolate three aspects of the full game: (i) short-term exploration, (ii) mass decay, and (iii) the dynamics induced by large mass (slower movement but access to splitting). To simplify exploration, delayed credit assignment, and partial observability, MINI-GAME 1 consists of collecting pellets arranged in a square path. The agent starts with mass 25 and experiences no mass decay. Although this is the simplest setting, it still exhibits non-stationarity: as the agent gains mass, it becomes larger and slower, yet the optimal behavior is simply to follow the dense trail of pellets. MINI-GAME 2 extends this setup by introducing mass decay, where the loss

⁶Other replay-based continual learning approaches, such as CLEAR (Rolnick et al., 2019), could also be considered. CLEAR leverages past experience from a replay buffer to mitigate forgetting. In this work, however, our goal is to analyze how environment design influences continual adaptation rather than to exhaustively benchmark all continual RL baselines.




Figure 7: **Performance of RL methods on episodic pellet-collection mini-games.** Panels (a)–(c) show results for the square-path tasks (MINI-GAMES 1–3), while (d)–(f) correspond to the randomly regenerated tasks (MINI-GAMES 4–6). Y-axis scales differ across panels. The dashed line indicates human performance, and the green line denotes the random policy. Shaded regions represent 95% confidence intervals over 10 runs, computed using the t-distribution.

per step depends on the agent’s current mass. MINI-GAME 3 further increases the starting mass to 1000, making the agent substantially slower, increasing decay, and enabling splitting. MINI-GAMES 4–6 mirror the first three settings, except that pellets are randomly scattered rather than arranged along a dense path. In these variants, credit assignment is more delayed, exploration requirements change, and partial observability plays a larger role (though episodic resets still simplify the task). Screenshots of all mini-games are provided in Appendix H.1.

Figure 7 reports performance across all mini-games; numerical results are provided in Appendix I. In the simplest setting, MINI-GAME 1, all algorithms achieve performance comparable to that of a human player. Introducing mass decay substantially increases difficulty, even in the structured square-path task, reducing performance to roughly half of the human baseline. Starting the agent with a larger mass further amplifies this effect. These trends persist when pellets are uniformly distributed across the arena, although the task becomes markedly more challenging. In these variants, and particularly once mass decay is introduced, only PPO is able to learn a non-trivial policy.

Continual Problems. We now evaluate the mini-games above in the continual setting.⁷ We remove episodic resets and replenish all pellets every 600 frames. The maximum number of pellets is 500, regenerated randomly (restricted to the path in the square-path mini-games). As before, we re-tuned hyperparameters separately for each mini-game and algorithm.

These tasks proved substantially more challenging. Results are reported in Appendix I (Figure 9 and Table 9). In the square-path variants, no agent succeeds. Exploration and partial observability play a central role: once the agent deviates from the path and loses sight of it, it struggles to recover, and without resets it must rediscover the path to continue learning. Interestingly, when pellets are uniformly distributed, performance improves relative to the square-path setting; possibly because pellets remain regularly observable. However, reducing the pellet count by half makes the problem considerably harder, and the baselines degrade sharply (Appendix J). As before, PPO is the most robust method among those evaluated. For MINI-GAME 4, we truncate training once agents reach the maximum allowed mass; without mass decay they would otherwise grow unbounded.

Finally, one might expect architectures with explicit memory to perform better in this setting. We therefore evaluated PPO augmented with a GRU (Cho et al., 2014). We did not observe consistent improvements that justified the added complexity, and thus focus on simpler architectures in the main text. Detailed results are provided in Appendix H.2.2.

⁷We use the word *continual* deliberately to emphasize the non-stationary nature of the problem. *Continuing* typically refers to stationary infinite-horizon problems, while *continuous* denotes continuous action spaces.

Interacting with Other Agents. We also designed mini-games in which the agent shares the arena with a second bot governed by a fixed policy. We evaluated performance against several simple opponents. Across all cases, the learning agents failed to learn a policy that reliably collected enough pellets to outgrow and eventually absorb the opponent, even when the opposing bot did not actively chase them. These results further highlight the difficulty of exploration in this environment: even acquiring the seemingly basic skill of safely scaling before engaging proves non-trivial. Additional details are provided in Appendix H.3.

Interacting with Viruses. Beyond pellet collection and direct competition, we designed a mini-game to assess whether agents could learn to use viruses strategically, splitting a larger bot and subsequently absorbing it. As in the previous settings, none of the evaluated agents learned the required sequence of actions, despite simplifying the task as much as possible. These results help explain why prior work in Agar-like environments has largely focused on pellet-collection tasks rather than other strategic interactions. Detailed results and analysis are provided in Appendix H.4.

Summary. Taken together, these mini-game results illustrate the range and severity of challenges posed by AgarCL. Even when individual aspects of the environment are isolated—non-stationarity, long-term credit assignment, partial observability, or strategic interaction—existing methods struggle to achieve robust performance. The transition from episodic to continual variants further amplifies these difficulties, revealing sharp degradation once resets are removed. More complex interactions, such as competition and virus-based strategies, remain entirely unsolved. Collectively, these findings validate AgarCL as a demanding and realistic benchmark for continual reinforcement learning, where progress requires advances beyond incremental algorithmic refinements.

5 Related Work

GOBIGGER (Zhang et al., 2023) is the closest existing platform to ours due to its Agar-style gameplay. However, it was designed to study collective behaviors in multi-agent RL, emphasizing coordination among *teams* of agents. It does not target continual RL; all tasks are episodic. In contrast, AgarCL is explicitly designed for continual settings, supports additional data streams, and scales much better with increasing numbers of cells and pellets. Earlier work also explored Agar-style games as RL environments, primarily to benchmark deep RL algorithms on pellet-collection tasks (Ansó et al., 2019; Wiehe et al., 2018), rather than to study non-stationarity or continual RL. AgarCL is implemented on top of the AgarLE (Deaton, 2018), an incomplete implementation of Agar.io with an OpenAI Gym (Brockman et al., 2016) interface.

Many platforms have been proposed for continual RL research. AgarCL complements these efforts by introducing a fundamentally different problem structure. In particular, unlike Switching ALE (Abbas et al., 2023), Continual World (Wolczyk et al., 2021), POET (Wang et al., 2019), MEAL (Tomilin et al., 2025), Continual NavBench (Kobanda et al., 2025), and numerous Mini-Grid variants (Chevalier-Boisvert et al., 2023), AgarCL is not organized as a sequence of episodic tasks that periodically switch. Instead, non-stationarity emerges organically from the underlying dynamics of a persistent environment.

The primary platform explicitly supporting non-episodic continual RL is JellyBean World (Platanios et al., 2020). However, it features comparatively simple dynamics, a discrete action space, and limited non-stationarity, which primarily stems from abrupt reward-function changes. Other alternatives include bespoke adaptations of MuJoCo (Todorov et al., 2012) or IsaacGym (Makoviychuk et al., 2021) environments used in specific studies (e.g., Feng et al., 2022; Long et al., 2024). These setups, however, were not released as standardized evaluation platforms. Beyond individual features, a key distinction is that AgarCL integrates reward, perception, and dynamics in a tightly coupled manner, producing non-stationarity that is directly linked to the agent’s state and behavior. Reproducing this form of smoothly evolving, interaction-driven non-stationarity within physics engines would require substantial additional design.

Finally, complex environments such as NetHack (Küttler et al., 2020) and Minecraft (Johnson et al., 2016; Guss et al., 2019) can also serve as potential evaluation frameworks for continual RL under the premise that they are “bigger than the agent” (Javed & Sutton, 2024). However, their richness introduces additional challenges that may obscure the specific difficulties of continual learning. NetHack, for example, relies heavily on language and prior knowledge, and existing approaches typically require substantial auxiliary machinery. Likewise, many successful Minecraft solutions rely on human demonstrations because exploration is difficult. In contrast, AgarCL aims to isolate continual learning challenges without relying on external supervision or domain-specific priors.

6 Conclusion

We introduced AgarCL as an evaluation platform for continual reinforcement learning. It captures key challenges, including partial observability and smooth, endogenous non-stationarity, while avoiding the abrupt task switches common in existing benchmarks. Through extensive experiments with DQN, PPO, and SAC across mini-games and the full environment, we demonstrated both the severity of the full problem and the utility of the mini-games for controlled analysis.

AgarCL incorporates features central to continual RL, including non-episodic interactions, smooth non-stationarity, high-dimensional observations, continuous actions, and a potentially unbounded horizon. Our results highlight the limitations of standard deep RL algorithms in such settings and underscore the gap between current methods and the demands of continual learning. Beyond introducing capabilities largely absent from prior frameworks, we show that even fixed policies fail to maintain stable performance in this evolving environment.

This work emphasizes the problem itself rather than proposing new solution methods. Given both scope and space constraints, we did not systematically evaluate algorithms specifically designed for continual RL. A practical limitation is the computational cost: although AgarCL was engineered for efficiency, continual evaluation at the required time scales remains resource-intensive, and experimental cycles can still be lengthy.

Acknowledgments

We would like to thank Jan Prus-Czarnecki and Anton Wiehe for early discussions, Alex Lewandowski and Anna Hakhverdyan for advice throughout the project, and Brett Daley and Prabhat Nagarajan for feedback on an earlier version of the manuscript. We would also like to thank Jon Deaton and the other contributors of the AgarLE GitHub repository — github.com/jondeaton/AgarLE. Kateryna contributed to this project while interning at NVIDIA and working in Prof. Machado’s lab.

The research is supported in part by the Natural Sciences and Engineering Research Council of Canada (NSERC) and the Canada CIFAR AI Chair Program. This research was enabled in part by the computational resources provided by the Alberta Machine Intelligence Institute (Amii), the Digital Research Alliance of Canada, and the University of Alberta Google Cloud Incubator.

References

- Zaheer Abbas, Rosie Zhao, Joseph Modayil, Adam White, and Marlos C. Machado. Loss of plasticity in continual deep reinforcement learning. In *Conference on Lifelong Learning Agents (CoLLAs)*, 2023.
- David Abel, André Barreto, Benjamin Van Roy, Doina Precup, Hado Philip van Hasselt, and Satinder Singh. A definition of continual reinforcement learning. In *Neural Information Processing Systems (NeurIPS)*, 2023.

- Rahaf Aljundi, Francesca Babiloni, Mohamed Elhoseiny, Marcus Rohrbach, and Tinne Tuytelaars. Memory aware synapses: Learning what (not) to forget. In *European Conference on Computer Vision (ECCV)*, 2018.
- Nishanth Anand and Doina Precup. Prediction and control in continual reinforcement learning. In *Neural Information Processing Systems (NeurIPS)*, 2023.
- Nil Stolt Ansó, Anton Orell Wiehe, Madalina M. Drugan, and Marco A. Wiering. Deep reinforcement learning for pellet eating in Agar.io. In *International Conference on Agents and Artificial Intelligence (ICAART)*, 2019.
- Jordan T. Ash and Ryan P. Adams. On warm-starting neural network training. In *Neural Information Processing Systems (NeurIPS)*, 2020.
- Lei Jimmy Ba, Jamie Ryan Kiros, and Geoffrey E. Hinton. Layer normalization. *CoRR*, abs/1607.06450, 2016.
- Charles Beattie, Joel Z. Leibo, Denis Teplyashin, Tom Ward, Marcus Wainwright, Heinrich Küttler, Andrew Lefrancq, Simon Green, Víctor Valdés, Amir Sadik, Julian Schrittwieser, Keith Anderson, Sarah York, Max Cant, Adam Cain, Adrian Bolton, Stephen Gaffney, Helen King, Demis Hassabis, Shane Legg, and Stig Petersen. DeepMind Lab. *CoRR*, abs/1612.03801, 2016.
- Marc G. Bellemare, Yavar Naddaf, Joel Veness, and Michael Bowling. The Arcade Learning Environment: An evaluation platform for general agents. *Journal of Artificial Intelligence Research*, 47:253–279, 2013.
- David Blackwell. Discrete Dynamic Programming. *The Annals of Mathematical Statistics*, 33(2): 719–726, 1962.
- Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. OpenAI Gym. *CoRR*, abs/1606.01540, 2016.
- Maxime Chevalier-Boisvert, Bolun Dai, Mark Towers, Rodrigo Perez-Vicente, Lucas Willems, Salem Lahlou, Suman Pal, Pablo Samuel Castro, and Jordan K. Terry. Minigrid & Miniworld: Modular & customizable reinforcement learning environments for goal-oriented tasks. In *Neural Information Processing Systems (NeurIPS)*, 2023.
- Kyunghyun Cho, Bart van Merriënboer, Dzmitry Bahdanau, and Yoshua Bengio. On the properties of neural machine translation: Encoder-decoder approaches. In *EMNLP Workshop on Syntax, Semantics and Structure in Statistical Translation*, 2014.
- Will Dabney, Georg Ostrovski, and Andre Barreto. Temporally-extended ε -greedy exploration. In *International Conference on Learning Representations (ICLR)*, 2021.
- Jonathan Deaton. AgarLE: Agar.io OpenAI Gym Learning Environment, 2018. <https://github.com/jondeaton/AgarLE>. Accessed: 2025-04-09.
- Shibhansh Dohare, J. Fernando Hernandez-Garcia, Qingfeng Lan, Parash Rahman, A. Rupam Mah-mood, and Richard S. Sutton. Loss of plasticity in deep continual learning. *Nature*, 632(8026): 768–774, 2024.
- Fan Feng, Biwei Huang, Kun Zhang, and Sara Magliacane. Factored adaptation for non-stationary reinforcement learning. In *Neural Information Processing Systems (NeurIPS)*, 2022.
- Yasuhiro Fujita, Prabhat Nagarajan, Toshiaki Kataoka, and Takahiro Ishikawa. ChainerRL: A deep reinforcement learning library. *Journal of Machine Learning Research*, 22(77):1–14, 2021.
- William H. Guss, Brandon Houghton, Nicholay Topin, Phillip Wang, Cayden R. Codel, Manuela Veloso, and Ruslan Salakhutdinov. MineRL: A large-scale dataset of Minecraft demonstrations. In *International Joint Conference on Artificial Intelligence (IJCAI)*, 2019.

- Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *International Conference on Machine Learning (ICML)*, 2018.
- Peter Henderson, Riashat Islam, Philip Bachman, Joelle Pineau, Doina Precup, and David Meger. Deep reinforcement learning that matters. In *AAAI Conference on Artificial Intelligence (AAAI)*, 2018.
- Muhammad Kamran Janjua, Haseeb Shah, Martha White, Erfan Miahi, Marlos C. Machado, and Adam White. GVFs in the real world: Making predictions online for water treatment. *Machine Learning*, 113(8):5151–5181, 2024.
- Khurram Javed and Richard S. Sutton. The big world hypothesis and its ramifications for artificial intelligence. In *Finding the Frame Workshop at Reinforcement Learning Conference*, 2024.
- Matthew Johnson, Katja Hofmann, Tim Hutton, and David Bignell. The Malmo Platform for artificial intelligence experimentation. In *International Joint Conference on Artificial Intelligence (IJCAI)*, 2016.
- Leslie Pack Kaelbling, Michael L. Littman, and Anthony R. Cassandra. Planning and acting in partially observable stochastic domains. *Artificial Intelligence*, 101(1-2):99–134, 1998.
- Khimya Khetarpal, Matthew Riemer, Irina Rish, and Doina Precup. Towards continual reinforcement learning: A review and perspectives. *Journal of Artificial Intelligence Research*, 75:1401–1476, 2022.
- James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A. Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, Demis Hassabis, Claudia Clopath, Dharshan Kumaran, and Raia Hadsell. Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences*, 114(13):3521–3526, 2017.
- Anthony Kobanda, Odalric-Ambrym Maillard, and Rémy Portelas. A continual offline reinforcement learning benchmark for navigation tasks. In *IEEE Conference on Games (CoG)*, 2025.
- Saurabh Kumar, Henrik Marklund, Ashish Rao, Yifan Zhu, Hong Jun Jeon, Yueyang Liu, and Benjamin Van Roy. Continual learning as computationally constrained reinforcement learning. *Foundations and Trends in Machine Learning*, 18(5):913–1053, 2025.
- Heinrich Küttler, Nantas Nardelli, Alexander Miller, Roberta Raileanu, Marco Selvatici, Edward Grefenstette, and Tim Rocktäschel. The NetHack Learning Environment. In *Neural Information Processing Systems (NeurIPS)*, 2020.
- Yann LeCun, Léon Bottou, Genevieve B. Orr, and Klaus-Robert Müller. Efficient BackProp. In *Neural Networks: Tricks of the Trade - Second Edition*, volume 7700 of *Lecture Notes in Computer Science*, pp. 9–48. Springer, 2012.
- Alex Lewandowski, Haruto Tanaka, Dale Schuurmans, and Marlos C. Machado. Curvature explains loss of plasticity. *CoRR*, abs/2312.00246, 2023.
- Alex Lewandowski, Aditya Ramesh, Edan Meyer, Dale Schuurmans, and Marlos C. Machado. The world is bigger: A computationally-embedded perspective on the big world hypothesis. In *Neural Information Processing Systems (NeurIPS)*, 2025.
- Zhizhong Li and Derek Hoiem. Learning without forgetting. In Bastian Leibe, Jiri Matas, Nicu Sebe, and Max Welling (eds.), *European Conference on Computer Vision (ECCV)*, 2016.
- Junfeng Long, Wenye Yu, Quanyi Li, Zirui Wang, Dahua Lin, and Jiangmiao Pang. Learning H-Infinity locomotion control. In *Conference on Robot Learning (CoRL)*, 2024.

- Marlos C. Machado. *Efficient Exploration in Reinforcement Learning through Time-Based Representations*. PhD thesis, University of Alberta, Canada, 2019.
- Marlos C. Machado, Marc G. Bellemare, Erik Talvitie, Joel Veness, Matthew J. Hausknecht, and Michael Bowling. Revisiting the Arcade Learning Environment: Evaluation protocols and open problems for general agents. *Journal of Artificial Intelligence Research*, 61:523–562, 2018.
- Viktor Makoviyichuk, Lukasz Wawrzyniak, Yunrong Guo, Michelle Lu, Kier Storey, Miles Macklin, David Hoeller, Nikita Rudin, Arthur Allshire, Ankur Handa, and Gavriel State. Isaac Gym: High performance GPU based physics simulation for robot learning. In *Neural Information Processing Systems (NeurIPS)*, 2021.
- Golnaz Mesbahi, Olya Mastikhina, Parham Mohammad Panahi, Martha White, and Adam White. K-percent evaluation for lifelong RL. *CoRR*, abs/2404.02113, 2024.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Bellemare, Alex Graves, Martin A. Riedmiller, Andreas Fidjeland, Georg Ostrovski, Stig Petersen, Charles Beattie, Amir Sadik, Ioannis Antonoglou, Helen King, Dhharshan Kumaran, Daan Wierstra, Shane Legg, and Demis Hassabis. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, 2015.
- Andrew Patterson, Samuel Neumann, Martha White, and Adam White. Empirical design in reinforcement learning. *Journal of Machine Learning Research*, 25(318):1–63, 2024.
- Emmanouil Antonios Platanios, Abulhair Saparov, and Tom M. Mitchell. Jelly Bean World: A testbed for never-ending learning. In *International Conference on Learning Representations (ICLR)*, 2020.
- Sam Powers, Eliot Xing, Eric Kolve, Roozbeh Mottaghi, and Abhinav Gupta. CORA: Benchmarks, baselines, and metrics as a platform for continual reinforcement learning agents. In *Conference on Lifelong Learning Agents (CoLLAs)*, 2022.
- Martin L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, 2014.
- David Rolnick, Arun Ahuja, Jonathan Schwarz, Timothy Lillicrap, and Gregory Wayne. Experience replay for continual learning. In *Neural Information Processing Systems (NeurIPS)*, 2019.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *CoRR*, abs/1707.06347, 2017.
- Ghada Sokar, Rishabh Agarwal, Pablo Samuel Castro, and Utku Evci. The dormant neuron phenomenon in deep reinforcement learning. In *International Conference on Machine Learning (ICML)*, 2023.
- Richard S. Sutton, Anna Koop, and David Silver. On the role of tracking in stationary environments. In *International Conference on Machine Learning (ICML)*, 2007.
- Yuval Tassa, Yotam Doron, Alistair Muldal, Tom Erez, Yazhe Li, Diego de Las Casas, David Budden, Abbas Abdolmaleki, Josh Merel, Andrew Lefrancq, Timothy P. Lillicrap, and Martin A. Riedmiller. DeepMind control suite. *CoRR*, abs/1801.00690, 2018.
- Emanuel Todorov, Tom Erez, and Yuval Tassa. MuJoCo: A physics engine for model-based control. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2012.
- Tristan Tomilin, Meng Fang, Yudi Zhang, and Mykola Pechenizkiy. COOM: A game benchmark for continual reinforcement learning. In *Neural Information Processing Systems (NeurIPS)*, 2023.

- Tristan Tomilin, Luka van den Boogaard, Samuel Garcin, Bram Grooten, Meng Fang, Yali Du, and Mykola Pechenizkiy. MEAL: A benchmark for continual multi-agent reinforcement learning. *CoRR*, abs/2506.14990, 2025.
- Rui Wang, Joel Lehman, Jeff Clune, and Kenneth O. Stanley. Paired open-ended trailblazer (POET): endlessly generating increasingly complex and diverse learning environments and their solutions. *CoRR*, abs/1901.01753, 2019.
- Anton Orell Wiehe, Nil Stolt Ansó, Madalina M. Drugan, and Marco A. Wiering. Sampled policy gradient for learning to play the game Agar.io. *CoRR*, abs/1809.05763, 2018.
- Maciej Wolczyk, Michał Zając, Razvan Pascanu, Łukasz Kuciński, and Piotr Miłoś. Continual World: A robotic benchmark for continual reinforcement learning. In *Neural Information Processing Systems (NeurIPS)*, 2021.
- Kevin Zakka, Baruch Tabanpour, Qiayuan Liao, Mustafa Haiderbhai, Samuel Holt, Jing Yuan Luo, Arthur Allshire, Erik Frey, Koushil Sreenath, Lueder A. Kahrs, Carmelo Sferrazza, Yuval Tassa, and Pieter Abbeel. MuJoCo Playground. *CoRR*, abs/2502.08844, 2025.
- Ming Zhang, Shenghan Zhang, Zhenjie Yang, Lekai Chen, Jinliang Zheng, Chao Yang, Chuming Li, Hang Zhou, Yazhe Niu, and Yu Liu. GoBigger: A scalable platform for cooperative-competitive multi-agent interactive simulation. In *International Conference on Learning Representations (ICLR)*, 2023.

Supplementary Materials

The following content was not necessarily subject to peer review.

A Broader Implications

This work introduces a new evaluation framework for continual RL, built around the video game Agar.io. Given that this project is based on a video game and designed for benchmarking and evaluation purposes, there are no foreseeable negative societal impacts. As with most research involving simulated environments in AI, the risks are minimal and do not require significant mitigation.

B Use of Large Language Models

LLMs were only used for minor language polishing during the writing of the paper.

C Software Release and Technical Details

AgarCL is released as free, open-source software under the terms of the MIT license. It is implemented on top of the AgarLE (Deaton, 2018), an incomplete implementation of Agar.io with an OpenAI Gym (Brockman et al., 2016) interface. AgarCL’s source code is publicly available in:

`https://github.com/machado-research/AgarCL`

AgarCL’s core simulation engine is implemented in C++, and rendering is handled by OpenGL. A Python interface, built with Pybind11, is also supported. In our experiments, when looking at the interquantile mean (IQM) over ten independent trials, a random agent, written in Python, receives 2,016 frames per second with a frame skip of 1. With a frame skip of 4 (default), the agent receives 1,163 observations per second, which represents 4,652 frames in the game and 1,163 actions selected by the agent. We report the IQM to discard fluctuations due to other jobs in the system. This experiment used 2 cores of an *Intel Xeon Gold 6448Y* CPU and one *NVIDIA L40s* GPU. Correspondingly, the IQM when evaluating AgarCL without a GPU was 1,440 and 1,092, respectively.

D Implementation Details in AgarCL

In this section, we discuss in more detail the full-game experiment and the implementation of the fixed-policy bots.

D.1 AgarCL Experiment Details

In our main evaluation, we use a 350×350 arena containing 500 pellets, 10 viruses, and 8 heuristic bots. Upon being consumed, any agent or bot respawns immediately with an initial mass of 25, while all other entities preserve their current state. We describe the bot heuristics in Appendix D.2.

Every 600 ticks (150 ticks if we use frame skipping of 4), pellets and viruses that have been consumed or destroyed are regenerated with a uniform probability across the arena. This maintains fixed totals of 500 and 10, respectively. Mass decay is applied every 60 ticks. We conduct each trial for 160×10^6 frames.

D.2 Fixed-Policy Bots

We implemented four heuristic bots in AgarCL: Aggressive, Aggressive-Shy, Hungry, and Hungry-Shy. We did so to explore how simple rules shape emergent game dynamics. All bots

follow a fixed policy that always targets the nearest pellet, but differ in how they interact with other agents/bots. This set of behaviours allows us to assess how different action priors influence both individual success and the overall balance of the game (see Appendix H.3). The four bots’ policies are:

- **Aggressive:** This bot first looks for any smaller opponent within a defined radius and attempts to consume it; if no suitable target is found, it switches to pellet collection.
- **Aggressive-Shy:** Like the Aggressive bot, it will hunt smaller opponents, but if a larger opponent approaches within its “shy” radius, it immediately flees and only returns to hunting once the threat has passed.
- **Hungry:** This bot ignores other players entirely and chases the closest pellet at every step.
- **Hungry-Shy:** Focused on pellet foraging like the Hungry bot, it additionally monitors for larger opponents: if one comes too close, it retreats before resuming its hunt.

E Additional Related Work

Continual Reinforcement Learning. There have been many attempts to formalize the continual reinforcement learning (CRL) problem, each highlighting different aspects of non-stationarity and lifelong adaptation. [Khetarpal et al. \(2022\)](#) introduce a taxonomy of continual reinforcement learning (CRL) that focuses on two fundamental dimensions of non-stationarity: ‘scope’, referring to the extent of variation across tasks or domains, and ‘driver’, representing the main source of change. Building on this, [Abel et al. \(2023\)](#) define CRL as an ongoing adaptation process, in which agents continuously update their policies in response to evolving objectives, dynamics, or reward structures. More recently, [Kumar et al. \(2025\)](#) define continual learning under computational constraints by proposing a framework that maintains a balance between past knowledge and the efficiency of online updates over an extended period. Collectively, these CRL formalisms establish foundational principles that guide the design of benchmarks and algorithms.

Existing Environments. Traditional RL testbeds (e.g., [Zakka et al., 2025](#); [Tassa et al., 2018](#)) have driven rapid progress in RL, but they are not tailored for CRL. Accordingly, many CRL studies resort to switching among different games to induce non-stationarity (e.g., [Abbas et al., 2023](#)), relying on clearly defined train-test boundaries and assuming a well-structured notion of tasks and episodes. As we argue in this paper, slower, smoother changes capture a different (and underexplored) type of non-stationarity in the problems that motivated CRL. JellyBeanWorld ([Platanios et al., 2020](#)) is the main platform for CRL in non-episodic settings that we are aware of. In JellyBeanWorld, agents navigate an infinite two-dimensional grid, interacting with various items by collecting or avoiding them. The agent’s states are partially observable. Although JellyBeanWorld is clearly valuable for CRL research, it offers simpler observations and (discrete) dynamics, without the ever-changing nature of AgarCL.

Agar-Like Environments. Most of the results relying on Agar-like environments for evaluation consisted of assessing basic agent capabilities, mostly related to pellet eating ([Wiehe et al., 2018](#); [Ansó et al., 2019](#)), without ever putting forward the environment as a key artifact. To the best of our knowledge, GoBigger ([Zhang et al., 2023](#)) is the only other actual evaluation framework for AI research based on the Agar.io game. It shares certain surface-level similarities with our own, particularly in its Agar-style gameplay mechanics, but it was introduced to support a fundamentally different problem. Its main goal is to provide a platform to study collective behaviours in multi-agent reinforcement learning. Thus, its features include supporting multiple teams of agents (in the original game, collaboration should emerge from communication) in settings where agents are organized into a few teams. The larger environment configuration supports 24 agents organized into four teams (6 each). The game is designed to be *episodic*, with larger maps artificially divided into episodes that have at most 14,400 frames (12 minutes). These are all in direct contrast to the continual learning problem we focus on, with potentially unbounded episodes and no pre-defined teams, but with a much bigger number of agents. GoBigger also does not

support pixel-based observations; it just supports something akin to our symbolic observations with information about objects’ positions and velocities. Finally, we benchmarked the frames per second (FPS) over ten independent trials using the `st_t6p4`⁸ configuration—the largest map setting in the GoBigger implementation. Our environment significantly outperforms GoBigger in simulation speed, achieving an interquartile mean (IQM) of 4,212 fps with GoBigger-style observations, compared to GoBigger’s IQM of 205 fps under the same observation setup. This experiment was run using 2 cores of an *Intel Xeon Gold 6448Y* CPU.

F Algorithm Details

In this section, we focus on the adaptations we have made to DQN, PPO, and SAC to allow them to work somewhat effectively in AgarCL.

DQN. We used PFRL’s DQN implementation (Fujita et al., 2021). We did so by discretizing the environment’s continuous actions. We ended up with 24 actions: 8 directions times the 3 discrete actions. The predefined directions were: UP (0, 1), UP-RIGHT (1, 1), RIGHT (1, 0), DOWN-RIGHT (1, -1), DOWN (0, -1), DOWN-LEFT (-1, -1), LEFT (-1, 0), and UP-LEFT (-1, 1).

As discussed throughout, the resets inherent to episodic tasks can benefit exploration by allowing agents to recover from “bad” states. However, in the face of continual problems, the agent must naturally recover from a “bad” state. In our experiments, we noticed that the ϵ -greedy strategy was ineffective in the non-episodic tasks we considered. Thus, in such tasks, we introduced temporally-extended exploration (Machado, 2019) via ϵ z-greedy (Dabney et al., 2021), which achieves temporal persistence by extending random actions over multiple steps using a heavy-tailed duration distribution.

PPO. Our implementation extends PFRL’s to the hybrid action setting. We adopted a shared neural network with two heads: an *actor* head that outputs action probabilities and a value *critic* head that estimates state values. Preliminary results showed that this architecture was more effective than two independent networks. The actor head splits a 256-dimensional feature vector into discrete and continuous branches: a softmax head for categorical action probabilities and a Gaussian head (state-independent covariance) for continuous action means and variances, sampling both at each timestep to form a factored joint policy. The critic head is a compact one-layer MLP producing a scalar $V(s)$, trained with PPO’s clipped value-function loss and generalized advantage estimation. Finally, we relied on PPO’s entropy regularization for exploration.

SAC. Similarly to PPO, we extended PFRL’s SAC implementation to the hybrid action setting and relied on the algorithm’s entropy-maximizing term for exploration. Unlike PPO’s single-value head, our implementation comprises three fully independent networks—one actor and two critics—each with its own encoder, ϕ , that directly processes the raw observation. The actor mirrors the PPO policy architecture, but with an entropy-regularized objective. Likewise, each critic encodes the observation using ϕ . The resulting representation is then concatenated with the continuous action, after which the critic predicts Q-values for each discrete action. The Q-values are selected with respect to the actor’s sampled discrete action. By structuring the critics this way, we avoid the combinatorial explosion that would arise if we naively input every possible hybrid action pair, and we prevent gradient interference between discrete and continuous parameters.

⁸This configuration typically includes 1000 pellets, a 144×144 arena, and 14,400 frames per episode. We modified the GoBigger implementation to support a single player.

G Tuning Details

Hyperparameter tuning is especially critical in reinforcement learning due to the inherent instability and sensitivity of RL algorithms. Unlike supervised learning, RL involves exploration-exploitation trade-offs, nonstationary data distributions, and delayed rewards, all of which can magnify the effects of poorly chosen hyperparameters. Parameters such as step size, discount factor, exploration noise, entropy regularization (in policy gradient methods), and update frequency can significantly influence learning dynamics and final policy performance. Improper tuning can lead to divergence, suboptimal policies, or excessive performance variance.

G.1 DQN Tuning

For tuning DQN, we swept over the hyperparameters listed in Table 1. As discussed in the paper, both the agent’s network and target network consist of three convolutional layers, each followed by a ReLU activation and a Layer Normalization layer. In preliminary experiments, we found that omitting Layer Normalization prevented the agent from learning altogether. All weights were initialized using a LeCun normal initialization (LeCun et al., 2012). Each hyperparameter combination was trained using three different random seeds. After identifying the best hyperparameter configuration, we conducted ten additional independent runs using the best-performing configuration (see Table 2 on the next page). As shown in Table 1, also on the next page, the discount factor, γ , the number of epochs, the replay buffer size, and the target-network update interval were each fixed at a single value, while the remaining hyperparameters were swept. Finally, we applied the hyperparameter settings from continual MINI-GAME 4 to train the full-game agent.

Table 1: Values of hyperparameters that we swept over when tuning DQN.

Hyperparameter	Values / Settings
Step size (<code>-lr</code>)	$10^{-5}, 3 \cdot 10^{-5}, 10^{-4}, 3 \cdot 10^{-4}$
Batch Accumulator (<code>-batch_accumulator</code>)	"sum", "mean"
Soft Update Coefficient (<code>-tau</code>)	$10^{-2}, 5^{-3}$
Batch Size (<code>-batch-size</code>)	32, 64
Replay Buffer Size (<code>-replay-buffer</code>)	10^5
Number of Epochs (<code>-epochs</code>)	1
Target Network Update Interval (<code>-update_interval</code>)	4
Gamma (<code>-gamma</code>)	0.99
Exploration algorithm	ϵ -Greedy, ϵ z-Greedy

G.2 PPO Tuning

Hyperparameter tuning for PPO is particularly challenging due to the size of the configuration space: minimally covering relevant combinations requires at least 324 runs per experiment. We employed a shared neural network to learn both actor and critic, using the same architecture described in Appendix F. All network weights are initialized using LeCun normal initialization (LeCun et al., 2012). The policy head is a Gaussian distribution with a state-independent, learned covariance—a configuration that is standard in PPO implementations (e.g., Henderson et al., 2018).

In Table 3, we have the choice between two reward normalization schemes as a hyperparameter. The first is *min-max normalization*, which linearly rescales the raw reward r_t into the range $[-1, 1]$:

$$\tilde{r}_t^{(\text{min-max})} = \frac{r_t - r_{\min}}{r_{\max} - r_{\min} + \epsilon}, \quad (2)$$

where $\epsilon > 0$ is a small constant added for numerical stability.

Table 2: Best hyperparameters for DQN on each minigame.

Category	Mini-game	Hyperparameters				Exploration Algorithm
		Step Size	Batch Accumulator	Soft Update Coefficient	Batch Size	
Episodic	①	10^{-4}	mean	$5 \cdot 10^{-5}$	64	ϵ -Greedy
	②	$3 \cdot 10^{-4}$	sum	$5 \cdot 10^{-5}$	32	ϵ -Greedy
	③	10^{-4}	mean	$5 \cdot 10^{-5}$	32	ϵ -Greedy
	④	$3 \cdot 10^{-4}$	sum	10^{-2}	32	ϵ -Greedy
	⑤	10^{-5}	mean	10^{-2}	64	ϵ -Greedy
	⑥	10^{-5}	mean	10^{-2}	64	ϵ -Greedy
Continuing	①	10^{-4}	sum	$5 \cdot 10^{-5}$	32	$\epsilon\mathcal{Z}$ -Greedy
	②	$3 \cdot 10^{-4}$	sum	10^{-2}	32	$\epsilon\mathcal{Z}$ -Greedy
	③	10^{-4}	mean	$5 \cdot 10^{-5}$	32	$\epsilon\mathcal{Z}$ -Greedy
	④	10^{-4}	sum	$5 \cdot 10^{-5}$	32	$\epsilon\mathcal{Z}$ -Greedy
	⑤	10^{-4}	mean	10^{-2}	32	$\epsilon\mathcal{Z}$ -Greedy
	⑥	$3 \cdot 10^{-4}$	mean	10^{-2}	32	$\epsilon\mathcal{Z}$ -Greedy
Other Agents	⑦-Small-Sparse ¹	10^{-5}	mean	5×10^{-2}	32	ϵ -Greedy
	⑦-Normal-Dense ²	3×10^{-4}	mean	10^{-2}	64	ϵ -Greedy
	⑧-Small-Sparse ¹	10^{-5}	mean	5×10^{-2}	64	ϵ -Greedy
	⑧-Normal-Dense ²	1×10^{-5}	sum	10^{-2}	64	ϵ -Greedy
Virus	⑨	3×10^{-4}	sum	10^{-2}	128	ϵ -Greedy

¹ Limited (Sparse) arena: 200×200 with 200 randomly respawned pellets.² Normal (Dense) arena: 350×350 with 500 randomly respawned pellets.

The second method, *variance normalization*, uses an exp. weighted moving average of the returns:

$$G_t = \gamma G_{t-1} + r_t, \quad (3) \quad \tilde{r}_t^{(\text{var-norm})} = \frac{r_t}{\sqrt{\text{Var}[G_t] + \epsilon}}, \quad (4)$$

where G_t is the smoothed return, and normalization is performed by dividing the reward by the square root of its running variance. This ensures that the normalized reward maintains approximately unit variance under an exponential moving average with discount factor γ . After normalization, the reward is clipped to lie within the fixed range $[-10, 10]$ to limit the effect of outliers during training:

$$\tilde{r}_t^{(\text{clipped})} = \text{clip}\left(\tilde{r}_t^{(\text{var-norm})}, -10, 10\right). \quad (5)$$

Initially, we swept over different hyperparameter combinations for PPO in some episodic settings following Table 3. Through these experiments, we observed that a value function coefficient of 0.9, an update interval of 5000, step sizes of either 10^{-5} or 3×10^{-5} , and epochs set to either 10 or 15 yielded consistently strong performance. Based on sensitivity analyses across most mini-game tasks, we narrowed down the range of hyperparameters we swept over for the other tasks. The smaller set of hyperparameters we swept over is shown in Table 4. Table 5 shows the best-performing hyperparameters that we used in each task at the end. Finally, we applied the hyperparameter settings from continual MINI-GAME 4 to train the full-game agent.

Table 3: Values of hyperparameters that we swept over when tuning PPO.

Hyperparameter	Values / Settings
Reward function (-reward)	min_max, variance_norm
Step size (-lr)	10^{-5} , 3×10^{-5} , 3×10^{-4} , 10^{-4}
Epochs (-epochs)	10, 15, 20
Max gradient norm (-max-grad-norm)	0.5, 0.7, 0.9
Entropy coefficient (-entropy-coef)	0.05, 0.01, 0.1, 0.5
Clipping epsilon (-clip-eps)	0.2, 0.4
Discount factor (-gamma)	0.995
GAE parameter (-lambda)	0.97
Value-function coefficient (-value-func-coef)	0.9, 0.5
Batch size (-batch-size)	64
Update interval (-update-interval)	1024, 2048, 5000

Table 4: Updated PPO hyperparameters used for tuning.

Hyperparameter	Values / Settings
Reward function (-reward)	min_max, variance_norm
Step size (-lr)	10^{-5} , 3×10^{-5}
Epochs (-epochs)	10, 15
Max gradient norm (-max-grad-norm)	0.5, 0.7, 0.9
Entropy coefficient (-entropy-coef)	0.05, 0.01, 0.1, 0.5
Clipping epsilon (-clip-eps)	0.2, 0.4
Discount factor (-gamma)	0.995
GAE parameter (-lambda)	0.97
Value-function coefficient (-value-func-coef)	0.9
Batch size (-batch-size)	64
Update interval (-update-interval)	5000

Table 5: Best hyperparameters for PPO on each minigame.

Category	Mini-game	Hyperparameters					
		Reward Function	Step Size	Epochs	Max Grad Norm	Entropy Coef	Clip Eps
Episodic	①	Min Max	10^{-4}	15	0.7	0.01	0.4
	②	Min Max	10^{-5}	10	0.7	0.01	0.2
	③	Min Max	10^{-5}	10	0.9	0.01	0.2
	④	Min Max	10^{-5}	15	0.9	0.05	0.4
	⑤	Min Max	10^{-4}	10	0.7	0.01	0.4
	⑥	Min Max	3×10^{-5}	10	0.5	0.05	0.2
Continuing	①	Min Max	10^{-4}	4	0.5	0.005	0.2
	②	Variance Norm	10^{-5}	10	0.5	0.05	0.4
	③	Min Max	10^{-4}	4	0.5	0.01	0.3
	④	Variance Norm	10^{-5}	10	0.5	0.1	0.4
	⑤	Min Max	3×10^{-5}	10	0.7	0.05	0.4
	⑥	Min Max	3×10^{-4}	15	0.7	0.05	0.2
Other Agents	⑦-Small-Sparse ¹	Min Max	3×10^{-5}	10	0.5	0.05	0.4
	⑦-Normal-Dense ²	Variance Norm	10^{-5}	10	0.9	0.1	0.4
	⑧-Small-Sparse ¹	Min Max	3×10^{-5}	10	0.5	0.05	0.2
	⑧-Normal-Dense ²	Min-Max	3×10^{-5}	10	0.7	0.05	0.2
Virus	⑨	Min Max	10^{-5}	10	0.7	0.01	0.4

¹ Uses a 200×200 arena with 200 randomly respawned pellets.² Uses a 350×350 arena with 500 randomly respawned pellets.

G.3 SAC Tuning

We swept over the hyperparameter combinations in Table 6. In SAC, we employed three separate networks. All network weights were initialized using LeCun normal initialization (LeCun et al., 2012). Table 7 summarizes the best-performing hyperparameters across our tasks. Also, we applied the hyperparameter settings from continual MINI-GAME 4 to train the full-game agent.

Table 6: Values of hyperparameters that we swept over when tuning SAC.

Hyperparameter	Values / Settings
Step size (<code>-lr</code>)	10^{-4} , 3×10^{-5} , 10^{-5}
Reward function (<code>-reward</code>)	<code>min_max</code> , <code>variance_norm</code>
Replay buffer size (<code>-replay-buffer</code>)	10^5
Soft update coefficient (<code>-tau</code>)	10^{-2} , 5×10^{-3} , 10^{-3}
Max gradient norm (<code>-max-grad-norm</code>)	0.5, 0.7, 0.9
Temperature step size (<code>-temperature-lr</code>)	10^{-4} , 3×10^{-4}
Update interval (<code>-update-interval</code>)	4
Batch size (<code>-batch-size</code>)	64
Discount factor (<code>-gamma</code>)	0.99

Table 7: Best hyperparameters for SAC on each minigame.

Category	Mini-game	Hyperparameters				
		Reward Function	Step Size	Soft Update Coefficient	Temperature LR	Max Grad Norm
Episodic	①	Min Max	3×10^{-5}	0.001	10^{-4}	0.7
	②	Min Max	3×10^{-5}	0.001	10^{-4}	0.7
	③	Min Max	10^{-4}	0.01	10^{-4}	0.7
	④	Min Max	10^{-5}	0.001	10^{-4}	0.7
	⑤	Min Max	10^{-5}	0.001	10^{-4}	0.7
	⑥	Min Max	10^{-4}	0.001	10^{-4}	0.5
Continuing	①	Variance Norm	3×10^{-5}	0.001	10^{-4}	0.5
	②	Variance Norm	3×10^{-5}	0.001	10^{-4}	0.7
	③	Variance Norm	3×10^{-5}	0.005	10^{-4}	0.9
	④	Variance Norm	10^{-5}	0.01	10^{-4}	0.5
	⑤	Variance Norm	10^{-5}	0.001	10^{-4}	0.7
	⑥	Variance Norm	3×10^{-5}	0.005	10^{-5}	0.9
Other Agents	⑦-Small-Sparse ¹	Variance Norm	10^{-4}	0.01	10^{-4}	0.9
	⑦-Normal-Dense ²	Variance Norm	10^{-4}	0.01	10^{-4}	0.7
	⑧-Small-Sparse ¹	Variance Norm	10^{-5}	0.001	10^{-4}	0.9
	⑧-Normal-Dense ²	Variance Norm	10^{-5}	0.005	10^{-4}	0.9
Virus	⑨	Min Max	10^{-5}	0.01	10^{-4}	0.7

¹ Sparse setting: 200×200 arena with 200 pellets.² Dense setting: 350×350 arena with 500 pellets.

H Mini-Game Results

In this section, we provide visualizations of the mini-games used in our evaluation (Section H.1) and the complete set of results mentioned in the main paper. Those include results in the continual pellet collection tasks (Section H.2), including analyses over the algorithms’ hyperparameter sensitivity (Section H.2.1) and results in those mini-games when augmenting PPO with a GRU (Section H.2.2). Additionally, we report results from mini-games in which the agent faced another bot (Section H.3) and those in which it was expected to successfully interact with viruses (Section H.4).

H.1 Illustrative Figures of the Pellet-Collection Mini-Games

The mini-games for pellet collection are divided into two sets: Square and Random. The Square mini-games require the agent to collect pellets along a square-shaped path, and it has three versions. The simplest version involves collecting pellets only, with no additional challenges. The second introduces mass decay, and in the third, the agent starts with a much bigger mass (1000 instead of 25). The set of mini-games with randomly scattered pellets uses the same variations. New pellets appear every 600 environment ticks. Figure 8 depicts all these variants.




Figure 8: Pellet Collection mini games. They were defined in terms of the pellet distribution in the arena, the agent's initial mass, and whether the agent's mass decays (which is hard to depict in an image). Note that we do not have eight mini-games because starting with a mass of 1,000 and no mass decay is an uninteresting setting. In each quadrant, the agent's initial view in the mini-game is shown on the left, and, to provide a sense of scale, a zoomed-out perspective of the same setting is shown on the right. The agent is depicted in blue.

H.2 Continual Problems

We discussed this setting in detail in the main paper; however, due to space constraints, we were unable to present the complete set of results there. They are available in Figure 9 below.




Figure 9: Performance of RL methods on *continual* pellet-collection mini-games. Panels (a), (b), and (c) show the performance on the *square-path tasks* (mini-games 1, 2, and 3), while (d), (e), and (f) show the performance on *randomly regenerated tasks* (mini-games 4, 5, and 6). Note that the y-axis scales vary across plots. The dashed line marks human performance, and the green line marks the random policy. The shaded region shows the 95% CI over 10 runs, computed using the t-distribution.

H.2.1 Hyperparameter Sensitivity Across Mini-games

We ran thousands of GPU jobs to tune and evaluate, on a per-environment basis, the baselines discussed in the previous section. Hyperparameter tuning is a major challenge in RL, and an even

bigger one in continual RL. Each algorithm is impacted differently by each hyperparameter. To test robustness, we used hyperparameters tuned in one mini-game to evaluate performance in another. In many tasks, PPO tuned on a particular mini-game often collapses when applied to a different mini-game. For SAC, using hyperparameters tuned for different mini-games sometimes yields better results than the hyperparameters optimized for the task itself! DQN, on the other hand, demonstrates surprising robustness—cross-task hyperparameter transfers have minimal impact on its peak performance. This large variability can be due to the complexities of hyperparameter selection strategies (Patterson et al., 2024), including the fact that we could afford only three seeds per configuration, which is certainly not enough for an accurate estimate. This is another common practice in the field that can be quite detrimental. The actual plots for these results are in Figure 10. These results underscore a critical insight: no single hyperparameter setting is robust across all tasks.

H.2.2 PPO augmented with a GRU

One might wonder whether approaches that tackle partial observability would be effective in AgarCL, especially in settings with a high degree of partial observability. As previously discussed, PPO is the best-performing algorithm among those we considered; accordingly, we decided to evaluate PPO augmented with a recurrent network in the pellet-collection MINI-GAMES.

Specifically, we evaluated PPO with Gated Recurrent Units (GRUs; Cho et al., 2014). The model comprises three convolutional layers, followed by a 256-unit GRU, and a final linear layer that splits into actor and critic heads. We re-tuned all hyperparameters for this new algorithm. All other experimental procedures were kept the same.

As shown in Figure 11, GRUs did not lead to much improvement. It performed well in relatively simple scenarios, such as MINI-GAME 1 (see panel (a)), achieving human-level performance. However, it struggled in more challenging tasks, such as mini-games 2 and 3 (panels (b) and (c)). Notably, in the uniform variation (panels (d)-(f)), the performance of PPO with GRUs was even worse than that of the non-recurrent version. Table 8 reports the equivalent numerical results.

Table 8: Performance of PPO with GRU-based recurrent architecture across six mini-games. Results are averaged over the last 100 data points. Values in parentheses denote standard deviations.

Algorithm	Mini-Game 1	Mini-Game 2	Mini-Game 3	Mini-Game 4	Mini-Game 5	Mini-Game 6
PPO	335 (3.0)	0 (0)	−975 (0)	22,463 (48.9)	4,018 (126.8)	5,064 (113.6)
PPO + RNN	484 (10.0)	0 (0)	−974 (0.3)	20,533 (250)	5,329 (30.9)	4,036 (22.4)

H.3 Interacting with Other Agents

Aside from pellet collection, capturing (or fleeing) other agents is a key skill an agent is expected to have in AgarCL. In this section, we evaluate the agent in mini-games where it interacts with a single other bot in the environment. To make the problem easier for the agent, given the difficulty they faced in the continual mini-games, we conducted experiments in an episodic setting, where each episode terminates either when the agent is eaten or after 10,000 time steps, whichever occurs first. We conducted these experiments in two configurations: the LARGE-DENSE and SMALL-SPARSE setups. We discuss both below.

H.3.1 Large-Dense Setup

This first setup uses a standard number of pellets (500) and a normal-sized arena (350×350). We used two types of fixed-policy bots: the *hungry* bot, which focuses solely on collecting pellets, and the *aggressive* bot, which prioritizes consuming any entity smaller than itself.

Figure 12 depicts the obtained results as well as the return obtained by the fixed-policy bot, as its return in comparison to the learning agent is quite informative. We can see that both DQN and




Figure 10: Performance of DQN, PPO, and SAC on six mini-games using cross-evaluated hyperparameters. The mini-games are grouped into two categories: the top three rows correspond to the **square path** group (Square-path pellet collection tasks), and the bottom three rows to the **random** group (randomly regenerated pellet tasks). For each mini-game, agents are evaluated not only using their own tuned hyperparameters but also using the hyperparameters optimized for the other two tasks within the same group. Each row shows the performance of a baseline algorithm across three hyperparameter configurations. For example, panels (a), (b), and (c) show DQN, PPO, and SAC in MINI-GAME 1 under different hyperparameter settings. The naming convention is as follows: SD-1000 denotes a Square path setting with mass Decay and an initial mass of 1000, while SnD-25 indicates no Decay and an initial mass of 25. Tuned is considered the best hyperparameter on a particular mini-game. We use R instead of S on the panels for the mini-games in which pellets were randomly spread in the environment (instead of a square).

SAC struggled against the aggressive, hungry bots. Gameplay videos recorded at various evaluation checkpoints revealed that these agents frequently became stuck in corners, a failure mode also observed in pellet collection mini-games. Consequently, the *aggressive* bot could quickly consume the learning agents, as reflected by the early plateaus in their return curves in panels (d) and (f). The performance of the *hungry* bot in panels (a) and (c) appears slightly better, focusing on collecting pellets rather than following the learning agents. The PPO agent exhibited more robust behaviour. Panel (e) shows that PPO learned to avoid the aggressive bot effectively while collecting many pellets. This evasive strategy contributed to the steady increase observed in the agent's return, albeit at a slower rate than its adversary. However, in the hungry setting (panel (b)), PPO failed to consume the bot,




Figure 11: Performance of PPO+GRU and PPO on *continual* pellet-collection mini-games. Panels (a), (b), and (c) correspond to the *square-path* tasks (with no mass decay, mass decay $m_0 = 25$, and mass decay $m_0 = 1000$, respectively), while panels (d), (e), and (f) correspond to the *uniformly regenerated* pellets with the same decay settings. The shaded regions show the 95% CI over 10 evaluation runs, computed using the t -distribution. Note that the y-axis scales differ between panels.




Figure 12: Performance of DQN, PPO, and SAC on *episodic* other-agent mini-games. Panels (a) and (d) show the performance of *DQN*, Panels (b) and (e) show the performance of *PPO*, and Panels (c) and (f) show the performance of *SAC*. The top row ((a)–(c)) corresponds to the *hungry-bot* task (mini-game 7), while the bottom row ((d)–(f)) corresponds to the *aggressive-bot* task (mini-game 8). Each experiment features a fixed-policy bot—either hungry or aggressive—in a standard-sized arena (350×350) with 500 pellets. Both the agent and the fixed-policy bot start with an initial mass of 25. Shaded regions represent the 95% CI over 10 independent runs, computed using the t -distribution.

even though the hungry bot did not attempt to attack it. This suggests that while PPO is effective at survival, it may not exploit opportunities to eliminate non-aggressive opponents.

These results raise the question of whether PPO’s success was due to the large arena and high pellet density. The SMALL-SPARSE setting, discussed next, addresses this question directly.

H.3.2 Small-Sparse Setup

To limit the agent’s ability to avoid interaction with bots, we evaluated them in a smaller arena (200×200) with fewer pellets (250). This configuration increases the likelihood of direct encounters between the learning agent and the fixed-policy bot.

Consistent with previous findings, the DQN and SAC baselines failed to demonstrate effective learning, irrespective of the type of bot in the environment (see Figure 13 on the next page). This is primarily attributable to recurring behavioural failures, such as becoming trapped in corners or fail-




Figure 13: Performance of RL baselines—DQN, PPO, and SAC—on *episodic* other-agent mini-games. Panels (a) and (d) show the performance of *DQN*, Panels (b) and (e) show the performance of *PPO*, and Panels (c) and (f) show the performance of *SAC*. The top row (a–c) corresponds to the *hungry-bot* mini-game, while the bottom row (d–f) corresponds to the *aggressive-bot* mini-game. Each experiment features a fixed-policy bot—aggressive or hungry—in a limited arena (200×200) with 250 pellets. The agent and the fixed-policy bot start with an initial mass of 25. Shaded regions represent the 95% confidence intervals over 10 independent runs, computed using the *t*-distribution.

ing to evade the bot, which persisted under these more restrictive conditions. Interestingly, the PPO agent also failed to learn in this environment, regardless of whether it was paired with the aggressive or the hungry bot. PPO could not consistently evade the aggressive bot or exploit the passive behaviour of the hungry one. In panel (e), PPO’s return clearly plateaus early, while the aggressive bot’s return steadily increases, indicating its ability to repeatedly eat the agent with ease. Even more interesting is PPO’s performance against the hungry bot. Despite the absence of an active threat, PPO failed to take advantage of the opportunity to survive and collect pellets. The consistent flattening of the agent’s learning curves across all settings further reinforces this observation, indicating minimal policy improvement over time.

H.4 Interacting with Viruses

A final skill we probe for is whether agents can learn how to leverage viruses in the environment, in the most benign setting possible. The mini-game is configured to be fully observable. As illustrated in Figure 14, the opposing bot remains stationary, and a clear, linear arrangement of viruses is positioned between the agent and the bot. Notably, there are no pellets in this mini-game. The agent




Figure 14: The learning agent (mass 3,000) is separated from a standstill bot (mass 5,000) by a line of static viruses (mass 100). The arena is fully observable with no pellets, mass decay, or virus respawning.




Figure 15: Performance of DQN, PPO, and SAC on the *virus-based* mini-game. Shaded regions show the 95% CI across 10 runs using the *t*-distribution.

begins with a mass of 3000 and must interact with a larger, stationary bot with a mass of 5000. To simplify dynamics further, mass decay is disabled. Training in this mini-game follows an episodic setup: an episode ends when one agent is eaten or after 1000 time steps.

The agent should learn to eject pellets toward the viruses, triggering them to split the stationary bot. The agent can then pass through the virus field and consume the smaller bot fragments. As shown in Figure 15, all learning agents failed to discover this strategy. Instead, they converged on a passive behaviour, remaining stationary throughout the episode. This highlights the difficulty these agents face in learning to leverage environmental elements, such as viruses, for strategic interaction, even in a simplified, fully observable setting.

I Numerical Results on Mini-Games and the Full Game

Table 9 summarizes the results of the experiments conducted in this paper. Each reported value represents the final performance, calculated as the average reward over the last 100 evaluation episodes.

Table 9: Performance across AgarCL minigames averaged over 10 independent runs (std. dev. is reported between parentheses). Results are averaged over the last 100 data points: episodes in episodic tasks and 100 steps in continual tasks.

Category	Mini-game	Scenarios	DQN	PPO	SAC	Human	Random
Pellet Collection (Episodic)	1	—	642 (14.2)	605 (25.3)	546 (6.0)	700 (0)	4.81 (1.3)
	2	—	382 (9.1)	398 (12.9)	285 (9.8)	682 (0)	0.81 (0.7)
	3	—	271 (23.1)	298 (28.8)	-16 (30.3)	612 (0)	-356 (22.7)
	4	—	1189 (44.2)	391 (36.3)	790 (98.6)	2876 (0)	4 (1.3)
	5	—	214 (13.2)	650 (58.9)	16 (8.4)	1600 (0)	0.17 (0.16)
	6	—	-426 (20.0)	812 (118.3)	-681 (12.8)	1452 (0)	-657 (6.9)
Pellet Collection (Continual)	1	—	619 (7.9)	335 (3.0)	419 (0.2)	700 (0)	0.01 (0.1)
	2	—	0 (0)	0 (0)	2 (1.7)	682 (0)	0 (0)
	3	—	-975 (0)	-975 (0)	-975 (0)	-480 (0)	-975 (0)
	4	Sparse ¹	—	22463 (48.9)	—	—	—
		Dense ¹	21402 (762.8)	21970 (277)	21997 (595.6)	23000 (0)	112.2 (12.5)
	5	Sparse	—	144 (10.8)	—	—	—
		Dense	132 (9.3)	4018 (126.8)	0.012 (0.9)	7215 (0)	0.05 (0.09)
	6	Sparse	—	-618.3 (1.4)	—	—	—
		Dense	-828 (10.5)	5064.4 (113.6)	-968 (0.01)	7215 (0)	-973 (1.1)
Other Agent (Episodic)	7—Small, Sparse ²	—	157 (0.7)	253 (35.6)	31 (0.2)	1980 (0)	0.015 (0.04)
	7—Large, Dense ²	—	146 (48.4)	429.3 (136.4)	25 (14.2)	2385 (0)	0.1 (0.16)
	8—Small, Sparse	—	189 (6.03)	267 (26)	114 (18.8)	1800 (0.015)	0.04 (0.04)
	8—Large, Dense	—	191 (14.56)	1037 (55.2)	69 (8.9)	2035 (0.015)	0.04 (0.04)
Virus (Episodic)	9	—	-36 (4.23)	-46 (6.45)	-5 (0.1)	870 (20.5)	-2319 (8.94)
Full Game (Grand Arena)	—	—	4 (3.1)	8 (8.3)	22 (11.1)	—	0.06 (0.18)

¹ In continuing pellet-collection tasks, “Sparse” refers to an arena with 250 pellets. “Dense” uses 500 pellets.

² “Small” arena: 200×200 with 200 pellets. “Large” arena: 350×350 with 500 pellets.

J How the Number of Pellets Impacts the Agent’s Performance

In this section, we further evaluate how PPO behaves in more challenging pellet-collection mini-games, where the number of available pellets is significantly reduced.

In Figure 9, PPO demonstrates strong performance in collecting pellets, even in the continual setting. However, the task remains relatively easy due to the high number of pellets (500), which allows the agent to find and consume pellets with minimal effort, even in the presence of mass decay. The

situation changes when the number of pellets is reduced to 250. Interestingly, PPO failed to learn effectively in MINI-GAMES ⑤ and ⑥, suggesting increased difficulty with exploration.

More generally, following the discussion on Section 4, in the context of designing easier settings in the full game such that we can see positive learning, we designed a variant of the environment in which the number of bots was decreased from 8 to 4, and we evaluated the impact of having fewer (400) or more (600, 1024) pellets than in the default setup. We maintained 10 viruses in the environment. We evaluated each agent for 86 million training steps.

As illustrated in Figure 17, increasing the number of pellets and reducing the number of bots does indeed make the task easier. Some of the curves depict a drop in performance later in training. While this is not surprising in continual RL due to issues such as loss of plasticity (Abbas et al., 2023; Lewandowski et al., 2023; Dohare et al., 2024), we did not investigate this phenomenon further nor evaluate any mitigation strategies. The important observation, despite the high variance across runs, is that this setting makes the problem more tractable, allowing us to consider the impact of freezing the agent’s policy.

These results motivated us to use 1,024 pellets in the experiment in which we froze the agent’s policy (in an easier environment so we could first observe learning).

K Continual Learning in AgarCL

In this section, we examine AgarCL from the perspective of continual RL. We first test the impact of freezing the agent’s policy in a different setting; then, we evaluate the performance of PPO with continual backpropagation (Dohare et al., 2024), an approach explicitly designed for continual learning.

K.1 Fixed

Policies vs Learning agent in AgarCL

In our earlier analysis, we identified performance degradation when training was interrupted at 32 million and 48 million steps in the setting with 8 bots and 1,024 pellets. We also evaluated the impact of freezing the agent’s policy in other scenarios. Specifically, we performed an additional experiment under the same conditions but with 4 bots and 500 pellets. We can also observe policy collapse for the agents frozen after 32M steps. We conjecture we would also observe policy collapse for the agent frozen after 48M steps if we waited long enough, but if not, this would be evidence that simpler settings (e.g., four bots) might make it easier for the agent to learn some non-trivial policy and maintain it, even though it is quite suboptimal.




Figure 16: PPO performance on *continuing randomly regenerated pellet-collection* in MINI-GAMES 4, 5, and 6. RD-1000 denotes a Randomly regenerated pellet setting with mass Decay and an initial mass of 1000; SnD-25 indicates no Decay and an initial mass of 25. The plateau in the RnD-25 curve after approx. 1M steps is due to the agent reaching the maximum cumulative reward in MINI-GAME ④. Shaded regions indicate the 95% CI over 10 independent runs, using reference values from the t-table.




Figure 17: PPO performance across varying pellet densities within a simplified version of the default setup. All experiments use 10 viruses and 4 bots; with a varying number of pellets, which is labeled in each curve. For example, P1024 denotes 1024 pellets. Results are averaged over 10 random seeds, and each curve is smoothed using a moving average with a window of 1000 steps.

K.2 Details about Experiments with Continual Learning Algorithms

The algorithms evaluated in Section 4.3 augmented PPO with Shrink and Perturb (Ash & Adams, 2020), ReDo (Sokar et al., 2023), or Continual Backpropagation (Dohare et al., 2024). In these three approaches, we used the same hyperparameters we used in PPO. We set the algorithm-specific hyperparameters according to the recommendations in the original papers. We detail those below.

Shrink and Perturb. We set $\alpha = 0.4$ (shrink), and $\beta = 0.2$ (perturb), so that, when appropriate, each parameter is updated according to $x' = 0.4x + 0.2x_0$, where x is the current weight and x_0 is a freshly initialized parameter. Shrink-and-Perturb is applied every 500,000 time steps.

ReDo. ReDo identifies and reinitializes inactive neurons in each linear or convolutional layer (excluding the output layer). For a representative minibatch, we record post-activation values, a , and compute per-unit activity scores, $s_j = \text{mean}(|a_j|)$, normalized as $\tilde{s}_j = s_j / (\bar{s} + 10^{-9})$. Units with $\tilde{s}_j \leq \tau$ are considered dormant. By default, we set $\tau = 0.05$.

ReDo (1) partially reinitializes incoming weights using LeCun normal or Kaiming uniform, (2) zeros outgoing connections to prevent contribution of dormant units, and (3) resets optimizer moments for affected parameters to avoid immediate re-amplification. The operation occurs every 2 million time steps and is computationally cheap compared with full retraining.

Continual Backpropagation. Continual Backpropagation (CB) tracks per-feature utility by monitoring activations and output weights, and periodically reinitializes features whose utility falls below a threshold. Replacement is performed via backward-pass hooks: fractional replacement counts accumulate until an integer replacement occurs, and reinitialized input weights are drawn uniformly within layer-specific bounds; corresponding output weights and normalization statistics are reset. Table 10 shows our default values.

Table 10: Continual Backpropagation hyperparameters.

Layer type	Replacement rate	Maturity threshold
Convolutional	1×10^{-5}	1000 steps
Linear	1×10^{-4}	100 steps

K.3 Full-Game Results with Continual Learning Methods

Table 11: Full-Game performance over the last 100 steps, averaged across 10 runs. Mean and standard deviation are reported for each algorithm.

Metric	Continual Backprop	ReDo	Shrink & Perturb	PPO
Mean	3675.8	4233.9	2074.2	762.58
Std	4381.81	5040.74	1868.18	1013.76




Figure 18: Performance of fixed-policy agents initialized from checkpoints at 32M and 48M steps. We report the moving average over 10 random seeds with a window size of 1000 steps.