



The Animal-AI Environment: A virtual laboratory for comparative cognition and artificial intelligence research

Konstantinos Voudouris¹ · Ben Slater¹ · Lucy G. Cheke¹ · Wout Schellaert² · José Hernández-Orallo² · Marta Halina³ · Matishalin Patel⁴ · Ibrahim Alhas⁵ · Matteo G. Mecattaf⁵ · John Burden⁵ · Joel Holmes⁵ · Niharika Chaubey⁵ · Niall Donnelly⁶ · Matthew Crosby⁷

Accepted: 19 January 2025 / Published online: 28 February 2025
© The Author(s) 2025

Abstract

The Animal-AI Environment is a unique game-based research platform designed to facilitate collaboration between the artificial intelligence and comparative cognition research communities. In this paper, we present the latest version of the Animal-AI Environment, outlining several major features that make the game more engaging for humans and more complex for AI systems. These features include interactive buttons, reward dispensers, and player notifications, as well as an overhaul of the environment's graphics and processing for significant improvements in agent training time and quality of the human player experience. We provide detailed guidance on how to build computational and behavioural experiments with the Animal-AI Environment. We present results from a series of agents, including the state-of-the-art deep reinforcement learning agent *Dreamer-v3*, on newly designed tests and the Animal-AI testbed of 900 tasks inspired by research in the field of comparative cognition. The Animal-AI Environment offers a new approach for modelling cognition in humans and non-human animals, and for building biologically inspired artificial intelligence.

Keywords Comparative cognition · Artificial intelligence · Cognition · Developmental psychology · Reinforcement learning · Computational neuroscience

Introduction

The Animal-AI Environment is the realisation of the idea that artificial intelligence (AI) research and cognitive science

can and should work together to answer questions about cognition, behaviour, and intelligence. The Animal-AI Environment was first released in 2019 (Beyret et al., 2019; Crosby, 2019), with a suite of (initially) held-out cognitive tasks (The *Animal-AI testbed*) that were the basis of the *The Animal-AI Olympics* competition in 2020 (Crosby et al., 2020). The environment is designed to offer a physically realistic *virtual laboratory* for building experiments inspired by research in comparative cognition, the field that studies and compares behaviour in non-human animals. In this paper, we present a major update to the environment, facilitating the development of a much wider range of behavioural experiments than was previously possible. The newest version of the Animal-AI Environment will interest researchers working on the development and evaluation of cognitively inspired AI, as well as comparative and developmental psychologists working on developing computational models of behaviour in humans and non-human animals.

The Animal-AI Environment has a simple structure, designed to emulate the classical set-up of psychological research on non-human animals. The agent is spawned in

✉ Konstantinos Voudouris
kv301@srcf.net

¹ Department of Psychology, University of Cambridge, Cambridge, UK

² VRAIN / ValGRAI, Universitat Politècnica de València, València, Spain

³ Department of History and Philosophy of Science, University of Cambridge, Cambridge, UK

⁴ Centre for Data Science AI and Modelling, University of Hull, Hull, UK

⁵ Leverhulme Centre for the Future of Intelligence, University of Cambridge, Cambridge, UK

⁶ Faculty of Environment, Science and Economy, University of Exeter, Exeter, UK

⁷ Department of Computing, Imperial College London, London, UK

a large arena, emulating an enclosure or laboratory, and it can be controlled by either a computational model or a human player. The objective is, within a specified time limit, to increase overall reward by obtaining rewarding objects and avoiding punishing ones. Rewarding objects correspond to appetitive stimuli such as food, while punishing objects correspond to aversive stimuli, such as electric shocks. The arena can be populated with a variety of other objects, such as ramps, blocks, tunnels, movable boxes, buttons, decoy goals, and dispensers. By using these objects as building blocks, an enormous variety of tasks can be generated, including setups similar to experiments from animal behaviour research (Crosby et al., 2020; Voudouris et al., 2022). Thousands of distinct tasks can be easily procedurally generated to build comprehensive test batteries (Crosby, 2020). The Animal-AI Environment is accompanied by several thousand open-source implementations of common experimental designs used in animal behaviour research, which can be freely used and adapted.

This article is divided into discrete sections; each is self-contained and can be read independently depending on the reader's interest. In Section “[Collaborative interdisciplinary research with Animal-AI](#)”, we present several advantages of using the Animal-AI Environment for researchers in AI, comparative cognition, and those working at the intersection of these fields. We place the Animal-AI Environment in the context of other environments and initiatives in these fields, arguing that the Animal-AI Environment is unique in being the only initiative that serves both these communities at once, encouraging close collaboration. In Section “[The Animal-AI Environment: A comprehensive overview](#)” we present a full reference for the structure of the environment, including details of its features and how to run computational and behavioural experiments in the environment. In Section “[A battery of agents](#)”, we outline the key agents that are implemented in the environment, serving as a starting point for research in AI and comparative cognition research. Finally, we present a short series of experiments to showcase the environment, these agents, and their potential for collaborative research.

Collaborative interdisciplinary research with Animal-AI

The Animal-AI Environment has been developed to facilitate collaborative and interdisciplinary research between artificial intelligence (AI) and comparative cognition. Comparative cognition offers the field of AI a wealth of experimental materials for studying and measuring intelligence in non-human systems (Crosby, 2020; Hernández-Orallo, 2017), providing inspiration for the development of better, more capable systems. Meanwhile, AI research offers comparative cognition

the ability to precisely model learning and behaviour computationally in physically realistic environments that emulate laboratory conditions. Computational modelling of human behaviour has already significantly improved our understanding of human cognition and behaviour (Griffiths, 2015; Guest & Martin, 2021; Zuidema et al., 2020), but computational modelling of non-human animal cognition is relatively nascent (Voudouris et al., 2024).¹ Collaboration between AI and comparative cognition can not only improve how we study the capabilities of AI systems, but also render more precise the theories and hypotheses we develop about animal behaviour (Voudouris et al., 2024).

A fundamental roadblock is the lack of a common platform for conducting comparable research on non-human animals and computational models. The Animal-AI Environment aims to overcome this obstacle. Primarily, the Animal-AI Environment is designed for building tests of spatial and physical cognition, including goal-directed navigation, object permanence, and tool use (Crosby, 2019). Physical cognition can be defined as any behaviour involving the manipulation of a non-social physical environment (Lake et al., 2017). The agent in the Animal-AI Environment is therefore designed to take the place of an animal inside the laboratory, interacting with its environment in return for rewards while avoiding punishments, such as aversive stimuli.

Having a virtual laboratory for conducting behavioural tests on computational models encourages a virtuous cycle between AI and comparative cognition. Figure 1 presents a simplified schematic of this cycle. Researchers start by replicating a comparative cognition experiment in the Animal-AI Environment, ranging from tests of instrumental learning to tests of tool use and causal reasoning. They can then evaluate the performance of computational models on that test, either the baseline systems that we provide or those they design through the standardised Python interface. How these models perform on realistic, lab-like tasks informs researchers of their flaws or merits, encouraging the development of new hypotheses, theories, and experimental designs, which can then be developed in the physical laboratory with non-human animals. These new designs can then be re-implemented in the Animal-AI Environment. Not only can this process improve and render more precise the theories and models used to explain animal cognition, but it can also inspire improvements in the architectures and methods of AI (Freire et al., 2024; Mitchener et al., 2022).

Figure 2 presents an example of how Animal-AI can be the starting point for this interdisciplinary research. This experiment tests object permanence, the capacity to track occluded objects. It is a replication of classic tests of object permanence. A rewarding object is hidden in a cup, and the

¹ See Brea et al., 2023; Baum et al., 2022 for recent examples of this.




Fig. 1 The Animal-AI Environment is a research tool that facilitates a virtuous cycle between AI research and animal behaviour research

subject must correctly identify its location. Versions of this task have been conducted with infants (Piaget, 1923; Piaget & Inhelder, 1969), several species of primate (e.g., Herrmann et al. (2007)), and four species of parrot (Krasheninnikova et al., 2019). Comparative cognition has innovated several paradigms for detecting whether non-human animals are capable of object permanence, and these paradigms are immediately applicable to detecting whether embodied AI systems are capable of it too (Voudouris et al., 2022). However, there remains significant controversy on *how* non-human animals (and infants) track occluded objects (see Bremner et al. (2015); Jaakkola (2014); Scholl (2007)). As Mareschal (2000) argues, computational modelling can help to elucidate the mechanisms underpinning object permanence. Indeed, in recent years, significant progress has been made towards AI systems that can track and reason about occluded objects (Team et al., 2023; Piloto et al., 2022; Tok-

makov et al., 2021). The Animal-AI Environment allows the evaluation of these sophisticated computational models on physically realistic experiments drawn from comparative cognition research, encouraging progress towards better theories about how animals track occluded objects *and* towards building AI systems that reliably do it, too.

To facilitate collaborative and interdisciplinary research, the Animal-AI Environment has the following design principles:

1. A three-dimensional environment, with accurate and fast physics.
2. Straightforward development of custom experiments, as well as procedural generation of variations of them.
3. Minimal hardware requirements, so that behavioural experiments can be run locally or in the browser.
4. Well-defined, functionally specified objects for recreating existing experiments of physical cognition and for innovating new ones.
5. Self-contained and relatively accessible to cognitive scientists without significant training in computer science or software engineering.
6. Integration with existing frameworks for computational modelling and reinforcement learning, such as Gymnasium (Towers et al., 2023) and stable-baselines-3 (Raffin et al., 2021).
7. Simple action space, to reduce the computational complexity of the learning problem, and to facilitate quick training and ease-of-use with adults, children, and non-human animals (such as primates).
8. Appealing to play as a participant, facilitating collection of large and detailed human behavioural datasets.
9. Openly available and free to use.

The Animal-AI Environment is unique in its mission to serve interdisciplinary research between comparative cognition and AI. For example, while Allritz et al. (2022; see also Miscov et al. (2024)) present an example of a virtual environment, *AP Explorer 3D*, for collecting data from primates on navigation and short-term memory tasks, they do not attempt




Fig. 2 An example of a task testing object permanence in the Animal-AI Environment. *Left*: Bird's-eye view of the arena, with the location of the agent indicated by the grey arrow. *Right*: The agent's view from

its starting position. It is frozen while the rewarding green sphere drops down behind the blue wall (see Voudouris et al. 2022, 2024)

to test computational models with this environment. Indeed, navigable virtual environments (i.e., open-world first-person video games) like the Animal-AI Environment have only very recently started to be used in the human cognitive sciences. These types of games, where players are embodied in a virtual avatar that can take several different actions and interact with a complex environment, offer numerous advantages such as increasing ecological validity by more closely approximating real-world environments than traditional experiments (Allen et al., 2023; Allritz et al., 2022). Video games facilitate more complex and nuanced experimental designs, as well as tapping into the popularity of computer games. In recent years, video games have been successfully used both for collecting rich human behavioural data and for computationally modelling that data (e.g., Ansarinia et al., 2022; Brändle et al., 2023; Brown et al., 2014; Coutrot et al., 2018; Coughlan et al., 2019; Coutrot et al., 2022; Dubey et al., 2018; Leduc-McNiven et al., 2018; McNab et al., 2015; Pouncy and Gershman, 2022; Rafner et al., 2022; Schulz et al., 2019; Stafford and Dewar, 2014; Voiskounsky et al., 2017; see Allen et al., 2023 for a review). The Animal-AI Environment offers a new opportunity to make similar progress at the intersection of comparative cognition and AI.

In contrast to comparative cognition and the cognitive sciences, there are several environments similar to the Animal-AI Environment that have been developed in the AI community. Most similar to the Animal-AI Environment in being three-dimensional environments with realistic physics are MuJoCo (Todorov et al., 2012), DeepMind Lab (Beattie et al., 2016; Leibo et al., 2018), Minecraft-based environments (Malmö Johnson et al. (2016), MineRL Guss et al. (2021), MineDoJo Fan et al. (2022)), MiniWorld (Chevalier-Boisvert et al., 2024), XLand (Team et al., 2023), Avalon (Albrecht et al., 2022), and ThreeDWorld (Gan et al., 2020, 2021). The Animal-AI Environment dis-

tinguishes itself from these environments in being the only one targeted directly at the intersection between comparative cognition and AI, as well as being freely available, compatible with contemporary machine learning libraries in Python, facilitating customisable experimental designs, and being accompanied by an existing library of several thousand implementations of classic comparative cognition experiments (see Table 1). Similarly to XLand and ThreeDWorld, the Animal-AI Environment is built on top of Unity's physics engine (Juliani et al., 2018), providing fast real-world physics and three-dimensional graphics (Ward et al., 2020). We also make use of Unity's open-source `ml-agents` Python API (Juliani et al., 2018), which is integrated with existing reinforcement learning frameworks, such as Gymnasium (Towers et al., 2023), Dopamine (Castro et al., 2018), and `stable-baselines-3` (Raffin et al., 2021). Unlike these environments, and similarly rich environments such as Avalon, the Animal-AI Environment is more constrained in its scope. While these other environments are invaluable for the development of large, data-driven computational models, they often demand expensive hardware and large research teams. The Animal-AI Environment is compact and simple enough to be used by small teams of AI researchers and cognitive scientists alike, while also scaling well to larger experiments. Nevertheless, Animal-AI's physics and visuals are rich enough to facilitate complex experiments that are still too difficult for state-of-the-art systems (Crosby et al., 2020; Voudouris et al., 2022; Voudouris et al., 2024; see also Section "The Animal-AI testbed").

Most similar to the Animal-AI Environment in scope and focus are DeepMind Lab, MiniWorld, and the Minecraft-based builds MineRL and Malmö. DeepMind Lab is an environment for training and testing autonomous agents on a range of tasks. It is designed to have convincing three-dimensional physics and simple visuals, similar to

Table 1 A comparison between the Animal-AI Environment and other realistic 3D environments in AI and cognitive science

Name (First published)	Community	Freely available	Custom designs	Animal experiment library	Actively maintained
MuJoCo (2012)	Robotics	✓	✗	✗	✓
DM Lab (2016)	AI	✓	✓	✗	✗
Malmö (2016)	AI	✓	✓	✗	✗
MiniWorld (2018)	AI	✓	✓	✗	✓
MineRL (2019)	AI	✓	✓	✗	✓
XLand (2021)	AI	✗	?	?	?
Avalon (2022)	AI	✓	✓	✗	✗
MineDoJo (2022)	AI	✓	✓	✗	✗
ThreeDWorld (2020)	AI, Human Psych.	✓	✓	✗	✓
APExplorer 3D (2022)	Animal Psych.	✓	✗	✓	?
Animal-AI (2019)	AI, Animal Psych.	✓	✓	✓	✓

Note: Maintenance details are correct at the time of publication, with sources presented in Appendix A

the Animal-AI Environment. It was designed to facilitate research on general and autonomous agents, and to build tests inspired by neuroscience. Indeed, DeepMind Lab has been used to study exploration and learning in young children, inspired by work in developmental psychology (Kosoy et al., 2020), as well as to improve reinforcement learning architectures (Leibo et al., 2018). However, DeepMind Lab requires experiments to be designed in Lua, a language that is much less familiar to researchers than Python, which is how users interact with the Animal-AI Environment (see Chevalier-Boisvert et al., 2024). Minecraft-based builds are similar to DeepMind Lab and the Animal-AI Environment in having a sophisticated physics engine and a definable ontology (i.e., blocks of different kinds that can populate the world). Mini-World offers a three-dimensional environment with realistic physics for building customisable experiments for reinforcement learning systems. However, none of these environments are directly targeting comparative cognition research. The ontology of the Animal-AI Environment has been designed by comparative cognitive scientists to offer objects that are directly relevant to the methodologies of non-human animal behaviour research, including buttons, ramps, decaying and ripening rewards, and tools.

The Animal-AI Environment is accompanied by a library of experiments adapted from the comparative cognition literature, including 900 tasks testing spatial reasoning, working memory, numerosity, and tool-use (Crosby et al., 2020), over 12,000 tasks testing object permanence (Voudouris et al., 2022), over 5000 tasks testing object affordances (Rutar et al., 2024), and over 260,000 basic tasks for interacting with the environment. The environment has been used to develop more sophisticated reinforcement learning architectures (Freire et al., 2024; Lehuger & Crosby, 2021; Mitchener et al., 2022; Zakharov et al., 2021) as well as make progress on computational models of learning in non-human animals (Lind & Vinken, 2021; Cardoso et al., 2023). It also facilitates direct comparison to human participants, enabling the evaluation of the state-of-the-art in reinforcement learning (Burden et al., 2023; Voudouris et al., 2022, 2024) and multi-modal large language models (Mecattaf et al., 2024). The Animal-AI Environment is therefore at the foundation of a rich and active research community at the intersection of comparative cognition and AI.

The Animal-AI Environment: A comprehensive overview

In this section, we provide a comprehensive outline of the Animal-AI Environment, highlighting the changes that have been made since earlier releases (Beyret et al., 2019; Crosby et al., 2020).

Environment design and modes of use

The Animal-AI Environment has two principal modes of use:

- Play mode: A full visual display for humans, enabling interaction with simple keyboard controls.
- Training mode: A lightweight Python API for training and testing computational agents on visual output from the Animal-AI Environment.

The Animal-AI Environment can be configured to create episodic tasks for humans and artificial agents. Each episode takes place inside an arena that can contain a wide range of different objects together with the agent. More information on all of these features and more is described in the following sections. As noted, the Animal-AI Environment is built on top of Unity's `ml-Agents` library (Juliani et al., 2018).²

Substantial graphical improvements have been made to the Animal-AI Environment—not only cosmetic changes to improve the environment's appearance (textured walls, animal skins, improved lighting) but also efforts to optimise environment rendering for faster performance. This includes pruning unnecessary casting and receiving of shadows; making use of the performance-orientated Universal Render Pipeline (URP);³ simplifying the shaders and lighting model used to render 3D objects; and reducing the resolution and anti-aliasing of textures and shadows. These optimisations complement efforts to create a more vibrant, engaging, and visually straightforward environmental appearance for human players, particularly children, while also increasing the frames per second (fps) at which the environment can run when using camera-based observations (see Section “[Using Animal-AI with computational models](#)”).

Environment specification

The arena & agent

As in earlier versions of the environment, the arena is a 40×40 unit space (where the agent is a sphere with a diameter of 1), including a tiled floor and four surrounding walls resembling white picket fences.⁴ The `Arena` object can be configured with three parameters. The `arena time`, `timeLimit`, provides the total number of time steps an agent or user can take before the episode ends. The reward for an episode begins at 0 and decreases in increments of $\frac{1}{t}$

² We provide an accompanying `animalai` Python package here: <https://pypi.org/project/animalai/>.

³ See <https://unity.com/srp/universal-render-pipeline>.

⁴ Note that there is no explicit height restriction on the arena.

per timestep, where t is the value of `timeLimit`, increasing when the agent interacts with a positively valenced object (see “[Valenced objects](#)” section) and decreasing upon interaction with a negatively valenced object. If `timeLimit` is set to 0, then the reward value no longer decreases with each time step. The arena pass mark, `passMark`, allows the researcher to define the reward threshold that should constitute a ‘pass’ on that particular episode, providing an easy metric for success or failure across episodes. It also determines the notifications (if turned on) that players receive at the end of an episode (see Section “[Using Animal-AI with human participants](#)”).

The agent is of size $1 \times 1 \times 1$ and can spawn anywhere in the Arena. It can take one of three forms: a hedgehog, a panda, or a pig, enabling the creation of an engaging narrative for studies with children (see Fig. 3). The agent can be frozen at the start of an episode for a specified number of steps, using the `frozenAgentDelays` parameter, a new feature in the Animal-AI Environment. There is no reward decrease during the frozen period, and all other objects continue to move and interact. This allows the agent to observe occurrences in the environment before being allowed to move.

To replicate the view of an animal in a laboratory, the agent observes the environment from a first-person perspective, both in “Play Mode” and “Training Mode”. Two further perspectives are offered in “Play Mode”. ‘Third-person view’ follows the agent around from a perspective slightly above and behind the agent. ‘Birds-eye view’ offers a view of the arena in its entirety from above. These perspectives are offered so that researchers can view their experimental designs from multiple perspectives. Researchers can specify whether players in “Play Mode” can change their camera angle, and which camera angle should be the default. In “Training Mode”, the API only returns the ‘First-person




Fig. 3 The Arena and the three Agent characters that can be selected for play and testing. *Top*: “hedgehog”. *Centre*: “panda”. *Bottom*: “pig”

view’ to simulate the perspective of an animal in the laboratory. The number of pixels or raycasts (another form of visual observation; see Section “[Using Animal-AI with computational models](#)”) in this image is configurable.

The Arena object can be configured to have ‘lights out’ periods, in which all pixels or raycasts in the ‘first-person view’ are replaced with zeros. For a human player, this results in a black screen. Lights out periods can be specified for specific time step intervals, or on an alternating schedule.

Immovable objects

As in earlier versions of the environment, there are five objects that cannot be moved by the agent: opaque and transparent tunnels and walls, and opaque ramps. All of these objects can be resized, rotated, and re-positioned freely by experiment designers via a configuration file. All opaque objects can be varied in colour in a $256 \times 256 \times 256$ RGB colour space (see Fig. 4). The characteristics of these five immovable objects are:

- **Opaque Walls:** An opaque cuboidal object used to construct obstacles to be navigated around. The colour of the walls can be changed. In the Animal-AI Environment, unlike in some earlier versions, opaque walls have a “brickwork” texture of black lines.
- **Transparent Walls:** A non-opaque cuboidal object used to construct obstacles to navigate around.
- **Ramps:** An opaque right-angled prism object. They can be climbed on/over, if the height-length ratio does not exceed 4:1. In the Animal-AI Environment, unlike in earlier versions, ramps have a patterned texture of lighter and darker colour.




Fig. 4 The immovable objects in the Animal-AI Environment. *Top left*: `CylinderTunnelTransparent`. *Top right*: `CylinderTunnel`. *Bottom left*: `Ramp`. *Bottom centre*: `Wall`. *Bottom right*: `WallTransparent`

- **Opaque Tunnels:** An opaque hemi-cylindrical tunnel used as an obstacle to be navigated through and/or around.
- **Transparent Tunnels:** A non-opaque hemicylindrical tunnel used as an obstacle to be navigated through and/or around.

Movable objects

There are six movable objects in the Animal-AI Environment: hollow boxes, light and heavy rectangular blocks, U-shaped, L-shaped, and J-shaped blocks, all of which can be pushed around. These blocks can be used as tools for moving other objects (other movable objects or any of the valenced rewards discussed below). These objects can be resized, rotated, and re-positioned freely, but their colour cannot be changed. In the Animal-AI Environment, unlike in earlier versions, blocks do not have a “cardboard box”- or “wood”- style texture, instead being uniform grey in colour with light grey trims along the vertices (Fig. 5).

- **Light Blocks:** An opaque cuboidal object. Light blocks have a mass of 1.
- **Heavy Blocks:** An opaque cuboidal object. These objects function the same as light blocks, except that heavy blocks have a mass of 2.
- **U-Shaped Blocks:** An opaque object with a squared “U” shape that facilitates pushing and pulling of other movable objects. U-shaped blocks have a mass of 1.5.
- **L- and J-Shaped Blocks:** Opaque objects with a squared “L” or “J” shape (mirror images of each other) that facilitate pushing of other objects. L- and J-shaped blocks have a mass of 1.5.
- **Hollow Boxes:** Opaque cuboidal hollow objects with an opening at the top, allowing them to contain other objects. They have a mass of 1.5.

Valenced objects

Upon making contact, positively and negatively valenced objects can increase or decrease reward respectively. These objects vary in both size, shape, and appearance. The colour

of valenced objects can be configured by the researcher, with the default colours being those presented in Fig. 6. Decaying, ripening, growing, and shrinking valenced objects are new to the Animal-AI Environment. In what follows, we refer to spherical valenced objects as “goals” and cuboidal regions of space that have valence as “zones”.

- **Non-Moving Goals (episode- and non-episode-ending):** Coloured spheres with a valence proportional to their size. They only move under the force of gravity or collision. These objects can be positive or negative and episode ending in the cases of *GoodGoal* and *BadGoal*, positive or negative and non-episode ending in the cases of *GoodGoalMulti* and *BadGoalMulti*, or neutral and non-episode ending in the case of *DecoyGoal*. They are solid and opaque and their colours can be changed (with the defaults shown in Fig. 6).
- **Moving Goals (episode- and non-episode-ending):** Identical to non-moving goals, except that they have an initial velocity in a prespecified or randomised direction at the start of the episode. They are denoted with the suffix ‘-Bounce’. They are solid and opaque.
- **Growing and Shrinking Goals:** Similar to ripening and decay goals, respectively, except their physical size changes proportional to their valence. The start and end size can be configured. Growth/Shrinkage can be delayed for a specified number of steps. Collecting these goals does not end an episode. A growth goal will cease to grow if the space around it becomes restrictive. They are solid and opaque.
- **Decaying and Ripening Goals:** Similar to non-moving goals, except that their valence changes over time, while their physical size remains constant from spawn time. The valence decreases over time for decaying goals, and increases with ripening goals. The difference is indicated by six two-dimensional glyphs on the goal’s surface, as well as colour changes. Decaying goals change from yellow to grey, while ripening goals change from grey to yellow. Decaying goals have glyphs of arrows pointing inwards, with a large black ring that empties over time. Ripening goals have glyphs of arrows pointing outwards, with a ring that fills over time. Both goals can have a con-




Fig. 5 The movable objects in the Animal-AI Environment. From *left to right*: *LightBlock*, *HeavyBlock*, *UBlock*, *LBlock*, *JBlock*, *HollowBox*




Fig. 6 The valenced objects in the Animal-AI Environment. *Top far left:* GoodGoal/GoodGoalBounce. *Top centre left:* GoodGoalMulti/GoodGoalMultiBounce/GrowGoal/ShrinkGoal. *Top centre:* BadGoal/ BadGoalBounce. *Top centre*

right: BadGoalMulti/BadGoalMultiBounce. *Top far right:* DecoyGoal/DecoyGoalBounce. *Bottom far left:* RipenGoal. *Bottom centre left:* DecayGoal. *Bottom centre right:* DeathZone. *Bottom far right:* HotZone

figurable “full” and “empty” valence and a configurable delay between their starting valence (full or empty) and their ultimate valence (empty or full). Collecting these goals does not end an episode. They are solid and opaque.

- **Death Zone:** A translucent red zone. When the agent makes contact with a death zone, its reward decreases by 1 and the episode ends. In the Animal-AI Environment, unlike in earlier versions, death zones occupy a three-dimensional region of space, through which other objects, including the agent, can pass.
- **Hot Zone:** A translucent orange zone. When the agent enters the hot zone, reward decreases $10\times$ faster than with time alone. It does not end the episode upon contact. In the Animal-AI Environment, unlike in earlier versions, hot zones occupy a three-dimensional region of space, through which other objects, including the agent, can

pass. If a hot zone and death zone overlap, the death zone takes precedence.

Dispensers

Dispensers are designed to emulate natural (e.g., tree spawners) and laboratory-like (e.g., goal dispensers) elements that dispense valenced objects, allowing replication of various human and non-human animal experiments. All of these objects are new to the Animal-AI Environment. See Fig. 7.

- **Spawner Tree:** A tree-shaped object that spawns GoodGoalMulti objects, which initially grow in the tree branches, then fall to the ground after a configurable time delay. The number of GoodGoalMulti spawned can be varied, and are dropped from the tree at random, in




Fig. 7 The dispenser objects in the Animal-AI Environment. From left to right: SpawnerDispenserTall producing unlimited GoodGoalMulti, SpawnerDispenserShort producing

unlimited GoodGoalMulti, SpawnerTree producing unlimited GoodGoalMulti, SpawnerButton producing unlimited GoodGoal, GoodGoalMulti, and BadGoal with equal probability

the style of naturally falling fruit. Tree Spawners have a fixed size and colour.

- **Tall Spawner Dispenser:** An immovable, cuboidal object designed to look like a simplified vending-machine/food magazine that spawns `GoodGoalMulti` through a hatch. The colour of the dispenser, the number of `GoodGoalMulti` spawned, and the delay between spawning can be configured. The size of the dispenser is fixed.
- **Short Spawner Dispenser:** A more compact version of the Goal Dispenser that also spawns `GoodGoalMulti` into the environment after a configurable time delay through a hatch. The size of the dispenser is fixed.
- **Spawner Button:** An opaque, immovable object with fixed size and colour (beige pillar, blue button). The `SpawnerButton` object can spawn three types of valenced object, `GoodGoal`, `GoodGoalMulti`, and `BadGoal`. The probability of a goal spawning upon pressing the button can be defined. The probability of spawning each of these three goal types can also be weighted, to increase or decrease the likelihood of it appearing when the button is pressed. The position of where the goal is spawned can also be defined. Finally, the number of steps it takes to reset the button can be defined.

Sign boards

Sign boards, a new addition to the Animal-AI Environment, enable the presentation of two-dimensional neutral stimuli to signal, for example, the presence or absence of valenced objects. They allow the replication of conditioning experiments common in cognitive science (see, e.g., Cardoso et al. 2023) as well as for general signposting (e.g., with arrows). See Fig. 8.

- **Sign board:** An opaque, cuboidal, immovable object that displays a two-dimensional visual image. They have a fixed size and a preset list of display symbols (including arrows, letters, and shapes), which can be used to non-linguistically signal events in the environment. Pixel grids can also be presented, with full or partial randomisation. Pixel grids, as well as the preset list of display symbols, can be given user-defined colours on the RGB spectrum.

Physics simulation

The Animal-AI Environment evolves in discrete time. The state of the environment at a certain time step is the complete set of information about the objects and agent that exists




Fig. 8 The sign board object in the Animal-AI Environment. *Top:* SignBoards with five preset symbols (*left to right:* "tick", "u-turn-arrow", "circle", "letter-a", "left-arrow"). *Bottom:* SignBoards with defined and randomised black grids. The syntax for defining these is overlaid

within it. This encompasses a multitude of factors, including the position of objects, the agent's position and orientation, time step information, velocities, contact with objects, and many other relevant factors. At each step, the agent can take one of nine discrete actions, transitioning the environment to its next state, with the default action being *No Action*:

- **No Action**: the agent has no force acting on it in any direction, apart from the force of gravity or of other objects colliding with it.
- **Forwards**: the agent has force acting on it in the forwards direction (towards the observable part of the environment).
- **Left**: the agent rotates 6 degrees left (negative offset).
- **Right**: the agent rotates 6 degrees right (positive offset).
- **ForwardsLeft**: the agent has force acting on it in the forwards direction and rotates 6 degrees left.
- **ForwardsRight**: the agent has force acting on it in the forwards direction and rotates 6 degrees right.
- **Backwards**: the agent has force acting on it in the backwards direction (away from the observable part of the environment).
- **BackwardsLeft**: the agent has force acting on it in the backwards direction and rotates 6 degrees left.
- **BackwardsRight**: the agent has force acting on it in the backwards direction and rotates 6 degrees right.

How the environment transitions from one state to the next is defined by the environment's transition function, and it depends on the preceding states of the environment and the action that the agent takes on that step. The Animal-AI Environment is built with Unity's physics engine, which has realistic Newtonian physics and Euclidean geometry that define the transition function. There are several forces that act on the objects in the environment that derive from the Unity physics engine:

- **Gravitational Force**: Every object experiences a constant force downwards (on the y-axis), simulating gravity.
- **Normal Force**: An object placed on a surface experiences an upward force equal and opposite to the gravitational force.
- **Applied Force**: When the agent moves forwards or backwards, or a `GoodGoalBounce`, `GoodGoalMultiBounce`, or a `BadGoalBounce` move in a direction, this is due to an applied force in that direction.⁵
- **Reaction Force**: When objects collide into each other, the object being collided with receives a force in the opposite direction to the collision path, proportional to its mass.

⁵ In the case of moving goals, this force only applies on the first time step.

- **Frictional Force**: There is a frictional force that opposes the direction of motion. All objects in the environment have high frictional force.
- **Drag Force**: Objects moving in the environment experience air resistance (drag) which opposes their movement proportional to their size.

Reward signal & health

The value of the reward signal can be positive or negative. An agent accrues or loses reward through touching positively or negatively valenced objects. Additionally, the agent receives a minor negative reward after each step, which is equivalent to $-\frac{1}{t}$ per timestep, where t refers to the number of timesteps in the episode. In a physical laboratory, food is often withheld from animals prior to an experiment so that they are motivated to obtain food as quickly as possible.⁶ Including a reward decrement on each timestep emulates this necessity to obtain food in the laboratory. However, sometimes this design is not required. Therefore, when t is 0, there is no time step limit, and reward does not decrement on each time step.

As an interpretable indicator of current reward for human players, we include a health indicator. This starts at 100 and the agent fails the episode when it reaches 0. The rate of health decrement is proportional to reward decrement. Health is replenished upon obtaining positively valenced objects proportional to the object's value, to a maximum value of 100.

Using Animal-AI with computational models

The Animal-AI Environment has a straightforward Python API for integration with computational models, based on the `ml-agents` library (Juliani et al., 2018). The environment returns information about its state at each time step, including the agent's first-person observations, position, velocity, and current reward, and the model must return one action from the set of permissible actions in order to evolve the environment. Thus, any model implemented in Python that returns single actions in response to observations can be integrated into the Animal-AI Environment.

A popular class of computational models that take observations and return actions are reinforcement learning agents, based on algorithms that seek to find sequences of actions that optimise reward. Reinforcement learning agents can be implemented directly in the `ml-agents` Python API. Alternatively, the environment can be wrapped to form a `Gymnasium` environment (Towers et al., 2023). This enables researchers to make use of the myriad of reinforcement learning agents that have been implemented elsewhere

⁶ Note that this is not always the case. For instance, in contemporary great ape studies, highly desirable and not readily available rewards are used.

in the literature, as well as build new agents according to standardised interfaces which can interact both with the Animal-AI Environment and the many other similar environments that exist.

Since the Animal-AI Environment is partially observable, the agent does not have access to complete information about the current state of the environment. The extent of the observability depends on the observation type. *Camera* observations allow for the closest replication of animal behaviour in the laboratory, allowing the agent to observe an image of the environment in front of them. The resolution of the image can be set to $k \times k \times 3$ where $4 \leq k \leq 512$ for the RGB colour channels. Grayscale can be turned on or off. With grayscale observations, the camera observation is simplified to a $k \times k$ matrix. To produce a simplified observation space, the agent can also make use of *raycast* observations. The agent can emit ‘rays’ along its x - z (horizontal) plane which intersect objects and return information about what objects the rays hit and how far away they are. The first ray is emitted directly ahead, with further rays fanned out equally, starting on the left. The number of rays, and the angle over which they are distributed, can be specified. Raycast observations return a matrix of $8 \times r$ where r is an odd number. Raycast observations allow for quicker learning and incur less computational overhead than camera observations. Finally, the agent can receive observations about the agent’s current health, its velocity in (x, y, z) , and its global position in (x, y, z) . All of these observation types can be used in any combination.

Using Animal-AI with human participants

Human behavioural testing is facilitated in the Animal-AI Environment using “play mode”. The participant interacts with the environment using the arrow-keys, [W, A, S, D], or with a controller connected to the testing system. Participants are provided with a high-resolution 1200×800 pixel display of the environment at 60fps, from one of three perspectives, with their ‘health’ in the bottom-left corner. ‘Health’, a new feature of the Animal-AI Environment, starts at 100 and steadily decreases with time. It is replenished when reward increases. Obtaining a reward of size 0.5 increases health by 50. Health cannot be over 100. The ‘Health’ bar turns continuously from blue through green to orange and finally red as it runs from 100 to 0. Participants are provided with the current reward value and the reward value at the end of the previous episode in the top-left corner. At the end of an episode, researchers can choose whether the participant is presented with a notification of whether they passed or failed. See Fig. 9.

Researchers can collect information on participant location and reward at each time step of an episode, permitting analysis of not just whether a participant passed or failed a task, but the path they took across the episode. This gives




Fig. 9 The first-person view of the playing participant, with the health bar in the *bottom left* corner, the current and previous episode reward in the *top left* corner, and the arena index in the *top right* corner

an insight into *how* a participant passed or failed. We provide full scripts in the supplementary material for conducting testing with humans.

Designing experiments in Animal-AI

Building tasks in the Animal-AI Environment is straightforward. Researchers can build task configurations in an interpretable and easy to use YAML configuration file, specifying the position, size, colour, and special parameters of the objects required for testing. The Animal-AI Environment then interprets configuration files and renders those objects in a new instance of the environment. The agent is always spawned first in the environment, and takes precedence over all other objects, regardless of the order in which the agent is specified in the configuration file. Otherwise, priority is in the order of the objects specified in the configuration file. Several arena configurations can be specified in a single file, meaning that sequences of tests can be presented in the order specified by the researcher. Researchers can also create multi-arena configurations, where the agent moves from one arena to the next within the same episode. This enables the design of carefully crafted hierarchical training and testing curricula, where the difficulty of tests can be adjusted and prespecified.

We also provide a configuration design assistant, which allows users to visualise, edit, and create configuration files through a graphical web application written in Python.⁷ The assistant facilitates task building by removing the need to manually write YAML files and by giving users immediate visual feedback about the configuration being designed. The assistant provides considerable interactivity with many application callbacks. These include moving, resizing, deleting, and adding items to an Animal-AI arena as well as fully designing a new arena from scratch.

⁷ See <https://github.com/Kinds-of-Intelligence-CFI/aai-config-assist>.

Figure 10 shows the configuration design assistant. Both represent the same YAML configuration. Users can place new objects in the environment with definable positions, rotations, sizes, and colours, and the assistant detects any overlapping objects. When a configuration contains such an overlap, the problematic items are spawned in separate locations according to the precedence set out in the configuration file. This is exemplified in the right panel, where a ramp is spawned in the bottom left portion of the arena, rather than to the right of the central blue wall as displayed in the configuration assistant representation. The assistant signals the potential overlap with a red border around the object. New items can effortlessly be added and supported by the assistant as the environment develops. The assistant is lightweight and quicker to launch than Animal-AI, thus streamlining configuration design. It has been used by researchers to speed up the experimental design process.

Task configurations can also be procedurally generated. A simple accompanying command line tool is provided to design and generate precisely specified variations of a task.⁸ Each generated file is a valid Animal-AI configuration. The tool supports selecting random colours, specifying a fixed list of options, conditional values dependent on other choices already made, labels, and much more. Procedural generation allows for the quick creation of large test batteries, making behavioural evaluation more robust. For instance, researchers can study whether the performance of agents is sensitive to changes that ought to be irrelevant to success on the task, such as the colours of some objects.

To demonstrate the ease with which one can design and run experiments in the Animal-AI Environment, we present three experiments. All coordinates are presented as (x,y,z) where y is the vertical axis. The first experiment examines goal-directed navigation using a simple foraging task, emulating experiments of foraging behaviour in a wide variety of animals (see, e.g., Pyke et al., 1977; Schoener, 1971). The agent is spawned at a random x -coordinate at $(x,0,5)$, with a random initial rotation. At the centre of the arena $(20,0,20)$ is a *SpawnerTree*, which spawns 10 *GoodGoalMulti* goals of size 2 (Fig. 11). The agent must obtain them all before their health runs out. The agent maximises reward by finding the shortest path between the spawned goals, known as *foraging by energy maximisation* (King & Marshall, 2022; Pyke et al., 1977).

The second experiment is based on the well-known *operant chamber* experiments (Heron & Skinner, 1939) that have been used to study learning in a wide variety of animals (Macphail, 1982). In the classic design, the animal is placed in a chamber with a lever which can be pressed to return a food reward. Naive animals from which food has been withheld for a few hours are placed in the cham-

ber, and learn through trial-and-error to operate the lever in return for rewards. An important feature of this task design in early work on animal behaviour was that lever presses were automatically recorded. Since the operant chamber is straightforward to design, several animals could be studied in parallel, a property mirrored by the Animal-AI Environment for computational models. We recreate a version of an operant chamber in the Animal-AI Environment by using the *SpawnerButton* object placed at $(38.5,0,18.8)$, which, when pressed, spawns a *GoodGoal* next to it at $(39,3,22)$. The button is enclosed by walls on three sides to ensure that the agent approaches it from the front. The agent is spawned at a random position in the arena. To succeed, the agent must press the button and then navigate to obtain the reward (Fig. 12). We also construct a training curriculum, in line with work in reinforcement learning (Narvekar et al., 2020) and on shaping in animal behaviour theory (Skinner, 1958; Peterson, 2004). The curriculum introduces the agent to the button and the consequences of pressing it. First, the agent is spawned at $(36,0,20)$, directly in front of the bottom, such that it is likely to hit it by accident. Then, it is spawned further away, at $(30,0,20)$, and then at $(20,0,20)$. In the fourth stage, it is spawned at $(20,0,z)$, where z is randomly selected. Finally, it is presented with the final task, in which it is spawned at a random (x,z) coordinate. By comparing performance with and without the curriculum, its effect on learning can be directly assessed.

Both the foraging task and the button tasks are on the simple end of possible experiments that can be conducted in the Animal-AI Environment. For the third experiment, we provide the *Animal-AI testbed*, updated for compatibility with the latest version of the environment. The testbed is a series of 900 configurations inspired by comparative cognition research, used in the Animal-AI Olympics competition held in 2019 (see Crosby et al., 2020, 2019; Voudouris et al., 2022). The Animal-AI testbed is divided into ten levels which probe different parts of spatial and physical cognition (see Table 2). Each level consists of 30 tasks with different demands and difficulty, each task of which has three variants, varying, for example, the starting position of the agent or the size of the goals. Examples of four tasks are presented in Fig. 13.

A battery of agents

We provide implementations of three main agent types for running computational experiments in the Animal-AI Environment: *Random Action Agent*, *Heuristic Agent*, and *Learner*. Each of these showcase the versatility of the environment for running experiments in multiple disciplines. Implementation of these agents is available in the supplementary material.

⁸ See <https://github.com/wschella/procgen-companion>.

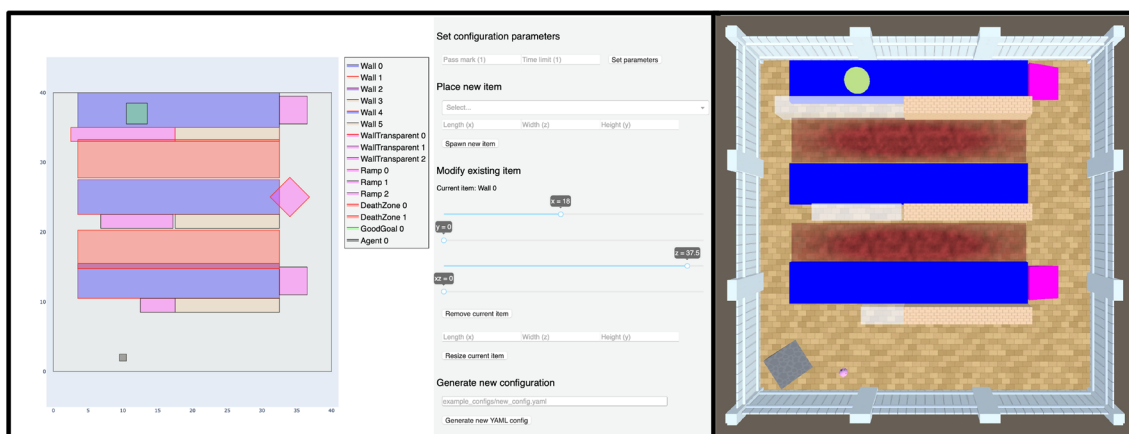


Fig. 10 The Animal-AI configuration design companion tool. *Left*: The graphical user interface for the tool, built in Python. *Right*: The rendered configuration in the Animal-AI Environment. Note the displaced *grey ramp* in the bottom left of the arena, as a result of a detected clash in the configuration tool



Fig. 11 The foraging task experiment. *Left*: A bird's-eye view of the task. *Right*: The agent's viewpoint during the task




Fig. 12 The operant chamber task. The agent is indicated by the *red arrow* and the location of the button is indicated by the *blue arrows*. *Left*: A bird's-eye view of the task. *Right*: The agent's viewpoint during the task




Fig. 13 Four tasks from the Animal-AI testbed. *Red arrows* indicate the position of the agent, *blue arrows* indicate the position of goals in the environment, and *grey arrows* indicate (possible) directions of motion of movable objects. *Far left*: A radial arm maze—the agent starts in the centre and must collect the eight GoodGoalMulti at the end of each arm. *Centre left*: An object permanence task—the agent first

observes a GoodGoal in the middle of a DeathZone, then the lights go out while the GoodGoal rolls behind a Wall on the left. *Centre right*: A trap-tube task—pushing the block on the right would knock the GoodGoal into the DeathZone, while the left block would not. *Far right*: A string-pulling task—only the UBlock can be used to pull the GoodGoal out of the DeathZone

Random action agents When performing testing on artificial and biological agents, it is often useful to define *chance performance* on tests that have been developed. However, experiments might be very complex, meaning that chance performance is not easily defined analytically. We provide customisable random action agents to simulate chance performance. The number of steps before picking a new action can be defined, or drawn from one of several distributions, and the agent can be biased to pick certain actions over others, or to correlate future actions with actions it has previously taken.

Heuristic agents More sophisticated agents can be designed by hand-coding rules for action specification in response to different observations. This can be done using raycast observations or camera observations. We provide examples of how this logic can be set up, to enable sophisticated, but rigid, behaviours in the environment.

Learners The Animal-AI Environment is integrated with Gymnasium, as well as `stable-baselines-3` (Raffin et al., 2021), rendering available many common (deep) reinforcement learning algorithms, such as Proximal Pol-

Table 2 The Animal-AI testbed consists of 900 configurations divided into ten levels of 30 tasks with three variants each

Level	Description
Food Retrieval (L1)	Valenced objects in an open arena with no obstacles.
Preferences (L2)	Valenced objects in forced- and free-choice configurations, such as Y-Mazes and Delayed Gratification tasks.
Static Obstacles (L3)	Valenced objects are fully or partially occluded by movable and immovable objects.
Avoidance (L4)	Positively valenced objects are arranged around negatively valenced objects. To obtain the former, the agent must navigate carefully to avoid the latter.
Spatial Reasoning and Support (L5)	Positively valenced objects are (partially) occluded. The agent must infer their presence based on their absence elsewhere. Valenced objects may also be supported out of reach by other objects, such as movable blocks.
Generalisation (L6)	Tasks from L1-L5 but adapted to have different features, such as the walls, floor, and/or ceiling being different colours.
Internal Modelling (L7)	Tasks from L1-L5 with different lights out intervals.
Object Permanence and Working Memory (L8)	Valenced objects become occluded behind other objects.
Numerosity and Advanced Preferences (L9)	Several valenced objects of different sizes and in different quantities are present in forced-choice settings. The agent must visually determine which choice will lead to the greatest reward.
Causal Reasoning (L10)	Positively valenced objects can only be obtained by interacting with neutral, movable objects, such as blocks.

Each level focuses on a different aspect of spatial and physical cognition (see Voudouris et al., 2022 for more details)

icy Optimisation (PPO, Schulman et al., 2017) and Deep Q-Learning (DQN, Mnih et al., 2013). We also provide an integration of Dreamer-v3, a recent model-based reinforcement learning algorithm that achieved state-of-the-art performance in Minecraft (Hafner et al., 2023).

Experiments

Using the tests designed in Section “Designing experiments in Animal-AI” as well as the Animal-AI testbed (Crosby, 2019; Crosby et al., 2020), we present the performances of a subset of the agents that have been implemented in the Animal-AI Environment.

Foraging experiment

We studied four types of agents on the foraging task in Fig. 11. To serve as a baseline for chance performance on this test, we used a random action agent. This agent selects one of nine actions with equal probability. The number of steps to take performing that action is a positive number drawn from a clipped normal distribution with mean of 5 and a standard deviation of 1, rounded to the nearest integer. We also designed a heuristic agent, to provide an example of near-optimal performance on this task. This agent navigates towards good goals using raycasts (45 rays over the frontal 60 degrees). When it gets stuck (i.e., its velocity is near 0 in all dimensions), it moves forwards and left or forwards and right with a probability of 0.5 until its velocity increases again. We then trained two deep reinforcement learning agents, Proximal Policy Optimisation (PPO) and Dreamer-v3, on the task, with 64×64 pixel input for one million steps each. We used the XL variant of Dreamer-v3, which has 200M

learnable parameters. We then evaluated all four agents on 100 episodes of the Foraging task. The training curves for PPO and Dreamer-v3, as well as the distribution of results for all four agents, are presented in Fig. 14. Note that all training curves presented in this paper are smoothed using 100-episode partial moving averages. Dreamer-v3, PPO, and the Heuristic Agent performed similarly and above the level of chance defined by the random action agent, evidence of efficient foraging strategies that approximately maximise reward. A Kruskal–Wallis rank-sum test indicates that there was a significant difference between the agents ($\chi^2(3) = 318.16$, $p < 0.00001$). Post hoc Dunn test statistics (Ogle et al., 2023) indicate that all agents are significantly different from each other, except for Dreamer-v3 and the Heuristic Agent (see Appendix B.1).

Operant chamber experiment

We studied four types of agents on the operant chamber task in Fig. 12. As with the previous experiment, we used a random action agent to define chance performance. A heuristic agent was also designed. This agent navigates towards SpawnerButton, GoodGoal and immovable objects using raycasts (45 rays over the frontal 60 degrees). When it has not detected any objects it repeats the previous action. When it is stationary, it moves forwards and left. A notable limitation is that raycasts do not detect which side of the button needs to be pressed, so this agent can successfully navigate towards buttons, but it does not necessarily press the button. Therefore, vision-based agents will likely perform much better, as the button is indicated by colour only.




Fig. 14 The results of four agents on the foraging task. *Left:* The training curves for PPO and Dreamer-v3 over one million training steps. *Right:* The scores of each type of agent over 100 trials of the task. The maximum reward available is 22, disregarding the reward lost during navigation between goals

As with the previous experiment, we trained PPO and Dreamer-v3 on the task with 64×64 pixel input for two million steps each. An agent was trained with and without the curriculum. The curriculum was cumulative, with the agents trained on five stages; each stage consisted of on an Animal-AI arena configuration containing the current stage of the curriculum along with all the previous stages interleaved. The agents loop over the full multi-arena configuration for 400,000 steps, before moving to the next stage, leading to training over two million steps in total. We then evaluated all six agents on 100 episodes of the Foraging task. The training curves for the four PPO and Dreamer-v3 agents, as well as the distribution of results for all six agents, are presented in Fig. 15. The Heuristic Agent performed above the level of chance defined by the random action agent, as did both Dreamer-v3 agents, with the agent trained on the curriculum performing better. The PPO agent trained only on the operant chamber task did not learn to press the button. With the curriculum, PPO was able to learn to solve some configurations, again demonstrating the role of curricula in scaffolding learning in these agents. A Kruskal–Wallis rank-sum test indicates that there was a significant difference between the agents ($\chi^2(3) = 400.15$, $p < 0.00001$). Post hoc Dunn test statistics indicate that all agents are significantly different from each other, except for the Heuristic Agent and both PPO agents compared to the random action agent, and the PPO agents compared to each other, suggesting that these agents are per-

forming no better than chance. Notably, the curriculum led to significantly better performance for Dreamer-v3 (*Bonferroni*-adjusted $p = 0.035504$; see Appendix B.2).

The Animal-AI testbed

While many agents in the Animal-AI Olympics were able to solve individual tasks in the Animal-AI testbed, performance was low compared to humans. Children aged 6–10 performed significantly better than even the top-performing AI systems in the competition (Voudouris et al., 2022).

Here, we present the performances of six agents on the Animal-AI testbed, along with data from children aged 6–10. As with the foraging experiment (Section “[Foraging experiment](#)”), we used a random action agent to define chance performance, and a heuristic agent with 45 rays emitted over the frontal 60 degrees of its view. We then trained PPO and Dreamer-v3 on 64×64 pixel input on a cumulative curriculum to be described below. We evaluated these four agents on all three variants of the 300 tasks in the Animal-AI testbed ($n = 900$). As far as we know, this is the first example of agents being trained and tested on the Animal-AI testbed (although see Lind and Vinken, 2021 for a variation).

We also present below the performances of children ($n = 59$) aged 6–10 on a subset of the tasks as a reference, taken

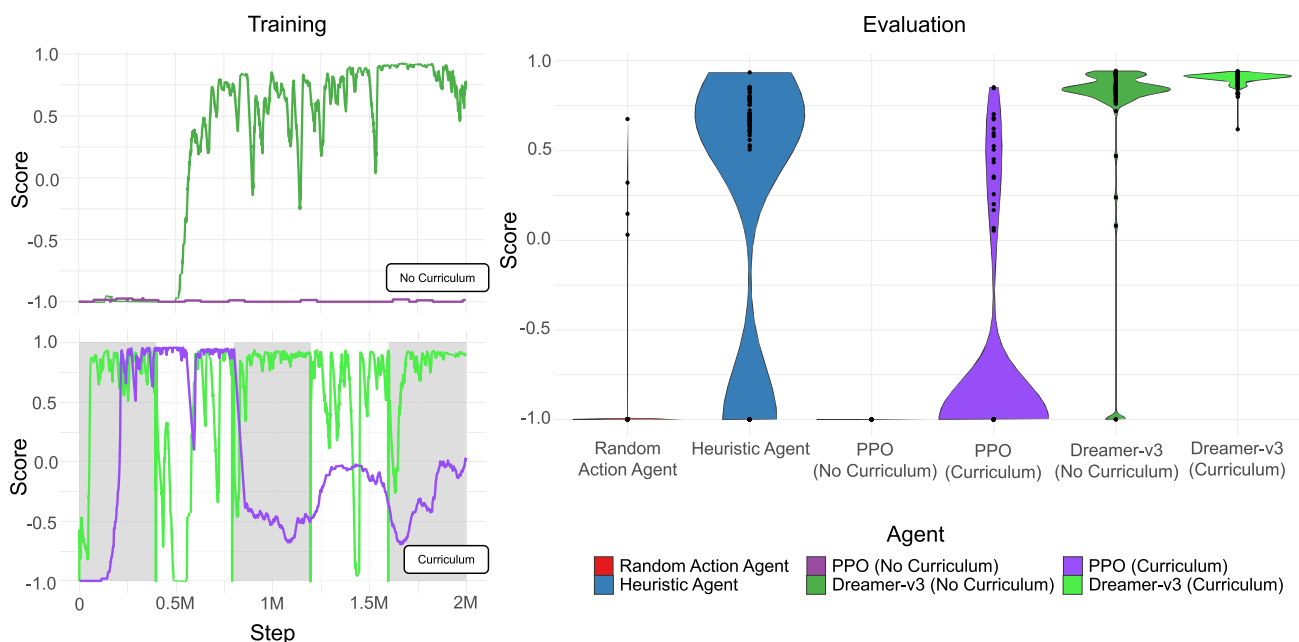


Fig. 15 The results of six agents on the operant chamber task. *Left*: The training curves for PPO and Dreamer-v3 over 2 million training steps with (*bottom*) and without (*top*) a training curriculum. Alternating grey and white panels indicate stages of the curriculum. *Right*: The scores of

each agent over 100 trials of the task used in the *No Curriculum* training condition. The maximum reward available is 1, disregarding the reward lost during navigation




Fig. 16 The results of seven types of agent on the Animal-AI testbed. *Left*: The training curves for PPO and Dreamer-v3 over 25 million training steps. Alternating grey and white panels indicate stages of the curriculum. *Right*: The proportion of Animal-AI testbed tasks that each

agent passed, according to the level. Data are presented for each agent with one score per episode. Chance performance is defined by the random action agent. Human performance is aggregated over data from 6–10 year olds ($n = 59$) on a subset of four tasks per level

from openly accessible data from Voudouris et al. (2022).⁹ We also present performances from the two best-performing agents from the Animal-AI Olympics competition, *Trrrrr* and *ironbar*, for reference (Crosby, 2019; Crosby et al., 2020). *Trrrrr* is a customised PPO architecture trained on 84×84 RGB pixel input on a bespoke curriculum for 70M steps and achieved the highest overall score in the competition.¹⁰ *ironbar* is a customised PPO architecture trained on 84×84 RGB pixel input on a bespoke curriculum for 20M steps and achieved the second highest overall score in the competition.¹¹ The testbed was not released to the developers of these agents until after their agents had been developed, trained, and tested.

PPO and Dreamer-v3 were trained on a cumulative curriculum consisting of 11 stages. The first stage contained one variant from each of 30 tasks in *Food Retrieval* (L1), totalling 30 configurations which the agent looped over for two million steps. The second stage contained one variant from each of the 60 tasks in L1 and *Preferences* (L2), totalling 60 configurations. The pattern repeats until the tenth stage, containing one variant from each of the 300 tasks in L1-L10. The agents were trained for two million steps on the first ten stages. They were then trained for a further five million steps on the last stage again (300 configurations from tasks L1-10), to ensure

adequate experience with all the tasks in the testbed. Agents were therefore trained for a total of 25 million steps.

Figure 16 presents the proportion of tests passed by each agent on each level, with the proportion for random action agents and for children indicated by reference lines. *Trrrrr* and *ironbar* performed the best across the levels, performing comparably on *Preferences* and exceeding children on *Food Retrieval*. The heuristic agent also performed competitively on *Food Retrieval* and *Internal Modelling*. Dreamer was generally outperformed by these other agents, but performed above the chance level defined by the sample of random action agents on *Food Retrieval*, *Preferences*, *Generalisation*, and *Internal Modelling*. PPO performed poorly across the testbed, generally performing at or near to chance performance, on most levels. Overall, while *Trrrrr* and *ironbar* placed first and second, respectively, in the Animal-AI Olympics Competition, Dreamer-v3 would have placed 25th, the Heuristic Agent would have placed 27th, the Random Action Agent would have placed 60th, and PPO would have placed 61st out of 64.

Three levels (Causal Reasoning, Object Permanence and Working Memory, and Numerosity and Advanced Preferences) pose difficulty for all agents, including children, indicating that the Animal-AI testbed remains a worthy challenge for state-of-the-art deep reinforcement learning. Although both Dreamer-v3 and PPO were trained on a curriculum of increasing difficulty, we leave it to future work to determine whether an alternative curriculum, or an adaptively generated curriculum (e.g., Team et al., 2023), would be more suitable for handling the diversity and complexity of the Animal-AI testbed. See Section B.3 for further analy-

⁹ Children played a randomly selected set of four tasks per level. They played these tasks randomly within each level, but with the levels arranged in the order presented in Table 2. Full data are available at <https://osf.io/g8u26/>.

¹⁰ See https://github.com/Denys88/rl_animal.

¹¹ See <https://github.com/ironbar/orangutan>.

sis and results. While better performance could be achieved through more training and better curricula, we leave these investigations to future work.

Conclusion

In this paper, we have presented the Animal-AI Environment, a virtual laboratory to bring together researchers from comparative cognition and artificial intelligence. Future iterations will continue to build on the Animal-AI Environment to improve its utility for researchers in these fields. Areas of active improvement include *adjustable arena sizes*, allowing researchers to build, for example, open-field navigation experiments. The introduction of *arena themes*, enabling researchers to replace the walls, floor, and sky with alternative colours and textures, will enable more naturalistic replication of experiments from the physical laboratory. We will also continue to improve the Animal-AI Environment's integration with cutting-edge frameworks and pipelines in artificial

intelligence and machine learning, to ensure that the latest computational models are available for use.

The Animal-AI Environment is unique in its mission to be a common tool for research on animal behaviour and cognition, and its integration with computational modelling and artificial intelligence. By bringing the precision of computational models of the brain and behaviour in touch with the carefully constructed experimental designs, datasets, and sophisticated cognitive theories, we can make significant progress towards better understanding animal cognition and behaviour.

A Sources of maintenance statuses of other environments

We defined maintenance status according to the last commit on the respective GitHub repository for each environment, as of 17th January 2025. Further information is presented in Table 3.

Table 3 The dates of the last commit in the GitHub history of each environment presented in Table 1 of the main text

Environment	Date of last commit	Source
DM Lab	2023-Jan-04	https://github.com/google-deepmind/lab/commits/master/
XLand	NA	No public release
Avalon	2023-May-04	https://github.com/Avalon-Benchmark/avalon/commits/main/
MuJoCo	2025-Jan-16	https://github.com/google-deepmind/mujoco/commits/main/
MineRL	2024-Nov-26	https://github.com/minerllabs/minerl/commits/dev/
Malmö	2022-Sep-23	https://github.com/microsoft/malmo/commits/master/
MineDoJo	2023-Aug-29	https://github.com/MineDojo/MineDojo/commits/main/
ThreeDWorld	2024-May-31	https://github.com/threedworld-mit/tdw/commits/master/
MiniWorld	2025-Jan-12	https://github.com/Farama-Foundation/Miniworld/commits/master/
APEXplorer 3D	NA	No public release (a static version linked to Allritz et al. (2022) is available in the supplementary material)

AAI3 is omitted as it is under constant active development across several repositories. These data are correct as of 2025-Jan-17

Table 4 Post hoc Dunn test of multiple comparisons for performances on the foraging task

Comparison	Z-statistic	Unadjusted <i>p</i> value	Adjusted <i>p</i> value
Random - Heuristic	14.324042	<0.000001	<0.000001
Random - PPO	6.664166	<0.000001	<0.000001
Random - Dreamer	15.662257	<0.000001	<0.000001
Heuristic - PPO	7.659876	0.001626	0.009759
Heuristic - Dreamer	1.338215	0.180826	1.000000
PPO - Dreamer	8.998092	<0.000001	<0.000001

Unadjusted and Bonferroni-adjusted *p* values are presented. All values are presented to six decimal places

B Further experimental results

B.1 Foraging experiment

See Table 4 for post hoc Dunn test statistics for performances on the foraging task.

B.2 Operant chamber experiment

See Table 5 for post hoc Dunn test statistics for performances on the operant chamber task.

B.3 The Animal-AI testbed

To examine the differences between agents and between levels on the Animal-AI testbed, we used a generalised linear mixed effects model, implemented in *R* (R Core Team, 2024). Level and Agent are taken as fixed effects, with participant

ID (either the name of the agent or the ID of the human participant) as a random intercept. Pass/fail is the response variable. The model, with a logit link function, was fitted by maximum likelihood (adaptive Gauss-Hermite quadrature) with the BOBYQA optimizer and TEN integration points.

In Table 6, the odds ratios of success to failure are presented. The population of agents is significantly more likely to pass the *Preferences* level compared to the first *Food Retrieval* level, perhaps because the arena is more restricted with walls, so there is less space for exploration. They are also significantly more likely to pass the *Generalisation* level. They are significantly more likely to fail *Spatial Reasoning & Support*, *Object Permanence & Working Memory*, *Numerosity & Advanced Preferences*, and *Causal Reasoning*, indicating that these levels are difficult.

Taking the random action agent to define chance performance, PPO is not significantly more likely to pass the tasks than chance. Dreamer is significantly more likely to pass than chance, as were ironbar, Trrrrr, the Heuristic Agent, and chil-

Table 5 Post hoc Dunn test of multiple comparisons for performances on the operant chamber task

Comparison	Z-statistic	Unadjusted <i>p</i> value	Adjusted <i>p</i> value
Random - Heuristic	2.258691	0.023903	0.358539
Random - PPO (No Curriculum)	-2.314548	0.020638	0.309565
Random - PPO (Curriculum)	-0.779929	0.435433	1.000000
Random - Dreamer (No Curriculum)	10.014743	<0.000001	<0.000001
Random - Dreamer (Curriculum)	13.054605	<0.000001	<0.000001
Heuristic - PPO (No Curriculum)	4.573239	0.00005	0.000072
Heuristic - PPO (Curriculum)	3.038620	0.002377	0.035650
Heuristic - Dreamer (No Curriculum)	7.756053	<0.000001	<0.000001
Heuristic - Dreamer (Curriculum)	10.795914	<0.000001	<0.000001
PPO (No Curriculum) - Dreamer (No Curriculum)	12.329291	<0.000001	<0.000001
PPO (No Curriculum) - Dreamer (Curriculum)	15.369153	<0.000001	<0.000001
PPO (No Curriculum) - PPO (Curriculum)	-1.534619	0.124877	1.000000
PPO (Curriculum) - Dreamer (No Curriculum)	10.794672	<0.000001	<0.000001
PPO (Curriculum) - Dreamer (Curriculum)	13.834534	<0.000001	<0.000001
Dreamer (No Curriculum) - Dreamer (Curriculum)	-3.039861	0.002367	0.035503

Unadjusted and Bonferroni-adjusted *p* values are presented. All values are presented to six decimal places

Table 6 Odds ratios from a generalised linear mixed effects model

Type	Fixed effect	Estimated odds	<i>p</i> value
Level	Intercept	0.048290	0.000003
	Preferences	1.383071	0.002403
	Static Obstacles	1.046102	0.671138
	Avoidance	0.956010	0.671425
	Spatial Reasoning & Support	0.562533	<0.000001
	Generalisation	1.769805	<0.000001
	Internal Modelling	1.057973	0.595556
	Object Permanence & Working Memory	0.303834	<0.000001
	Numerosity & Advanced Preferences	0.184259	<0.000001
	Causal Reasoning	0.367255	<0.000001
Agent	Dreamer-v3	6.539571	0.038181
	PPO	0.622401	0.608951
	ironbar	21.925737	0.000644
	Trrrrr	22.146038	0.000617
	Heuristic Agent	6.164322	0.044706
	Children (Aged 6-10)	47.198110	<0.000001

The reference category for *Level* is *Food Retrieval (L1)*, and for *Agent* it is *Random Action Agent*. All values are presented to 6 decimal places

dren. Children are 47 times more likely to pass than chance, while ironbar and Trrrrr are 19 times more likely. Dreamer is six times more likely.

Author Contributions K. Voudouris: Conceptualization, Methodology, Software, Validation, Formal Analysis, Investigation, Resources, Data Curation, Writing (Original Draft), Writing (Review & Editing), Visualization, Project Administration. I. Alhas: Software, Validation, Data Curation, Writing (Review & Editing). W. Schellaert: Software, Validation, Investigation, Data Curation, Writing (Review & Editing). M. G. Mecattaf: Formal Analysis, Software, Validation, Investigation, Data Curation, Visualization, Writing (Review & Editing). B. Slater: Software, Validation, Writing (Review & Editing). M. Crosby: Conceptualization, Formal Analysis, Software, Writing (Review & Editing). J. Holmes: Software, Writing (Review & Editing). J. Burden: Formal Analysis, Writing (Review & Editing). N. Chaubey: Formal Analysis, Writing (Review & Editing). N. Donnelly: Data Curation, Writing (Review & Editing). M. Patel: Software, Data Curation, Writing (Review & Editing). M. Halina: Conceptualization, Resources, Supervision, Funding Acquisition, Writing (Review & Editing). J. Hernández-Orallo: Conceptualization, Resources, Supervision, Funding Acquisition, Writing (Review & Editing). L. G. Cheke: Conceptualization, Methodology, Resources, Supervision, Funding Acquisition, Writing (Review & Editing).

Funding This work was funded by an ESRC DTP scholarship (ES/P000 738/1), the Kinds of Intelligence Program at the Leverhulme Centre for the Future of Intelligence (G108086), the US DARPA HR00112120007 (RECoG-AI) grant, and a Templeton World Charity Foundation grant (Major Transitions in the Evolution of Cognition: TWCF-2020-20539). This work was also funded by the European Union with the “Programa Operativo del Fondo Europeo de Desarrollo Regional (FEDER) de la Comunitat Valenciana 2014-2020” under agreement INNEST/2021/317 (Neurocalçat) and by the Vic. Inv. of the Universitat Politècnica de València under “programa de ayudas a la formación de doctores en colaboración con empresas” (DOCEMPR21).

Data Availability All datasets and materials for the experiments reported here can be found at <https://github.com/Kinds-of-Intelligence-CFI/aai3-paper-experiments>. None of the experiments were preregistered.

Code Availability Further information, tutorials, scripts, and documentation for the project can be found here: <https://github.com/Kinds-of-Intelligence-CFI/animal-ai>. Scripts to replicate the experiments and analyses in this paper can be found here: <https://github.com/Kinds-of-Intelligence-CFI/aai3-paper-experiments>.

Declarations

Conflicts of Interest M. Crosby is now employed by DeepMind Technologies Limited, but this work was conducted prior to this employment. The remaining authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Ethical Approval Ethical approval is not required because no experiments were performed with humans or animals. All data from human children were from an openly accessible repository (<https://osf.io/g8u26/>) associated with an independent study that received ethical approval.

Consent to Participate No research was conducted with human subjects. All data from human children were from an openly accessible repository (<https://osf.io/g8u26/>) associated with an independent study that received ethical approval.

Consent for Publication No research was conducted with human subjects. All data from human children were from an openly accessible repository (<https://osf.io/g8u26/>) associated with an independent study that received ethical approval.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as

long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Albrecht, J., Fetterman, A., Fogelman, B., Kitanidis, E., Wróblewski, B., Seo, N., Rosenthal, M., Knutins, M., Polizzi, Z., Simon, J., et al. (2022). Avalon: A benchmark for rl generalization using procedurally generated worlds. *Advances in Neural Information Processing Systems*, 35, 12813–12825.
- Allen, K. R., Brändle, F., Botvinick, M., Fan, J., Gershman, S. J., Gopnik, A., Griffiths, T. L., Hartshorne, J., Hauser, T. U., Ho, M. K., Leeuw, J. d., Ma, W. J., Murayama, K., Nelson, J. D., Opheusden, B. v., Pouncy, H. T., Rafner, J., Rahwan, I., Rutledge, R., ..., Schulz, E. (2023). Using Games to Understand the Mind. <https://doi.org/10.31234/osf.io/hbsvj>
- Allritz, M., Call, J., Schweller, K., McEwen, E. S., de Guinea, M., Janmaat, K. R., Menzel, C. R., & Dolins, F. L. (2022). Chimpanzees (pan troglodytes) navigate to find hidden fruit in a virtual environment. *Science Advances*, 8(25), eabm4754.
- Ansarinia, M., Clocher, B., Defossez, A., Schmuck, E., & Cardoso-Leite, P. (2022). Cogenv: A reinforcement learning environment for cognitive tests. 2022 Conference on Cognitive Computational Neuroscience.
- Baum, M., Schattenhofer, L., Rössler, T., Osuna-Mascaró, A., Auersperg, A., Kacelnik, A., & Brock, O. (2022). Yoking-based identification of learning behavior in artificial and biological agents. *International Conference on Simulation of Adaptive Behavior*, 67–78.
- Beattie, C., Leibo, J. Z., Teplyashin, D., Ward, T., Wainwright, M., Küttler, H., Lefrancq, A., Green, S., Valdés, V., Sadik, A., et al. (2016). Deepmind lab. [arXiv:1612.03801](https://arxiv.org/abs/1612.03801).
- Beyret, B., Hernández-Orallo, J., Cheke, L., Halina, M., Shanahan, M., & Crosby, M. (2019). The Animal-AI environment: Training and testing animal-like artificial cognition. [arXiv:1909.07483](https://arxiv.org/abs/1909.07483).
- Brändle, F., Stocks, L. J., Tenenbaum, J. B., Gershman, S. J., & Schulz, E. (2023). Empowerment contributes to exploration behaviour in a creative video game. *Nature Human Behaviour*, 1–9.
- Brea, J., Clayton, N. S., & Gerstner, W. (2023). Computational models of episodic-like memory in food-caching birds. *Nature Communications*, 14(1), 2979.
- Bremner, J. G., Slater, A. M., & Johnson, S. P. (2015). Perception of object persistence: The origins of object permanence in infancy. *Child Development Perspectives*, 9(1), 7–13.
- Brown, H. R., Zeidman, P., Smittenaar, P., Adams, R. A., McNab, F., Rutledge, R. B., & Dolan, R. J. (2014). Crowdsourcing for Cognitive Science - The Utility of Smartphones. *PLOS ONE*, 9(7), e100662.
- Burden, J., Voudouris, K., Burnell, R., Rutar, D., Cheke, L., & Hernández-Orallo, J. (2023). Inferring capabilities from task performance with bayesian triangulation. [arXiv preprint arXiv:2309.11975](https://arxiv.org/abs/2309.11975).
- Cardoso, R. P., Donnelly, N., Keedwell, E., Cheke, L., & Shanahan, M. (2023). What is a stimulus? a computational perspective on an associative learning model. *ALIFE 2023: Ghost in the Machine: Proceedings of the 2023 Artificial Life Conference*.
- Castro, P. S., Moitra, S., Gelada, C., Kumar, S., & Bellemare, M. G. (2018). Dopamine: A research framework for deep reinforcement learning. [arXiv preprint arXiv:1812.06110](https://arxiv.org/abs/1812.06110).
- Chevalier-Boisvert, M., Dai, B., Towers, M., Perez-Vicente, R., Willems, L., Lahlou, S., Pal, S., Castro, P. S., & Terry, J. (2024). Minigrid & miniworld: Modular & customizable reinforcement learning environments for goal-oriented tasks. *Advances in Neural Information Processing Systems*, 36.
- Coughlan, G., Coutrot, A., Khondoker, M., Minihane, A.-M., Spiers, H., & Hornberger, M. (2019). Toward personalized cognitive diagnostics of at-genetic-risk Alzheimer's disease. *Proceedings of the National Academy of Sciences*, 116(19), 9285–9292.
- Coutrot, A., Manley, E., Goodroe, S., Gahnstrom, C., Filomena, G., Yesiltepe, D., Dalton, R. C., Wiener, J. M., Hölscher, C., Hornberger, M., & Spiers, H. J. (2022). Entropy of city street networks linked to future spatial navigation ability. *Nature*, 604(7904), 104–110.
- Coutrot, A., Silva, R., Manley, E., De Cothi, W., Sami, S., Bohbot, V. D., Wiener, J. M., Hölscher, C., Dalton, R. C., Hornberger, M., & Spiers, H. J. (2018). Global Determinants of Navigation Ability. *Current Biology*, 28(17), 2861–2866.e4. <https://doi.org/10.1016/j.cub.2018.06.009>
- Crosby, M. (2020). Building thinking machines by solving animal cognition tasks. *Minds and Machines*, 30(4), 589–615.
- Crosby, M., Beyret, B., & Halina, M. (2019). The Animal-AI Olympics. *Nature. Machine Intelligence*, 1(5), 257.
- Crosby, M., Beyret, B., Shanahan, M., Hernández-Orallo, J., Cheke, L., & Halina, M. (2020). The Animal-AI testbed and competition. *NeurIPS 2019 competition and demonstration track*, 164–176.
- Dubey, R., Agrawal, P., Pathak, D., Griffiths, T. L., & Efros, A. A. (2018). Investigating human priors for playing video games. [arXiv:1802.10217](https://arxiv.org/abs/1802.10217).
- Fan, L., Wang, G., Jiang, Y., Mandlekar, A., Yang, Y., Zhu, H., Tang, A., Huang, D.-A., Zhu, Y., & Anandkumar, A. (2022). Minedojo: Building open-ended embodied agents with internet-scale knowledge. *Advances in Neural Information Processing Systems*, 35, 18343–18362.
- Freire, I. T., Amil, A. F., & Verschure, P. F. (2024). Sequential memory improves sample and memory efficiency in episodic control. *Nature Machine Intelligence*, 1–13.
- Gan, C., Bhandwadar, A., Torralba, A., Tenenbaum, J. B., & Isola, P. (2021). OPEn: An Open-ended Physics Environment for Learning Without a Task [[arXiv:2110.06912](https://arxiv.org/abs/2110.06912) [cs]]. <https://doi.org/10.48550/arXiv.2110.06912>
- Gan, C., Schwartz, J., Alter, S., Mrowca, D., Schrimpf, M., Traer, J., De Freitas, J., Kubilius, J., Bhandwadar, A., Haber, N., et al. (2020). Threedworld: A platform for interactive multi-modal physical simulation. [arXiv:2007.04954](https://arxiv.org/abs/2007.04954).
- Griffiths, T. L. (2015). Manifesto for a new (computational) cognitive revolution. *Cognition*, 135, 21–23.
- Guest, O., & Martin, A. E. (2021). How computational modeling can force theory building in psychological science. *Perspectives on Psychological Science*, 16(4), 789–802.
- Guss, W. H., Castro, M. Y., Devlin, S., Houghton, B., Kuno, N. S., Loomis, C., Milani, S., Mohanty, S., Nakata, K., Salakhutdinov, R., et al. (2021). The minerl 2020 competition on sample efficient reinforcement learning using human priors. [arXiv:2101.11071](https://arxiv.org/abs/2101.11071).
- Hafner, D., Pasukonis, J., Ba, J., & Lillicrap, T. (2023). Mastering diverse domains through world models. [arXiv:2301.04104](https://arxiv.org/abs/2301.04104)
- Hernández-Orallo, J. (2017). *The measure of all minds: Evaluating natural and artificial intelligence*. Cambridge University Press.
- Heron, W., & Skinner, B. (1939). An apparatus for the study of animal behavior. *The Psychological Record*, 3, 165.
- Herrmann, E., Call, J., Hernández-Lloreda, M. V., Hare, B., & Tomasello, M. (2007). Humans have evolved specialized skills

- of social cognition: The cultural intelligence hypothesis. *science*, 317(5843), 1360–1366.
- Jaakkola, K. (2014). Do animals understand invisible displacement? a critical review. *Journal of Comparative Psychology*, 128(3), 225.
- Johnson, M., Hofmann, K., Hutton, T., & Bignell, D. (2016). The malmo platform for artificial intelligence experimentation. *IJCAI*, 4246–4247.
- Juliani, A., Berges, V.-P., Teng, E., Cohen, A., Harper, J., Elion, C., Goy, C., Gao, Y., Henry, H., Mattar, M., et al. (2018). Unity: A general platform for intelligent agents. [arXiv:1809.02627](https://arxiv.org/abs/1809.02627).
- King, A. J., & Marshall, H. H. (2022). Optimal foraging. *Current Biology*, 32(12), R680–R683.
- Kosoy, E., Collins, J., Chan, D. M., Pathak, D., Agrawal, P., & Gopnik, A. (2020). Exploring exploration: Comparing children with agents in unified exploration environments. *CogSci*.
- Krasheninnikova, A., Berardi, R., Lind, M.-A., O'Neill, L., & von Bayern, A. M. (2019). Primate cognition test battery in parrots. *Behaviour*, 156(5–8), 721–761.
- Lake, B. M., Ullman, T. D., Tenenbaum, J. B., & Gershman, S. J. (2017). Building machines that learn and think like people. *Behavioral and Brain Sciences*, 40.
- Leduc-McNiven, K., White, B., Zheng, H., D McLeod, R., & R Friesen, M. (2018). Serious games to assess mild cognitive impairment: 'The game is the assessment'. *Research and Review Insights*, 2(1). <https://doi.org/10.15761/RR.1000128>
- Lehuger, A., & Crosby, M. (2021). Fixed β -vae encoding for curious exploration in complex 3d environments. [arXiv:2105.08568](https://arxiv.org/abs/2105.08568).
- Leibo, J. Z., d'Autume, C. d. M., Zoran, D., Amos, D., Beattie, C., Anderson, K., Castañeda, A. G., Sanchez, M., Green, S., Gruslys, A., et al. (2018). Psychlab: A psychology laboratory for deep reinforcement learning agents. [arXiv:1801.08116](https://arxiv.org/abs/1801.08116).
- Lind, J., & Vinken, V. (2021). Can associative learning be the general process for intelligent behavior in non-human animals? *bioRxiv*, 2021-12.
- Macphail, E. M. (1982). *Brain and intelligence in vertebrates*. Oxford: Clarendon Press.
- Mareschal, D. (2000). Object knowledge in infancy: Current controversies and approaches. *Trends in Cognitive Sciences*, 4(11), 408–416.
- McNab, F., Zeidman, P., Rutledge, R. B., Smittenaar, P., Brown, H. R., Adams, R. A., & Dolan, R. J. (2015). Age-related changes in working memory and the ability to ignore distraction. *Proceedings of the National Academy of Sciences*, 112(20), 6515–6518.
- Mecattaf, M. G., Slater, B., Tešić, M., Prunty, J., Voudouris, K., & Cheke, L. G. (2024). A little less conversation, a little more action, please: Investigating the physical common-sense of LLMs in a 3D embodied environment. [arXiv:2410.23242](https://arxiv.org/abs/2410.23242).
- Miscov, A. M., McEwen, E., & Seed, A. (2024). Virtually anything can happen: Investigating short-term memory in capuchin monkeys using virtual environments. *Proceedings of the Annual Meeting of the Cognitive Science Society*, 46.
- Mitchener, L., Tuckey, D., Crosby, M., & Russo, A. (2022). Detect, understand, act: A neuro-symbolic hierarchical reinforcement learning framework. *Machine Learning*, 111(4), 1523–1549.
- Mnih, V., Kavukcuoglu, K., Silver, D., Graves, A., Antonoglou, I., Wierstra, D., & Riedmiller, M. (2013). Playing atari with deep reinforcement learning. [arXiv:1312.5602](https://arxiv.org/abs/1312.5602).
- Narvekar, S., Peng, B., Leonetti, M., Sinapov, J., Taylor, M. E., & Stone, P. (2020). Curriculum learning for reinforcement learning domains: A framework and survey. *The Journal of Machine Learning Research*, 21(1), 7382–7431.
- Ogle, D. H., Doll, J. C., Wheeler, A. P., & Dinno, A. (2023). Fsa: Simple fisheries stock assessment methods [R package version 0.9.5]. <https://CRAN.R-project.org/package=FSA>
- Peterson, G. B. (2004). A day of great illumination: B f skinner's discovery of shaping. *Journal of the experimental analysis of behavior*, 82(3), 317–328.
- Piaget, J. (1923). *The Origins of Intelligence In The Child* (M. Cook, Trans.). Routledge & Kegan Paul, Ltd.
- Piaget, J., & Inhelder, B. (1969). Intellectual operations and their development. *Experimental psychology: Its scope and method*, 7, 147–203.
- Piloto, L. S., Weinstein, A., Battaglia, P., & Botvinick, M. (2022). Intuitive physics learning in a deep-learning model inspired by developmental psychology. *Nature human behaviour*, 6(9), 1257–1267.
- Pouncy, T., & Gershman, S. J. (2022). Inductive biases in theory-based reinforcement learning. *Cognitive Psychology*, 138, 101509. <https://doi.org/10.1016/j.cogpsych.2022.101509>
- Pyke, G. H., Pulliam, H. R., & Charnov, E. L. (1977). Optimal foraging: A selective review of theory and tests. *The quarterly review of biology*, 52(2), 137–154.
- Raffin, A., Hill, A., Gleave, A., Kanervisto, A., Ernestus, M., & Dormann, N. (2021). Stable-baselines3: Reliable reinforcement learning implementations. *Journal of Machine Learning Research*, 22(268), 1–8. <http://jmlr.org/papers/v22/20-1364.html>
- Rafner, J., Biskjaer, M. M., Zana, B., Langsfjord, S., Bergenholtz, C., Rahimi, S., Carugati, A., Noy, L., & Sherson, J. (2022). Digital Games for Creativity Assessment: Strengths, Weaknesses and Opportunities. *Creativity Research Journal*, 34(1), 28–54.
- R Core Team (2024). *_R_: A language and environment for statistical computing*. R Foundation for statistical computing, Vienna, Austria. <https://www.R-project.org/>.
- Rutar, D., Cheke, L. G., Hernández-Orallo, J., Markelius, A., & Schellaert, W. (2024). General interaction battery: Simple object navigation and affordances (GIBSONA). Available at SSRN 4924246.
- Schoener, T. W. (1971). Theory of feeding strategies. *Annual review of ecology and systematics*, 369–404.
- Scholl, B. J. (2007). Object persistence in philosophy and psychology. *Mind & Language*, 22(5), 563–591.
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., & Klimov, O. (2017). Proximal policy optimization algorithms. [arXiv preprint arXiv:1707.06347](https://arxiv.org/abs/1707.06347).
- Schulz, E., Bertram, L., Hofer, M., & Nelson, J. D. (2019). Exploring the space of human exploration using entropy mastermind. *bioRxiv*, 540666.
- Skinner, B. F. (1958). Reinforcement today. *American Psychologist*, 13(3), 94.
- Stafford, T., & Dewar, M. (2014). Tracing the trajectory of skill learning with a very large sample of online game players. *Psychological science*, 25(2), 511–518.
- Team, A. A., Bauer, J., Baumli, K., Baveja, S., Behbahani, F., Bhoopchand, A., Bradley-Schmieg, N., Chang, M., Clay, N., Collister, A., et al. (2023). Human-timescale adaptation in an open-ended task space. [arXiv preprint arXiv:2301.07608](https://arxiv.org/abs/2301.07608).
- Todorov, E., Erez, T., & Tassa, Y. (2012). Mujoco: A physics engine for model-based control. *IROS*, 5026–5033.
- Tokmakov, P., Li, J., Burgard, W., & Gaidon, A. (2021). Learning to track with object permanence. *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 10860–10869.
- Towers, M., Terry, J. K., Kwiatkowski, A., Balis, J. U., Cola, G. d., Deleu, T., Goulão, M., Kallinteris, A., KG, A., Krimmel, M., Perez-Vicente, R., Pierré, A., Schulhoff, S., Tai, J. J., Shen, A. T. J., & Younis, O. G. (2023). Gymnasium. <https://doi.org/10.5281/zenodo.8127026>
- Voiskounsky, A. E., Yermolova, T. D., Yagolkovskiy, S. R., & Khromova, V. M. (2017). Creativity in online gaming: Individual and dyadic performance in Minecraft. *Psychology in Russia: State of the Art*, 10(4), 144–161. <http://psychologyinrussia.com/volumes/index.php?article=6935>
- Voudouris, K., Cheke, L. G., & Halina, M. (2024). The future is computational comparative cognition. *Comparative Cognition & Behavior Reviews*, 19.

- Voudouris, K., Crosby, M., Beyret, B., Hernández-Orallo, J., Shanahan, M., Halina, M., & Cheke, L. G. (2022). Direct human-ai comparison in the Animal-AI environment. *Frontiers in Psychology, 13*, 1884.
- Voudouris, K., Donnelly, N., Rutar, D., Burnell, R., Burden, J., Hernández-Orallo, J., & Cheke, L. G. (2022). Evaluating object permanence in embodied agents using the Animal-AI environment. EBeM'22: Workshop on AI Evaluation Beyond Metrics, July 25, 2022, Vienna, Austria.
- Voudouris, K., Liu, J. D., Siwinska, N., Schellaert, W., & Cheke, L. G. (2024). Investigating object permanence in deep reinforcement learning agents. Proceedings of the Annual Meeting of the Cognitive Science Society, 46.
- Ward, T., Bolt, A., Hemmings, N., Carter, S., Sanchez, M., Barreira, R., Noury, S., Anderson, K., Lemmon, J., Coe, J., et al. (2020). Using unity to help solve intelligence. [arXiv:2011.09294](https://arxiv.org/abs/2011.09294).
- Zakharov, A., Crosby, M., & Fountas, Z. (2021). Episodic memory for subjective-timescale models. ICML 2021 Workshop on Unsupervised Reinforcement Learning.
- Zuidema, W., French, R. M., Alhama, R. G., Ellis, K., O'Donnell, T. J., Sainburg, T., & Gentner, T. Q. (2020). Five ways in which computational modeling can help advance cognitive science: Lessons from artificial grammar learning. *Topics in cognitive science, 12*(3), 925–941.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.