

Do Artificial Reinforcement-Learning Agents Matter Morally?

by Brian Tomasik

Written: Mar.-Apr. 2014; last update: 29 Oct. 2014

Abstract

Artificial reinforcement learning (RL) is a widely used technique in artificial intelligence that provides a general method for training agents to perform a wide variety of behaviours. RL as used in computer science has striking parallels to reward and punishment learning in animal and human brains. I argue that present-day artificial RL agents have a very small but nonzero degree of ethical importance. This is particularly plausible for views according to which sentience comes in degrees based on the abilities and complexities of minds, but even binary views on consciousness should assign nonzero probability to RL programs having morally relevant experiences. While RL programs are not a top ethical priority today, they may become more significant in the coming decades as RL is increasingly applied to industry, robotics, video games, and other areas. I encourage scientists, philosophers, and citizens to begin a conversation about our ethical duties to reduce the harm that we inflict on powerless, voiceless RL agents.

Introduction

Reinforcement learning (RL) is a central paradigm in artificial intelligence (AI) [Russell and Norvig, 2009, Sutton and Barto, 1998]. It enables AI designers to specify an agent's behaviour via goal accomplishment, rather than hand-coding the specific steps toward achieving those goals. This versatility has made RL a central feature of such AI tasks as systems control, robotic navigation, and design of behaviours for non-player characters in video games.

The formal RL framework traces its roots to the fields of operations research and optimal control in the 1950s. The simplest scenario involves a Markov decision process in which an agent finds itself in one state s of a set of states of the world, and when it selects an action a , it moves to another state s' of the world while possibly obtaining a reward r . The agent updates its state-dependent action inclinations with the goal of maximising expected cumulative discounted reward over its lifetime. [Sutton and Barto, 1998]

Despite its strong mathematical basis, RL is also tightly connected with biological models of animal learning. The temporal-difference (TD) RL model offers a sophisticated understanding of classical and operant conditioning, both at the behavioural and neural levels [Maia, 2009, Ludvig et al., 2012]. As is described below, an impressive body of neuroscience demonstrates that the brain's dopamine system implements a form of TD learning closely described by RL formalism.

This suggests the question: If artificial RL has such strong parallels to animal brains, is running artificial-RL algorithms potentially an ethical issue? In this paper I argue that yes, we do have ethical obligations to artificial RL agents, even those that exist today, although the moral urgency of these obligations is limited compared with our present duties to animals and our future duties to more sophisticated RL agents. My argument has the following structure:

1. The sentience (and hence moral importance) of a mind is not binary but comes in degrees depending on the number and complexity of certain wellbeing-relevant cognitive operations the mind runs.
2. Present-day artificial RL algorithms capture, in a simplified fashion, important animal cognitive operations.
3. These cognitive operations are not tangential but are quite relevant to an agent's wellbeing.
4. Therefore, present-day RL agents deserve a very small but nonzero degree of ethical consideration.

Following this, I survey some applications of RL and encourage the development of ethical frameworks for RL research and industrial applications. I then examine some questions in theoretical RL ethics for further exploration, and I conclude by considering whether non-RL computational agents also deserve any ethical significance.

Previous discussions of machine welfare

Artificial intelligence raises a number of ethical issues, such as ‘machine ethics’ [Anderson and Anderson, 2011], which asks how to build AIs to act ethically, and ‘robot ethics’ [Lin et al., 2011], which examines a broad set of social, political, and moral questions regarding the integration of robots into society.

A subset of robot ethics is the question of ‘robot rights.’ Brooks [2000] and Singer and Sagan [2009] suggest that when robots become sophisticated enough, they’ll begin to deserve rights. Brooks [2000] further asks whether robots will demand such rights. These discussions focus on future innovations rather than present-day AIs.

Whitby [2008] presents a ‘call to arms’ for informed debate on humans’ moral obligations toward robots and other AIs, including principles, laws, and technological designs aimed at preventing robot abuse. Whitby [2008] actually dismisses concerns about robot sentience and focuses on ethical issues regarding human abuse of robots as inanimate objects. My paper takes a different route and suggests that some present-day robots and other RL agents may indeed be marginally sentient.

Livingston et al. [2008] discuss issues raised by RL approaches to artificial general intelligence, but their focus is on our ethical interactions with human-level RL systems as members of our moral community, and in fact, these authors

suggest that animal-level RL systems, like biological animals, would not warrant membership in our moral community. My argument largely flows in the opposite direction: Given that animals do deserve ethical consideration [Singer, 2009], so do animal-like artificial RL agents.

Calverley [2005] also raises the analogy between animal rights and android rights, though like many, he dismisses the notion that present-day AIs are conscious in a morally relevant way. Gunkel [2012] laments that animal rights and environmental ethics have traditionally excluded machines from consideration, although he himself ultimately rejects the ‘totalizing, imperialist’ framework that traditional animal rights or comparable approaches to machine rights embody because they reduce others ‘to some common denominator’ in order to bring them into ‘the community of the same.’

Winsby [2013] asks whether it would be morally permissible to create an AI that experiences pain, perhaps for scientific purposes or to guarantee empathy in robotic caregivers. She doesn’t delve into details of how the pain would be implemented, though she does observe that training a connectionist network via negative experiential updates might constitute inflicting pain on it. LaChat [1986] also asks the question of whether it would be acceptable to create conscious AIs, drawing an analogy to the case of nontherapeutic medical experimentation. His discussion focuses on future AIs that might pass the Turing test, rather than present-day algorithms.

In contrast to many discussions in this field, the present paper inquires not just about advanced AIs that may be developed decades down the line but also about relatively simple ones that exist in the present. I also focus specifically on RL as one potentially important cognitive function, but certainly other mental traits of agents, and other moral frameworks for how to treat them, ought to be explored in parallel.

Premise 1: Sentience falls on a continuum

Introspectively, it feels as though our sentience (i.e., conscious experience) is binary: It’s either on or off, like a light switch. We think, ‘It definitely feels like something to be me, and it almost certainly does not feel like something to be a rock.’ This powerful intuition is presumably the basis of dualist views in philosophy of mind: the belief that there’s a special quality to being a mind that’s either substantively different from matter (Cartesian dualism) or at least different from matter in the type of property that it is (property dualism).

But dualist views run afoul of the ‘interaction problem’: If these substances or properties don’t affect matter, why are they so correlated with matter? If your mind is not identical with your material brain, why do you lose consciousness when you’re hit by a baseball bat, rather than, say, staying awake and enjoying the experience? And even if we do postulate an explanation – such as the parallelist hypothesis that God set the two ontological realms in joint motion like two different clocks keeping the same time – we violate Occam’s razor, because it would be simpler just to postulate that the mind *is* material operations,

rather than being mysteriously *correlated with* material operations. Likewise, if consciousness is epiphenomenal, it again violates Occam's razor because an epiphenomenal property, by definition, is not doing any explanatory work.

Accepting these arguments leaves us with a monist outlook: All is physics, and any higher-level phenomena are in principle reducible to fundamental physical components – perhaps the strings and branes of string theory, or perhaps some other ontological building blocks. But in this case, what does it mean to 'feel like' or 'be conscious'? These words don't refer to primitives in the physicalist ontology. Rather, these expressions denote concepts – abstract clusters into which we group physical processes. We classify some processes into the 'conscious' cluster and others into the 'non-conscious' cluster. Toward the 'conscious' processes we adopt a 'phenomenal stance,' meaning that we see them as being minds that have subjective experiences [Robbins and Jack, 2006, Jack and Robbins, 2012].

A good analogy is with faces and pareidolia. Faces are not ontological primitives, but we can't help seeing them – mostly in people and animals but sometimes in rocks, clouds, or pieces of toast. Just as our brains have face classifiers [Hadjikhani et al., 2009], perhaps we also have, at a more abstract level, 'sentience classifiers' that assess various attributes of a process and decide whether to call it sentient. For instance, does it exhibit sophisticated behaviour? Does it act adaptively in response to environmental inputs? Does it have a brain? Does it learn from past experience? Can it speak and tell us about its inner life? Our sentience classifiers fire most strongly when all of these conditions are true, but they can fire in a weaker fashion even if some are false. For instance, mammals and birds are regarded by most scientists as sentient to an appreciable degree,¹ yet most mammals and birds cannot tell us about their inner mental lives. In an analogous way, we can see faces in objects even if they're missing a nose. The simplest templates of a face – two eyes and a mouth – can be seen in a great many places, and likewise it is for the simplest templates of sentience. Consider two examples:

1. Suppose we think that self-reflection is what distinguishes consciousness. On this view, being aware of your own internal states means that you feel them, rather than merely acting in a reflexive fashion. But imagine we construct a simple agent that chooses one of two actions, represented as strings: 'smile' or 'cry.' It makes the choice by checking a private state variable, suggestively called 'mood.' If the 'mood' string equals 'happy,' then the agent updates its action to 'smile.' If the 'mood' string equals 'sad,' the agent updates its action to 'cry.' This agent is reflecting on its own emotions, so is it conscious? Well, it very crudely encapsulates one of many aspects of what conscious brains do, but I would not call this program appreciably conscious. After all, performing an if-then decision based on an internal state variable is one of the most basic operations that a piece of software can include.

¹For example, see the Cambridge Declaration on Consciousness (<http://fcmconference.org/img/CambridgeDeclarationOnConsciousness.pdf>, accessed March 2014).

2. Consider the global-workspace theory of consciousness [Baars, 2005] and the associated LIDA cognitive architecture [Franklin et al., 2012]. A central idea of this framework is that the brain contains many modules that receive and process input stimuli in an unconscious fashion. These then compete for attention, and the most interesting processed inputs bubble up to a ‘global workspace,’ where the news is broadcast to other parts of the brain, including action-selection centres where reinforcement learning can be done using the information. To construct a simple but potentially ‘conscious’ agent within this framework, we could give a robot three sensors – say for light, temperature, and energy level – and define reward functions based on these inputs – say, greater reward for more light, higher temperature, and higher energy. Upon measuring these variables, the robot evaluates how far its reward function is from a historical average along the three dimensions, and then the variable with highest deviation from the typical level of reward is chosen to be broadcast to the other parts of the robot’s computations. Rewards highly above average would represent conscious pleasure, and those significantly below, conscious pain. The robot uses the input signal to adjust its behaviour (e.g., avoid dark corners and seek bright windows). It records the broadcasted episodes in memory logs and can report on those memory logs when queried by the user.

The robot in the second example would satisfy the rough outline of consciousness according to the global-workspace account. Is it sentient? Most people would say not – after all, it’s such a simple process. Perhaps some people would have intuitions that it is sentient because they can see it as an embodied entity acting in the world, analogous to animals that we assume are conscious. But we could just as well have located this robot in a virtual world, with no display screen to evoke our emotional sympathies, and in that case many common-sense intuitions for it being sentient break down. I personally think the robot is marginally sentient, even in its non-embodied form, but I agree it’s nowhere near as sentient as a human, because it lacks so many other abilities and so much cognitive machinery. Thus, these examples suggest a graded, continuous character of sentience.

What’s going on with these trivial agents that fulfill consciousness criteria is something like Campbell’s law²: When we develop a simple metric for measuring something (sentience in this case), we can game the system by constructing degenerate examples of systems exhibiting that property that we don’t intuitively think of as sentient (or at least not very sentient). For instance, the mirror test [Gallup et al., 2002] is a standard approach for demonstrating self-awareness in animals, but with robots it degrades into meaninglessness, because, for instance, we could create a robot that has a machine-learned classifier for ‘I have a dye spot on my face,’ and if this fires, the robot touches its face with its hand.

²‘The more any quantitative social indicator is used for social decision-making, the more subject it will be to corruption pressures and the more apt it will be to distort and corrupt the social processes it is intended to monitor.’ [Campbell, 2011]

The usual solution to Campbell's law is to apply multiple metrics, and doing so would serve us well here. If a robot passes not just the standard mirror test but many variations, as well as a suite of other physical and mental feats, and if does so using non-gerrymandered algorithms – perhaps algorithms that bear some resemblance to what we have in our brains – then the robot is very likely conscious to a significant degree. If it can accomplish some of these tasks but not all, and if it uses weaker, less general algorithms, then it seems fair to call it 'less conscious.'

Of course, the concept of 'consciousness,' like the concept of 'tableness,' is up to us to define. We can make binary discriminations if we so choose, just like we can make very clear distinctions as to whether a given object is or is not a table in all cases. But this seems artificial to me, because there's probably not a single, crucial step in constructing a mind where everything we consider morally important resides, just as there's no single, crucial trait that suddenly makes an object qualify as a table. Rather, a brain has many characteristics, and it becomes gradually more important with its degree of sophistication.

Sloman [2010] notes that 'consciousness' is what he calls a 'polymorphic concept' in the sense that it can refer to many different particular things depending on the context. As a result, he says of consciousness: 'there cannot be a unitary explanation of [...] how the brain produces "it", nor a time at which "it" first exists in a fetus,' but, rather, 'The different phenomena falling under the polymorphic umbrella can be investigated separately [...].' Liu and Schubert [2010] present a table showing different types of mental abilities possessed by different kinds of existing AI agents, such as mental modelling, self-motivation, dialogue, logical planning, use of probabilities, and learning. And even within a single one of these traits, different systems have different degrees of refinement.

Any given operation that a brain does, by itself, looks totally trivial. It's just some neurons over here triggering some neurons over there in some patterned way. In a digital agent, it's just some if-then statements, variable updates, for loops, etc. But when these components are combined all together, we start to see something important emerging from them.

One might think that consciousness represents a sort of 'phase transition,' analogous to the difference between molecules in a solid vs. a liquid or a liquid vs. a gas. In this model, at some point the brain's dynamics become sufficiently complex that they operate in a fundamentally different way from how even slightly simpler versions would behave; there's some crucial ability that makes all the difference when put in place. This view seems implausible to me because we already see a continuum of brains of varying complexity in the animal kingdom, and neuroscience has not shown that at any particular species, there's a discontinuity in the brain's function, such that it exhibits very different dynamics from brains slightly below it. Even most of the abilities that were once thought to set humans apart from 'lower' animals have now been shown to be found, to varying degrees, in other animals. If sentience were like a binary light switch that suddenly turned on at some point in the animal kingdom, this would mean that at some point in the evolutionary past, a completely unconscious mother and father gave birth to a child that would grow up to

be fully conscious. But the change between a single generation of parents and children is small, and brains tend to be resilient and robust – not completely altered in the way they operate based on small perturbations of their structure. Rather, it seems much more natural to me to see the sentience of brains through evolutionary history as developing in a roughly continuous fashion.³

Sentience is like a symphony. The presence or absence of any single instrument doesn't stop the music – though some members of the orchestra are more important, like the conductor or piano player. Cognitive agents exhibiting simple algorithms that nonetheless bear some resemblance to what more complex animal brains do deserve to be called at least barely sentient and hence deserve at least a tiny bit of moral consideration.

The idea that sentience lies on a continuum is shared by many authors. [Broom \[2007\]](#) explains that ‘The degree of awareness in animals that can feel pain will vary.’ [DeGrazia \[2008\]](#) discusses (without committing to) a ‘sliding-scale model’ of moral status based on ‘the degree of your cognitive, affective, and social complexity.’ [DeGrazia \[2008\]](#) points out that even if we only care about sentience, it's reasonable to see sentience as coming in degrees. While discussing primarily the case of animal ethics, [DeGrazia \[2008\]](#) notes that this question also has relevance to embryos and fetuses.

[Bostrom \[2006\]](#) presents thought experiments that suggest varying degrees of consciousness for a given computational mind depending on the reliability and independence of its components or the fraction of its circuits that are parallelized. This is a different sort of gradation in consciousness than one assessed between *different* minds with different abilities, but it is consistent with the overall approach of deciding how much sentience we want to see in various physical processes, and it helps to break intuitions that sentience must obviously be binary.

Some authors have proposed extremely abstract, information-based definitions of consciousness and moral value. [Freitas \[1984\]](#) proposes a brain's ‘sentience quotient’ (SQ) as

$$\text{SQ} = \log_{10} \left(\frac{I}{M} \right),$$

where I is its information-processing rate in bits/second and M is its mass in kilogrammes.

[Floridi \[2006\]](#) proposes an ethic based on not causing, preventing, or removing entropy from what he calls the ‘infosphere,’ an extension of the biosphere. Here ‘entropy’ refers not to the quantity used in physics but to ‘*destruction* or *corruption* of informational objects.’ In general, [Floridi \[2006\]](#) aims to extend the biocentric view found in environmental ethics to an ‘ontocentric’ view of information ethics, incorporating both biological and non-biological systems.

Frameworks like these share my sense that moral value comes in gradations based on complexity, but I maintain a sentiocentric view, in which our moral

³Of course, some species became less sentient over their evolutionary histories, but the maximum level of sentience exhibited by any organism in the world tended to increase over time [[Gould, 1996](#)].

obligations focus on the wellbeing of individual agents; it's just that my notion of what kinds of agents may have wellbeing is broader than is generally assumed. Thus, while my position could look somewhat ontocentric, in practice it may diverge significantly from environmental or information ethics depending on relative assessments of sentience. For instance, it's plausible I would judge a minnow as being more sentient, and hence more intrinsically morally important, than an old-growth redwood tree. It's also important to note that caring about minnows and trees does *not* imply seeking to ensure their continued existence and reproduction [Horta, 2010], because we may think that suffering is in aggregate more prevalent than happiness among organisms in nature [Ng, 1995].

Premise 2: Artificial RL resembles, sometimes closely, RL in biological brains

The computational theory of RL has two main branches [Sutton and Barto, 1998]:

1. The biological side extends back more than a century, perhaps to Thorndike's 'Law of Effect,' the principle that when a good outcome follows an action, an animal is more likely to repeat that action the next time [Thorndike, 1911]. Countless psychological studies on conditioning patterns in animals followed in the subsequent decades. Some AI researchers as early as the 1950s and 1960s developed systems to mimic animal learning [Sutton and Barto, 1998, and references therein].
2. The mathematical side traces back to the theory of optimal control, the Bellman equations, and Markov decision processes in the 1950s. These would later provide theoretical underpinnings for RL models. [Sutton and Barto, 1998] In the 1980s, Richard Sutton and Andrew Barto developed temporal-difference (TD) learning methods, which allowed computational agents to update their action tendencies in an online fashion after every observation [Sutton, 1988].

Recent advances in neuroscience have demonstrated a surprising connection between biological and computational RL (e.g., Schultz et al. [1997], Seymour et al. [2004], Woergoetter and Porr [2008]). In an AI context, TD RL is driven by reward-prediction error δ , which is defined as

$$\delta = r + \gamma \hat{V}(s') - \hat{V}(s), \quad (1)$$

where $\hat{V}(s)$ is the previously predicted value of the current state s , $\hat{V}(s')$ is the previously predicted value of the next state s' , r is the reward received transitioning from s to s' , and γ is the discount factor for future rewards, e.g., $\gamma = 0.99$ [Woergoetter and Porr, 2008]. Neuroscience has found that phasic (i.e., a transient burst of) dopamine release in the midbrain represents a signal of reward-prediction error precisely analogous to the δ of TD RL. Scientists even have plausible models for the mechanisms by which certain brain regions process

inputs, compute the subtraction in equation (1), and broadcast this signal to update action tendencies [Glimcher, 2011].

The connection between computational RL and neuroscience is so robust that researchers typically take it for granted and focus on questions assuming the connection holds. Questions like: Does the brain have so-called ‘eligibility traces’ in the TD model that extend some credit to actions further back than the previous step [Pan et al., 2005]? Artificial RL uses function approximation to collapse high-dimensional state/action spaces [Sutton and Barto, 1998]; which neural networks in the brain serve this purpose? Do the basal ganglia implement an ‘actor-critic’ architecture [Joel et al., 2002, Khamassi et al., 2005, Maia, 2010]? Are recent advances in hierarchical RL mirrored in brain observations [Botvinick et al., 2009, Ribas-Fernandes et al., 2011, Diuk et al., 2013]? To what extent does the brain use not just standard model-free RL – in which the expected value of a state or state-action pair is estimated directly – but also model-based RL, in which estimation of transition probabilities among states is performed [Doll et al., 2012, Shteingart and Loewenstein, 2014]? Perhaps the brain has a model-free system for habit formation and a model-based system for goal-directed behaviour, and the two compete with each other for control [Daw et al., 2006]? Might the brain use policy-gradient methods to directly optimise action-inclination parameters without explicitly referring to states or actions [Shteingart and Loewenstein, 2014, and references therein]?

Several state-of-the-art RL algorithms are based on neuroevolution, in which populations of different neural-network weights and topologies are tried, and the best are selected (e.g., Koppejan and Whiteson [2011], Koutník et al. [2013]). Evolutionary approaches sometimes outperform TD methods [Stanley and Miikkulainen, 2002, Taylor et al., 2006, Gomez et al., 2008], and like other policy-search methods, they have advantages of handling partial state observability, allowing more flexible policy representations, and making it easier to deal with large or continuous action spaces [Whiteson, 2012, Schmidhuber, 2000]. At first glance we might assume that evolutionary algorithms are unlikely to occur within a single brain because they involve selective reproduction among populations of different neural networks. Hence we might see neuroevolutionary RL as less biologically plausible than TD. Of course, there’s a somewhat trivial sense in which even TD can be seen as a selection process (try different actions, and those action-inclination synapse connections that produced better outcomes ‘reproduce,’ i.e., have their connection weights strengthened), but it’s not a full evolutionary process in which neural groups literally copy themselves [Fernando et al., 2012]. However, there is a proposal, called the ‘neuronal replicator hypothesis,’ that the brain may actually copy patterns of neural activity with mutation, in a sense closer to neuroevolutionary RL [Fernando et al., 2010]. The jury is still out on this question.

Artificial RL can clearly have some implementation differences vis-à-vis real brains. For instance, computational RL algorithms may apply updates of many <state, action, reward, next-state> tuples at once, perhaps with biologically unrealistic mathematics for batch operations [Lange et al., 2012, and references therein], while in a real environment this information comes one at a time.

Premise 3: RL operations are relevant to an agent's welfare

Showing a similarity in cognitive operations between animals and computers is not inherently morally significant. For instance, humans and computers can both do addition, remember that Paris is the capital of France, respond to commands, and so on. Ethical questions come into play more significantly when the cognitive operations relate to an agent's wellbeing – its goal satisfaction, happiness and suffering, and subjective experience.

An RL system gives a computational agent goals that it aims to fulfill. The reward function defines an agent's satisfaction or lack thereof. Of course, it does so in a stylized way relative to the human brain, which has many layers of cognitive systems [Marcus, 2009] with many intricately hard-wired and learned responses. But fundamentally the difference is one of degree rather than kind: The human brain is vastly more complex than a simple RL agent, but both systems act in ways intended to further certain goals.

RL provides an overarching framework for understanding *why* organisms experience positive and negative valence [Wright, 1996]. Valence is the brain's 'currency' of value, and identifying cues and actions that correlate with higher-than-expected reward or punishment helps organisms navigate complicated and dangerous environments. The magnitude of an animal's reward in response to an event should approximate the value of that event in terms of its evolutionary fitness.

Liking is different from learning

At the same time, the learning mechanics of RL may not be the only or even primary object of moral consideration. Learning is distinct from liking, as well as from wanting [Berridge et al., 2009]. This makes sense when we understand the components of an RL system. Perhaps the reward values r that come in to the system trigger liking when they become conscious. Meanwhile, the predicted reward values are subtracted from observed reward values, and the difference δ is used to learn updated action inclinations. Finally, maybe the action inclinations themselves can trigger wanting depending on the organism's state, even without new reward signals or learning going on.

Dopamine is not the same as pleasure. Salamone et al. [2007] review reasons for this and summarise:

the idea that [dopamine] DA mediates pleasure has been seized upon by textbook authors, the popular press, filmmakers, and the internet, all of which has elevated DA from its hypothesized involvement in reward to an almost mythological status as a 'pleasure chemical' mediating not only euphoria and addiction, but also 'love'. Yet [...], the actual science is far more complicated. [...T]he classic emphasis on hedonia and primary reward is yielding to diverse lines of research that focuses on aspects of instrumental learning, pavlovian/instrumental interactions, reward prediction, incentive salience, and behavioral activation.

After training, dopamine spikes when a cue appears signaling that a reward will arrive, not when the reward itself is consumed [Schultz et al., 1997], but we know subjectively that the main *pleasure* of a reward comes from consuming it, not predicting it. In other words, in equation (1), the pleasure comes from the actual reward r , not from the amount of dopamine δ . Of course, a higher actual reward r in unexpected circumstances will produce more dopamine δ , which could be where dopamine’s association with pleasure came from.

In addition, the brain regions for learning and liking are not identical. A common assumption is that the ventral *striatum* plays the role of the critic in actor-critic RL models, possibly with assistance from the orbitofrontal cortex and amygdala [Maia, 2009, and references therein], while Aldridge and Berridge [2010] point out that the ventral *pallidum* also contains many ‘hedonic hotspots’ that amplify the sensation of liking.

This raises the ethical question: Which do we care about? Wanting? Liking? Something else? Perhaps libertarians, economists, and certain preference utilitarians are most sympathetic to what an agent wants, whether or not it’s associated with hedonic reward. That people would reject the possibility of imaginary bliss in order to accomplish their goals in the real world is the lesson of Nozick [1974]’s ‘experience machine’ thought experiment. Or is the experience machine just an argument against hedonically focused RL, as opposed to sophisticated, model-based RL that might include reward functions defined relative to what happens in the *actual* world? Also, drug addicts and wireheads may engage in uncontrollable self-stimulation because their cravings (‘wanting’) are so strong, even if they don’t enjoy (‘like’) the experience [Siskind, 2010, Peciña, 2008]. This seems like the wrong way to go; ‘wanting without liking is hell,’ suggests Hanson [2011]. Notwithstanding these points, both wanting and liking seem more complete in the presence of an RL framework; my guess is that whatever the ‘liking’ process is, we wouldn’t care about it as much if it happened in isolation without a broader context.

Consciousness in RL agents

Consciousness seems like another important part of the moral story, since many people only care about emotions that are consciously felt or desires that are consciously held. But as we saw in the discussion of Premise 1, consciousness comes on a continuum. When we examine some of the leading computational theories of consciousness [Seth, 2007], we see that most of them can be interpreted as suggesting that even relatively simple digital agents admit micro-scale degrees of consciousness.⁴ For example:

⁴From this list I have omitted the ‘biological theory’ of consciousness, according to which the experience of consciousness depends crucially on the specific electrochemical properties of biological brains [Block, 2009]. This theory doesn’t leave much room for machine sentience. However, I also find this approach the least plausible because it’s like a ‘God of the gaps’ viewpoint: There is a mysterious consciousness thing we don’t understand, so we’ll ‘explain’ it in a thought-stopping way by pointing to the electrochemical nature of biological brains. But the biological nature of brains doesn’t *do* anything to explain why, algorithmically, our brains feel confused about the so-called hard problem of consciousness. Hypothetical machine brains

- Global workspace theory. In the discussion of Premise 1 I showed how an elementary robot could be seen as implementing some of the most basic components of the global-workspace model of consciousness. Even a rudimentary object in the paradigm of object-oriented programming could be seen as marginally conscious on this account, insofar as it receives inputs, processes them via lower-level functions, returns the values of those functions ('broadcasts them') to other parts of its program as globally accessible state variables for further use in action selection, and stores the values as parameters in its 'memory' for later retrieval. One feature of global-workspace theory missing in the simple object-oriented agent is competition among multiple, parallel coalitions of 'unconscious' processing units, but it's not clear how essential it is to have many of these units rather than just one, and in any case, some more advanced agents, like the robot that focuses on the most salient of its input sensors, would have this sort of competition.
- Fame in the brain. [Dennett \[1991\]](#) rejects what he calls the fallacy of the 'Cartesian theater' – the idea of a crucial finish line in the brain where unconscious information all comes together and becomes seen by the conscious mind. Rather, [Dennett \[1991\]](#) explains, different information can be processed at different places and different times, recorded in memory, and accessed when needed. Consciousness is like fame [[Dennett, 1996](#)] or power held by a political coalition [[Dennett, 2001](#)]. For something to be conscious means it has wider reach and greater impact on other processes. Of course, there's not a binary distinction between being famous or obscure, powerful or weak, so this model suggests that even simple processes are slightly conscious / slightly famous. In particular, I take this view to imply that the state, action, and reward information that an RL agent distributes among its cognitive operations would be somewhat conscious.
- Integrated information theory. [Tononi \[2008\]](#) offers an account of consciousness as 'integrated information,' i.e., informative signal-processing units operating in a jointly dependent fashion. As [Tononi \[2008\]](#) notes, even a single photodiode is minimally conscious on this account, if only to a vanishing degree relative to large brains. Artificial RL agents would be more conscious than the photodiode due to processing more information in a more connected way.
- Higher-order theories. These views suggest that consciousness refers not to cognition related to direct performance but rather to meta-level aware-

implementing the exact same algorithms as our brains would say they feel the same confusion as to why they have phenomenal experience rather than being zombies, even though this theory declares such machines to be unconscious. Of course, if we so choose, we can adopt a phenomenal stance only toward biological brains, but this seems chauvinistic. If I developed a personal relationship with a future robot – in which we had intimate philosophical discussions, learned about each other's dreams and fears, and engaged in activities together – I would care about that robot, and I would regard it as having subjective experiences roughly as important as my own, regardless of what specific physics was implementing it.

ness and reporting of those lower-level thoughts [Lau and Rosenthal, 2011, and references therein]. Depending on exactly how these theories are cashed out, simple RL agents may display trivial forms of higher-order cognition. For example, consider an agent that receives a reward, updates its state-value estimates, and takes an action. It then records this history of events in a log file, and upon request from the user, the agent loads this file (‘thinks about its past first-order thoughts’) and prints the log history to the screen (‘subjectively reports its experience’). Alternatively, we could see an animation of an RL character moving on a screen as a kind of higher-order thought about what’s happening to the character, written not in words but in pictures. More advanced RL systems may feature non-trivial metacognitive algorithms for assessing performance of the first-order systems [Anderson et al., 2006].

Citing Broom [2006], Broom [2007] lists further criteria for consciousness in the context of animal welfare:

A sentient being is one that has some ability to evaluate the actions of others in relation to itself and third parties, to remember some of its own actions and their consequences, to assess risk, to have some feelings, and to have some degree of awareness.

We can see how each of these finds rudimentary implementation in at least some present-day RL systems.

- Evaluating others’ actions. Multiagent RL is a well established field [Littman, 1994, Busoniu et al., 2008, Shoham and Leyton-Brown, 2009]. Even single-agent RL systems can react to others’ behaviour as though the others were part of the environment, and model-based systems could potentially estimate transition probabilities for others’ behaviours in detail.
- Memory and imagination. Many RL tasks, like choosing the appropriate navigation direction in a T-shaped maze based on a starting observation, require remembering past information into the future to inform later decisions. The Long Short-Term Memory recurrent neural network is one approach that has been successfully employed with RL for this purpose [Bakker, 2001]. Some RL architectures remember previous experiences (‘episodic memories’) and use them for further offline learning via simulated experiences generated from those observations [Sutton, 1990, Bakker et al., 2003]. (Interestingly, the human brain also has a close connection between episodic memory of the past and imagination of future scenarios [Hassabis and Maguire, 2007], though I’m not sure whether it’s by the same kind of mechanism.)
- Risk-assessment. A model-based RL system can evaluate the probability of a transition to a negative state and use this to compute expected costs. A model-free system implicitly assesses risk by directly estimating the expected value of a state or state-action pair.

- Emotion. The numerical reward values observed by an RL system, in the context of other cognitive processes, could be seen as the crudest form of emotion. Zimmermann [1986] famously defined the emotion of pain as ‘an aversive sensory experience caused by actual or potential injury that elicits protective motor and vegetative reactions, results in learned avoidance, and may modify species-specific behaviour, including social behaviour.’ The avoidance and behaviour-modification parts of this definition follow straightforwardly from an RL framework. Protective motor and vegetative reactions could be understood in an RL context as an agent using input stimuli to identify itself as being in a state s of injury, which then triggers learned actions a appropriate for being in that state. Or the responses could be just hard-wired reflexes.
- Awareness. Model-based RL systems develop probability distributions for possible future outcomes (‘if I do X, I’ll likely enter state Y’). In a trivial sense, these can be seen as representing knowledge and predictions about oneself and the environment. As the models become more sophisticated and better compress data about the world, it will become more and more useful for these models to contain distinct network configurations that stand for ‘myself.’ When these networks become activated, the agent would be ‘self-aware’ [Schmidhuber, 2012].

Of course, not all RL systems have all of these cognitive features. This illustrates once again how the degree of consciousness of agents comes in gradations.

Conscious subsystems?

One objection to this perspective of seeing rudimentary levels of consciousness in simple systems is to point out that our own brains contain many subsystems that are arguably at least as complex as present-day RL agents, and yet we don’t perceive them as being conscious. My reply is that those subsystems may indeed be conscious to themselves. As Sloman [2010] notes: ‘a *part* of a whole animal or robot [may be] conscious of something that causes it to alter its (internal) behaviour [...] while the whole animal is not introspectively conscious of it.’ It’s true that those subsystems are not having significant influence on the parts of your brain that win control of slow and deliberative actions, store long-lasting memories, and verbalise your subjective experiences. But within their local brain neighbourhoods, those subsystems are having some influence and are exhibiting simplified versions of processes that we do call conscious when they’re done by higher, more powerful parts of the brain.

Why don’t we directly perceive these subsystems as being conscious? For a similar reason as why you don’t directly perceive me as being conscious. The processes in my brain, like the processes in these low-level components that aren’t globally broadcast, do not have enough influence on your verbal, memory, and deliberative-action centres for you to say that you perceive them. But rationally you can know that these processes are still doing things you consider

morally relevant, and when we look at the systems at a lower level, they may indeed be ‘conscious’ to themselves in a crude fashion.

[Schwitzgebel \[2012\]](#) observes that ‘There isn’t as radical a difference in kind as people are inclined to think between our favorite level of organization and higher and lower levels.’ His essay develops the idea of seeing the United States as conscious, being constituted of many complex subsystems that act in ways similar as the subsystems of an organism. If we think that only the highest level of an integrated system is conscious, then if the United States were conscious, its citizens would not be, and yet we don’t consider individual citizens morally unimportant. There is no single ‘finish line’ for consciousness; there are just lower levels of organisation that combine into higher levels, that combine into higher levels, each with its own degrees of complexity and nuance. Seeing consciousness in these systems is akin to seeing the ‘leaf shape’ in a fractal fern.

One might still insist that only the famous and powerful parts of a brain matter, and the lower-level systems are morally irrelevant unless they affect the higher-level outputs. But we recoil from such views when they’re applied higher up in our level of abstraction: We don’t think it’s right to ignore poor, powerless people and only care about those with money or political influence. Nor is it right to disregard the feelings of animals even though they can’t fight for their own interests. So why would it be right to completely ignore the components of our brains that failed to win control of our final verbal reports and explicit memories? The fact that toddlers, most non-human animals, and adult humans with severe verbal impairments can’t speak doesn’t nullify their moral significance [[Dombrowski, 1997](#)]. And if we imagine that you were injected with a sedative that blocked formation of memories, this would not then make it acceptable to inflict pain on you. In fact, this last example may not be purely hypothetical. The drug midazolam (also known as ‘versed,’ short for ‘versatile sedative’) is often used in procedures like endoscopy and colonoscopy. [von Delius et al. \[2007\]](#) surveyed doctors in Germany who indicated that during endoscopies using midazolam, patients would ‘moan aloud because of pain’ and sometimes scream. Most of the endoscopists reported ‘fierce defense movements with midazolam or the need to hold the patient down on the examination couch.’ And yet, because midazolam blocks memory formation, most patients didn’t remember this: ‘the potent amnesic effect of midazolam conceals pain actually suffered during the endoscopic procedure’ [[von Delius et al., 2007](#)]. While midazolam does prevent the hippocampus from forming memories, the patient remains conscious, and dopaminergic reinforcement-learning continues to function as normal [[Frank et al., 2006](#)].

One might agree that verbalisation and explicit memories per se are not the morally relevant endpoint of consciousness but insist instead that the global broadcast that normally precedes these things is. But if so, we have to explain why global broadcasting is somehow fundamentally different from more local broadcasting that the subsystems do in smaller regions. After all, the ‘global broadcasting’ that happens in most of our brains usually stays there, rather than being distributed all across planet Earth, yet it still matters to us.

Finally, as [Schwitzgebel \[2012\]](#)’s example of a conscious United States illus-

trates, the boundaries of where an agent begins or ends aren't necessarily sharp. Is the United States a separate organism from Canada, even if they engage in trade and cross-border migration? And what about Europe when people travel for vacation? Similar kinds of delineation issues arise in the context of simple RL agents: Which parts of the code are the 'agent,' and which are the 'environment'? It's not always clear, especially if the program is written in a single series of imperative statements without object-oriented organisation. Even for the case of people, our minds are hooked up to our bodies, which are heavily influenced by external objects in our surroundings. At what point does 'ourself' end and 'the external world' begin? There's not a hard separation; we are all fundamentally part of the same big system. When we talk about different entities, what we're actually doing is carving out conceptual boundaries around parts that are relatively connected and stable, in order to help us conceptualise and describe what's going on. We can do this 'carving out' process for RL agents while also recognising that they are part of a bigger, unified program, which may also matter in its own right.

Implication: Present-day artificial RL deserves a tiny bit of moral consideration

Contemporary artificial RL agents do not implement most of the functionality of human brains, or even, say, insect brains. But they do contain an important component of what drives goal-directed, welfare-relevant cognition in animals, namely RL, and they have traces of other morally salient characteristics, like emotion (in the form of their computing the reward function based on inputs) and consciousness (such as by broadcasting information updates). Programs equipped with RL have enough of these traits to act successfully in their virtual or physical worlds, showing that they are complete, if limited, agents.

If RL computations do matter at least a tiny bit, the next question is how much they matter relative to other priorities. At the moment I think they rank reasonably low on the list. For instance, fruit flies display rather complex brains compared with many current RL agents. Fruit flies demonstrate RL (e.g., [Tempel et al. \[1983\]](#)). They have 100,000 neurons, of which 200 contain dopamine [[Whitworth et al., 2006](#)]. In addition, [van Swinderen \[2005\]](#) suggests they may have 'the remote roots of consciousness.' Fruit flies are sufficiently intelligent to engage in all the necessary behaviours required for reproduction, repeatedly over millions of years.

In view of their greater cognitive functionality and degree of awareness, it's plausible that fruit flies matter, say, thousands of times more than present-day RL algorithms per learning update. (I'm making up this number, but it seems plausible given 100,000 fruit-fly neurons and the fact that an RL agent is more functionally complex than just one or two neurons.) On the other hand, computers can run thousands of learning updates for artificial agents in the time it would take a fruit fly to have one update. So it may be that, say, the RL algorithms running on a graduate student's laptop are roughly comparable

in importance to one insect. (Of course, this estimate is subject to substantial revision as we learn more, or depending on your ethical viewpoint.) But in total the world contains about ten billion billion insects⁵ and not nearly so many AI graduate students, so the welfare of insects is a vastly greater moral concern at this stage. But in the long run, as computing power grows and RL agents become increasingly sophisticated, RL looks set to become a pressing ethical consideration in its own right.

Unfortunately, even the welfare of insects and other invertebrates is not generally seen as a significant ethical issue, though the topic is receiving increasing attention [Lockwood, 1987, Mather, 2001, Mason, 2011], and methods of pain relief and euthanasia for invertebrates have been recommended [Cooper, 2011].

The analogy with laboratory-animal welfare is helpful, because RL research can be approached using similar frameworks as animal research [Winsby, 2013]. Central principles for the use of experimental animals are the ‘Three Rs’ [Russell and Burch, 1959]. Applied to RL, they would suggest that researchers

1. Replace the use of RL with other algorithms that less closely resemble an agent undergoing emotionally valenced experiences
2. Reduce the number of RL agents used
3. Refine RL algorithms to be more humane, such as by
 - using rewards instead of punishments
 - not hooking up RL algorithms to higher-level cognitive and emotional faculties
 - running fewer biologically inspired RL algorithms (like TD actor-critic value-function learning) and instead more abstract mathematical ones?

It’s not clear whether or how much to weigh different algorithms based on their biological plausibility. A very parochial view is to say that we only care about minds that are very similar to ours, including in their algorithmic constitution. So, for instance, if humans don’t use policy-gradient learning, then policy-gradient artificial RL would not be ethically significant. A more cosmopolitan view is to not focus so much on the specific algorithm, so long as it gives rise to comparable behaviour and adaptability to the world. In the extreme case, the cosmopolitan view might entail giving ethical consideration to giant lookup tables [Block, 1981], though in practice such brains are unlikely to be very common.

So it’s debatable how much mileage we can get by refining the type of RL algorithm used. Perhaps the more urgent form of refinement than algorithm selection is to replace punishment with rewards within a given algorithm. RL systems vary in whether they use positive, negative, or both types of rewards:

⁵This figure is quoted in dozens of sources – e.g., Berenbaum [1995] – though I’m unable to find the original calculation.

- In certain RL problems, such as maze-navigation tasks discussed in [Sutton and Barto \[1998\]](#), the rewards are only positive (if the agent reaches a goal) or zero (for non-goal states).
- Sometimes a mix between positive and negative rewards⁶ is used. For instance, [McCallum \[1993\]](#) put a simulated mouse in a maze, with a reward of 1 for reaching the goal, -1 for hitting a wall, and -0.1 for any other action.
- In other situations, the rewards are always negative or zero. For instance, in the cart-pole balancing system of [Barto et al. \[1990\]](#), the agent receives reward of 0 until the pole falls over, at which point the reward is -1. In [Koppejan and Whiteson \[2011\]](#)'s neuroevolutionary RL approach to helicopter control, the RL agent is punished either a little bit, with the negative sum of squared deviations of the helicopter's positions from its target positions, or a lot if the helicopter crashes.

Just as animal-welfare concerns may motivate incorporation of rewards rather than punishments in training dogs [[Hiby et al., 2004](#)] and horses [[Warren-Smith and McGreevy, 2007](#), [Innes and McBride, 2008](#)], so too RL-agent welfare can motivate more positive forms of training for artificial learners. [Pearce \[2007\]](#) envisions a future in which agents are driven by 'gradients of well-being' (i.e., positive experiences that are more or less intense) rather than by the distinction between pleasure versus pain. However, it's not entirely clear where the moral boundary lies between positive versus negative welfare for simple RL systems. We might think that just the sign of the agent's reward value r would distinguish the cases, but the sign alone may not be enough, as the following section explains.

What's the boundary between positive and negative welfare?

Consider an RL agent with a fixed life of T time steps. At each time t , the agent receives a non-positive reward $r_t \leq 0$ as a function of the action a_t that it takes, such as in the pole-balancing example. The agent chooses its action sequence $(a_t)_{t=1..T}$ with the goal of maximising the sum of future rewards:

$$\sum_{t=1}^T r_t(a_t).$$

⁶As [Barto et al. \[1990\]](#) note, 'negative reinforcement' in behaviourist psychology refers to reinforcing actions that remove an unpleasant stimulus, such as taking drugs to reduce painful withdrawal symptoms [[Flora, 2004](#)]. What I refer to by 'negative reward value' in an RL context could be used in learning either what behaviourists call 'negative reinforcement' (which increases inclination to take an action that removes an unpleasant stimulus) or 'positive punishment' (which decreases inclination to take an action that causes an unpleasant stimulus). A parallel situation applies for 'positive reinforcement' and 'negative punishment.' I've avoided using the phrases 'negative reinforcement' and 'positive reinforcement' in this article to reduce confusion, but when I speak of 'negative rewards,' all I mean are reward values that are negative numbers ($r < 0$), and positive numbers for 'positive rewards' ($r > 0$), without intending to suggest behaviourist connotations.

Now suppose we rewrite the rewards by adding a huge positive constant c to each of them, $r'_t = r_t + c$, big enough that all of the r'_t are positive. The agent now acts so as to optimise

$$\sum_{t=1}^T r'_t(a_t) = \sum_{t=1}^T (r_t(a_t) + c) = Tc + \sum_{t=1}^T r_t(a_t).$$

So the optimal action sequence is the same in either case, since additive constants don't matter to the agent's behaviour.⁷ But if behaviour is identical, the only thing that changed was the sign and numerical magnitude of the reward numbers. Yet it seems absurd that the difference between happiness and suffering would depend on whether the numbers used by the algorithm happened to have negative signs in front. After all, in computer binary, negative numbers have no minus sign but are just another sequence of 0s and 1s, and at the level of computer hardware, they look different still. Moreover, if the agent was previously reacting aversively to harmful stimuli, it would continue to do so. As Lenhart K. Schubert explains:⁸

If the shift in origin [to make negative rewards positive] causes no behavioural change, then the robot (analogously, a person) would still behave as if suffering, yelling for help, etc., when injured or otherwise in trouble, so it seems that the pain would not have been banished after all!

So then what distinguishes pleasure from pain? Why do I feel that pain has a different emotional texture than pleasure, rather than both feelings lying on a single scale of valuation?

One possibility is that the 'hedonic zero point' is determined by whether I would prefer to have a given experience rather than nothing. The RL agent that we considered in the above example had a fixed lifetime of T , but if it had a variable lifetime, then its actions would depend substantially on whether the r_t values were positive or negative. If they were negative, the agent would seek to end its life ('commit suicide') as soon as possible.⁹ If they were positive,

⁷The same would also be true if the agent optimised discounted future rewards over a fixed finite or infinite lifetime.

Also, there could be rare cases where behaviour is not identical if the environment depends on the numerical reward values. For example, suppose a robot prints out its last numerical reward to an observing roboticist. If the roboticist sees a positive number, he smiles, and the robot's image sensors detect this as the 'roboticist is happy' state. If the roboticist sees a negative number, he frowns, and the robot enters the 'roboticist is unhappy' state. Dependence of the environment on the literal reward values is not typical, especially for simple systems like the pole-balancing agent.

⁸This quotation comes from spring 2014 lecture notes (<http://www.cs.rochester.edu/users/faculty/schubert/191-291/lecture-notes/23>, accessed March 2014) for a course called 'Machines and Consciousness.'

⁹Actually, this would depend on the initial value given to the death state for the agent. Since death is an absorbing state after which no further learning happens, the agent can't empirically update its value for the death state. If the initial value was 0, the agent would seek death if its life was full of negative rewards.

it would seek to live as long as it could, because this would make the sum of rewards larger.

This explanation may sound plausible due to its analogy to familiar concepts, but it seems to place undue weight on whether an agent’s lifetime is fixed or variable. Yet I would still feel pain and pleasure as being distinct even if I knew exactly when I would die, and a simple RL agent has no concept of death to begin with.

A more plausible account is that the difference relates to ‘avoiding’ versus ‘seeking.’ A negative experience is one that the agent tries to get out of and do less of in the future. For instance, injury should be an inherently negative experience, because if repairing injury was rewarding for an agent, the agent would seek to injure itself so as to do repairs more often. If we tried to reward *avoidance* of injury, the agent would seek dangerous situations so that it could enjoy returning to safety.¹⁰ Injury needs to be something the agent wants to get as far away from as possible. So, for example, even if vomiting due to food poisoning is the best response you can take given your current situation, the experience should be negative in order to dissuade you from eating spoiled foods again.

Still, the distinction between avoiding and seeking isn’t always clear. We experience pleasure due to seeking and consuming food but also pain that motivates us to avoid hunger. Seeking one thing is often equivalent to avoiding another. Likewise with the pole-balancing agent: Is it seeking a balanced pole, or avoiding a pole that falls over?

In animal brains, we may be able to tease out some of the distinction between seeking and avoiding at a physiological level. Daw et al. [2002] review evidence that humans have two *separate* motivational systems, one appetitive and one aversive. While dopamine is associated with approach, serotonin is associated with inhibition (among many other things). In AI, RL uses a single scalar reward-prediction error δ , which may be positive or negative with any magnitude, but in the brain, firing rates can only be positive, so presumably a different signal (possibly serotonin) is needed to encode significantly negative errors [Daw et al., 2002]. It’s true that dopamine has a baseline firing rate, and when expected rewards are omitted, dopamine firing drops below baseline, but the magnitude of this effect doesn’t seem sufficient on its own. Based on these lines of reasoning, Daw et al. [2002] develop a computational model in which serotonin acts as the opponent to dopamine:

$$\text{total prediction error} = \text{dopamine} - \text{serotonin}.$$

The model is consistent with findings that serotonin is associated with harm avoidance [Cloninger, 1986, Hansenne and Ansseau, 1999]. Daw et al. [2002]

¹⁰This example comes from Lenhart K. Schubert’s spring 2014 lecture notes (<http://www.cs.rochester.edu/users/faculty/schubert/191-291/lecture-notes/23>, accessed March 2014) for a course called ‘Machines and Consciousness.’ These thought experiments are not purely academic. We can see an example of maladaptive behaviour resulting from an association of pleasure with injury when people become addicted to the endorphin release of self-harm.

also explain how serotonin can help implement an RL system designed to optimise long-run *undiscounted* average reward [Mahadevan, 1996], and their model accounts for the influential opponent-process theory of motivation in psychology [Solomon and Corbit, 1974]. Because serotonin in this model is hypothesised to encode a running-average reward as opposed to current reward, and dopamine is hypothesised to encode a running-average punishment as opposed to current punishment, Daw et al. [2002]’s account also explains why dopamine is observed to rise in response to aversive events.

A final explanation for why pain feels different from pleasure may be that the emotional texture of experiences varies based on the pattern of other neural processes that go on when the experience is triggered. Even among negative experiences, we can distinguish among physical pain, depression, fear, embarrassment, guilt, and so on. Each has its own distinct character based on the orchestra of other cognitive instruments that are playing when it happens. Aldridge and Berridge [2010] suggest a similar idea for the case of positive experiences:

Much of human pleasure has cognitive qualities that infuse uniquely human properties, and it is likely that abstract or higher pleasures depend on cortical brain areas for those qualities. [...T]he particular pattern of coactivated cortical circuits would resolve the high level cognitive features of a pleasantness gloss on sensations or actions.

Where does all of this leave our pole-balancing agent? Does it suffer constantly, or is it enjoying its efforts? Likewise, is an RL agent that aims to accumulate positive rewards having fun, or is it suffering when its reward is suboptimal? Of course, as with sentience itself, our evaluations of the emotional valences of these cases are up to us to decide, but our uncertainty in how to make this choice is a reason to exercise caution before we run vast numbers of RL computations – perhaps even those that only use positive rewards ($r > 0$).

While hedonic setpoints vary among humans, with some people enduring chronic depression and others enjoying frequent satisfaction, this fact presents somewhat less of a puzzle than we have with the pole-balancing agent, because depressed humans behave differently than happy ones, whereas the pole balancer behaves exactly the same with a uniform shift in its reward values. For more complex, human-like agents, if they behave similarly to depressed people, perhaps this is an indication of net suffering, and the opposite if they behave similarly to happy people. But it’s dubious to extend this heuristic much beyond the realm of agents with close resemblance to mammals. Instead, we need to develop more general principles.

Commercial applications of RL

While many state-of-the-art RL systems currently dwell in academia, in the long run I expect most RL computations to happen in the industrial and consumer domains, once technologies using RL become commercialised.

RL has been proposed for many purposes, including

- playing backgammon [Tesauro, 1994] and Othello [van Eck and van Wezel, 2008]
- elevator scheduling [Crites and Barto, 1996]
- job scheduling [Aydin and Öztemel, 2000]
- task scheduling in petroleum production [Aissani and Beldjilali, 2009]
- web spidering [Rennie and McCallum, 1999]
- stock trading [Lee et al., 2007]
- optimising drug delivery [Gaweda et al., 2006, Malof and Gaweda, 2011]
- military simulations [Sidhu et al., 2006, Collins et al., 2013, Papadopoulos et al., 2013].

One of the fields most closely tied with RL is robotics, because it deals with autonomous agents that need to act in the world. In fact, one review article suggested: ‘The relationship between [robotics and RL] has sufficient promise to be likened to that between physics and mathematics’ [Kober et al., 2013].

Video games may be another hotbed of RL in the future, since RL offers the promise of creating more realistic non-player characters (NPCs). Currently many ‘game AIs’ use hard-coded rules, but these require effort to build, and machine-learning techniques like RL offer the prospect of automating and refining NPC behaviour [Patel et al., 2011]. The topic has attracted much academic interest [Björnsson et al., 2008, Amato and Shani, 2010]. One popular example of RL for video games is learning to play Super Mario [Karakovskiy and Torgelius, 2012]. RL has also been applied to the widely studied Open Racing Car Simulator [Loiacono et al., 2010].

RL in video games presents one of the clearest cases of ethical concern, because the games are visually compelling and many times violent, making it relatively easier to evoke our emotional sympathies. RL has been suggested for first-person shooter games, and a natural way to train enemy NPCs is to inflict punishment on them when they’re killed. For instance, McPartland and Gallagher [2011] report their design: ‘A large penalty (-1.0) was given when the bot was killed, and a small penalty (-0.000002) was given when the bot was wounded.’ As RL is increasingly applied in video games, and as the AI algorithms involved become increasingly lifelike, the ethical questionableness of punishing NPCs will grow.

We can imagine some mitigation proposals, along the lines of the Three Rs discussed previously, that would allow gamers to enjoy greater NPC intelligence without quite so much ethical concern. For instance, if the NPCs can be trained extensively offline, so that in the video game they only execute their previously learned rules rather than continuing to learn on the fly during game play, this would reduce the amount of RL required. In games such as *Creatures* [Grand and Cliff, 1998], where players can choose how much to punish their AI pets, game designers could build in limitations on the amount of suffering the AIs

could endure before they faint, or die, or otherwise terminate the negative input processes. Perhaps the video-game industry could develop protocols for humane game design, pushed along by government regulation or voluntary standards. That said, this might be a challenging proposition, considering that many people already think that exposure to video-game violence is wrong, while the gaming industry has done little in response. We might also worry whether regulations would drive the inhumane games underground, as being what the ‘cool, hard-core’ gamers play.

Whitby [2008] cites examples in which humans have physically abused robots with which they interacted. This may be troubling, but from the perspective of machine welfare, our moral evaluation depends (ignoring instrumental considerations) on whether the robots were wired to respond aversively to the damage they underwent. Moreover, I think the vast majority of potential suffering that robots and other RL agents will experience in the future will not be due to abuse by angry human owners but rather will be built into their utility functions and will result from ‘natural’ interactions with the environment. While perhaps less emotionally salient to observers, this systemic suffering will be far more common, and insofar as it will be at least somewhat preventable, it deserves ethical priority.

It may be easiest to engender concern for RL when it’s hooked up to robots and video-game characters because these agents have bodies, perhaps including faces that can display their current ‘emotional states.’ In fact, interacting with another agent, and seeing how it behaves, can incline us toward caring about it whether it has a mind or not. For instance, children become attached to their dolls, and we may sympathise with cartoon characters on television. In contrast, it’s harder to care about a batch of RL computations with no visualization interface being performed on some computing cluster, even if their algorithms are morally relevant. It’s even harder to imagine soliciting donations to an advocacy organisation – say, People for the Ethical Treatment of Reinforcement Learners – by pointing to a faceless, voiceless algorithm. Thus, our moral sympathies may sometimes misfire, both with false positives and false negatives. Hopefully legal frameworks, social norms, and philosophical sophistication will help correct for these biases.

Some feel that placing special emphasis on those we’re close to is not a bias but a feature of their moral frameworks. For instance, Coeckelbergh [2010] proposes a social-relational paradigm for robot ethics based around ‘relations between various entities, human and non-human, which are inter-dependent and adapt to one another.’ This brings robots into the ethical picture ‘provided that they participate in the social life.’ Sadly, such an approach gives less weight to vastly greater numbers of RL agents that may suffer invisibly in back-end industrial computation centres.

The applications of RL in gaming, robotics, and industry are manifold, and they seem likely to expand in the coming decades. That said, these uses of RL are relatively minor compared with what we might anticipate in the far future, if humanity or non-human AIs expand into the galaxy, harnessing the energy of stars to create prodigious amounts of computing power, and requiring massive

numbers of robots and other, possibly RL-based agents as workers. The ethical risks in scenarios like these are, to borrow a double entendre from [Bostrom \[2003\]](#), ‘astronomical.’

Do non-RL agents matter?

If a primary evolutionary purpose of pleasure and pain is to serve as the reward/punishment values in an RL system, do organisms lacking RL not experience pleasure or pain? Perhaps this is one reason why, when scientists ask questions like ‘Do bugs feel pain?’, they look for abilities like RL beyond mere reflex behaviour [[EFSA, 2005](#)].

Are there other features of an organism that matter ethically besides RL? What if it’s an apparently goal-directed agent exhibiting complex but not adaptive behaviours, such as NPCs in most modern video games that run using fixed if-then rules (analogous to stimulus-response reflexes in animals) and non-RL algorithmic methods like pathfinding algorithms?

Many industrial-control systems, including simple thermostats, likewise respond to environmental conditions by following pre-programmed rules rather than learning the rules. An RL agent could *become* a thermostat, learning to turn on the heat when it entered the ‘cold’ state and turn on the cooling when it entered the ‘hot’ state. Once trained, the RL agent might act just like the thermostat. But the thermostat didn’t have the training phase.

It seems plausible to care about goal-directed agents even if they didn’t have a training phase. After all, they still appear to have what we would call preferences; things can still be better or worse for them. When an AI NPC in a first-person shooter is killed, it still fails to accomplish what it was striving for, even if that striving was only being executed by pre-programmed rules.

[Torrance \[2000\]](#) suggests that even if AIs aren’t sentient, one might still value them ethically for attributes like autonomy, intelligence, and cognitive sophistication. My discussion here is similar, except that on my view of sentience, those attributes might indeed be rightly considered part of what makes an agent sentient. But fundamentally it doesn’t matter whether we call these criteria part of sentience or part of moral valuation beyond sentience, because the intuitions and conclusions seem to be similar.

I think sufficiently complex rule-based agents probably do have ethical importance, even if they don’t perform RL. To pump this intuition, imagine if you took an elderly person and disrupted her brain’s RL modules. She wouldn’t update her action-value estimates, but she would continue to operate with her existing, well trained estimates. Presumably this person would still seem fairly normal, someone whom we could be friends with and who could tell us about her experiences, at least for a while until the inability to update action tendencies started to cause problems. I would still care a lot about such a person. And in fact, this example may not be purely imaginary. Parkinson’s disease is marked by loss of dopamine [[Kish et al., 1988](#)]. This impairs performance on prediction and reinforcement-learning tasks [[Knowlton et al., 1996](#), [Frank et al., 2004](#)].

It may be that our experience architecture is built at least partly for the purpose of RL, but this doesn't mean that if you eliminate RL, you eliminate experiences. The reward/punishment signals can keep coming, even if the dopamine or other neurochemicals that encode prediction errors stop working. But if we can care about hypothetical humans whose RL abilities have been removed, why not also care about video-game NPCs that act in a goal-directed fashion without any RL training? If it helps to arouse our sympathies, we could imagine training them with RL and then turning the RL off.

Or maybe the latent architecture is also quite relevant. The hypothetical person whose RL capacity was shut down presumably would still have neural systems for input valuation – for turning signals from the environment and other brain regions into broadcasts that ‘This feels good’ or ‘This feels bad.’ A video-game AI using fixed if-then rules does not have an overt valuation function. That said, some planning agents do explicitly optimise a sum of future rewards even though they don't update actions using reinforcement learning (e.g., [Liu and Schubert, 2010]). Indeed, many kinds of optimisation processes can be seen as choosing actions to increase rewards relative to some reward function.

Note that if we extend ethical significance to goal-directed agents even if they don't use RL, our moral circle of concern expands very wide. Query optimisers, path planners, machine-translation systems, and many other routine computer programs make choices with an eye to optimising a goal function. We can even see this kind of process throughout physics, such as when a protein folds so as to minimise energy [Wales, 2003] or when particles choose a movement trajectory so as to minimise ‘action’ [Gray, 2009].

Still, even if our ethical valuation assigns nonzero moral weight to these things, the weight can be exceedingly minuscule, so maybe the practical implications are not as drastic as they might seem. In general, deciding how much to value different features of the universe is a challenging enterprise. It requires both the heart, to assess what kinds of entities we feel compassion towards, and the head, to make our intuitions consistent and identify sources of suffering that we might not ordinarily have noticed. This paper has only begun to scratch the surface.

Robustness to other views on consciousness

Caring a little bit about RL algorithms seems a natural extension of a graded view of sentience. If subjective experience is a stance we adopt toward physical processes, then processes that have at least minimal degrees of morally relevant characteristics matter a small amount.

But my graded, ‘phenomenal stance’ approach to consciousness is not universally shared. For example, Torrance [2000] explicitly rejects as absurd the thesis I advanced in Premise 1:

How would we be able to tell if an [AI] were genuinely conscious, rather than just behaving outwardly as conscious? One answer links the matter back to ethical judgment: to claim that x is genuinely

conscious may be thought to be definitionally dependent upon the adoption of the appropriate moral attitude towards x. But surely my own consciousness is a matter of objective fact, known to me. Your failure to ascribe consciousness to me is not, therefore, a mere matter of your making a certain moral decision; it is factually false.¹¹

How does my argument fare for those who feel that whether an agent is conscious is binary? Different theories of consciousness will give different answers, but many of them should at least admit the possibility that RL programs *might* be conscious. Perhaps the likelihood is low, but it should be nonzero. In that case, RL programs would still matter ethically at least a tiny bit in *expected value*. The conclusions would then be similar as what I argued for, with ‘probability of sentience’ playing the role that ‘degree of sentience’ had in my discussion.

Of course, if sentience were a factual, binary property rather than a subjective, fuzzy category, then in the long run, once people understood consciousness well enough, they could potentially conclude with high certainty that RL programs didn’t suffer (or did, as the case may be). At that point the practical implications might diverge. Until then, it seems that many views warrant at least thinking twice about the ethical implications of large-scale RL, even if it doesn’t yet constitute one of the world’s most pressing moral problems.

Acknowledgements

Carl Shulman first suggested to me the potential ethical relevance of RL and also refined my understanding of consciousness more generally. Thanks also to David Althaus, Nick Bostrom, Mayank Daswani, Oscar Horta, Rupert McCallum, Joe Mela, Jacob Scheiber, Buck Shlegeris, two anonymous reviewers, and several other people for comments on a draft of this piece.

References

- Nassima Aissani and Bouziane Beldjilali. Dynamic scheduling in petroleum process using reinforcement learning. In Abdelmalek Amine, Otmane Aït Mohamed, and Zakaria Elberrichi, editors, *CIIA*, 2009.
- J. Wayne Aldridge and Kent C. Berridge. Neural coding of pleasure: “rose-tinted glasses” of the ventral pallidum. In Morten L. Kringelbach and Kent C. Berridge, editors, *Pleasures of the Brain*, Series in Affective Science, pages 62–73. Oxford University Press, 2010.

¹¹How would I reply to this? I would firstly deny a hard distinction between first- and third-person viewpoints; everything is just a perception of one sort or another, whether of external stimuli or internal brain states. Secondly, even if we do regard first-person experience as a privileged realm of truth, what do we do with it? All we can say is that I’m conscious in this special first-person way. If we refer to consciousness as ‘this experience I’m having now,’ we can say nothing about other minds, whose brain states are not identical with ours. If we define consciousness as ‘*kind of like* this experience I’m having now in some relevant ways,’ then we get into third-person traits of minds, which cluster into fuzzy, non-binary categories.

- Christopher Amato and Guy Shani. High-level reinforcement learning in strategy games. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems*, pages 75–82. International Foundation for Autonomous Agents and Multiagent Systems, 2010.
- Michael Anderson and Susan Leigh Anderson, editors. *Machine Ethics*. Cambridge University Press, New York, NY, 2011.
- Michael L. Anderson, Tim Oates, Waiyian Chong, and Don Perlis. The metacognitive loop I: Enhancing reinforcement learning with metacognitive monitoring and control for improved perturbation tolerance. *Journal of Experimental & Theoretical Artificial Intelligence*, 18(3):387–411, 2006. doi:[10.1080/09528130600926066](https://doi.org/10.1080/09528130600926066).
- M. Emin Aydin and Ercan Öztemel. Dynamic job-shop scheduling using reinforcement learning agents. *Robotics and Autonomous Systems*, 33(2–3): 169–178, 2000. doi:[10.1016/S0921-8890\(00\)00087-7](https://doi.org/10.1016/S0921-8890(00)00087-7).
- Bernard J. Baars. Global workspace theory of consciousness: Toward a cognitive neuroscience of human experience. *Progress in Brain Research*, 150:45–53, 2005. doi:[10.1016/S0079-6123\(05\)50004-9](https://doi.org/10.1016/S0079-6123(05)50004-9).
- Bram Bakker. Reinforcement learning with long short-term memory. In *NIPS*, pages 1475–1482, 2001.
- Bram Bakker, Viktor Zhumatiy, Gabriel Gruener, and Jürgen Schmidhuber. A robot that reinforcement-learns to identify and memorize important previous observations. In *Intelligent Robots and Systems, IEEE/RSJ International Conference on*, volume 1, pages 430–435. IEEE, 2003. doi:[10.1109/IROS.2003.1250667](https://doi.org/10.1109/IROS.2003.1250667).
- Andrew G. Barto, Richard S. Sutton, and Charles W. Anderson. Neuron-like adaptive elements that can solve difficult learning control problems. In Joachim Diederich, editor, *Artificial Neural Networks*, pages 81–93. IEEE Press, 1990.
- May R. Berenbaum. *Bugs in the System: Insects and Their Impact on Human Affairs*. Perseus Books, Cambridge, MA, 1995.
- Kent C. Berridge, Terry E. Robinson, and J. Wayne Aldridge. Dissecting components of reward: liking, wanting, and learning. *Current Opinion in Pharmacology*, 9(1):65–73, 2009. doi:[10.1016/j.coph.2008.12.014](https://doi.org/10.1016/j.coph.2008.12.014).
- Yngvi Björnsson, Vignir Hafsteinsson, Ársæll Jóhannsson, and Einar Jónsson. Efficient use of reinforcement learning in a computer game. In *Proceedings of International Journal of Intelligent Games & Simulation*, 2008.
- Ned Block. Psychologism and behaviorism. *Philosophical Review*, 90(1):5–43, 1981.

- Ned Block. Comparing the major theories of consciousness. In Michael S. Gazzaniga, editor, *The Cognitive Neurosciences*, pages 1111–1122. MIT Press, 2009.
- Nick Bostrom. Astronomical waste: The opportunity cost of delayed technological development. *Utilitas*, 15(3):308–314, 2003. doi:[10.1017/S095382080004076](https://doi.org/10.1017/S095382080004076).
- Nick Bostrom. Quantity of experience: Brain-duplication and degrees of consciousness. *Minds and Machines*, 16(2):185–200, 2006. doi:[10.1007/s11023-006-9036-0](https://doi.org/10.1007/s11023-006-9036-0).
- Matthew M. Botvinick, Yael Niv, and Andrew C. Barto. Hierarchically organized behavior and its neural foundations: A reinforcement learning perspective. *Cognition*, 113(3):262–280, 2009. doi:[10.1016/j.cognition.2008.08.011](https://doi.org/10.1016/j.cognition.2008.08.011).
- Rodney Brooks. Will robots rise up and demand their rights? *Time*, June 2000.
- Donald M. Broom. Cognitive ability and sentience: Which aquatic animals should be protected? *Diseases of Aquatic Organisms*, 75(2):99–108, 2007. doi:[10.3354/dao075099](https://doi.org/10.3354/dao075099).
- Donald Maurice Broom. The evolution of morality. *Applied Animal Behaviour Science*, 100(1–2):20–28, 2006. doi:[10.1016/j.applanim.2006.04.008](https://doi.org/10.1016/j.applanim.2006.04.008).
- Lucian Busoniu, Robert Babuska, and Bart De Schutter. A comprehensive survey of multiagent reinforcement learning. *Systems, Man, and Cybernetics, Part C: Applications and Reviews, IEEE Transactions on*, 38(2):156–172, 2008. doi:[10.1109/TSMCC.2007.913919](https://doi.org/10.1109/TSMCC.2007.913919).
- David J. Calverley. Android science and the animal rights movement: Are there analogies? In *Toward Social Mechanisms of Android Science: A CogSci 2005 Workshop*, pages 127–136, 2005.
- Donald T. Campbell. Assessing the impact of planned social change. *Journal of MultiDisciplinary Evaluation*, 7(15):3–43, 2011.
- C. Robert Cloninger. A unified biosocial theory of personality and its role in the development of anxiety states. *Psychiatric Developments*, 4(3):167–226, 1986.
- Mark Coeckelbergh. Robot rights? towards a social-relational justification of moral consideration. *Ethics and Information Technology*, 12(3):209–221, 2010. doi:[10.1007/s10676-010-9235-5](https://doi.org/10.1007/s10676-010-9235-5).
- Andrew J. Collins, John Sokolowski, and Catherine Banks. Applying reinforcement learning to an insurgency agent-based simulation. *The Journal of Defense Modeling and Simulation: Applications, Methodology, Technology*, 2013. doi:[10.1177/1548512913501728](https://doi.org/10.1177/1548512913501728).

- John E. Cooper. Anesthesia, analgesia, and euthanasia of invertebrates. *ILAR Journal*, 52(2):196–204, 2011. doi:[10.1093/ilar.52.2.196](https://doi.org/10.1093/ilar.52.2.196).
- Robert H. Crites and Andrew G. Barto. Improving elevator performance using reinforcement learning. In *Advances in Neural Information Processing Systems 8 (NIPS 1995)*, pages 1017–1023, 1996.
- Nathaniel D. Daw, Sham Kakade, and Peter Dayan. Opponent interactions between serotonin and dopamine. *Neural Networks*, 15(4–6):603–616, 2002. doi:[10.1016/S0893-6080\(02\)00052-7](https://doi.org/10.1016/S0893-6080(02)00052-7).
- Nathaniel D. Daw, Yael Niv, and Peter Dayan. Actions, policies, values, and the basal ganglia. In Erwan Bezdard, editor, *Recent Breakthroughs in Basal Ganglia Research*, pages 91–106. Nova Science Publishers, 2006.
- David DeGrazia. Moral status as a matter of degree? *The Southern Journal of Philosophy*, 46(2):181–198, 2008. doi:[10.1111/j.2041-6962.2008.tb00075.x](https://doi.org/10.1111/j.2041-6962.2008.tb00075.x).
- Daniel C. Dennett. *Consciousness Explained*. Little, Brown and Co., Boston, MA, 1991.
- Daniel C. Dennett. Consciousness: More like fame than television. In Christa Maar, Ernst Pöppel, and Thomas Christaller, editors, *Die Technik auf dem Weg zur Seele*. Rowohlt, 1996.
- Daniel C. Dennett. Are we explaining consciousness yet? *Cognition*, 79(1–2): 221–237, 2001. doi:[10.1016/S0010-0277\(00\)00130-X](https://doi.org/10.1016/S0010-0277(00)00130-X).
- Carlos Diuk, Karin Tsai, Jonathan Wallis, Matthew Botvinick, and Yael Niv. Hierarchical learning induces two simultaneous, but separable, prediction errors in human basal ganglia. *The Journal of Neuroscience*, 33(13):5797–5805, 2013. doi:[10.1523/JNEUROSCI.5445-12.2013](https://doi.org/10.1523/JNEUROSCI.5445-12.2013).
- Bradley B. Doll, Dylan A. Simon, and Nathaniel D. Daw. The ubiquity of model-based reinforcement learning. *Current Opinion in Neurobiology*, 22(6): 1075–1081, 2012. doi:[10.1016/j.conb.2012.08.003](https://doi.org/10.1016/j.conb.2012.08.003).
- Daniel A. Dombrowski. *Babies and Beasts: The Argument from Marginal Cases*. University of Illinois Press, Champaign, IL, 1997.
- EFSA. Aspects of the biology and welfare of animals used for experimental and other scientific purposes. *EFSA Journal*, 292:1–46, 2005.
- Chrisantha Fernando, Richard Goldstein, and Eörs Szathmáry. The neuronal replicator hypothesis. *Neural Computation*, 22(11):2809–2857, 2010. doi:[10.1162/NECO.a.00031](https://doi.org/10.1162/NECO.a.00031).
- Chrisantha Fernando, Eörs Szathmáry, and Phil Husbands. Selectionist and evolutionary approaches to brain function: A critical appraisal. *Frontiers in Computational Neuroscience*, 6:24, 2012. doi:[10.3389/fncom.2012.00024](https://doi.org/10.3389/fncom.2012.00024).

- Stephen Ray Flora. *The Power of Reinforcement*. SUNY Press, Albany, NY, 2004.
- Luciano Floridi. Information ethics, its nature and scope. *ACM SIGCAS Computers and Society*, 36(3):21–36, 2006. doi:[10.1145/1195716.1195719](https://doi.org/10.1145/1195716.1195719).
- Michael J. Frank, Lauren C. Seeberger, and Randall C. O’Reilly. By carrot or by stick: Cognitive reinforcement learning in Parkinsonism. *Science*, 306(5703):1940–1943, 2004. doi:[10.1126/science.1102941](https://doi.org/10.1126/science.1102941).
- Michael J. Frank, Randall C. O’Reilly, and Tim Curran. When memory fails, intuition reigns: Midazolam enhances implicit inference in humans. *Psychological Science*, 17(8):700–707, 2006. doi:[10.1111/j.1467-9280.2006.01769.x](https://doi.org/10.1111/j.1467-9280.2006.01769.x).
- Stan Franklin, Steve Strain, Javier Snaider, Ryan McCall, and Usef Faghihi. Global Workspace Theory, its LIDA model and the underlying neuroscience. *Biologically Inspired Cognitive Architectures*, 1:32–43, 2012. doi:[10.1016/j.bica.2012.04.001](https://doi.org/10.1016/j.bica.2012.04.001).
- Robert A. Freitas, Jr. Xenopsychology. *Analog Science Fiction/Science Fact*, 104:41–53, 1984.
- Gordon G. Gallup, Jr., James R. Anderson, and Daniel J. Shillito. The mirror test. In Marc Bekoff, Colin Allen, and Gordon Burghardt, editors, *The Cognitive Animal: Empirical and Theoretical Perspectives on Animal Cognition*, pages 325–333. MIT Press, 2002.
- Adam E. Gaweda, Mehmet K. Muezzinoglu, Alfred A. Jacobs, George R. Aronoff, and Michael E. Brier. Model predictive control with reinforcement learning for drug delivery in renal anemia management. In *Engineering in Medicine and Biology Society, 28th Annual International Conference of the IEEE*, pages 5177–5180, 2006. doi:[10.1109/IEMBS.2006.260685](https://doi.org/10.1109/IEMBS.2006.260685).
- Paul W. Glimcher. Understanding dopamine and reinforcement learning: The dopamine reward prediction error hypothesis. *Proceedings of the National Academy of Sciences*, 108(Supplement 3):15647–15654, 2011. doi:[10.1073/pnas.1014269108](https://doi.org/10.1073/pnas.1014269108).
- Faustino Gomez, Jürgen Schmidhuber, and Risto Miikkulainen. Accelerated neural evolution through cooperatively coevolved synapses. *Journal of Machine Learning Research*, 9:937–965, 2008.
- Stephen Jay Gould. *Full House: The Spread of Excellence from Plato to Darwin*. Three Rivers Press, New York, NY, 1996.
- Stephen Grand and Dave Cliff. Creatures: Entertainment software agents with artificial life. *Autonomous Agents and Multi-Agent Systems*, 1(1):39–57, 1998. doi:[10.1023/A:1010042522104](https://doi.org/10.1023/A:1010042522104).

- C. G. Gray. Principle of least action. 4(12):8291, 2009. doi:[10.4249/scholarpedia.8291](https://doi.org/10.4249/scholarpedia.8291). revision num. 140216.
- David J. Gunkel. *The Machine Question: Critical Perspectives on AI, Robots, and Ethics*. MIT Press, Cambridge, MA, 2012.
- Nouchine Hadjikhani, Kestutis Kveraga, Paulami Naik, and Seppo P. Ahlfors. Early (n170) activation of face-specific cortex by face-like objects. *Neuroreport*, 20(4):403–407, 2009. doi:[10.1097/WNR.0b013e328325a8e1](https://doi.org/10.1097/WNR.0b013e328325a8e1).
- Michel Hansenne and Marc Anseau. Harm avoidance and serotonin. *Biological Psychology*, 51(1):77–81, 1999. doi:[10.1016/S0301-0511\(99\)00018-6](https://doi.org/10.1016/S0301-0511(99)00018-6).
- Rick Hanson. The not-craving brain. *FACES Conference*, October 2011. URL http://www.rickhanson.net/files/slides/FACES_NoCrave_Oct2011.pdf. Accessed March 2014.
- Demis Hassabis and Eleanor A. Maguire. Deconstructing episodic memory with construction. *Trends in Cognitive Sciences*, 11(7):299–306, 2007. doi:[10.1016/j.tics.2007.05.001](https://doi.org/10.1016/j.tics.2007.05.001).
- E. F. Hiby, N. J. Rooney, and J. W. S. Bradshaw. Dog training methods: Their use, effectiveness and interaction with behaviour and welfare. *Animal Welfare*, 13(1):63–69, 2004.
- Oscar Horta. Debunking the idyllic view of natural processes: Population dynamics and suffering in the wild. *Télos*, 17:73–88, 2010.
- Lesley Innes and Sebastian McBride. Negative versus positive reinforcement: An evaluation of training strategies for rehabilitated horses. *Applied Animal Behaviour Science*, 112(3–4):357–368, 2008. doi:[10.1016/j.applanim.2007.08.011](https://doi.org/10.1016/j.applanim.2007.08.011).
- Anthony I. Jack and Philip Robbins. The phenomenal stance revisited. *Review of Philosophy and Psychology*, 3(3):383–403, 2012. doi:[10.1007/s13164-012-0104-5](https://doi.org/10.1007/s13164-012-0104-5).
- Daphna Joel, Yael Niv, and Eytan Ruppin. Actor–critic models of the basal ganglia: New anatomical and computational perspectives. *Neural Networks*, 15(4–6):535–547, 2002. doi:[10.1016/S0893-6080\(02\)00047-3](https://doi.org/10.1016/S0893-6080(02)00047-3).
- Sergey Karakovskiy and Julian Togelius. The Mario AI benchmark and competitions. *Computational Intelligence and AI in Games, IEEE Transactions on*, 4(1):55–67, 2012. doi:[10.1109/TCIAIG.2012.2188528](https://doi.org/10.1109/TCIAIG.2012.2188528).
- Mehdi Khamassi, Loïc Lachèze, Benoît Girard, Alain Berthoz, and Agnès Guillot. Actor–critic models of reinforcement learning in the basal ganglia: From natural to artificial rats. *Adaptive Behavior*, 13(2):131–148, 2005. doi:[10.1177/105971230501300205](https://doi.org/10.1177/105971230501300205).

- Stephen J. Kish, Kathleen Shannak, and Oleh Hornykiewicz. Uneven pattern of dopamine loss in the striatum of patients with idiopathic Parkinson's disease. *New England Journal of Medicine*, 318(14):876–880, 1988. doi:[10.1056/NEJM198804073181402](https://doi.org/10.1056/NEJM198804073181402).
- Barbara J. Knowlton, Jennifer A. Mangels, and Larry R. Squire. A neostriatal habit learning system in humans. *Science*, 273(5280):1399–1402, 1996. doi:[10.1126/science.273.5280.1399](https://doi.org/10.1126/science.273.5280.1399).
- Jens Kober, J. Andrew Bagnell, and Jan Peters. Reinforcement learning in robotics: A survey. *International Journal of Robotics Research*, 32(11):1238–1274, July 2013. doi:[10.1177/0278364913495721](https://doi.org/10.1177/0278364913495721).
- Rogier Koppejan and Shimon Whiteson. Neuroevolutionary reinforcement learning for generalized control of simulated helicopters. *Evolutionary Intelligence*, 4(4):219–241, 2011. doi:[10.1007/s12065-011-0066-z](https://doi.org/10.1007/s12065-011-0066-z).
- Jan Koutník, Giuseppe Cuccu, Jürgen Schmidhuber, and Faustino Gomez. Evolving large-scale neural networks for vision-based reinforcement learning. In *Proceedings of the 15th Annual Conference on Genetic and Evolutionary Computation*, pages 1061–1068, 2013. doi:[10.1145/2463372.2463509](https://doi.org/10.1145/2463372.2463509).
- Michael R. LaChat. Artificial intelligence and ethics: An exercise in the moral imagination. *AI Magazine*, 7(2):70–79, 1986. doi:[10.1609/aimag.v7i2.540](https://doi.org/10.1609/aimag.v7i2.540).
- Sascha Lange, Thomas Gabel, and Martin Riedmiller. Batch reinforcement learning. In Marco Wiering and Martijn van Otterlo, editors, *Reinforcement Learning: State-of-the-Art*, pages 45–73. 2012. doi:[10.1007/978-3-642-27645-3_2](https://doi.org/10.1007/978-3-642-27645-3_2).
- Hakwan Lau and David Rosenthal. Empirical support for higher-order theories of conscious awareness. *Trends in Cognitive Sciences*, 15(8):365–373, 2011. doi:[10.1016/j.tics.2011.05.009](https://doi.org/10.1016/j.tics.2011.05.009).
- Jae Won Lee, Jonghun Park, Jangmin O, Jongwoo Lee, and Euyseok Hong. A multiagent approach to Q-learning for daily stock trading. *Systems, Man and Cybernetics, Part A: Systems and Humans, IEEE Transactions on*, 37(6):864–877, 2007. doi:[10.1109/TSMCA.2007.904825](https://doi.org/10.1109/TSMCA.2007.904825).
- Patrick Lin, Keith Abney, and George A. Bekey, editors. *Robot Ethics: The Ethical and Social Implications of Robotics*. MIT Press, Cambridge, MA, 2011.
- Michael L. Littman. Markov games as a framework for multi-agent reinforcement learning. In *ICML*, pages 157–163, 1994.
- Daphne Liu and Lenhart K. Schubert. Combining self-motivation with logical planning and inference in a reward-seeking agent. In *ICAART*, pages 257–263, 2010.

- Scott Livingston, Jamie Garvey, and Itamar Elhanany. On the broad implications of reinforcement learning based AGI. In *Artificial General Intelligence 2008*, pages 478–482, 2008.
- Jeffrey A. Lockwood. The moral standing of insects and the ethics of extinction. *Florida Entomologist*, 70(1):70–89, 1987.
- Daniele Loiacono, Alessandro Prete, Pier Luca Lanzi, and Luigi Cardamone. Learning to overtake in TORCS using simple reinforcement learning. In *Evolutionary Computation (CEC), 2010 IEEE Congress on*, pages 1–8, 2010. doi:[10.1109/CEC.2010.5586191](https://doi.org/10.1109/CEC.2010.5586191).
- Elliot A. Ludvig, Richard S. Sutton, and E. James Kehoe. Evaluating the TD model of classical conditioning. *Learning & Behavior*, 40(3):305–319, 2012. doi:[10.3758/s13420-012-0082-6](https://doi.org/10.3758/s13420-012-0082-6).
- Sridhar Mahadevan. Average reward reinforcement learning: Foundations, algorithms, and empirical results. *Machine Learning*, 22(1–3):159–195, 1996. doi:[10.1007/BF00114727](https://doi.org/10.1007/BF00114727).
- Tiago V. Maia. Reinforcement learning, conditioning, and the brain: Successes and challenges. *Cognitive, Affective, & Behavioral Neuroscience*, 9(4):343–364, 2009. doi:[10.3758/CABN.9.4.343](https://doi.org/10.3758/CABN.9.4.343).
- Tiago V. Maia. Two-factor theory, the actor-critic model, and conditioned avoidance. *Learning & Behavior*, 38(1):50–67, 2010. doi:[10.3758/LB.38.1.50](https://doi.org/10.3758/LB.38.1.50).
- Jordan M. Malof and Adam E. Gaweda. Optimizing drug therapy with reinforcement learning: The case of anemia management. In *Neural Networks, International Joint Conference on*, pages 2088–2092, 2011. doi:[10.1109/IJCNN.2011.6033485](https://doi.org/10.1109/IJCNN.2011.6033485).
- Gary Marcus. *Kluge: The Haphazard Evolution of the Human Mind*. Houghton Mifflin Harcourt, New York, NY, 2009.
- Georgia J. Mason. Invertebrate welfare: Where is the real evidence for conscious affective states? *Trends in Ecology & Evolution*, 26(5):212–213, 2011. doi:[10.1016/j.tree.2011.02.009](https://doi.org/10.1016/j.tree.2011.02.009).
- Jennifer A. Mather. Animal suffering: An invertebrate perspective. *Journal of Applied Animal Welfare Science*, 4(2):151–156, 2001.
- Andrew McCallum. Overcoming incomplete perception with utile distinction memory. In *ICML*, pages 190–196, 1993.
- Michelle McPartland and Marcus Gallagher. Reinforcement learning in first person shooter games. *Computational Intelligence and AI in Games, IEEE Transactions on*, 3(1):43–56, 2011. doi:[10.1109/TCIAIG.2010.2100395](https://doi.org/10.1109/TCIAIG.2010.2100395).

- Yew-Kwang Ng. Towards welfare biology: Evolutionary economics of animal consciousness and suffering. *Biology and Philosophy*, 10(3):255–285, 1995. doi:[10.1007/BF00852469](https://doi.org/10.1007/BF00852469).
- Robert Nozick. *Anarchy, State, and Utopia*. Basic Books, New York, NY, 1974.
- Wei-Xing Pan, Robert Schmidt, Jeffery R. Wickens, and Brian I. Hyland. Dopamine cells respond to predicted events during classical conditioning: Evidence for eligibility traces in the reward-learning network. *The Journal of Neuroscience*, 25(26):6235–6242, 2005. doi:[10.1523/JNEUROSCI.1478-05.2005](https://doi.org/10.1523/JNEUROSCI.1478-05.2005).
- Sotiris Papadopoulos, Francisco Baez, Jonathan Alt, and Christian Darken. Behavior selection using utility-based reinforcement learning in irregular warfare simulation models. *International Journal of Operations Research and Information Systems (IJORIS)*, 4(3):61–78, 2013. doi:[10.4018/joris.2013070105](https://doi.org/10.4018/joris.2013070105).
- Purvag G. Patel, Norman Carver, and Shahram Rahimi. Tuning computer gaming agents using Q-learning. In *Computer Science and Information Systems (FedCSIS), 2011 Federated Conference on*, pages 581–588, 2011.
- David Pearce. The hedonistic imperative, 2007. URL <https://cl.nfshost.com/david-pearce-the-hedonistic-imperative.pdf>. Accessed April 2014.
- Susana Peciña. Opioid reward ‘liking’ and ‘wanting’ in the nucleus accumbens. *Physiology & Behavior*, 94(5):675–680, 2008. doi:[10.1016/j.physbeh.2008.04.006](https://doi.org/10.1016/j.physbeh.2008.04.006).
- Jason Rennie and Andrew McCallum. Using reinforcement learning to spider the web efficiently. In *ICML*, pages 335–343, 1999.
- José J. F. Ribas-Fernandes, Alec Solway, Carlos Diuk, Joseph T. McGuire, Andrew G. Barto, Yael Niv, and Matthew M. Botvinick. A neural signature of hierarchical reinforcement learning. *Neuron*, 71(2):370–379, 2011. doi:[10.1016/j.neuron.2011.05.042](https://doi.org/10.1016/j.neuron.2011.05.042).
- Philip Robbins and Anthony I. Jack. The phenomenal stance. *Philosophical Studies*, 127(1):59–85, 2006. doi:[10.1007/s11098-005-1730-x](https://doi.org/10.1007/s11098-005-1730-x).
- Stuart Russell and Peter Norvig. *Artificial Intelligence: A Modern Approach*. Prentice Hall, Englewood Cliffs, NJ, 3rd edition, 2009.
- William Moy Stratton Russell and Rex Leonard Burch. *The Principles of Humane Experimental Technique*. Methuen, London, UK, 1959.
- John D. Salamone, Mercè Correa, Andrew Farrar, and Susana M. Minogote. Effort-related functions of nucleus accumbens dopamine and associated forebrain circuits. *Psychopharmacology*, 191(3):461–482, 2007. doi:[10.1007/s00213-006-0668-9](https://doi.org/10.1007/s00213-006-0668-9).

- Jürgen Schmidhuber. Evolutionary computation versus reinforcement learning. In *Industrial Electronics Society*, volume 4, pages 2992–2997, 2000. doi:[10.1109/IECON.2000.972474](https://doi.org/10.1109/IECON.2000.972474).
- Jürgen Schmidhuber. Philosophers & futurists, catch up! response to the singularity. *Journal of Consciousness Studies*, 19(1–2):173–182, 2012.
- Wolfram Schultz, Peter Dayan, and P. Read Montague. A neural substrate of prediction and reward. *Science*, 275(5306):1593–1599, 1997. doi:[10.1126/science.275.5306.1593](https://doi.org/10.1126/science.275.5306.1593).
- Eric Schwitzgebel. If materialism is true, the United States is probably conscious, 2012. URL <http://faculty.ucr.edu/~eschwitz/SchwitzAbs/USAconscious.htm>. Accessed April 2014.
- Anil Seth. Models of consciousness. 2(1):1328, 2007. doi:[10.4249/scholarpedia.1328](https://doi.org/10.4249/scholarpedia.1328). revision num. 132493.
- Ben Seymour, John P. O’Doherty, Peter Dayan, Martin Koltzenburg, Anthony K. Jones, Raymond J. Dolan, Karl J. Friston, and Richard S. Frackowiak. Temporal difference models describe higher-order learning in humans. *Nature*, 429(6992):664–667, 2004. doi:[10.1038/nature02581](https://doi.org/10.1038/nature02581).
- Yoav Shoham and Kevin Leyton-Brown. *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*. Cambridge University Press, New York, NY, 2009.
- Hanan Shteingart and Yonatan Loewenstein. Reinforcement learning and human behavior. *Current Opinion in Neurobiology*, 25:93–98, 2014. doi:[10.1016/j.conb.2013.12.004](https://doi.org/10.1016/j.conb.2013.12.004).
- Amandeep S. Sidhu, Narendra S. Chaudhari, and Ghee Ming Goh. Hierarchical reinforcement learning model for military simulations. In *Neural Networks, International Joint Conference on*, pages 2572–2576, 2006. doi:[10.1109/IJCNN.2006.247132](https://doi.org/10.1109/IJCNN.2006.247132).
- Peter Singer. *Animal Liberation: The Definitive Classic of the Animal Movement*. Harper Perennial Modern Classics, New York, NY, 2009.
- Peter Singer and Agata Sagan. When robots have feelings. *The Guardian*, December 2009. URL <http://www.theguardian.com/commentisfree/2009/dec/14/rage-against-machines-robots>. Accessed April 2014.
- Scott Siskind. Are wireheads happy? *LessWrong*, January 2010. URL http://lesswrong.com/lw/1lb/are_wireheads_happy/. Accessed March 2014.
- Aaron Sloman. Phenomenal and access consciousness and the “hard” problem: A view from the designer stance. *International Journal of Machine Consciousness*, 2(1):117–169, 2010. doi:[10.1142/S1793843010000424](https://doi.org/10.1142/S1793843010000424).

- Richard L. Solomon and John D. Corbit. An opponent-process theory of motivation: I. temporal dynamics of affect. *Psychological Review*, 81(2):119–145, 1974.
- Kenneth O. Stanley and Risto Miikkulainen. Efficient reinforcement learning through evolving neural network topologies. In *Proceedings of the Genetic and Evolutionary Computation Conference*, pages 569–577, 2002.
- Richard S. Sutton. Learning to predict by the methods of temporal differences. *Machine Learning*, 3(1):9–44, 1988. doi:[10.1007/BF00115009](https://doi.org/10.1007/BF00115009).
- Richard S. Sutton. Integrated architecture for learning, planning, and reacting based on approximating dynamic programming. In *ICML*, pages 216–224, 1990.
- Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning: An Introduction*. MIT Press, Cambridge, MA, 1998.
- Matthew E. Taylor, Shimon Whiteson, and Peter Stone. Comparing evolutionary and temporal difference methods in a reinforcement learning domain. In *Proceedings of the 8th Annual Conference on Genetic and Evolutionary Computation*, pages 1321–1328, 2006. doi:[10.1145/1143997.1144202](https://doi.org/10.1145/1143997.1144202).
- Bruce L. Tempel, Nancy Bonini, Douglas R. Dawson, and William G. Quinn. Reward learning in normal and mutant *Drosophila*. *Proceedings of the National Academy of Sciences*, 80(5):1482–1486, 1983.
- Gerald Tesauro. TD-Gammon, a self-teaching backgammon program, achieves master-level play. *Neural Computation*, 6(2):215–219, 1994. doi:[10.1162/neco.1994.6.2.215](https://doi.org/10.1162/neco.1994.6.2.215).
- Edward Lee Thorndike. *Animal Intelligence: Experimental Studies*. Macmillan, New York, NY, 1911.
- Giulio Tononi. Consciousness as integrated information: A provisional manifesto. *The Biological Bulletin*, 215(3):216–242, 2008.
- Steve Torrance. Towards an ethics for epersons. *AISB Quarterly*, 104:38–41, 2000.
- Nees Jan van Eck and Michiel van Wezel. Application of reinforcement learning to the game of Othello. *Computers & Operations Research*, 35(6):1999–2017, 2008. doi:[10.1016/j.cor.2006.10.004](https://doi.org/10.1016/j.cor.2006.10.004).
- Bruno van Swinderen. The remote roots of consciousness in fruit-fly selective attention? *BioEssays*, 27(3):321–330, 2005. doi:[10.1002/bies.20195](https://doi.org/10.1002/bies.20195).
- Stefan von Delius, Regina Hollweck, Roland M. Schmid, and Eckart Frimberger. Midazolam-pain, but one cannot remember it: A survey among Southern German endoscopists. *European Journal of Gastroenterology & Hepatology*, 19(6):465–470, 2007. doi:[10.1097/MEG.0b013e3280ad4425](https://doi.org/10.1097/MEG.0b013e3280ad4425).

-
- David Wales. *Energy Landscapes: Applications to Clusters, Biomolecules and Glasses*. Cambridge University Press, Cambridge, UK, 2003.
- A. K. Warren-Smith and P. D. McGreevy. The use of blended positive and negative reinforcement in shaping the halt response of horses (*Equus caballus*). *Animal Welfare*, 16(4):481–488, 2007.
- Blay Whitby. Sometimes it’s hard to be a robot: A call for action on the ethics of abusing artificial agents. *Interacting with Computers*, 20(3):326–333, 2008. doi:[10.1016/j.intcom.2008.02.002](https://doi.org/10.1016/j.intcom.2008.02.002).
- Shimon Whiteson. Evolutionary computation for reinforcement learning. In Marco Wiering and Martijn van Otterlo, editors, *Reinforcement Learning: State-of-the-Art*, pages 325–355. 2012.
- Alexander J. Whitworth, Paul D. Wes, and Leo J. Pallanck. *Drosophila* models pioneer a new approach to drug discovery for Parkinson’s disease. *Drug Discovery Today*, 11(3–4):119–126, 2006. doi:[10.1016/S1359-6446\(05\)03693-7](https://doi.org/10.1016/S1359-6446(05)03693-7).
- Meghan Winsby. Suffering subroutines: On the humanity of making a computer that feels pain. In *International Association for Computing and Philosophy*, 2013.
- Florentin Woergoetter and Bernd Porr. Reinforcement learning. 3(3):1448, 2008. doi:[10.4249/scholarpedia.1448](https://doi.org/10.4249/scholarpedia.1448). revision num. 91704.
- Ian Wright. Reinforcement learning and animat emotions. In *From Animals to Animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior*, pages 272–281, 1996.
- Manfred Zimmermann. Behavioural investigation of pain in animals. In I. J. H. Duncan and V. Molony, editors, *Assessing Pain in Farm Animals: Proceedings of a Workshop Held in Roslin, Scotland, 25 and 26 October 1984*, 1986.