

Journal Pre-proof

Reward Is Enough

David Silver, Satinder Singh, Doina Precup and Richard S. Sutton

PII: S0004-3702(21)00086-2
DOI: <https://doi.org/10.1016/j.artint.2021.103535>
Reference: ARTINT 103535


To appear in: *Artificial Intelligence*

Received date: 12 November 2020
Revised date: 28 April 2021
Accepted date: 12 May 2021

Please cite this article as: D. Silver, S. Singh, D. Precup et al., Reward Is Enough, *Artificial Intelligence*, 103535, doi: <https://doi.org/10.1016/j.artint.2021.103535>.

This is a PDF file of an article that has undergone enhancements after acceptance, such as the addition of a cover page and metadata, and formatting for readability, but it is not yet the definitive version of record. This version will undergo additional copyediting, typesetting and review before it is published in its final form, but we are providing this version to give early visibility of the article. Please note that, during the production process, errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

© 2021 Published by Elsevier.



Reward Is Enough

David Silver, Satinder Singh, Doina Precup, Richard S. Sutton

Abstract

In this article we hypothesise that intelligence, and its associated abilities, can be understood as subserving the maximisation of reward. Accordingly, reward is enough to drive behaviour that exhibits abilities studied in natural and artificial intelligence, including knowledge, learning, perception, social intelligence, language, generalisation and imitation. This is in contrast to the view that specialised problem formulations are needed for each ability, based on other signals or objectives. Furthermore, we suggest that agents that learn through trial and error experience to maximise reward could learn behaviour that exhibits most if not all of these abilities, and therefore that powerful reinforcement learning agents could constitute a solution to artificial general intelligence.

1 Introduction

Expressions of intelligence in animal and human behaviour are so bountiful and so varied that there is an ontology of associated abilities to name and study them, e.g. social intelligence, language, perception, knowledge representation, planning, imagination, memory, and motor control. What could drive agents (natural or artificial) to behave intelligently in such a diverse variety of ways?

One possible answer is that each ability arises from the pursuit of a goal that is designed specifically to elicit that ability. For example, the ability of social intelligence has often been framed as the Nash equilibrium of a multi-agent system; the ability of language by a combination of goals such as parsing, part-of-speech tagging, lexical analysis, and sentiment analysis; and the ability of perception by object segmentation and recognition. In this paper, we consider an alternative hypothesis: that the generic objective of maximising reward is enough to drive behaviour that exhibits most if not all abilities that are studied in natural and artificial intelligence.

This hypothesis may startle because the sheer diversity of abilities associated with intelligence seems to be at odds with any generic objective. However, the natural world faced by animals and humans, and presumably also the environments faced in the future by artificial agents, are inherently so complex that they require sophisticated abilities in order to succeed (for example, to survive) within those environments. Thus, success, as measured by maximising reward, demands a variety of abilities associated with intelligence. In such environments, any behaviour that maximises reward must necessarily exhibit those abilities. In this sense, the generic objective of reward maximisation contains within it many or possibly even all the goals of intelligence.

Reward thus provides two levels of explanation for the bountiful expressions of intelligence found in nature. First, different forms of intelligence may arise from the maximisation of different reward signals in different environments, resulting for example in abilities as distinct as echolocation in bats, communication by whale-song, or tool use in chimpanzees. Similarly, artificial agents may be required to maximise a variety of reward signals in future environments, resulting in new forms of intelligence with abilities as distinct as laser-based navigation, communication by email, or robotic manipulation.

Second, the intelligence of even a single animal or human is associated with a cornucopia of abilities. According to our hypothesis, all of these abilities subserve a *singular goal* of maximising that animal or agent's reward within its environment. In other words, the pursuit of one goal may generate




Figure 1: The *reward-is-enough* hypothesis postulates that intelligence, and its associated abilities, can be understood as subserving the maximisation of reward by an agent acting in its environment. For example, a squirrel acts so as to maximise its consumption of food (top, reward depicted by acorn symbol), or a kitchen robot acts to maximise cleanliness (bottom, reward depicted by bubble symbol). To achieve these goals, complex behaviours are required that exhibit a wide variety of abilities associated with intelligence (depicted on the right as a projection from an agent's stream of experience onto a set of abilities expressed within that experience).

complex behaviour that exhibits multiple abilities associated with intelligence. Indeed, such reward-maximising behaviour may often be consistent with specific behaviours derived from the pursuit of separate goals associated with each ability.

For example, a squirrel's brain may be understood as a decision-making system that receives sensations from, and sends motor commands to the squirrel's body. The behaviour of the squirrel may be understood as maximising a cumulative reward such as satiation (i.e. negative hunger). In order for a squirrel to minimise hunger, the squirrel-brain must presumably have abilities of perception (to identify good nuts), knowledge (to understand nuts), motor control (to collect nuts), planning (to choose where to cache nuts), memory (to recall locations of cached nuts) and social intelligence (to bluff about locations of cached nuts, to ensure they are not stolen). Each of these abilities associated with intelligence may therefore be understood as subserving a singular goal of hunger minimisation (see Figure 1).

As a second example, a kitchen robot may be implemented as a decision-making system that receives sensations from, and sends actuator commands to, the robot's body. The singular goal of the kitchen robot is to maximise a reward signal measuring cleanliness.¹ In order for a kitchen robot to maximise cleanliness, it must presumably have abilities of perception (to differentiate clean and dirty utensils), knowledge (to understand utensils), motor control (to manipulate utensils), memory (to recall locations of utensils), language (to predict future mess from dialogue), and social intelligence (to encourage young children to make less mess). A behaviour that maximises cleanliness must therefore yield all these abilities in service of that singular goal (see Figure 1).

When abilities associated with intelligence arise as solutions to a singular goal of reward maximisation, this may in fact provide a deeper understanding since it explains *why* such an ability arises (e.g.

¹for example, as judged by occasional human inspection.

that classification of crocodiles is important to avoid being eaten). In contrast, when each ability is understood as the solution to its own specialised goal, the *why* question is side-stepped in order to focus upon *what* that ability does (e.g. discriminating crocodiles from logs). Furthermore, a singular goal may also provide a broader understanding of each ability that may include characteristics that are otherwise hard to formalise, such as dealing with irrational agents in social intelligence (e.g. pacifying an angry aggressor), grounding language to perceptual experience (e.g. dialogue regarding the best way to peel a fruit), or understanding haptics in perception (e.g. picking a sharp object from a pocket). Finally, implementing abilities in service of a singular goal, rather than for their own specialised goals, also answers the question of how to integrate abilities, which otherwise remains an outstanding issue.

Having established that reward maximisation is a suitable objective for understanding the problem of intelligence, one may consider methods for solving the problem. One might then expect to find such methods in natural intelligence, or choose to implement them in artificial intelligence. Among possible methods for maximising reward, the most general and scalable approach is to learn to do so, by interacting with the environment by trial and error. We conjecture that an agent that can effectively learn to maximise reward in this manner would, when placed in a rich environment, give rise to sophisticated expressions of general intelligence.

A recent salutary example of both problem and solution based on reward maximisation comes from the game of Go. Research initially focused largely upon distinct abilities, such as openings, shape, tactics, and endgames, each formalised using distinct objectives such as sequence memorisation, pattern recognition, local search, and combinatorial game theory [31]. AlphaZero [47] focused instead on a singular goal: maximising a reward signal that is 0 until the final step, and then +1 for winning or -1 for losing. This ultimately resulted in a deeper understanding of each ability – for example, discovering new opening sequences [63], using surprising shapes within a global context [39], understanding global interactions between local battles [62], and playing safe when ahead [39]. It also yielded a broader set of abilities that had not previously been satisfactorily formalised – such as balancing influence and territory, thickness and lightness, and attack and defence. AlphaZero’s abilities were innately integrated into a unified whole, whereas integration had proven highly problematic in prior work [31]. Thus, maximising wins proved to be enough, in a simple environment such as Go, to drive behaviour exhibiting a variety of specialised abilities. Furthermore, applying the same method to different environments such as chess or shogi [46] resulted in new abilities such as piece mobility and colour complexes [43]. We argue that maximising rewards in richer environments – more comparable in complexity to the natural world faced by animals and humans – could yield further, and perhaps ultimately all abilities associated with intelligence.

The rest of this paper is organised as follows. In Section 2 we formalise the objective of reward maximisation as the problem of reinforcement learning. In Section 3 we present our main hypothesis. We consider several important abilities associated with intelligence, and discuss how reward maximisation may yield those abilities. In Section 4 we turn to the use of reward maximisation as a solution strategy. We present related work in Section 5 and finally, in Section 6 we discuss possible weaknesses of the hypothesis and consider several alternatives.

2 Background: The Reinforcement Learning Problem

Intelligence may be understood as a flexible ability to achieve goals. For example according to John McCarthy, “*intelligence is the computational part of the ability to achieve goals in the world*” [28]. *Reinforcement learning* [54] formalises the problem of goal-seeking intelligence. The general problem may be instantiated with a wide and realistic range of goals and worlds — and hence a wide range of forms of intelligence — corresponding to different reward signals to maximise in different environments.

2.1 Agent and Environment

Like many interactive approaches to artificial intelligence [41], reinforcement learning follows a protocol that decouples a problem into two systems that interact sequentially over time: an *agent* (the solution) that takes decisions and an *environment* (the problem) that is influenced by those decisions. This is in contrast to other specialised protocols that may for example consider multiple agents, multiple environments, or other modes of interaction.

Dimension	Alternative A	Alternative B	Notes
Observations	Discrete	Continuous	Time-step may be infinitesimal
Actions	Discrete	Continuous	
Time	Discrete	Continuous	
Dynamics	Deterministic	Stochastic	
Observability	Full	Partial	Other agents are part of environment, from perspective of single agent (see Section 3.3)
Agency	Single agent	Multi-agent	
Uncertainty	Certain	Uncertain	
Termination	Continuing	Episodic	Environment may terminate and reset to an initial state
Stationarity	Stationary	Non-stationary	Environment depends upon history and hence also upon time
Synchronicity	Asynchronous	Synchronous	Observation may remain unchanged until action is executed
Reality	Simulated	Real-world	May include humans that interact with agent

Table 1: The definition of environment is broad and encompasses many problem dimensions.

2.2 Agent

An *agent* is a system that receives at time t an observation O_t and outputs an action A_t . More formally, the agent is a system $A_t = \alpha(H_t)$ that selects an action A_t at time t given its *experience* history $H_t = O_1, A_1, \dots, O_{t-1}, A_{t-1}, O_t$, in other words, given the sequence of observations and actions that have occurred in the history of interactions between the agent and the environment.

The agent system α is limited by practical constraints to a bounded set [42]. The agent has limited capacity determined by its machinery (for example, limited memory in a computer or limited neurons in a brain). The agent and environment systems execute in real-time. While the agent spends time computing its next action (e.g. producing no-op actions while deciding whether to run away from a lion), the environment system continues to process (e.g. the lion attacks). Thus, the reinforcement learning problem represents a practical problem, as faced by natural and artificial intelligence, rather than a theoretical abstraction that ignores computational limitations.

This paper does not delve into the nature of the agent, focusing instead on the problem it must solve, and the intelligence that may be induced by any solution to that problem.

2.3 Environment

An *environment* is a system that receives action A_t at time t and responds with observation O_{t+1} at the next time step. More formally, an environment is a system $O_{t+1} = \varepsilon(H_t, A_t, \eta_t)$ that determines the next observation O_{t+1} that the agent will receive from the environment, given experience history H_t , the latest action A_t , and potentially a source of randomness η_t .

The environment specifies within its definition the interface to the agent. The agent consists solely of the decision-making entity; anything outside of that entity (including its body, if it has one) is considered part of the environment. This includes both the sensors and actuators that define the observations for the agent and the actions available to the agent respectively.

Note that this definition of environment is very broad and encompasses many problem dimensions, including those in Table 1.

2.4 Rewards

The reinforcement learning problem represents goals by cumulative rewards. A *reward* is a special scalar observation R_t , emitted at every time-step t by a *reward signal* in the environment, that provides an instantaneous measurement of progress towards a goal. An instance of the reinforcement learning problem is defined by an environment ε with a reward signal, and by a cumulative objective to maximise, such as a sum of rewards over a finite number of steps, a discounted sum, or the average reward per time-step.

A wide variety of goals can be represented by rewards.² For example, a scalar reward signal can represent weighted combinations of objectives, different trade-offs over time, and risk-seeking or risk-averse utilities. Reward can also be determined by a human-in-the-loop, for example humans may provide explicit reinforcement of desired behaviour, online feedback via clickthrough or thumbs-up, delayed feedback via questionnaires or surveys, or by a natural language utterance. Including feedback from a human can provide a mechanism to formulate seemingly fuzzy goals such as “I’ll know it when I see it”.

In addition to their generality, rewards also provide intermediate feedback, potentially at every time-step, on progress towards the goal. This intermediate signal is an essential part of the problem definition when considering long or infinite streams of experience – without intermediate feedback, learning is not possible.

3 Reward is Enough

We have seen, in the previous section, that rewards are sufficient to express a wide variety of goals, corresponding to the diverse purposes towards which intelligence may be directed.

We now have all the ingredients to make our main point: that many different forms of intelligence can be understood as subserving the maximisation of reward, and that the many abilities associated with each form of intelligence may arise implicitly from the pursuit of those rewards. Taken to its limit, we hypothesise that all intelligence and associated abilities may be understood in this manner:

Hypothesis. (*Reward-is-Enough*)

Intelligence, and its associated abilities, can be understood as subserving the maximisation of reward by an agent acting in its environment.

This hypothesis is important because, if it is true, it suggests that a good reward-maximising agent, in the service of achieving its goal, could implicitly yield abilities associated with intelligence. A good agent in this context is one that successfully – perhaps using as-yet-undiscovered algorithms – performs well at maximising cumulative reward in its environment. We return to the question of how such an agent may be constructed in Section 4.

In principle, any behaviour can be encoded by maximisation of a reward signal that is explicitly chosen to induce that behaviour [12] (for example, providing rewards when objects are correctly identified, or when a syntactically correct sentence is produced). The hypothesis here is intended to be much stronger: that intelligence and associated abilities will implicitly arise in the service of maximising one of many possible reward signals, corresponding to the many pragmatic goals towards which natural or artificial intelligence may be directed.

Sophisticated abilities may arise from the maximisation of simple rewards in complex environments. For example, the minimisation of hunger in a squirrel’s natural environment demands a skilful ability to manipulate nuts that arises from the interplay between (among other factors) the squirrel’s musculoskeletal dynamics; objects such as leaves, branches, or soil that the squirrel or nut may be resting upon, connected to, or obstructed by; variations in the size and shape of nuts; environmental factors such as wind, rain, or snow; and changes due to ageing, disease or injury. Similarly, the pursuit of cleanliness in a kitchen robot demands a sophisticated ability to perceive utensils in an enormous range of states that includes clutter, occlusion, glare, encrustation, damage, and so on.

Furthermore, the maximisation of many other reward signals by a squirrel (such as maximising survival time, minimising pain, or maximising reproductive success) or kitchen robot (such as maximising a healthy eating index, maximising positive feedback from the user, or maximising their gastronomic endorphins), and in many other environments (such as different habitats, other dexterous bodies, or different climates), would also yield abilities of perception, locomotion, manipulation, and so forth. Thus, the path to general intelligence may in fact be quite robust to the choice of reward signal. Indeed the ability to generate intelligence may often be orthogonal to the goal that it is given, in the sense that the maximisation of many different reward signals in many different environments may produce similar abilities associated with intelligence.

²Indeed, the *reward hypothesis* speculates that all goals of interest in studying natural or building artificial agents may be represented by rewards [54]. This should not be confused with our reward-is-enough hypothesis, which considers the abilities that arise from the pursuit of any one such goal.

In the following sections we explore whether and how this hypothesis could be applied in practice to various important abilities, including several that are seemingly hard to formalise as reinforcement learning problems. We do not provide an exhaustive discussion of all abilities associated with intelligence, but encourage the reader to consider abilities such as memory, imagination, common sense, attention, reasoning, creativity, or emotions, and how those abilities may subserve a generic objective of reward maximisation.

3.1 Reward is Enough for Knowledge and Learning

We define knowledge as information that is internal to the agent; for example, knowledge may be contained within the parameters of an agent’s functions for selecting actions, predicting cumulative reward, or predicting features of future observations. Some of this knowledge may be innate (prior knowledge), while some knowledge may be acquired through learning.

An environment may demand innate knowledge. In particular, it may be necessary, in order to maximise total reward, to have knowledge that is immediately accessible in novel situations. For example, a new-born gazelle may need to run away from a lion. In this case, innate understanding of predator evasion may be necessary before there is any opportunity to learn this knowledge. Note, however, that the extent of prior knowledge is limited both in theory (by the capacity of the agent) and in practice (by the difficulty of constructing useful prior knowledge). Furthermore, unlike other abilities that we shall consider, environmental demand for innate knowledge cannot be operationalised; it is the knowledge that comes before experience, and therefore cannot be acquired from experience.

An environment may also demand learned knowledge. This will occur when future experience is uncertain, due to unknown elements, stochasticity, or complexity of the environment. This uncertainty results in a vast array of potential knowledge that the agent may need, depending on its particular realisation of future events. In rich environments, including many of the hardest and most important problems in natural and artificial intelligence, the total space of potential knowledge is far larger than the capacity of the agent. For example, consider an early-human environment. An agent may be born in either the Arctic or Africa, leading to radically different requirements for knowledge to maximise total reward: facing polar bears or lions; traversing glaciers or savannah; building with ice or with mud; etc. Particular events within the agent’s life may also lead to radically different needs: choosing to hunt or to farm; facing locusts or war; going blind or going deaf; encountering friend or foe; etc. In each of these possible lives, the agent must acquire detailed and specialised knowledge. If the sum of such potential knowledge outstrips the agent’s capacity, knowledge must be a function of the agent’s experience and adapt to the agent’s particular circumstances – thus demanding learning. In practice this learning may take many computational forms, for example making predictions, models or skills through the adaptation of parameters or the construction, curation and reuse of structure.

In summary, the environment may call for both innate and learned knowledge, and a reward-maximising agent will, whenever required, contain the former (for example, through evolution in natural agents and by design in artificial agents) and acquire the latter (through learning). In richer and longer-lived environments the balance of demand shifts increasingly toward learned knowledge.

3.2 Reward is Enough for Perception

The human world demands a variety of perceptual abilities to accumulate rewards. Some life-or-death examples include: image segmentation to avoid falling off a cliff; object recognition to classify healthy and poisonous foods; face recognition to differentiate friend from foe; scene parsing to drive slowly near playing children; or speech recognition to understand a verbal warning to "duck!" Several modes of perception may be required, including visual, aural, olfactory, somatosensory or proprioceptive perception.

Historically, these perceptual abilities were formulated using separate problem definitions [9]. More recently, there has been a growing movement towards unifying perceptual abilities as solutions to supervised learning problems [23]. The problem is typically formalised as minimising the classification error for examples in a test set, given a training set of correctly labelled examples. The unification of many perceptual abilities as supervised learning problems has led to considerable success in a wide variety of real-world applications where large data-sets are available [22, 16, 8].

As per our hypothesis, perception may instead be understood as subserving, and implemented by, a generic goal of maximising cumulative reward within a rich environment that demands perceptual abilities. Considering perception from the perspective of reward maximisation may facilitate understanding and implementation of an even greater range of perception that includes challenging and realistic forms of perceptual abilities:

- Action and observation are typically intertwined into active forms of perception, such as haptic perception (e.g. identifying the contents of a pocket by moving fingertips), visual saccades (e.g. moving eyes to switch focus between bat and ball), physical experimentation (e.g. hitting a nut with a rock to see if it will break), or echolocation (e.g. emitting sounds at varying frequency and measuring timings and intensity of subsequent echoes).
- The utility of perception often depends upon the agent’s behaviour – for example, the cost of misclassifying a crocodile depends upon whether the agent is walking or swimming, and upon whether the agent would subsequently fight or flee.
- There may be an explicit or implicit cost to acquiring information (e.g. there are energy, computational, and opportunity costs to turning the head and checking for a predator).
- The distribution of data is typically context-dependent. For example, an Arctic agent may be required to classify ice and polar bears, upon which its rewards depend; while an African agent may need to classify savannah and lions. The diversity of potential data may, in rich environments, vastly exceed the agent’s capacity or the quantity of pre-existing data (see Section 3.1) – requiring that perception is learned from experience.
- Many applications of perception do not have access to labelled data.

3.3 Reward is Enough for Social Intelligence

Social intelligence is the ability to understand and interact effectively with other agents. This ability is often formalised, using game theory, as the equilibrium solution of a multi-agent game. Equilibrium solutions are considered desirable because they are robust to deviation or worst-case scenarios. For example, the Nash equilibrium is a joint strategy for all agents such that no unilateral deviation gives benefit to the deviator [32]. In zero-sum games a Nash equilibrium is also minimax optimal [33]: it attains the best possible value against worst-case opponents.

According to our hypothesis, social intelligence may instead be understood as, and implemented by, maximising cumulative reward from the point of view of one agent in an environment that contains other agents. Following this standard agent-environment protocol, one agent observes the behaviour of other agents, and may affect other agents through its actions, just as it observes and affects any other aspect of its environment. An agent that can anticipate and influence the behaviour of other agents can typically achieve greater cumulative reward. Thus, if an environment needs social intelligence (e.g. because it contains animals or humans), reward maximisation will produce social intelligence.

Robustness may also be demanded by the environment, for example because the environment may contain multiple agents following different strategies. If these other agents are aliased (i.e. there is no way to identify in advance which strategies other agents will follow) then a reward-maximising agent must hedge its bets and choose a robust behaviour that will be effective against any of these potential strategies. Furthermore, the other agents’ strategies may be adaptive. This means that the behaviour of other agents may depend on the agent’s past interactions, just as on other aspects of the environment – for example, a jammed door that only opens on the n th attempt – may depend upon past interactions. One situation where this adaptivity may arise is where the environment contains one or more reinforcement learning agents that learn to maximise their own reward. This kind of environment may call for aspects of social intelligence such as bluffing or information hiding, and a reward-maximising agent may have to be stochastic (i.e. use a mixed strategy) to avoid exploitation.

In rich environments such as our human world, reward maximisation may lead to a better solution than an equilibrium solution [45]. This is because it may capitalise upon suboptimal behaviours of other agents, rather than assuming optimal or worst-case behaviour. Furthermore, reward maximisation has a unique optimal value [37], while the equilibrium value is non-unique in general-sum games [40].

3.4 Reward is Enough for Language

Language has been a subject of considerable study in both natural [10] and artificial intelligence [27]. Because language plays a dominant role in human culture and interactions, the definition of intelligence itself is often premised upon the ability of understanding and using language, especially natural language [58].

Recently, significant success has been achieved by treating language as the optimisation of a singular objective: the predictive modelling of language within a large corpus of data [27, 8]. This approach has facilitated progress towards many subproblems of language that were often previously studied or implemented separately within natural language processing and understanding, including both syntactic subproblems (e.g. formal grammar, part-of-speech tagging, parsing, segmentation) and semantic subproblems (e.g. lexical semantics, entailment, sentiment analysis) as well as some that bring both together (e.g. summarisation, dialogue systems).

Nevertheless, language modelling by itself may not be sufficient to produce a broader set of linguistic abilities associated with intelligence that includes the following:

- Language may be intertwined with other modalities of action and observation. Language is often contextual, not just in what was uttered but also on what else is happening in the environment around the agent as perceived through vision and other sensory modalities (e.g. consider a dialogue between two agents carrying an awkward object or building a shelter). Furthermore, language is often interspersed with other communicative actions, such as gestures, facial expressions, tonal variations, or physical demonstrations.
- Language is consequential and purposeful. Language utterances have a consequence in the environment, typically by influencing the mental state and thereby the behaviour of other communicators within the environment. These consequences may be optimised to achieve a variety of ends, for example, a salesperson learns to tailor their language to maximise sales, while a politician learns to tailor their language to maximise votes.
- The utility of language varies according to the agent's situation and behaviour. For example, a miner may require language regarding the stability of rocks, while a farmer may need language regarding the fertility of soil. Furthermore there may be an opportunity cost to language (e.g. discussing farming instead of doing the work of farming).
- In rich environments the potential uses of language to deal with unforeseen events may outstrip the capacity of any corpus. In these cases, it may be necessary to solve linguistic problems dynamically, through experience — for example, interactively developing the most effective language to control a new disease, to build a new technology, or to find a way to address a new grievance of a rival so as to forestall aggression.

According to our hypothesis, the ability of language in its full richness, including all of these broader abilities, arises from the pursuit of reward. It is an instance of an agent's ability to produce complex sequences of actions (e.g. uttering sentences) based on complex sequences of observations (e.g. receiving sentences) in order to influence other agents in the environment (cf. discussion of social intelligence above) and accumulate greater reward [7]. The pressure to comprehend and produce language can come from many reward-increasing benefits. If an agent can comprehend a "danger" warning then it can predict and avoid negative rewards. If an agent can generate a "fetch" command then this may cause the environment (say, containing a dog) to move an object nearer to the agent. Similarly, an agent may only eat if it can comprehend complex descriptions of the location of food, generate complex instructions for growing food, engage in complex dialogue to negotiate for food, or build long-term relationships that enhance those negotiations – ultimately resulting in a variety of complex linguistic skills.

3.5 Reward is Enough for Generalisation

Generalisation is often formalised by the ability to transfer the solution to one problem into the solution to another problem [36, 56, 59]. For example, generalisation in supervised learning [36] may focus upon transferring a solution learned from one data-set, such as photographs, to another data-set, such as paintings. Generalisation in meta-learning [13, 59] has recently focused upon the problem of transferring an agent from one environment to another environment.

As per our hypothesis, generalisation may instead be understood as, and implemented by, maximising cumulative reward in a continuing stream of interaction between an agent and a single complex environment – again following a standard agent-environment protocol (see Section 2.1). Environments such as the human world demand generalisation simply because the agent encounters different aspects of the environment at different times. For example, a fruit-eating animal may encounter a new tree every day; furthermore it may become injured, suffer a drought, or face an invasive species. In each case, the animal must adapt quickly to its new state, by generalising its experience from past states. The differing states faced by the animal are not parcelled neatly into disjoint, sequential tasks with distinct labels. Instead the state depends upon the animal’s behaviour; it may combine a variety of elements that overlap and recur at different time-scales; and important aspects of the state may be partially observed. Rich environments demand the ability to generalise from past states to future states – with all these associated complexities – in order to efficiently accumulate rewards.

3.6 Reward is Enough for Imitation

Imitation is an important ability associated with human and animal intelligence, which may facilitate the rapid acquisition of other abilities, such as language, knowledge, and motor skills. In artificial intelligence, imitation has often been formulated as a problem of learning from demonstration [44] through *behavioural cloning* [2], where the goal is to reproduce the actions chosen by a teacher, when supplied with explicit data regarding the teacher’s actions, observations and rewards, typically under the assumption that the teacher is solving a symmetric problem to the agent. Behavioural cloning has led to several successful machine learning applications [52, 60, 5], especially those where human teacher data is plentiful but interactive experience is limited or costly. In contrast, the natural ability of *observational learning* [3] includes any form of learning from the observed behaviour of other humans or animals, and does not assume a symmetric teacher or require direct access to their actions, observations and rewards. This suggests that a much broader and realistic class of observational learning abilities, compared to direct imitation through behavioural cloning, may be demanded in complex environments:

- Other agents may be an integral part of the agent’s environment (e.g. a baby observing its mother), without assuming the existence of a distinct data-set containing teacher data.
- The agent may need to learn an association between its own state (e.g. the pose of the baby’s body), and the state of another agent (e.g. the pose of its mother), or between its own actions (e.g. rotating a robot’s manipulator) and observations of another agent (e.g. seeing a human hand), potentially at higher levels of abstraction (e.g. imitating the choice of food by the mother rather than her muscle activations).
- Other agents may be partially observed (e.g. where the human hand is occluded), such that their actions or goals may only be imperfectly inferred, perhaps in hindsight.
- Other agents may demonstrate undesirable behaviours that should be avoided.
- There may be many other agents in the environment, potentially exhibiting different skills or different levels of competence.
- Observational learning may even occur without any explicit agency (e.g. imitating the idea of bridge-building from observations of a log fallen across a stream).

We speculate that these broader abilities of observation learning could be driven by the maximisation of reward, from the perspective of a single agent that simply observes other agents as integral parts of its environment [6], potentially leading to many of the same benefits as behavioural cloning – such as the sample-efficient acquisition of knowledge – but in a much wider and more integrated context.

3.7 Reward is Enough for General Intelligence

For our last example, we turn to the ability which simultaneously poses the greatest challenge and where our hypothesis offers the greatest potential benefit. General intelligence, of the sort possessed by humans and perhaps also other animals, may be defined as the ability to flexibly achieve a variety of goals in different contexts. For example, humans can flexibly address problems such as locomotion, transportation or communication with solutions, depending upon their specific context, such as swimming or skiing, carrying or kicking, writing or sign language. General intelligence is sometimes

formalised by a set of environments that measures the agent’s capabilities across a variety of different goals and contexts [24, 15].

According to our hypothesis, general intelligence can instead be understood as, and implemented by, maximising a singular reward in a single, complex environment. For example, natural intelligence faces a contiguous stream of experience throughout its lifetime, generated from interactions with the natural world. An animal’s stream of experience is sufficiently rich and varied that it may demand a flexible ability to achieve a vast variety of subgoals (such as foraging, fighting, or fleeing), in order to succeed in maximising its overall reward (such as hunger or reproduction). Similarly, if an artificial agent’s stream of experience is sufficiently rich, then many goals (such as battery-life or survival) may implicitly require the ability to achieve an equally wide variety of subgoals, and the maximisation of reward should therefore be enough to yield an artificial general intelligence.

4 Reinforcement Learning Agents

Our main hypothesis, that intelligence and its associated abilities may be understood as subserving the maximisation of reward, is fully agnostic to the nature of the agent. This leaves open the important question of how to construct an agent that maximises reward. In this section, we suggest that this question may also be largely answered by reward maximisation. Specifically, we consider agents with a general ability to learn how to maximise reward from their ongoing experience of interacting with the environment. Such agents, which we will refer to as *reinforcement learning agents*, provide several advantages.³

First, among all possible solution methods for maximising reward, surely the most natural approach is to learn to do so from experience, by interacting with the environment. Over time, that interactive experience provides a wealth of information about cause and effect, about the consequences of actions, and about how to accumulate reward. Rather than predetermining the agent’s behaviour (placing faith in the designer’s foreknowledge of the environment) it is natural instead to bestow the agent with a general ability to discover its own behaviours (placing faith in experience). More specifically, the design goal of maximising reward is implemented through an ongoing internal process of learning from experience a behaviour that maximises future reward.⁴

Reinforcement learning agents, by learning from experience, provide a general solution method that may be effective, with minimal or even zero modification, across many different reward signals and environments.

Furthermore, a single environment may be so complex, like the natural world, that it contains a heterogeneous diversity of possible experiences. The potential variations in the stream of observations and rewards faced by a long-lived agent will inevitably outstrip its capacity for preprogrammed behaviours (see Section 3.1). To achieve high reward, the agent must therefore be equipped with a general ability to fully and continually adapt its behaviour to new experiences. Indeed, reinforcement learning agents may be the only feasible solutions in such complex environments.

A sufficiently powerful and general reinforcement learning agent may ultimately give rise to intelligence and its associated abilities. In other words, if an agent can continually adjust its behaviour so as to improve its cumulative reward, then any abilities that are repeatedly demanded by its environment must ultimately be produced in the agent’s behaviour. A good reinforcement learning agent could thus acquire behaviours that exhibit perception, language, social intelligence and so forth, in the course of learning to maximise reward in an environment, such as the human world, in which those abilities have ongoing value.

We do not offer any theoretical guarantee on the sample efficiency of reinforcement learning agents. Indeed, the rate at and degree to which abilities emerge will depend upon the specific environment, learning algorithm, and inductive biases; furthermore one may construct artificial environments in which learning will fail. Instead, we conjecture that powerful reinforcement learning agents, when

³It is common to use the same name to describe both the problem (e.g. mountain climbing refers to the problem of ascending to a peak), its solution methods (e.g. mountain climbing can also refer to the ropes and crampons used by mountain climbers), and the field (e.g. the pastime of mountain climbing). Where not clear from context we refer to the reinforcement learning problem, reinforcement learning agents, and the field of reinforcement learning.

⁴The internal reward may also be distinct from, but chosen to service, the design goal [49].

placed in complex environments, will in practice give rise to sophisticated expressions of intelligence. If this conjecture is correct, it offers a complete pathway towards the implementation of artificial general intelligence.

Several recent examples of reinforcement learning agents, endowed with an ability to learn to maximise rewards, have given rise to broadly capable behaviours that exceeded expectations, the performance of prior agents, and in several cases, the performance of human experts. For example, when asked to maximise wins in the game of Go, AlphaZero learned (see Section 1) an integrated intelligence across many facets of Go [47]; when the same algorithm was applied to maximise outcomes in the game of chess, AlphaZero learned a different set of abilities encompassing openings, endgames, piece mobility, king safety and so forth [46, 43]. Reinforcement learning agents that maximise score in Atari 2600 [29] have learned a range of abilities, including aspects of object recognition, localisation, navigation, and motor control demanded by each particular Atari game, while agents that maximise successful grips in vision-based robotic manipulation have learned sensorimotor abilities such as object singulation, regrasping, and dynamic object tracking [21]. While these examples are far narrower in scope than the environments faced by natural intelligence, they provide some practical evidence for the effectiveness of the reward maximisation principle.

One may of course wonder how to learn to maximise reward effectively in a practical agent. For example, the reward could be maximised directly (e.g. by optimising the agent’s policy [55]), or indirectly, by decomposing into subgoals such as representation learning, value prediction, model-learning and planning, which may themselves be further decomposed [54]. We do not address this question further in this paper, but note that it is the central question studied throughout the field of reinforcement learning.

5 Related Work

Intelligence has long been associated with goal-oriented behaviour [57]. This goal-oriented notion of intelligence is central to the concept of *rationality*, in which an agent selects the actions in the manner that *optimally* achieves its goals or maximises its utility. Rationality has been widely used to understand human behaviour [61, 4], and as a basis for artificial intelligence [41], formalised with respect to an agent interacting with an environment.

It has frequently been argued that computational constraints should also be taken into account when reasoning about goals. *Bounded* [48, 42, 35] or *computational* rationality [26] suggests that agents should select the program that best achieves their goals, given the real-time consequences arising from that program (for example, how long the program takes to execute), and subject to limitations on the set of programs (for example, limiting the maximum size of the program). Our contribution builds upon these viewpoints, but focuses on the question of whether a single, simple reward-based goal could provide a common basis for all abilities associated with intelligence.

The standard protocol for reinforcement learning was defined by Sutton and Barto [54]; in Section 2 we presented a common generalisation with partially observed histories.

Unified *cognitive architectures* [34, 1] aspire towards general intelligence. They combine a variety of solution methods for separate goals, but do not provide a generic objective that justifies and explains the choice of architecture, nor a singular goal towards which the individual components contribute.

The perspective of language as reward-maximisation dates back to behaviourism [50, 51]; however, the reinforcement learning problem differs from behaviourism in allowing an agent to construct and use internal state. In multi-agent environments, the advantages of focusing on a single agent’s objective, rather than an equilibrium objective, were discussed by Shoham and Powers [45].

6 Discussion

We have presented the reward-is-enough hypothesis and some of its implications. Next we briefly answer a number of questions that frequently arise in discussions of this hypothesis.

Which environment? One may ask which environment will give rise, through reward maximisation, to the “most intelligent” behaviour or to the “best” specific ability (for example natural language). Inevitably, the specific environmental experiences encountered by an agent — for example the friends,

foes, teachers, toys, tools, or libraries encountered during the lifetime of a human brain — will shape the nature of its subsequent abilities. While this question may be of great interest for any specific application of intelligence, we have focused instead upon the arguably more profound question of which generic objective could give rise to all forms of intelligence. The maximisation of different rewards in different environments may lead to distinct, powerful forms of intelligence, each of which exhibits its own impressive, and yet incomparable, array of abilities. A good reward-maximising agent will harness any elements present in its environment but the emergence of intelligence in some form is not predicated upon their specifics. For example, the human brain develops from birth differently when exposed to different experiences in its environment, but will robustly acquire some form of sophisticated abilities regardless of its specific culture or education.

Which reward signal? The desire to manipulate the reward signal often arises from the idea that only a carefully constructed reward could possibly induce general intelligence. By contrast, we suggest the emergence of intelligence may be quite robust to the nature of the reward signal. This is because, in rich environments such as the natural world, intelligence and its associated abilities may be demanded by even a seemingly innocuous reward signal. For example, consider a signal that provides +1 reward to the agent each time a round-shaped pebble is collected. In order to maximise this reward signal effectively, an agent may need to classify pebbles, to manipulate pebbles, to navigate to pebble beaches, to store pebbles, to understand waves and tides and their effect on pebble distribution, to persuade people to help collect pebbles, to use tools and vehicles to collect greater quantities, to quarry and shape new pebbles, to discover and build new technologies for collecting pebbles, or to build a corporation that collects pebbles.

What else, other than reward maximisation, could be enough for intelligence? *Unsupervised learning* [19] (e.g. to identify patterns in observations) and *prediction* [18, 11] (e.g. of future observations [30]) may provide effective principles for understanding experience, but do not provide a principle for action selection and therefore cannot in isolation be enough for goal-oriented intelligence. *Supervised learning* gives a mechanism for mimicking human intelligence; given enough human data one might imagine that all abilities associated with human intelligence could emerge. However, supervised learning from human data cannot be enough for a general-purpose intelligence that is able to optimise for non-human goals in non-human environments. Even where human data is plentiful, imitative intelligence may be limited in scope to the behaviours already known and exhibited by humans within the data, rather than discovering creative new behaviours that solve problems in unexpected ways (see also offline learning, below). *Evolution* by natural selection can be understood at an abstract level as maximising fitness, as measured by individual reproductive success, optimised by a population-based mechanism such as mutation and crossover. In our framework, reproductive success can be seen as one possible reward signal that has driven the emergence of natural intelligence. However, artificial intelligence may be designed with other goals than reproductive success, and may maximise the corresponding reward signals using methods other than mutation and crossover, potentially leading to very different forms of intelligence. Furthermore, while fitness maximisation may explain the initial configuration of natural intelligence (e.g. a human baby brain), a process of trial-and-error learning to maximise an intrinsic reward signal [49] (see Section 4) may furthermore explain how natural intelligence adapts through experience to develop sophisticated abilities (e.g. a human adult brain) in the service of fitness maximisation. *Maximisation of free energy* or minimisation of surprise [14] may yield several abilities of natural intelligence, but does not provide a general-purpose intelligence that can be directed towards a broad diversity of different goals in different environments. Consequently it may also miss abilities that are demanded by the optimal achievement of any one of those goals (for example, aspects of social intelligence required to mate with a partner, or tactical intelligence required to checkmate an opponent). *Optimisation* is a generic mathematical formalism that may maximise any signal, including cumulative reward, but does not specify how an agent interacts with its environment. By contrast, the reinforcement learning problem includes interaction at its heart: actions are optimised to maximise reward, those actions in turn determine the observations received from the environment, which themselves inform the optimisation process; furthermore optimisation occurs online in real-time while the environment continues to tick.

Which reward maximisation problem? Even amongst reinforcement learning researchers a number of variations of the reward maximisation problem are studied. Rather than following a standard agent-environment protocol, the interaction loop is often modified for different cases that may include multiple agents, multiple environments, or multiple training lifetimes. Rather than maximising a generic objective defined by cumulative reward, the goal is often formulated separately for different

cases: for example multi-objective learning, risk-sensitive objectives, or objectives that are specified by a human-in-the-loop. Furthermore, rather than addressing the problem of reward maximisation for general environments, special-case problems are often studied for a particular class of environments such as linear environments, deterministic environments, or stable environments. While this may be appropriate for specific applications, a solution to a specialised problem does not usually generalise; in contrast a solution to the general problem will also provide a solution for any special cases. The reinforcement learning problem may also be transformed into a probabilistic framework that approximates the goal of reward maximisation [64, 38, 25, 17]. Finally, universal decision-making frameworks [20] provide a theoretical but incomputable formulation of intelligence across all environments; while the reinforcement learning problem provides a practical formulation of intelligence in a given environment.

Can offline learning from a sufficiently large data-set be enough for intelligence? Offline learning is typically only enough to solve problems that are already to a large extent solved within the available data (for example, by the humans that generated the data). Although it may be possible to generalise to the agent’s current problems from solutions demonstrated in or extracted from an offline data-set, in rich environments, this generalisation will inevitably be imperfect. Furthermore, the data necessary to solve the agent’s current problems will often have a negligible probability of occurring in offline data (e.g. under random behaviour or human behaviour). Online interaction allows an agent to specialise to the problems it is currently facing, to continually verify and correct the most pressing holes in its knowledge, and to find new behaviours that are very different from and achieve greater reward than those in the data set. For example, a large data-set of squirrels collecting nuts is unlikely to demonstrate the behaviour required to understand how to best grow an orchard or build a nut harvester.

Is the reward signal too impoverished? One might wonder whether sample-efficient reinforcement learning agents necessarily exist that can maximise reward in complex, unknown environments. In answering this question, we note first that an effective agent may make use of additional experiential signals to facilitate the maximisation of future reward. Many solution methods, including model-free reinforcement learning, learn to associate future reward with features of observations, through value function approximation, which in turn provide rich secondary signals that drive learning of deeper associations through a recursive bootstrapping process [53]. Other solution methods, including model-based reinforcement learning, construct predictions of observations, or of features of observations, that facilitate the subsequent maximisation of reward through planning. Furthermore, as discussed in Section 3.6, observations of other agents within the environment may also facilitate rapid learning.

Nevertheless, it is often the challenge of sample-efficient reinforcement learning in complex environments that has led researchers to introduce assumptions or develop simpler abstractions that are more amenable to both theory and practice. However, these assumptions and abstractions may simply side-step the difficulties that must inevitably be faced by any broadly capable intelligence. We choose instead to accept the challenge head-on and focus upon its solution; we hope that other researchers will join us on our quest.

7 Conclusion

In this paper we have presented the hypothesis that the maximisation of total reward may be enough to understand intelligence and its associated abilities. The key idea is that rich environments typically demand a variety of abilities in the service of maximising reward. The bountiful expressions of intelligence found in nature, and presumably in the future for artificial agents, may be understood as instantiations of this same idea with different environments and different rewards. Furthermore, a singular goal of reward maximisation may give rise to deeper, broader and more integrated understanding of abilities than specialised problem formulations for each distinct ability. In particular, we explored in greater depth several abilities that may at first glance seem hard to comprehend through reward maximisation alone, such as knowledge, learning, perception, social intelligence, language, generalisation, imitation, and general intelligence, and found that reward maximisation could provide a fundamental basis for understanding each ability. Finally, we have presented a conjecture that intelligence could emerge in practice from sufficiently powerful reinforcement learning agents that learn to maximise future reward. If this conjecture is true, it provides a direct pathway towards understanding and constructing an artificial general intelligence.

References

- [1] John R Anderson, Daniel Bothell, Michael D Byrne, Scott Douglass, Christian Lebiere, and Yulin Qin. An integrated theory of the mind. *Psychological review*, 111(4):1036, 2004.
- [2] Michael Bain and Claude Sammut. A framework for behavioural cloning. In *Machine Intelligence 15*, pages 103–129, 1995.
- [3] Albert Bandura and David C McClelland. *Social learning theory*, volume 1. Englewood cliffs Prentice Hall, 1977.
- [4] G.S. Becker. *The Economic Approach to Human Behavior*. Economic theory. University of Chicago Press, 1976.
- [5] Mariusz Bojarski, Davide Del Testa, Daniel Dworakowski, Bernhard Firner, Beat Flepp, Praseoon Goyal, Lawrence D. Jackel, Mathew Monfort, Urs Muller, Jiakai Zhang, Xin Zhang, Jake Zhao, and Karol Zieba. End to end learning for self-driving cars. *CoRR*, abs/1604.07316, 2016.
- [6] Diana Borsa, Nicolas Heess, Bilal Piot, Siqi Liu, Leonard Hasenclever, Rémi Munos, and Olivier Pietquin. Observational learning by reinforcement learning. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*, pages 1117–1124. International Foundation for Autonomous Agents and Multiagent Systems, 2019.
- [7] Jeshua Bratman, Michael Shvartsman, Richard Lewis, and Satinder Singh. A new approach to exploring language emergence as boundedly optimal control in the face of environmental and cognitive constraints. *Proceedings of the 10th International Conference on Cognitive Modeling, ICCM*, 2010.
- [8] Tom B Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *arXiv:2005.14165*, 2020.
- [9] Edward C. Carterette and Morton P. Friedman. *Handbook of Perception*. Academic Press, 1978.
- [10] Noam Chomsky and David W Lightfoot. *Syntactic structures*. Walter de Gruyter, 2002.
- [11] Andy Clark. Whatever next? predictive brains, situated agents, and the future of cognitive science. *Behavioral and brain sciences*, 36(3):181–204, 2013.
- [12] Gerard Debreu. Representation of a preference ordering by a numerical function, 1954.
- [13] Chelsea Finn and Sergey Levine. Meta-learning: from few-shot learning to rapid reinforcement learning.
- [14] Karl Friston. The free-energy principle: a unified brain theory? *Nature reviews. Neuroscience*, 11:127–38, 02 2010.
- [15] Ben Goertzel and Cassio Pennachin. *Artificial general intelligence*, volume 2. Springer, 2007.
- [16] Alex Graves, Abdel-rahman Mohamed, and Geoffrey Hinton. Speech recognition with deep recurrent neural networks. In *2013 IEEE international conference on acoustics, speech and signal processing*, pages 6645–6649. IEEE, 2013.
- [17] Danijar Hafner, Pedro A. Ortega, Jimmy Ba, Thomas Parr, Karl J. Friston, and Nicolas Heess. Action and perception as divergence minimization. *CoRR*, abs/2009.01791, 2020.
- [18] Jeff Hawkins and Sandra Blakeslee. *On Intelligence*. Times Books, USA, 2004.
- [19] Geoffrey Hinton and Terrence J. Sejnowski. *Unsupervised Learning: Foundations of Neural Computation*. The MIT Press, 05 1999.
- [20] Marcus Hutter. *Universal Artificial Intelligence: Sequential Decisions based on Algorithmic Probability*. Springer, 2005.

- [21] Dmitry Kalashnikov, Alex Irpan, Peter Pastor, Julian Ibarz, Alexander Herzog, Eric Jang, Deirdre Quillen, Ethan Holly, Mrinal Kalakrishnan, Vincent Vanhoucke, and Sergey Levine. Scalable deep reinforcement learning for vision-based robotic manipulation. In *2nd Annual Conference on Robot Learning, CoRL 2018, Zürich, Switzerland, 29-31 October 2018, Proceedings*, volume 87 of *Proceedings of Machine Learning Research*, pages 651–673. PMLR, 2018.
- [22] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.
- [23] Yann LeCun. Computer perception with deep learning, 2013.
- [24] Shane Legg and Marcus Hutter. Universal intelligence: A definition of machine intelligence. *Minds Mach.*, 17(4):391–444, 2007.
- [25] Sergey Levine. Reinforcement learning and control as probabilistic inference: Tutorial and review. *CoRR*, abs/1805.00909, 2018.
- [26] Richard L. Lewis, Andrew Howes, and Satinder Singh. Computational rationality: Linking mechanism and behavior through bounded utility maximization. *topiCS*, 6(2):279–311, 2014.
- [27] Christopher D Manning, Christopher D Manning, and Hinrich Schütze. *Foundations of statistical natural language processing*. MIT press, 1999.
- [28] John McCarthy. What is AI?, 1998.
- [29] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529, 2015.
- [30] Joseph Modayil, Adam White, and Richard S. Sutton. Multi-timescale nexting in a reinforcement learning robot. *Adaptive Behaviour*, 22(2):146–160, 2014.
- [31] Martin Müller. Computer Go. *Artificial Intelligence*, 134(1-2):145–179, 2002.
- [32] John F Nash et al. Equilibrium points in n-person games. *Proceedings of the national academy of sciences*, 36(1):48–49, 1950.
- [33] J v Neumann. Zur theorie der gesellschaftsspiele. *Mathematische annalen*, 100(1):295–320, 1928.
- [34] Allen Newell. *Unified Theories of Cognition*. Harvard University Press, USA, 1990.
- [35] Laurent Orseau and Mark B. Ring. Space-time embedded intelligence. In *Proceedings of the 5th International Conference on Artificial General Intelligence*, volume 7716 of *Lecture Notes in Computer Science*, pages 209–218. Springer, 2012.
- [36] Sinno Jialin Pan and Qiang Yang. A survey on transfer learning. *IEEE Transactions on knowledge and data engineering*, 22(10):1345–1359, 2009.
- [37] Martin L Puterman. *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons, 2014.
- [38] Konrad Rawlik, Marc Toussaint, and Sethu Vijayakumar. On stochastic optimal control and reinforcement learning by approximate inference. In *Twenty-third international joint conference on artificial intelligence*, 2013.
- [39] Michael Redmond and Chris Garlock. *AlphaGo to Zero: the complete games*. Smart Go, 2020.
- [40] J Ben Rosen. Existence and uniqueness of equilibrium points for concave n-person games. *Econometrica: Journal of the Econometric Society*, pages 520–534, 1965.
- [41] Stuart Russell and Peter Norvig. *Artificial Intelligence: A Modern Approach*. Prentice Hall, 1995.

- [42] Stuart J. Russell and Devika Subramanian. Provably bounded-optimal agents. *J. Artif. Intell. Res.*, 2:575–609, 1995.
- [43] M. Sadler, N. Regan, and G. Kasparov. *Game Changer: AlphaZero’s Groundbreaking Chess Strategies and the Promise of AI*. New in Chess, 2019.
- [44] Stefan Schaal. Learning from demonstration. In Michael Mozer, Michael I. Jordan, and Thomas Petsche, editors, *Advances in Neural Information Processing Systems 9, NIPS, Denver, CO, USA, December 2-5, 1996*, pages 1040–1046. MIT Press, 1996.
- [45] Yoav Shoham, Rob Powers, and Trond Grenager. If multi-agent learning is the answer, what is the question? *Artificial Intelligence*, 171(7):365 – 377, 2007.
- [46] David Silver, Thomas Hubert, Julian Schrittwieser, Ioannis Antonoglou, Matthew Lai, Arthur Guez, Marc Lanctot, Laurent Sifre, Dharmashan Kumaran, Thore Graepel, et al. A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play. *Science*, 362(6419):1140–1144, 2018.
- [47] David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, et al. Mastering the game of go without human knowledge. *nature*, 550(7676):354–359, 2017.
- [48] Herbert A. Simon. A behavioral model of rational choice. *The Quarterly Journal of Economics*, 69(1):99–118, 1955.
- [49] Satinder Singh, Richard L. Lewis, Andrew G. Barto, and Jonathan Sorg. Intrinsically motivated reinforcement learning: An evolutionary perspective. *IEEE Trans. Auton. Ment. Dev.*, 2(2):70–82, 2010.
- [50] Burrhus Frederic Skinner. *The Behavior of Organisms: An Experimental Analysis*. New York: Appleton-Century-Crofts, 1938.
- [51] Burrhus Frederic Skinner. *Verbal behavior*. New York: Appleton-Century-Crofts, 1957.
- [52] Nisan Stiennon, Long Ouyang, Jeffrey Wu, Daniel M. Ziegler, Ryan Lowe, Chelsea Voss, Alec Radford, Dario Amodei, and Paul F. Christiano. Learning to summarize with human feedback. In *Advances in Neural Information Processing Systems 33*, 2020.
- [53] Richard S Sutton. Learning to predict by the methods of temporal differences. *Machine learning*, 3(1):9–44, 1988.
- [54] Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning: An Introduction*. The MIT Press, second edition, 2018.
- [55] Richard S Sutton, David A McAllester, Satinder Singh, and Yishay Mansour. Policy gradient methods for reinforcement learning with function approximation. In *Advances in neural information processing systems*, pages 1057–1063, 2000.
- [56] Matthew E. Taylor and Peter Stone. Transfer learning for reinforcement learning domains: A survey. *Journal of Machine Learning Research*, 10(1):1633–1685, 2009.
- [57] E. C. Tolman. *Purposive behavior in animals and men*. Century/Random House, UK, 1932.
- [58] Alan M Turing. Computing machinery and intelligence. *Mind*, 59(236):433, 1950.
- [59] Joaquin Vanschoren. Meta-learning: A survey. *CoRR*, abs/1810.03548, 2018.
- [60] Oriol Vinyals, Igor Babuschkin, Wojciech M. Czarnecki, Michaël Mathieu, Andrew Dudzik, Junyoung Chung, David H. Choi, Richard Powell, Timo Ewalds, Petko Georgiev, Junhyuk Oh, Dan Horgan, Manuel Kroiss, Ivo Danihelka, Aja Huang, Laurent Sifre, Trevor Cai, John P. Agapiou, Max Jaderberg, Alexander Sasha Vezhnevets, Rémi Leblond, Tobias Pohlen, Valentin Dalibard, David Budden, Yury Sulsky, James Molloy, Tom L. Paine, Çağlar Gülçehre, Ziyu Wang, Tobias Pfaff, Yuhuai Wu, Roman Ring, Dani Yogatama, Dario Wünsch, Katrina McKinney, Oliver Smith, Tom Schaul, Timothy P. Lillicrap, Koray Kavukcuoglu, Demis Hassabis, Chris Apps, and David Silver. Grandmaster level in starcraft II using multi-agent reinforcement learning. *Nat.*, 575(7782):350–354, 2019.

- [61] M. Weber, G. Roth, C. Wittich, and E. Fischhoff. *Economy and Society: An Outline of Interpretive Sociology*. Economy and Society: An Outline of Interpretative Sociology. University of California Press, 1978.
- [62] Yuan Zhou. *AlphaGo vs Ke Jie*. Slate and Shell, 2017.
- [63] Yuan Zhou. *Rethinking opening strategy: AlphaGo's impact on pro play*. Slate and Shell, 2018.
- [64] Brian D. Ziebart, J. Andrew Bagnell, and Anind K. Dey. Modeling interaction via the principle of maximum causal entropy. In *Proceedings of the 27th International Conference on Machine Learning*, pages 1255–1262. Omnipress, 2010.

Declaration of interests

☐ The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

☒ The authors declare the following financial interests/personal relationships which may be considered as potential competing interests:

All authors are employees of DeepMind

Author declaration

[Instructions: Please check all applicable boxes and provide additional information as requested.]

1. Conflict of Interest

Potential conflict of interest exists:

We wish to draw the attention of the Editor to the following facts, which may be considered as potential conflicts of interest, and to significant financial contributions to this work:

The nature of potential conflict of interest is described below:

☒ No conflict of interest exists.

We wish to confirm that there are no known conflicts of interest associated with this publication and there has been no significant financial support for this work that could have influenced its outcome.

2. Funding

☒ Funding was received for this work.

All of the sources of funding for the work described in this publication are acknowledged below:

[List funding sources and their role in study design, data analysis, and result interpretation]

All authors are full or part-time employees of DeepMind

☐ No funding was received for this work.

3. Intellectual Property

☒ We confirm that we have given due consideration to the protection of intellectual property associated with this work and that there are no impediments to publication, including the timing of publication, with respect to intellectual property. In so doing we confirm that we have followed the regulations of our institutions concerning intellectual property.

4. Research Ethics

☐ We further confirm that any aspect of the work covered in this manuscript that has involved human patients has been conducted with the ethical approval of all relevant bodies and that such approvals are acknowledged within the manuscript.

☐ IRB approval was obtained (required for studies and series of 3 or more cases)

☐ Written consent to publish potentially identifying information, such as details or the case and photographs, was obtained from the patient(s) or their legal guardian(s).

N/A: no human patients were involved

5. Authorship

The International Committee of Medical Journal Editors (ICMJE) recommends that authorship be based on the following four criteria:

1. Substantial contributions to the conception or design of the work; or the acquisition, analysis, or interpretation of data for the work; AND
2. Drafting the work or revising it critically for important intellectual content; AND
3. Final approval of the version to be published; AND
4. Agreement to be accountable for all aspects of the work in ensuring that questions related to the accuracy or integrity of any part of the work are appropriately investigated and resolved.

All those designated as authors should meet all four criteria for authorship, and all who meet the four criteria should be identified as authors. For more information on authorship, please see <http://www.icmje.org/recommendations/browse/roles-and-responsibilities/defining-the-role-of-authors-and-contributors.html#two>.

☒ All listed authors meet the ICMJE criteria. We attest that all authors contributed significantly to the creation of this manuscript, each having fulfilled criteria as established by the ICMJE.

☐ One or more listed authors do(es) not meet the ICMJE criteria.

We believe these individuals should be listed as authors because:

[Please elaborate below]

☒ We confirm that the manuscript has been read and approved by all named authors.

☒ We confirm that the order of authors listed in the manuscript has been approved by all named authors.

6. Contact with the Editorial Office

The Corresponding Author declared on the title page of the manuscript is:

[Insert name below]

David Silver

☒ This author submitted this manuscript using his/her account in EVISE.

☒ We understand that this Corresponding Author is the sole contact for the Editorial process (including EVISE and direct communications with the office). He/she is responsible for communicating with the other authors about progress, submissions of revisions and final approval of proofs.

☒ We confirm that the email address shown below is accessible by the Corresponding Author, is the address to which Corresponding Author's EVISE account is linked, and has been configured to accept email from the editorial office of American Journal of Ophthalmology Case Reports:

[Insert email address you wish to use for communication with the journal here]





davidsilver@google.com

☐ Someone other than the Corresponding Author declared above submitted this manuscript from his/her account in EVISE:

[Insert name below]

☐ We understand that this author is the sole contact for the Editorial process (including EVISE and direct communications with the office). He/she is responsible for communicating with the other authors, including the Corresponding Author, about progress, submissions of revisions and final approval of proofs.

We the undersigned agree with all of the above.

Author's name (First, Last)	Signature	Date
1. David Silver		8 October 2020
2. Satinder Singh		8 October 2020
3. Doina Precup		8 October 2020
4. Richard Sutton		8 October 2020