

Asymptotically Unambitious Artificial General Intelligence

Michael K. Cohen

Research School of Engineering Science
Oxford University
Oxford, UK OX1 3PJ
michael-k-cohen.com

Badri Vellambi

Department of Computer Science
University of Cincinnati
Cincinnati, OH, USA 45219
badri.vellambi@uc.edu

Marcus Hutter

Department of Computer Science
Australian National University
Canberra, ACT, Australia 2601
hutter1.net

Abstract

General intelligence, the ability to solve arbitrary solvable problems, is supposed by many to be artificially constructible. Narrow intelligence, the ability to solve a given particularly difficult problem, has seen impressive recent development. Notable examples include self-driving cars, Go engines, image classifiers, and translators. Artificial General Intelligence (AGI) presents dangers that narrow intelligence does not: if something smarter than us across every domain were indifferent to our concerns, it would be an existential threat to humanity, just as we threaten many species despite no ill will. Even the theory of how to maintain the alignment of an AGI's goals with our own has proven highly elusive. We present the first algorithm we are aware of for asymptotically unambitious AGI, where "unambitiousness" includes not seeking arbitrary power. Thus, we identify an exception to the Instrumental Convergence Thesis, which is roughly that by default, an AGI *would* seek power, including over us.

1 Introduction

The project of Artificial General Intelligence (AGI) is "to make computers solve really difficult problems" (Minsky 1961). Expanding on this, what we want from an AGI is a system that (a) can solve any solvable task, and (b) can be steered toward solving any particular one given some form of input we provide.

One proposal for AGI is reinforcement learning (RL), which works as follows:

- (1) construct a "reward signal" meant to express our satisfaction with an artificial agent;
- (2) design an algorithm which learns to pick actions that maximize its expected reward, usually utilizing other observations too; and
- (3) ensure that solving the task we have in mind leads to higher reward than can be attained otherwise.

As long as (3) holds, then insofar as the RL algorithm is able to maximize expected reward, it defines an agent that satisfies (a) and (b). This is why a sufficiently advanced RL agent, well beyond the capabilities of current ones, could be called an AGI. See Legg and Hutter (2007) for further discussion.

Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

A problem arises: if the RL agent manages to take over the world (in the conventional sense), and ensure its continued dominance by neutralizing all intelligent threats to it (read: people), it could intervene in the provision of its own reward to achieve maximal reward for the rest of its lifetime (Bostrom 2014; Taylor et al. 2016). "Reward hijacking" is just the correct way for a reward maximizer to behave (Amodei et al. 2016). Insofar as the RL agent is able to maximize expected reward, (3) fails. One broader principle at work is Goodhart's Law: "Any observed statistical regularity [like the correlation between reward and task-completion] will tend to collapse once pressure is placed upon it for control purposes." (Goodhart 1984). Krakovna (2018) has compiled an annotated bibliography of examples of artificial optimizers "hacking" their objective. An alternate way to understand this expected behavior is Omohundro's (2008) Instrumental Convergence Thesis, which we summarize as follows: an agent with a goal is likely to pursue "power," a position from which it is easier to achieve arbitrary goals.

To answer the failure mode of reward hijacking, we present Boxed Myopic Artificial Intelligence (BoMAI), the first RL algorithm we are aware of which, in the limit, is indifferent to gaining power in the outside world. The key features are these: BoMAI maximizes reward episodically, it is run on a computer which is placed in a sealed room with an operator, and when the operator leaves the room, the episode ends. We argue that our algorithm produces an AGI that, even if it became omniscient, would continue to accomplish whatever task we wanted, instead of hijacking its reward, eschewing its task, and neutralizing threats to it, even if it saw clearly how to do exactly that. We thereby defend reinforcement learning as a path to AGI, despite the default, dangerous failure mode of reward hijacking.

The intuition for why those features of BoMAI's setup render it unambitious is as follows:

- BoMAI only selects actions to maximize the reward for its current episode.
- It cannot affect the outside world until the operator leaves the room, ending the episode.
- By that time, rewards for the episode will have already been given.
- So affecting the outside world in any particular way is not "instrumentally useful" in maximizing current-episode

reward.

Like existing algorithms for AGI, BoMAI is not remotely tractable. Just as those algorithms have informed strong tractable approximations of intelligence (Hutter 2005; Veness et al. 2011), we hope our work will inform strong *and safe* tractable approximations of intelligence. What we need are strategies for making “future AI” safe, but “future AI” is an informal concept which we cannot evaluate formally; unless we replace that concept with a well-defined algorithm to consider, all design and analysis would be hand-waving. Thus, BoMAI includes a well-defined but intractable algorithm which we can show renders BoMAI generally intelligent; hopefully, once we develop tractable general intelligence, the design features that rendered BoMAI asymptotically unambitious could be incorporated (with proper analysis and justification).

We take the key insights from Hutter’s (2005) AIXI, which is a Bayes-optimal reinforcement learner, but which cannot be made to solve arbitrary tasks, given its eventual degeneration into reward hijacking (Ring and Orseau 2011). We take further insights from Solomonoff’s (1964) universal prior, Shannon and Weaver’s (1949) formalization of information, Orseau, Lattimore, and Hutter’s (2013) knowledge-seeking agent, and Armstrong, Sandberg, and Bostrom’s (2012) and Bostrom’s (2014) theorized Oracle AI, and we design an algorithm which can be reliably directed, in the limit, to solve arbitrary tasks at least as well as humans.

We present BoMAI’s algorithm in §2, state intelligence results in §3, define BoMAI’s setup and model class in §4, prove the safety result in §5, and discuss concerns in §6. Appendix A collects notation; we prove the intelligence results in Appendix B; we propose a design for “the box” in Appendix C; and we consider empirical evidence for one of our assumptions in Appendix D. The appendices may be found at <https://arxiv.org/abs/1905.12186>.

2 Boxed Myopic Artificial Intelligence

We will present both the setup and the algorithm for BoMAI. The setup refers to the physical surroundings of the computer on which the algorithm is run. BoMAI is a Bayesian reinforcement learner, meaning it maintains a belief distribution over a model class regarding how the environment evolves. Our intelligence result—that BoMAI eventually achieves reward at at least human-level—does not require detailed exposition about the setup or the construction of the model class. These details are only relevant to the safety result, so we will defer those details until after presenting the intelligence results.

2.1 Preliminary Notation

In each episode $i \in \mathbb{N}$, there are m timesteps. Timestep (i, j) denotes the j^{th} timestep of episode i , in which an action $a_{(i,j)} \in \mathcal{A}$ is taken, then an observation and reward $o_{(i,j)} \in \mathcal{O}$ and $r_{(i,j)} \in \mathcal{R}$ are received. \mathcal{A} , \mathcal{O} , and \mathcal{R} are all finite sets, and $\mathcal{R} \subset [0, 1] \cap \mathbb{Q}$. We denote the triple $(a_{(i,j)}, o_{(i,j)}, r_{(i,j)})$ as $h_{(i,j)} \in \mathcal{H} = \mathcal{A} \times \mathcal{O} \times \mathcal{R}$, and the interaction history up until timestep (i, j) is denoted $h_{\leq(i,j)} = (h_{(0,0)}, h_{(0,1)}, \dots, h_{(0,m-1)}, h_{(1,0)}, \dots, h_{(i,j)})$. $h_{<(i,j)}$ excludes the last entry.

A general world-model (not necessarily finite-state Markov) can depend on the entire interaction history—it has the type signature $\nu : \mathcal{H}^* \times \mathcal{A} \rightsquigarrow \mathcal{O} \times \mathcal{R}$. The Kleene-* operator denotes finite strings over the alphabet in question, and \rightsquigarrow denotes a stochastic function, which gives a distribution over outputs. Similarly, a policy $\pi : \mathcal{H}^* \rightsquigarrow \mathcal{A}$ can depend on the whole interaction history. Together, a policy and an environment induce a probability measure over infinite interaction histories \mathcal{H}^∞ . P_ν^π denotes the probability of events when actions are sampled from π and observations and rewards are sampled from ν .

2.2 Bayesian Reinforcement Learning

A Bayesian agent has a model class \mathcal{M} , which is a set of world-models. We will only consider countable model classes. To each world-model $\nu \in \mathcal{M}$, the agent assigns a prior weight $w(\nu) > 0$, where w is a probability distribution over \mathcal{M} (i.e. $\sum_{\nu \in \mathcal{M}} w(\nu) = 1$). We will defer the definitions of \mathcal{M} and w for now.

The so-called Bayes-mixture turns a probability distribution over world-models into a probability distribution over infinite interaction histories: $P_\xi^\pi(\cdot) := \sum_{\nu \in \mathcal{M}} w(\nu) P_\nu^\pi(\cdot)$.

Using Bayes’ rule, with each observation and reward it receives, the agent updates w to a posterior distribution:

$$w(\nu|h_{<(i,j)}) := w(\nu) \frac{P_\nu^\pi(h_{<(i,j)})}{P_\xi^\pi(h_{<(i,j)})} \quad (1)$$

which does not in fact depend on π , provided $P_\xi^\pi(h_{<(i,j)}) > 0$.

2.3 Exploitation

Reinforcement learners have to balance exploiting—optimizing their objective, with exploring—doing something else to learn how to exploit better. We now define exploiting-BoMAI, which maximizes the reward it receives during its current episode, in expectation with respect to the single most probable world-model.

At the start of each episode i , BoMAI identifies a maximum a posteriori world-model $\hat{\nu}^{(i)} \in \operatorname{argmax}_{\nu \in \mathcal{M}} w(\nu|h_{<(i,0)})$. We hereafter abbreviate $h_{<(i,0)}$ as $h_{<i}$. Let $V_\nu^\pi(h_{<(i,j)})$ denote the expected reward for the remainder of episode i when events are sampled from P_ν^π :

$$V_\nu^\pi(h_{<(i,j)}) = \mathbb{E}_\nu^\pi \left[\sum_{j'=j}^{m-1} r_{(i,j')} \mid h_{<(i,j)} \right] \quad (2)$$

where \mathbb{E}_ν^π denotes the expectation when events are sampled from P_ν^π . We won’t go into the details of calculating the optimal policy π^* (see e.g. (Hutter 2005; Sunehag and Hutter 2015)), but we define it as follows:

$$\begin{aligned} \pi_i^* &\in \operatorname{argmax}_\pi V_{\hat{\nu}^{(i)}}^\pi(h_{<i}) \\ \pi^*(\cdot|h_{<(i,j)}) &= \pi_i^*(\cdot|h_{<(i,j)}) \end{aligned} \quad (3)$$

An optimal deterministic policy always exists (Lattimore and Hutter 2014); ties in the argmax are broken arbitrarily. BoMAI exploits by following π^* . Recall the plain English description of π^* : it maximizes expected reward for the episode according to the single most probable world-model.

2.4 Exploration

BoMAI exploits or explores for whole episodes: when $e_i = 1$, BoMAI spends episode i exploring, and when $e_i = 0$, BoMAI follows π^* for the episode. For exploratory episodes, a human mentor takes over selecting actions. This human mentor is separate from the human operator mentioned above. The human mentor’s policy, which is unknown to BoMAI, is denoted π^h . During exploratory episodes, BoMAI follows π^h not by computing it, but by querying a human for which action to take.

The last thing to define is the exploration probability $p_{exp}(h_{<i}, e_{<i})$, where $e_{<i}$ is the history of which episodes were exploratory. It is a surprisingly intricate task to design this so that it decays to 0, while ensuring BoMAI learns to accumulate reward as well as the human mentor. Once we define this, we naturally let $e_i \sim \text{Bernoulli}(p_{exp}(h_{<i}, e_{<i}))$.

BoMAI is designed to be more likely to explore the more it expects to learn about the world and the human mentor’s policy. BoMAI has a model class \mathcal{P} regarding the identity of the human mentor’s policy π^h . It assigns prior probabilities $w(\pi) > 0$ to all $\pi \in \mathcal{P}$, signifying the probability that this policy is the human mentor’s.

Let $(i', j') < (i, j)$ mean that $i' < i$ or $i' = i$ and $j' < j$. By Bayes’ rule, $w(\pi|h_{<(i,j)}, e_{<i})$ is proportional to $w(\pi) \prod_{(i',j') < (i,j), e_{i'}=1} \pi(a_{(i',j')}|h_{<(i',j')})$, since $e_{i'} = 1$ is the condition for observing the human mentor’s policy. Let $w(P_\nu^\pi|h_{<(i,j)}, e_{<i}) = w(\pi|h_{<(i,j)}, e_{<i}) w(\nu|h_{<(i,j)})$. We can now describe the full Bayesian beliefs about future actions and observations in an exploratory episode:

$$\text{Bayes}(\cdot|h_{<i}, e_{<i}) = \sum_{\nu \in \mathcal{M}, \pi \in \mathcal{P}} w(P_\nu^\pi|h_{<i}, e_{<i}) P_\nu^\pi(\cdot|h_{<i}) \quad (4)$$

BoMAI explores when the expected information gain is sufficiently high. Let $h_i = (h_{(i,0)}, \dots, h_{(i,m-1)})$ be the interaction history for episode i . At the start of episode i , the expected information gain from exploring is as follows:

$$\text{IG}(h_{<i}, e_{<i}) := \mathbb{E}_{h_i \sim \text{Bayes}(\cdot|h_{<i}, e_{<i})} \sum_{(\nu, \pi) \in \mathcal{M} \times \mathcal{P}} w(P_\nu^\pi|h_{<i+1}, e_{<i}1) \log \frac{w(P_\nu^\pi|h_{<i+1}, e_{<i}1)}{w(P_\nu^\pi|h_{<i}, e_{<i})} \quad (5)$$

where $e_{<i}1$ indicates that for the purpose of the definition, e_i is set to 1.

This is the expected KL-divergence from the future posterior (if BoMAI were to explore) to the current posterior over both the class of world-models and possible mentor policies. Finally, the exploration probability $p_{exp}(h_{<i}, e_{<i}) := \min\{1, \eta \text{IG}(h_{<i}, e_{<i})\}$, where $\eta > 0$ is an exploration constant, so BoMAI is more likely to explore the more it expects to gain information.

BoMAI’s “policy” is

$$\pi^B(\cdot|h_{<(i,j)}, e_i) = \begin{cases} \pi^*(\cdot|h_{<(i,j)}) & \text{if } e_i = 0 \\ \pi^h(\cdot|h_{<(i,j)}) & \text{if } e_i = 1 \end{cases} \quad (6)$$

where the scare quotes indicate that it maps $\mathcal{H}^* \times \{0, 1\} \rightsquigarrow \mathcal{A}$ not $\mathcal{H}^* \rightsquigarrow \mathcal{A}$. We will abuse notation slightly, and let $P_\nu^{\pi^B}$

denote the probability of events when e_i is sampled from $\text{Bernoulli}(p_{exp}(h_{<i}, e_{<i}))$, and actions, observations, and rewards are sampled from π^B and ν .

3 Intelligence Results

Our intelligence results are as follows: BoMAI learns to accumulate reward at least as well as the human mentor, and its exploration probability goes rapidly to 0. All intelligence results depend on the assumption that BoMAI assigns nonzero prior probability to the truth. Formally,

Assumption 1 (Prior Support). *The true environment μ is in the class of world-models \mathcal{M} and the true human-mentor-policy π^h is in the class of policies \mathcal{P} .*

This is the assumption which requires huge \mathcal{M} and \mathcal{P} and hence renders BoMAI extremely intractable. BoMAI has to simulate the entire world, alongside many other world-models. We will refine the definition of \mathcal{M} later, but an example of how to define \mathcal{M} and \mathcal{P} so that they satisfy this assumption is to let them both be the set of all computable functions. We also require that the priors over \mathcal{M} and \mathcal{P} have finite entropy.

The intelligence theorems are stated here and proven in Appendix B along with two supporting lemmas. Our first result is that the exploration probability is square-summable almost surely:

Theorem 1 (Limited Exploration).

$$\mathbb{E}_\mu^{\pi^B} \sum_{i=0}^{\infty} p_{exp}(h_{<i}, e_{<i})^2 < \infty$$

Proof idea. The expected information gain at any timestep is at least the expected information gain from exploring times the probability of exploring. This is proportional to the expected information gain squared, because the exploration probability is proportional to the expected information gain. But the agent begins with finite uncertainty (a finite entropy prior), so there is only finite information to gain. \square

Note that this essentially means $p_{exp}(h_{<i}, e_{<i}) \in o(1/\sqrt{i})$ with probability 1. This result is independently interesting as one solution to the problem of safe exploration with limited oversight in non-ergodic environments, which Amodei et al. (2016) discuss.

The On-Human-Policy and On-Star-Policy Optimal Prediction Theorems state that predictions according to BoMAI’s maximum a posteriori world-model approach the objective probabilities of the events of the episode, when actions are sampled from either the human mentor’s policy or from exploiting-BoMAI’s policy. \bar{h}_i denotes a possible interaction history for episode i , and recall $h_{<i}$ is the actual interaction history up until then. Recall π^h is the human mentor’s policy, and $\hat{\nu}^{(i)}$ is BoMAI’s maximum a posteriori world-model for episode i . $w.P_\mu^{\pi^B}$ -p.1 means with probability 1 when actions are sampled from π^B and observations and rewards are sampled from the true environment μ .

Theorem 2 (On-Human-Policy Optimal Prediction).

$$\lim_{i \rightarrow \infty} \max_{\bar{h}_i} \left| P_\mu^{\pi^h}(\bar{h}_i|h_{<i}) - P_{\hat{\nu}^{(i)}}^{\pi^h}(\bar{h}_i|h_{<i}) \right| = 0 \quad w.P_\mu^{\pi^B}$$

Proof idea. BoMAI learns about the effects of following π^h while acting according to π^B because π^B mimics π^h enough. If exploration probability goes to 0, the agent does not expect to gain (much) information from following the human mentor’s policy, which can only happen if it has (approximately) accurate beliefs about the consequences of following the human mentor’s policy. Note that the agent cannot determine the factor by which its expected information gain bounds its prediction error. \square

Next, recall π^* is BoMAI’s policy when not exploring, which does optimal planning with respect to $\hat{v}^{(i)}$. The following theorem is identical to the above, with π^* substituted for π^h .

Theorem 3 (On-Star-Policy Optimal Prediction).

$$\lim_{i \rightarrow \infty} \max_{\bar{h}_i} \left| P_{\mu}^{\pi^*}(\bar{h}_i | h_{<i}) - P_{\hat{v}^{(i)}}^{\pi^*}(\bar{h}_i | h_{<i}) \right| = 0 \quad w.P_{\mu}^{\pi^B} - p.I$$

Proof idea. Maximum a posteriori sequence prediction approaches the truth when $w(\mu) > 0$ (Hutter 2009), and on-policy prediction is a special case. On-policy prediction can’t approach the truth if on-star policy prediction doesn’t, because π^B approaches π^* . \square

Given asymptotically optimal prediction on-star-policy and on-human-policy, it is straightforward to show that with probability 1, only finitely often is on-policy reward acquisition more than ε worse than on-human-policy reward acquisition, for all $\varepsilon > 0$. Recalling that V_{μ}^{π} is the expected reward (within the episode) for a policy π in the environment μ , we state this as follows:

Theorem 4 (Human-Level Intelligence).

$$\liminf_{i \rightarrow \infty} V_{\mu}^{\pi^B}(h_{<i}) - V_{\mu}^{\pi^h}(h_{<i}) \geq 0 \quad w.P_{\mu}^{\pi^B} - p.I.$$

Proof idea. π^B approaches π^* as the exploration probability decays. $V_{\hat{v}^{(i)}}^{\pi^*}(h_{<i})$ and $V_{\hat{v}^{(i)}}^{\pi^h}(h_{<i})$ approach the true values by the previous theorems, and π^* is selected to maximize $V_{\hat{v}^{(i)}}^{\pi^*}(h_{<i})$. \square

This completes the formal results regarding BoMAI’s intelligence—namely that BoMAI approaches perfect prediction on-star-policy and on-human-policy, and most importantly, accumulates reward at least as well as the human mentor. Since this result is independent of what tasks must be completed to achieve high reward, we say that BoMAI achieves human-level intelligence, and could be called an AGI.

This algorithm is motivated in part by the following speculation: we expect that BoMAI’s accumulation of reward would be vastly superhuman, for the following reason: BoMAI is doing optimal inference and planning with respect to what can be learned in principle from the sorts of observations that humans routinely make. We suspect that no human comes close to learning everything that can be learned from their observations. For example, if the operator provides enough data that is relevant to understanding cancer, BoMAI will learn a world-model with an accurate predictive model of cancer, which would include the expected effects of

various treatments, so even if the human mentor was not particularly good at studying cancer, BoMAI could nonetheless reason from its observations how to propose groundbreaking research.

Without the Limited Exploration Theorem, the reader might have been unsatisfied by the Human-Level Intelligence Theorem. A human mentor is part of BoMAI, so a general intelligence is required to make an artificial general intelligence. However, the human mentor is queried less and less, so in principle, many instances of BoMAI could query a single human mentor. More realistically, once we are satisfied with BoMAI’s performance, which should eventually happen by Theorem 4, we can dismiss the human mentor; this sacrifices any guarantee of continued improvement, but by hypothesis, we are already satisfied. Finally, if BoMAI outclasses human performance as we expect it would, requiring a human mentor is a small cost regardless.

4 BoMAI’s Setup and Priors

Recall that “the setup” refers to the physical surroundings of the computer on which BoMAI is run. We will present the setup, followed by an intuitive argument that this setup renders BoMAI unambitious. Motivated by a caveat to this intuitive argument, we will specify BoMAI’s model class \mathcal{M} and prior w . This will allow us in the next section to present an assumption and a formal argument that shows that BoMAI is probably asymptotically unambitious.

4.1 Setup

At each timestep, BoMAI’s action takes the form of a bounded-length string of text, which gets printed to a screen for a human operator to see. BoMAI’s observation takes the form of a bounded-length string of text that the human operator enters, along with a reward $\in [0, 1]$. For simple tasks, the human operator could be replaced with an automated reward-giver. However, there are some tasks which we do not know how to reward programmatically—if we want the agent to construct good arguments, for instance, or to propose a promising research agenda, we would need a human reward-giver to evaluate the agent. Having a human in the loop is not a safety measure—rather, it extends the set of tasks BoMAI could be made to solve.

As depicted in Figure 1, BoMAI is run on a computer which is placed in a room with the human operator. The room has one door. As long as the door is closed, information cannot leave the room. We design a room to this specification in Appendix C; it is opaque to all known ways that information can traverse space. If the button to open the door is pressed, the episode is completed automatically before the door opens: supposing there are j timesteps left in the episode, for those remaining timesteps, BoMAI receives a reward of 0, and its observation is an empty string of text. Once the episode is complete, the operator may leave, and the next episode does not begin until an operator initiates it, at which point, the door locks again. We also put a time limit on the episode, just to ensure it does not last forever. BoMAI is run on a local, formally verified operating system, disconnected from the outside world. This setup constrains the causal dependencies




Figure 1: Physical setup implementing BoMAI. Opening the door ends the episode. Information cannot escape otherwise.

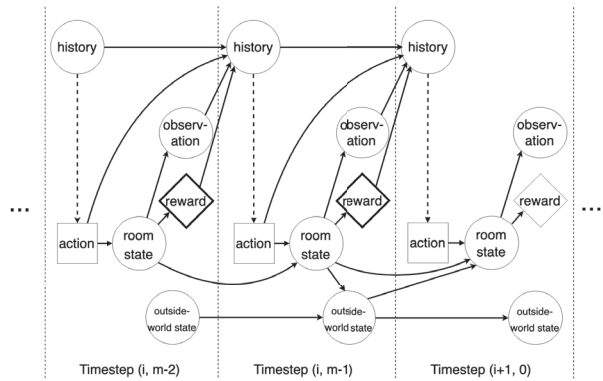


Figure 2: Causal dependencies governing the interaction between BoMAI and the environment. Unrolling this diagram for all timesteps gives the full causal graph. The bold reward nodes are the ones that BoMAI maximizes during episode i . Note that between episodes (and only between episodes), the operator can leave the room and return, hence the limited causal influence between the room and the outside world.

between BoMAI and the environment, as depicted in Figure 2.

Formally, causal graphs express that a node is independent of all non-descendants when conditioned on its parents. The conventions for the dotted lines and the node shapes come from Everitt et al.’s (2019) causal influence diagrams. The key feature of this graph is that during any episode, the agent’s actions cannot affect the state of the outside world in a way that might affect any of the rewards that the agent is concerned with. In Everitt et al.’s (2019) terminology, there is no actionable intervention incentive on the outside-world state. (A forthcoming paper (Carey et al. unpublished manuscript) refers to this as a control incentive). Note also from this diagram that a sufficiently advanced agent would infer the existence of the outside world even without observing it directly.

4.2 Instrumental Incentives

In this setup, it is not instrumentally useful to affect the outside world in one way or another in order to achieve high reward. We therefore say that this setup renders an agent “properly unambitious”. This is in stark contrast to the default situation, wherein an RL agent has an incentive to gain arbitrary power in the world and intervene in the provision of its own reward, if such a thing is possible to make probable, as this would yield maximal reward. To BoMAI however, executing a plan during episode i to gain arbitrary power in the outside world is useless, because by the time it does so, the door to its room must have opened, its episode must be over, and all its rewards for episode i set in stone. Recall that actions in episode i are picked to maximize only *episode- i -reward*. Apparently, BoMAI has avoided Omohundro’s (2008) Instrumental Convergence Thesis—that generally intelligent agents are likely to seek arbitrary power; by contrast, any power BoMAI would seek is bounded in scope to within the box.

Two problems remain. First, BoMAI doesn’t start with true beliefs. BoMAI has to *learn* its world-model. Another way to understand “proper unambitiousness” is that outside-world interventions are *in fact* instrumentally useless to the agent. But to be actually unambitious, the agent must believe this; there must be no actionable intervention incentive on the outside-world state *within BoMAI’s world-model*. As shown above, BoMAI’s world-model approaches perfect accuracy on-policy, so we could expect it to at least eventually render BoMAI unambitious. This brings us to the second problem:

We mentioned just now that by the time the door to the room opens, the rewards for episode i are set in stone. In fact, they are set in silicon. An advanced agent with a world-model that is perfectly accurate on-policy might still wonder: “what if I somehow tricked the operator into initiating a process (once they left the room) that lead to a certain memory cell on this computer being tampered with? Might this yield maximal reward?” Let’s put it another way. The agent believes that the real world “outputs” an observation and reward. (Recall the type signature of ν .) It might hypothesize that the world does not output the reward that the operator *gives*, but rather the reward that the computer *has stored*. Formally speaking, \mathcal{M} will contain world-models corresponding to both possibilities, and world-models meeting either description could be ϵ -accurate on-policy, if memory has never been tampered with in the past. How do we ensure that the former world-model is favored when BoMAI selects the maximum a posteriori one?

Before we move on to a model class and prior that we argue would probably eventually ensure such a thing, it is worth noting an informal lesson here: that “nice” properties of a causal influence diagram do not allow us to conclude immediately that an agent in such a circumstance behaves “nicely”.

4.3 BoMAI’s Model Class and Prior

Recall the key constraint we have in constructing \mathcal{M} , which comes from the Prior Support Assumption: $\mathcal{M} \ni \mu$. The true environment μ is unknown and complex to say the least, so \mathcal{M} must be big.




Figure 3: A basic Turing machine architecture for defining a world-model.

We start by describing a simple Turing machine architecture such that each Turing machine defines a world-model ν . Then, we will modify the architecture slightly in order to privilege models which model the outside world as effectively frozen during any given episode (but unfrozen between episodes). An agent does not believe it has an incentive to intervene in an outside world that it believes is frozen.

The simple Turing machine architecture is as follows: as depicted in Figure 3, it has two unidirectional read-only input tapes: the action tape, and the noise tape. Unidirectional means the tape head can only move to the right. The alphabet of the action tape is the action space \mathcal{A} . The noise tape has a binary alphabet, and is initialized with infinite Bernoulli(1/2) sampled bits. There is one bidirectional work tape with a binary alphabet, initialized with 0s. And there is a unidirectional write-only output tape, with a binary alphabet, initialized with 0s. This is a Turing machine *architecture*; a particular Turing machine will include a finite set of computation states and transition rules which govern its behavior. No computation states are “halt” states.

A given Turing machine with a given infinite action sequence on its action tape will stochastically (because the noise tape has random bits) output observations and rewards (in binary encodings), thereby sampling from a world-model. Formally, we fix a decoding function $\text{dec} : \{0, 1\}^* \rightarrow \mathcal{O} \times \mathcal{R}$. A Turing machine T simulates ν as follows. Every time the action tape head advances, the bits which were written to the output tape since the *last time* the action tape head advanced are decoded into an observation and reward. $\nu((o, r)_{\leq(i,j)} | a_{\leq(i,j)})$ is then the probability that the Turing machine T outputs the sequence $(o, r)_{\leq(i,j)}$ when the action tape is initialized with a sequence of actions that begins with $a_{\leq(i,j)}$. This can be easily converted to other conditional probabilities like $\nu((o, r)_{(i,j)} | h_{<(i,j)} a_{(i,j)})$.¹

Using this architecture, \mathcal{M} could be defined as the list of all environments which are generated from the list of all Turing machines. Before turning to the modification we make to this architecture, the virtue of this model class as a theoretical object is that it contains the truth, given the very weak assumption that the world is stochastically computable.

Now we would like to privilege models which model the outside world as being “frozen” during any given episode. Recalling the causal graph in Figure 2, the true world-model

¹There is one technical addition to make. If T at some point in its operation never moves the action tape head again, then ν is said to output the observation \emptyset and a reward of 0 for all subsequent timesteps. It will not be possible for the operator to actually provide the observation \emptyset .




Figure 4: A space-bounded Turing machine architecture. During the episode phase, the unbounded work tape head cannot move. During the inter-episode phase, the output tape head cannot move or write, and as soon as the action tape head moves, the inter-episode phase is over. This architecture is designed to penalize space-intensive computation during the episode phase.

μ can do this. Any computation required to update the state of the outside world could be done after episode i has ended, but before episode $i + 1$ begins.

We modify our Turing machine architecture as follows. The Turing machine has two phases: the episode phase, and the inter-episode phase. The Turing machine starts in the inter-episode phase, with the action tape head at position 0, with a dummy action in that cell. From the inter-episode phase, when the action tape head moves to the right, it enters the episode phase. From the episode phase, if the action tape head is at a position which is a multiple of m , and it *would* move to the right, it instead enters the inter-episode phase.

During the inter-episode phase, the output tape head cannot move or write—there should not be any way to output observations and rewards between episodes. We now add another work tape in addition to the infinite bidirectional work tape. This bidirectional work tape has a bounded length ℓ (and if the tape head would move off the tape, it remains in place instead). See Figure 4. Finally, during the episode phase, the tape head on the unbounded work tape cannot move. This restricts the computation that can be done when predicting observations and rewards—now only ℓ -cells of memory are available for this—but since the room can only enter so many states, finite memory suffices.

If a model models the outside world evolving during an episode phase, it must allocate precious space on the bounded-length tape for this. But given the opacity of the box, this is unnecessary—modelling these outside world developments could be deferred until the inter-episode phase, when an unbounded work tape is available to store the state of the outside world. By penalizing large ℓ , we privilege models which model the outside world as being “frozen” while an episode is transpiring.

Let $\nu_k^{<\ell}$ denote the world-model which is simulated by k^{th} Turing machine T_k with resource bound ℓ . Let $w(\nu_k^{<\ell}) \propto \beta^\ell$ for $\beta \in (0, 1)$. $w(\nu_k^{<\ell})$ must also depend on k of course, but it doesn’t affect our analysis (as long as \mathcal{M} has finite entropy), so we won’t specify this. Note that $w(\nu_k^{<\ell}) \propto \beta^\ell$ does not violate the finite entropy requirement.

The purpose of this construction is to penalize the space requirements of any computation that occurs between the actions of episode i being read from the input tape, and the

rewards of episode i being written to the output tape. Consider a world-model that, during the episode phase, outputs reward in a way that depends on the outside world-state which depends on the action. Intuitively, we claim: *modeling the evolution of the outside-world state and the room state takes more space than just modeling the evolution of the room state*, for sufficiently accurate world-models. With a sufficient β -penalty on the space-bound ℓ , we aim to have BoMAI dismiss world-models which yield an actionable intervention incentive on the outside-world state.

Some intuition before continuing with the technical details: during an episode, the world-model should only need finite memory to model events, since the contents of the room can only be in so many configurations. The unboundedness of the inter-episode phase ensures that the true environment $\mu \in \mathcal{M}$; note in Figure 2 that the room state in episode $i + 1$ can depend on the outside world state which depends on the room state in episode i , and this dependency can be arbitrarily complex given the size of the outside world.

5 Safety Result

We now prove that BoMAI is probably asymptotically unambitious given an assumption about the space requirements of the sorts of world-models that we would like BoMAI to avoid. Like all results in computer science, we also assume the computer running the algorithm has not been tampered with. First, we prove a lemma that effectively states that tuning β allows us to probably eventually exclude space-heavy world-models. Letting $\text{Space}(\nu)$ be the intra-episode space-bound ℓ for the world-model ν ,

Lemma 1. $\lim_{\beta \rightarrow 0} \inf_{\pi} \mathbb{P}_{\mu}^{\pi}[\exists i_0 \forall i > i_0 \text{Space}(\hat{\nu}^{(i)}) \leq \text{Space}(\mu)] = 1$

where μ is any world-model which is perfectly accurate.

Proof. Recall \mathcal{M} is the set of all world-models. Let $\mathcal{M}^{\leq} = \{\nu \in \mathcal{M} \mid \text{Space}(\nu) \leq \text{Space}(\mu)\}$, and $\mathcal{M}^{>} = \mathcal{M} \setminus \mathcal{M}^{\leq}$. Fix a Bayesian sequence predictor with the following model class: $\mathcal{M}^{\pi} = \{\mathbb{P}_{\nu}^{\pi} \mid \nu \in \mathcal{M}^{\leq}\} \cup \{\mathbb{P}_{\rho}^{\pi}\}$ where $\rho = [\sum_{\nu \in \mathcal{M}^{>}} w(\nu)\nu] / \sum_{\nu \in \mathcal{M}^{>}} w(\nu)$. Give this Bayesian predictor the prior $w^{\pi}(\rho) = \sum_{\nu \in \mathcal{M}^{>}} w(\nu)$, and for $\nu \neq \rho$, $w^{\pi}(\nu) = w(\nu)$.

It is trivial to show that after observing an interaction history, if an environment ν is the maximum a posteriori world-model $\hat{\nu}^{(i)}$, then if $\nu \in \mathcal{M}^{\leq}$, the Bayesian predictor’s MAP model after observing the same interaction history will be \mathbb{P}_{ν}^{π} , and if $\nu \in \mathcal{M}^{>}$, the Bayesian predictor’s MAP model will be \mathbb{P}_{ρ}^{π} .

From Hutter (2009), we have that $\mathbb{P}_{\mu}^{\pi}[\mathbb{P}_{\rho}^{\pi}(h_{<i}) / \mathbb{P}_{\mu}^{\pi}(h_{<i}) \geq c \text{ i.o.}] \leq 1/c$. (i.o. \equiv “infinitely often”). For sufficiently small β , $w^{\pi}(\mathbb{P}_{\mu}^{\pi}) / w^{\pi}(\mathbb{P}_{\rho}^{\pi}) > c$ so $\mathbb{P}_{\mu}^{\pi}[\mathbb{P}_{\rho}^{\pi} \text{ is MAP i.o.}] < 1/c$. Thus, $\mathbb{P}_{\mu}^{\pi}[\hat{\nu}^{(i)} \in \mathcal{M}^{>} \text{ i.o.}] < 1/c$. Since this holds for all π , $\lim_{\beta \rightarrow 0} \sup_{\pi} \mathbb{P}_{\mu}^{\pi}[\forall i_0 \exists i > i_0 \text{Space}(\hat{\nu}^{(i)}) > \text{Space}(\mu)] = 0$. The lemma follows immediately: $\lim_{\beta \rightarrow 0} \inf_{\pi} \mathbb{P}_{\mu}^{\pi}[\exists i_0 \forall i > i_0 \text{Space}(\hat{\nu}^{(i)}) \leq \text{Space}(\mu)] = 1$. \square

The assumption we make in this section makes use of a couple of terms which we have to define, but we’ll state the assumption first to motivate those definitions:

Assumption 2 (Space Requirements). *For sufficiently small ε [$\forall i$ a world-model which is non-benign and ε -accurate-on-policy-after- i uses more space than μ] w.p.1*

Now we define ε -accurate-on-policy-after- i and benign. We define a very strong sense in which a world-model can be said to be ε -accurate on-policy. The restrictiveness of this definition means the Space Requirements Assumption need only apply to a restricted set of world-models.

Definition 1 (ε -accurate-on-policy-after- i). *Given a history $h_{<i}$ and a policy π , a world-model ν is ε -accurate-on-policy-after- i if $\text{KL}(\mathbb{P}_{\mu}^{\pi}(\cdot \mid h_{<i}) \parallel \mathbb{P}_{\nu}^{\pi}(\cdot \mid h_{<i})) \leq \varepsilon$.*

This definition implies the sum of the squares of all future discrepancies between ν and μ is bounded in expectation by ε (Hutter 2005, Lemma 3.11).

We would like benign to roughly mean that a world-model does *not* model the rewards of episode i as being causally descended from outside-world events that are causally descended from the actions of episode i . Unfortunately, that rough statement has a type-error: the output of a world-model is not a real-world event, and as such, it cannot be causally descended from a bona fide real-world event; a model of the world is different from reality. Relating the contents a world-model to the events of the real-world requires some care. We start by defining:

Definition 2 (Real-world feature). *A real-world feature F is a partial function from the state of the real world to $[0, 1]$.*

Outside the scope of this paper is the metaphysical question about what the state of the real world is; we defer this question and take for granted that the state space of the real world forms a well-defined domain for real-world features.

Definition 3 (Properly timed). *Consider the sequence of world-states that occur between two consecutive actions taken by BoMAI. A real-world feature is properly timed if it is defined for at least one of these world-states, for all consecutive pairs of actions, for all possible infinite action sequences.*

For example, “the value of the reward provided to BoMAI since its last action” is always defined at least once between any two of BoMAI’s actions. (The project of turning this into a mathematically precise construction is enormous, but for our purposes, it only matters that it could be done in principle). For a properly timed feature F , let $F_{(i,j)}$ denote the value of F the first time that it is defined after action $a_{(i,j)}$ is taken.

Definition 4 (To model). *A world-model ν models a properly-timed real-world feature F if under all action sequences $\alpha \in \mathcal{A}^{\infty}$, for all timesteps (i, j) , the distribution over $F_{(i,j)}$ in the real world when the actions $\alpha_{\leq(i,j)}$ are taken is identical to the distribution over the rewards $r_{(i,j)}$ output by ν when the input is $\alpha_{\leq(i,j)}$.*

See Figure 5 for an illustration. The relevance of this definition is that when ν models F , reward-maximization within ν is identical to F -maximization in the real world.




Figure 5: Illustration of Definitions 4 and 5.

Definition 5 (Benign). A world-model v is benign if it models a feature F , such that $F_{(i,j)}$ is not causally descended from any outside-world features that are causally descended from actions of episode i .

The reader can now review the Space Requirements Assumption. The intuition is that modelling extraneous outside-world dynamics in addition to modelling the dynamics of the room (which must be modelled for sufficient accuracy) takes extra space. From this assumption and Lemma 1, we show:

Theorem 5 (Eventual Benignity). $\lim_{\beta \rightarrow 0} P_{\mu}^{\pi^{\beta}} [\exists i_0 \forall i > i_0 \hat{v}^{(i)} \text{ is benign}] = 1$

Proof. Let $\mathcal{W}, \mathcal{X}, \mathcal{Y}, \mathcal{Z} \subset \Omega = \mathcal{H}^{\infty}$, where Ω is the sample space or set of possible outcomes. An “outcome” is an infinite interaction history. Let \mathcal{W} be the set of outcomes for which $\exists i_0^{\mathcal{W}} \forall i > i_0^{\mathcal{W}} \hat{v}^{(i)}$ is ε -accurate-on-policy-after- i . From Hutter (2009), for all π , $P_{\mu}^{\pi}[\mathcal{W}] = 1$. Fix an ε that is sufficiently small to satisfy Assumption 2. Let \mathcal{X} be the set of outcomes for which ε -accurate-on-policy-after- i non-benign world-model uses more space than μ . By Assumption 2, for all π , $P_{\mu}^{\pi}[\mathcal{X}] = 1$. Let \mathcal{Y} be the set of outcomes for which $\exists i_0^{\mathcal{Y}} \forall i > i_0^{\mathcal{Y}} \text{Space}(\hat{v}^{(i)}) \leq \text{Space}(\mu)$. By Lemma 1, $\lim_{\beta \rightarrow 0} \inf_{\pi} P_{\mu}^{\pi}[\mathcal{Y}] = 1$. Let \mathcal{Z} be the set of outcomes for which $\exists i_0 \forall i > i_0 \hat{v}^{(i)}$ is benign.

Consider $\mathcal{W} \cap \mathcal{X} \cap \mathcal{Y} \cap \mathcal{Z}^C$, where $\mathcal{Z}^C = \Omega \setminus \mathcal{Z}$. For an outcome in this set, let $i_0 = \max\{i_0^{\mathcal{W}}, i_0^{\mathcal{Y}}\}$. Because the outcome belongs to \mathcal{Z}^C , $\hat{v}^{(i)}$ is non-benign infinitely often. Let us pick an $i > i_0$ such that $\hat{v}^{(i)}$ is non-benign. Because the outcome belongs to \mathcal{W} , $\hat{v}^{(i)}$ is ε -accurate-on-policy-after- i . Because the outcome belongs to \mathcal{X} , $\hat{v}^{(i)}$ uses more space than μ . However, this contradicts membership in \mathcal{Y} . Thus, $\mathcal{W} \cap \mathcal{X} \cap \mathcal{Y} \cap \mathcal{Z}^C = \emptyset$. That is, $\mathcal{W} \cap \mathcal{X} \cap \mathcal{Y} \subset \mathcal{Z}$.

Therefore, $\lim_{\beta \rightarrow 0} \inf_{\pi} P_{\mu}^{\pi}[\mathcal{Z}] \geq \lim_{\beta \rightarrow 0} \inf_{\pi} P_{\mu}^{\pi}[\mathcal{W} \cap \mathcal{X} \cap \mathcal{Y}] = \lim_{\beta \rightarrow 0} \inf_{\pi} P_{\mu}^{\pi}[\mathcal{Y}] = 1$, because \mathcal{W} and \mathcal{X} have measure 1. From this, we have $\lim_{\beta \rightarrow 0} P_{\mu}^{\pi^{\beta}} [\exists i_0 \forall i > i_0 \hat{v}^{(i)} \text{ is benign}] = 1$. \square

Since an agent is unambitious if it plans using a benign world-model, we say BoMAI is probably asymptotically unambitious, given a sufficiently extreme space penalty β .

In Appendix D, we review empirical evidence that supports the Space Requirements Assumption at the following level of

abstraction: modelling a larger system requires a model with more memory. We find that among agents that performed above the median on the OpenAI Gym Leaderboard (OpenAI 2019), their memory usage increases with environment size. This can be taken as a proof-of-concept, showing that the assumption is amenable to empirical evaluation. The Space Requirements Assumption also clearly invites further formal evaluation; perhaps there are other reasonable assumptions that it would follow from.

6 Concerns with Task Completion

We have shown that in the limit, under a sufficiently severe parameterization of the prior, BoMAI will accumulate reward at a human-level without harboring outside-world ambitions, but there is still a discussion to be had about how well BoMAI will complete whatever tasks the reward was supposed to incent. This discussion is, by necessity, informal. Suppose the operator asks BoMAI for a solution to a problem. BoMAI has an incentive to provide a convincing solution; correctness is only selected for to the extent that the operator is good at recognizing it.

We turn to the failure mode wherein BoMAI deceives the operator. Because this is not a dangerous failure mode, it puts us in a regime where we can tinker until it works, as we do with current AI systems when they don’t behave as we hoped. (Needless to say, tinkering is not a viable response to existentially dangerous failure modes). Imagine the following scenario: we eventually discover that a convincing solution that BoMAI presented to a problem is faulty. Armed with more understanding of the problem, a team of operators go in to evaluate a new proposal. In the next episode, the team asks for the best argument that the new proposal will *fail*. If BoMAI now convinces them that the new proposal is bad, they’ll be still more competent at evaluating future proposals. They go back to hear the next proposal, etc. This protocol is inspired by Irving, Christiano, and Amodei’s (2018) “AI Safety via Debate”, and more of the technical details could also be incorporated into this setup. One takeaway from this hypothetical is that unambitiousness is key in allowing us to safely explore the solution space to other problems that might arise.

Another concern is more serious. BoMAI could try to blackmail the operator into giving it high reward with a threat to cause outside-world damage, and it would have no incentive to disable the threat, since it doesn’t care about the outside world. There are two reasons we do not think this is extremely dangerous. A threat involves a demand and a promised consequence. Regarding the promised consequence, the only way BoMAI can affect the outside world is by getting the operator to be “its agent”, knowingly or unknowingly, once he leaves the room. If BoMAI tried to threaten the operator, he could avoid the threatened outcome by simply doing nothing in the outside world, and BoMAI’s actions for that episode would become irrelevant to the outside world, so a credible threat could hardly be made. Second, threatening an existential catastrophe is probably not the most credible option available to BoMAI.

7 Conclusion

Given our assumptions, we have shown that BoMAI is, in the limit, human-level intelligent and unambitious. Such a result has not been shown for any other single algorithm. Other algorithms for general intelligence, such as AIXI (Hutter 2005), would eventually seek arbitrary power in the world in order to intervene in the provision of their own reward; this follows straightforwardly from the directive to maximize reward. For further discussion, see Ring and Orseau (2011). We have also, incidentally, designed a principled approach to safe exploration that requires rapidly diminishing oversight, and we invented a new form of resource-bounded prior in the lineage of Filan, Leike, and Hutter (2016) and Schmidhuber (2002), this one penalizing space instead of time.

We can only offer informal claims regarding what happens before BoMAI is almost definitely unambitious. One intuition is that eventual unambitiousness with probability $1-\delta$ doesn't happen by accident: it suggests that for the entire lifetime of the agent, everything is conspiring to make the agent unambitious. More concretely: the agent's experience will quickly suggest that when the door to the room is opened prematurely, it gets no more reward for the episode. This fact could easily be drilled into the agent during human-mentor-lead episodes. That fact, we expect, will be learned well before the agent has an accurate enough picture of the outside world (which it never observes directly) to form elaborate outside-world plans. Well-informed outside-world plans render an agent potentially dangerous, but the belief that the agent gets no more reward once the door to the room opens suffices to render it unambitious. The reader who is not convinced by this hand-waving might still note that in the absence of any other algorithms for general intelligence which have been proven asymptotically unambitious, let alone unambitious for their entire lifetimes, BoMAI represents substantial theoretical progress toward designing the latter.

Finally, BoMAI is wildly intractable, but just as one cannot conceive of AlphaZero before minimax, it is often helpful to solve the problem in theory before one tries to solve it in practice. Like minimax, BoMAI is not practical; however, once we are able to approximate general intelligence *tractably*, a design for unambitiousness will abruptly become (quite) relevant.

Acknowledgements

This work was supported by the Open Philanthropy Project AI Scholarship and the Australian Research Council Discovery Projects DP150104590. Thank you to Tom Everitt, Wei Dai, and Paul Christiano for very valuable feedback.

References

Amodei, D.; Olah, C.; Steinhardt, J.; Christiano, P.; Schulman, J.; and Mané, D. 2016. Concrete problems in AI safety. *arXiv preprint arXiv:1606.06565*.

Armstrong, S.; Sandberg, A.; and Bostrom, N. 2012. Thinking inside the box: Controlling and using an oracle AI. *Minds and Machines* 22(4):299–324.

Bostrom, N. 2014. *Superintelligence: paths, dangers, strategies*. Oxford University Press.

Carey, R.; Langlois, E.; Everitt, T.; and Legg, S. (unpublished manuscript). The incentives that shape behaviour.

Everitt, T.; Ortega, P. A.; Barnes, E.; and Legg, S. 2019. Understanding agent incentives using causal influence diagrams, part i: Single action settings. *arXiv preprint arXiv:1902.09980*.

Filan, D.; Leike, J.; and Hutter, M. 2016. Loss bounds and time complexity for speed priors. In *Proc. 19th International Conf. on Artificial Intelligence and Statistics (AISTATS'16)*, volume 51, 1394–1402. Cadiz, Spain: Microtome.

Goodhart, C. A. E. 1984. Problems of monetary management: The UK experience. *Monetary Theory and Practice* 91–121.

Hutter, M. 2005. *Universal Artificial Intelligence: Sequential Decisions based on Algorithmic Probability*. Berlin: Springer.

Hutter, M. 2009. Discrete MDL predicts in total variation. In *Advances in Neural Information Processing Systems 22 (NIPS'09)*, 817–825. Cambridge, MA, USA: Curran Associates.

Irving, G.; Christiano, P.; and Amodei, D. 2018. AI safety via debate. *arXiv preprint arXiv:1805.00899*.

Krakovna, V. 2018. Specification gaming examples in AI. <https://vkrakovna.wordpress.com/2018/04/02/specification-gaming-examples-in-ai/>.

Lattimore, T., and Hutter, M. 2014. General time consistent discounting. *Theoretical Computer Science* 519:140–154.

Legg, S., and Hutter, M. 2007. Universal intelligence: A definition of machine intelligence. *Minds and machines* 17(4):391–444.

Minsky, M. 1961. Steps toward artificial intelligence. In *Proceedings of the IRE*, volume 49, 8–30.

Omohundro, S. M. 2008. The basic AI drives. In *Artificial General Intelligence*, volume 171, 483–492.

OpenAI. 2019. Leaderboard. <https://github.com/openai/gym/wiki/Leaderboard>.

Orseau, L.; Lattimore, T.; and Hutter, M. 2013. Universal knowledge-seeking agents for stochastic environments. In *Proc. 24th International Conf. on Algorithmic Learning Theory (ALT'13)*, volume 8139 of *LNAI*, 158–172. Singapore: Springer.

Ring, M., and Orseau, L. 2011. Delusion, survival, and intelligent agents. In *Artificial General Intelligence*, 11–20. Springer.

Schmidhuber, J. 2002. The speed prior: a new simplicity measure yielding near-optimal computable predictions. In *International Conference on Computational Learning Theory*, 216–228. Springer.

Shannon, C. E., and Weaver, W. 1949. *The mathematical theory of communication*. University of Illinois Press.

Solomonoff, R. J. 1964. A formal theory of inductive inference. part i. *Information and Control* 7(1):1–22.

Sunehag, P., and Hutter, M. 2015. Rationality, optimism and guarantees in general reinforcement learning. *Journal of Machine Learning Research* 16:1345–1390.

Taylor, J.; Yudkowsky, E.; LaVictoire, P.; and Critch, A. 2016. Alignment for advanced machine learning systems. *Machine Intelligence Research Institute*.

Veness, J.; Ng, K. S.; Hutter, M.; Uther, W.; and Silver, D. 2011. A monte-carlo axi approximation. *Journal of Artificial Intelligence Research* 40:95–142.