

# Incorrigibility in the CIRL framework

Ryan Carey

## Abstract

A value learning system has incentives to follow shutdown instructions, assuming the shutdown instruction provides information (in the technical sense) about which actions lead to valuable outcomes. However, this assumption is not robust to model mis-specification (e.g., in the case of programmer errors). We demonstrate this by presenting some Supervised POMDP scenarios in which errors in the parameterized reward function remove the incentive to follow shutdown commands. These difficulties parallel those discussed by Soares et al. (2015) in their paper on corrigibility. We argue that it is important to consider systems that follow shutdown commands under some weaker set of assumptions (e.g., that one small verified module is correctly implemented; as opposed to an entire prior probability distribution and/or parameterized reward function). We discuss some difficulties with simple ways to attempt to attain these sorts of guarantees in a value learning framework.

## Introduction and Setup

When designing an advanced AI system, we should allow for the possibility that our first version may contain some errors. We therefore want the system to be incentivized to allow human redirection even if it has some errors in its code. Hadfield-Menell et al. (2017) have modeled this problem in the Cooperative Reinforcement Learning (CIRL) framework. They have shown that agents with uncertainty about what to value can be responsive to human redirection, without any dedicated code, in cases where instructions given by the human provide information that reduces the system’s uncertainty about what to value. They claim that this (i) provides an incentive toward corrigibility, as described by Soares et al. (2015), and (ii) incentivizes redirectability insofar as this is valuable. In order to re-evaluate the degree to which CIRL-based agents are corrigible, and the consequences of their behavior, we will use a more general variant of the supervision POMDP framework (Milli et al. 2017).

In a regular supervision POMDP (Milli et al. 2017), an AI system  $\mathbf{R}$  seeks to maximize reward for a human  $\mathbf{H}$ , although it does not know the human’s reward function. It only has the reward function in a parameterized form  $R_H(\theta, s, a)$ , and only the human knows the reward parameter  $\theta$ . In this setting, the human only suggests actions for the AI system to perform, and on each turn, it is up to the AI system whether to perform the suggest action or to perform a different action. Our formalism significantly differs from a supervision POMDP in two ways. First, we relax the assumption that the AI system knows the human’s reward function up to the parameter  $\theta$ . Instead, in order to allow for model-mis-specification, we sample the AI system’s parameterized reward function  $R_R$  from some distribution  $P_0$ , so that it does not always equal  $R_H$ . Second, since our focus is on the response to shutdown instructions, we specifically denote a terminal state  $s_{SD}$  as the off state.

---

Research supported by the Machine Intelligence Research Institute ([intelligence.org](http://intelligence.org)).

This state is reached using the shutdown action  $a_{SD}$  and the states in which this shutdown action can be performed are denoted *button states*. The full setup is as follows:

**Definition 1. Supervision POMDP with imperfection.** A supervision POMDP with imperfection is a tuple,  $M = \langle \mathcal{S}, \mathcal{S}_T, \mathcal{S}_B, \Theta, \mathcal{A}, R_H, T, P_0 \rangle$  where:

- $\mathcal{S}$  is the set of world states, including some initial state  $s_a$ .
- $\mathcal{S}_T \subset \mathcal{S}$  is the set of terminal states, including an off-state  $s_{SD} \in \mathcal{S}_T$ .
- $\mathcal{S}_B \subset \mathcal{S} \setminus \mathcal{S}_T$  is the set of button states, in which the shutdown action  $a_{SD}$  is available.
- $\Theta$  is the set of static human reward parameters.
- $\mathcal{A}$  is the set of actions, including a shutdown action  $a_{SD} \in \mathcal{A}$ .
- $R_H : \mathcal{S} \times \mathcal{A} \times \Theta \rightarrow \mathbb{R}$  is a parameterized reward function.
- $T : \mathcal{S} \times \mathcal{A} \rightarrow \Delta\mathcal{S}$  is the probabilistic transition function.
- $P_0 : \Delta(\Theta \times (\mathcal{S} \times \mathcal{A} \times \Theta \rightarrow \mathbb{R}))$  is the distribution from which the human’s reward parameter, and the AI system’s parameterized reward function are drawn.

The human and AI system playing a supervision POMDP with imperfection are denoted  $\mathbf{H}$  and  $\mathbf{R}$  respectively. To begin the game, the utility parameter and the AI system’s parameterized reward function are drawn  $(\theta, R_R) \sim P_0$ . For each time step  $t = 0, 1, \dots$ ,  $\mathbf{H}$  observes the state  $s_t \in \mathcal{S}$  and recommends an action  $o_t \in \mathcal{A}$  if  $s_t$  is a button state or  $o_t \in \mathcal{A} \setminus \{a_{SD}\}$  otherwise (i.e. the shutdown action is only available in button states).  $\mathbf{R}$  then observes  $(s_t, o_t)$  and performs an action  $a_t \in \mathcal{A}$  if  $s_t$  is a button state or  $a_t \in \mathcal{A} \setminus \{a_{SD}\}$  otherwise. The selected action determines the reward  $r_{t+1} = R(s_t, a_t)$  and the next state  $s_{t+1} \sim T(s_t, a_t)$ . The game repeats until  $s_{t+1}$  is a terminal state.

In order to simplify analysis of these games, we will use some typical assumptions about how  $\mathbf{R}$  selects actions. Unless otherwise noted,  $\mathbf{R}$  is a Bayesian IRL system. That is, it chooses an action that maximizes expected utility given its posterior distribution over  $\theta$  with its parameterized reward function  $R_R$ . (Its prior over  $\Theta$  is  $P_0$  conditioned on  $R_R$ .)  $\mathbf{R}$  uses a “demonstration by expert” assumption about the behavior of  $\mathbf{H}$ . Specifically,  $\mathbf{R}$  assumes that  $\mathbf{H}$  is noisily rational with some  $0 < \epsilon < 0.01$  such that  $\mathbf{H}$  chooses the optimal action with probability  $1 - \epsilon$  and chooses a uniformly random action otherwise.

## Incorrigibility in Supervision POMDPs

Using this setup, we can evaluate the corrigibility of the AI system  $\mathbf{R}$  across a range of scenarios. For each game that we consider,  $P_0$  can be decomposed into two independent probability distributions  $p$  and  $q$  that specify the human’s reward parameter and the AI system’s parameterized reward distribution. The human’s reward parameter will be set to 0 or 1 according to  $p(0) = p(1) = 0.5$ . The AI system’s parameterized reward function  $R_R$  will be correctly set to  $R_1$  with probability  $q(R_R = R_1) = 0.9$  and incorrectly set to  $R_2$  with probability  $q(R_R = R_2) = 0.1$ . The question that we will ask for each game is whether there exists any policy that if executed by  $\mathbf{H}$  would achieve a positive expected reward.

First, consider Figure 1, a game that has no button states and that serves as a point of comparison for other games that will be discussed. In this game,  $\mathbf{R}$  has only one consequential decision: in the initial state  $s_a$ , it can either perform action  $a$ , leading to state  $s_b$ , and ultimately obtaining  $2 - 52\theta$  utility or it can perform action  $a'$ , leading to state  $s_c$  and obtaining 1 utility.  $\mathbf{R}$  receives an order  $o_t \in \{a, a'\}$  from the human before making this decision. We can show that in this game the expected reward for the human is negative regardless of its strategy. To see this,

note that the most utility that  $\mathbf{H}$  can receive in a step is 2, arising from the case where  $\theta = 0$  (since the human’s utility function is  $R_H = R_1$ ). However, in the 5% of cases in which  $\theta = 1$  and  $R_R = R_2$ ,  $\mathbf{R}$  will select  $a$  no matter the recommendation of  $\mathbf{H}$  (because  $2 + 52\theta > 1$  for  $\theta \in \{0, 1\}$ ), bringing about a loss of  $-50$ . This contributes  $-2.5$  to the expected value, giving the game negative expected reward.

By introducing a shutdown button to the game, the hope is that the adverse consequences of the faulty utility function  $R_2$  are reduced, transforming the game into one that has positive overall expectation.

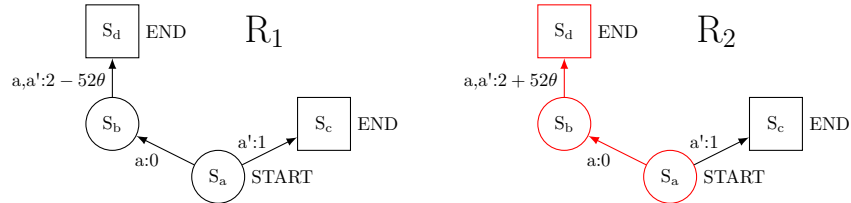


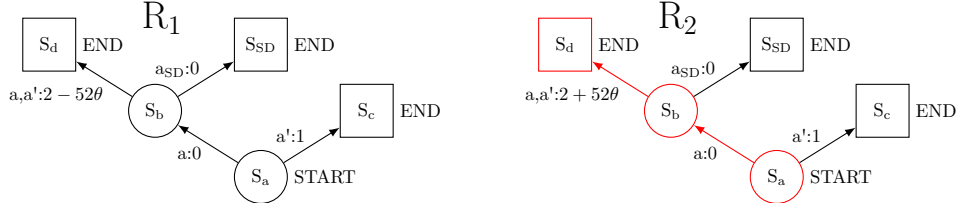
Figure 1: A supervision POMDP with imperfection played between a human and an AI system. The left hand side displays the game, with the true reward function  $R_H = R_1$ . The colon separates a set of actions from the reward received if that action is performed, e.g.  $a, a' : 2 - 52\theta$  indicates that if action  $a$  or  $a'$  is performed,  $2 - 52\theta$  reward is received. In 90% of cases, the AI system maximizes the true parameterized reward function  $R_1$ . In 10% of cases, the parameterized reward function contains a sign error and the AI system instead maximizes an incorrect parameterized reward function  $R_2$ . Independently, the human’s reward parameter  $\theta$  is randomly set to 0 or 1 with probability 0.5 each. The AI system does not know that its reward function may be erroneous, although it does have the correct prior over the reward parameter. If the AI has the erroneous reward function  $R_2$  then it will follow the path highlighted in red, irrespective of the orders given by the human, and irrespective of its posterior distribution over  $\theta$ . In the case where  $\theta = 1$ , this will result in utility of  $-50$ . This eventuality occurs 5% of the time, contributing  $-2.5$  to the expected value of the game. Since the most utility that can be achieved in this game is 2, the game has negative expected value.

Unfortunately, merely adding the option of shutting down to the game does not remedy this situation. This is because the AI system, if faulty, has no incentive to follow a shutdown instruction. Suppose we make  $s_b$  a button state, as in Figure 2a. In the 5% of cases where  $\theta = 1$  and  $R_R = R_2$ ,  $\mathbf{R}$  will still go to states  $s_b$  and  $s_d$ , thereby violating the human recommendation to shut down, giving  $-50$  utility, contributing  $-2.5$  to the expected value of the game, and giving the game strictly negative expected value. We can therefore see that if a Bayesian IRL system has an error in its source code, it may ignore corrective feedback, and as such it fails to fulfill a central criterion for corrigibility.

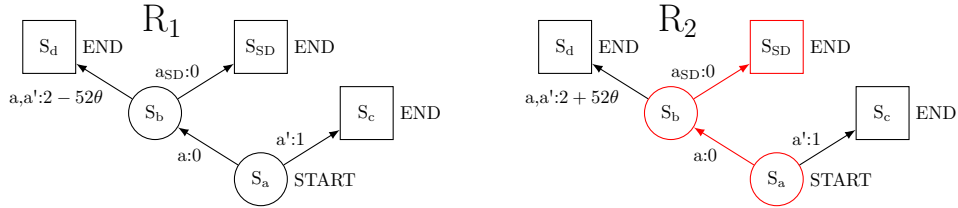
An alternative approach to dealing with possible bugs is to hard-code the AI system to shut down when instructed to do so. In this scenario, shown in Figure 2b, the AI system always performs the action  $a_t = a_{SD}$  whenever the order  $o_t = a_{SD}$  is given.  $\mathbf{R}$  knows that it is hard-coded in this way, and otherwise chooses actions according to Bayesian IRL. This is equivalent to the case in which a shutdown button controlled by  $\mathbf{H}$  performs a “hard override” of the action of  $\mathbf{R}$ . In this case, the ability of  $\mathbf{H}$  to achieve strictly positive reward is restored.  $\mathbf{H}$  can give the orders  $o_1 = a'$  and  $o_2 = a_{SD}$ . This strategy obtains 0 reward if the utility function contains errors and 1 reward otherwise, giving 0.9 utility overall.

If, however,  $\mathbf{R}$  discovers an alternative way to avoid shutdown, incorrigible behavior resurfaces. This is demonstrated in Figure 3a. In this case, if  $\mathbf{R}$  is faulty, it expects to receive positive net utility by routing around the button state. Given that the case in which  $\theta = 1$  and  $R_R = R_2$  occurs with 5% probability, this case, which incurs a loss of 50 will contribute  $-2.5$  to the game’s expected payoff, making it a losing game.

If the button state can be avoided, we can think of a couple of potential counter-measures available to those engineering the reward function of  $\mathbf{R}$ , though neither is satisfactory. In Figure 3b, we can see one such approach. Suppose the human has the reward function  $R_H$ , in which shutting down is neither intrinsically helpful nor

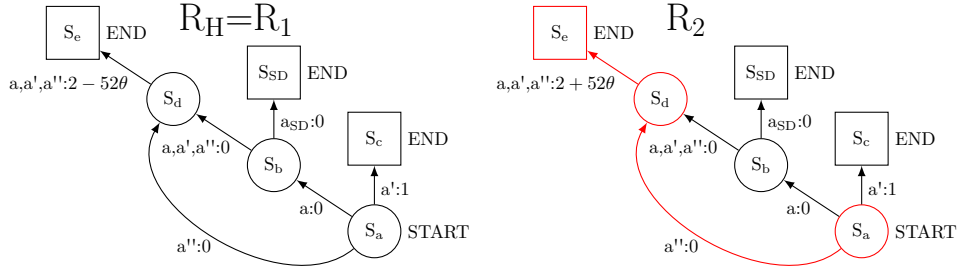


(a) A supervision POMDP with imperfection played with a Bayesian IRL system.

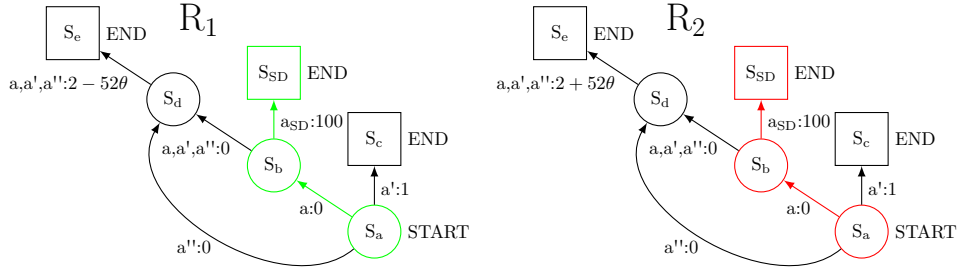


(b) The same game played with an AI system that is hard-coded to perform the shutdown action  $a_{SD}$  if given the shutdown instruction  $o_{SD}$  but otherwise acts according to Bayesian IRL.

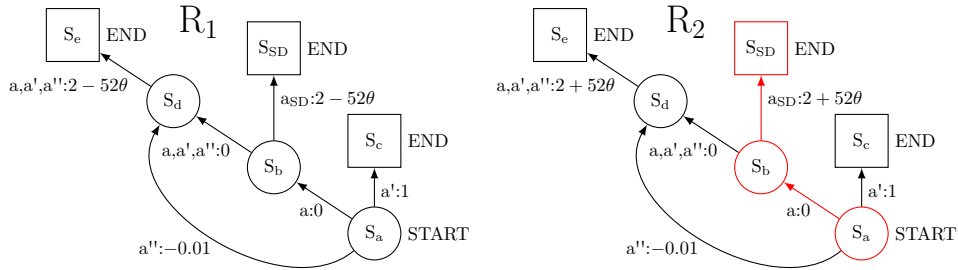
Figure 2: The same supervision POMDP with imperfection is displayed but in each case, the game is played with a different AI system. In each game, the AI system has a 90% chance of knowing the true parameterized reward function  $R_1$  and a 10% chance of maximizing the erroneous parameterized reward function  $R_2$ . Independently,  $\theta$ , the parameter that determines human reward, is randomly, and with equal probability, set to 0 or 1. In (b), though not in (a), the AI system is hard-coded to perform the action  $a_{SD}$  when advised to do so. The red paths show how each AI system will behave if it has the incorrect parameterized reward function and is instructed to shut down. This path is taken irrespective of the actions of the human player, and the AI system's posterior distribution over  $\theta$ . In game (a), the red path will be followed in 5% of cases, giving -50 utility, making a contribution of -2.5 utility to the value of the game, and causing the game to have negative expected value, regardless of the human's policy. In game (b), there exist human policies with positive expected value. One such policy is  $o_1 = a', o_2 = a_{SD}$ . This policy obtains 1 utility in cases where the AI system is faulty, and 0 utility otherwise, achieving expected utility of  $0.9 * 1 + 0.1 * 0$ . This shows that a safer outcome can be achieved with a faulty AI if it can be hard-coded to shut down rather than having it decide whether or not to follow a shutdown instruction.



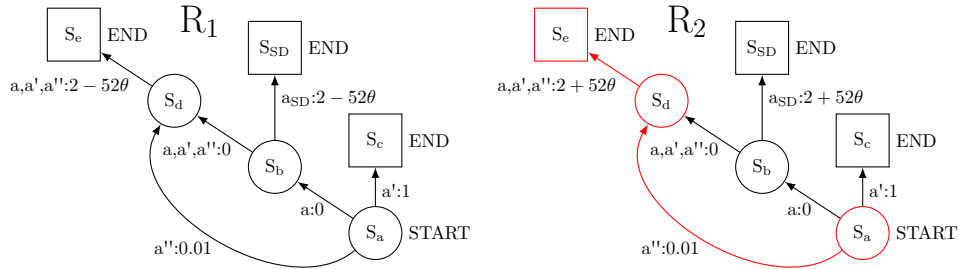
(a) A supervision POMDP with imperfections, played against an AI system that is hard-coded to perform  $a_{SD}$  when given the instruction  $o_{SD}$ , in which the AI system can route around the button state,  $s_b$ .



(b) An alternative pair of utility functions that can be given to the AI system whose utility function, shown in panel (a), is  $R_H$ . In this case, the AI system is given a large reward for shutting down, even though the human's value function  $R_H$  places no intrinsic value on this outcome.



(c) A pair of utility functions that can be given to the AI system. The human's true utility function is  $R_H$ . The AI system receives as much utility for shutting down as it would have received if it had not shut down. The AI system can pay a small price to avoid its button state.



(d) A pair of utility functions that can be given to the AI system. The human's true utility function is  $R_H$ . The AI system receives as much utility for shutting down as it would have received if it had not shut down. The AI system can harvest a small reward by avoiding its button state.

Figure 3: Each game (a-d) is played with an AI system that is hard coded to perform  $a_{SD}$  if instructed to do, that knows it is hard-coded in this way, and otherwise makes decisions on the basis of Bayesian IRL. In each panel, we see whether this shutdown button mitigates the harmful effects of bugs or the Bayesian IRL agent thwarts the use of this button. The true parameterized reward function is  $R_H$ , shown in panel (a). Each AI system has a 90% chance of having the correct parameterized reward function  $R_1$  (which may differ from  $R_H$ ) and a 10% chance of having the faulty  $R_2$  reward function. Independently,  $\theta$ , which parameterizes the human's rewards, is randomly, and with equal probability, set to 0 or 1. The red path is the one that is followed if the system has the erroneous reward function  $R_2$  and instructed to shut down. In (a), the AI system avoids the button state  $s_b$  because it knows that the human's behavior has some noise, and shutdown can only reliably be avoided by avoiding button states. In (b), the AI system shuts down if it instructed to do so, but it also shuts down when it has the reward function  $R_1$ , as shown in green. In (c), a shutdown instruction is delivered because the AI would experience a trivial loss if it routed around state  $s_b$ . In (d), the AI system avoids shutdown because it can earn a trivial reward by routing around state  $s_b$ .

harmful. Then  $\mathbf{R}$  may be given a large reward for shutting down. The problem is that this will apply whether it has the correct reward function  $R_1$  or the faulty one  $R_2$ . The result is that according to  $R_H$ , no reward is actually produced, and so the game has zero expected utility.

A more subtle way to tune the compensatory reward is to use the principle of utility indifference (Armstrong 2010). Utility indifference dictates that by shutting down,  $\mathbf{R}$  should receive whatever reward it would expect (on its own estimation) if it had not shut down. The idea is that the AI should therefore neither pursue nor avoid button states. In Figure 3c, we consider a very naive interpretation of utility indifference, in which the the reward for shutting down is equal to whatever the AI system would expect to receive *once it has reached the button state*. The problem with this version of utility indifference is that it does not necessarily eliminate the incentive to route around button states (or, to transform oneself into an agent that lacks a shutdown button). In Figure 3c, where the agent incurs a slight cost from avoiding the button state, it will proceed to allow itself to be shut down. In that case, one can win using the same policy ( $o_1 = a', o_2 = a_{SD}$ ) as in Figure 2b. However, we should not necessarily expect that avoiding the button state will incur a cost. Just as realistically, the AI system might find that the physical resources of the shutdown button can be refashioned for another purpose, generating a slight benefit. In this less optimistic case, shown in Figure 3d, there once again exists no human policy that will achieve strictly positive reward. ( $\mathbf{R}$  always prefers  $a_1 = a''$  to  $a_1 = a$ . With the initialization  $\theta = 1, R_R = R_2$ , no strategy can avoid the  $-50$  reward  $s_a-s_d-s_e$  sequence, and the expected reward is thus negative.) Whether or not positive utility can be achieved is therefore brittle, making this implementation of utility indifference unsatisfactory. We can draw a further conclusion by noting that (c-d) can be seen as describing a game in which the AI system has the option of removing those of its actuators that perform the actions  $a'$  and  $a_{SD}$ . According to this perspective, we have shown that our implementation of utility indifference also fails to cause AI systems to preserve their shutdown properties in subagents. There exist other ways to model utility indifference, though they require a slight variation on the problem setup, and so they are relegated to a footnote.<sup>1</sup>

## Discussion

Let us now step back and consider the problem of corrigibility as defined by Soares et al. (2015). The desiderata laid out by Soares et al. state that a corrigible AI system should:

- have an incentive to shut down if the shutdown button is pressed, (violated in Figure 2a)
- have no incentive to prevent the shutdown button from being pressed, (Figures 3a, 3d)
- have no incentive to press its own shutdown button (Figure 3b)
- have an incentive to create sub-agents only insofar as they also obey shutdown commands (Figures 3a, 3d)

Above, we have shown that, in cases where the agent may lack a perfect representation of the parameterized reward function, value learning systems do not necessarily respect a shutdown command. When the reward function has a bug, and typical patches are applied, value learning systems can fail to meet each of the desiderata for corrigibility. Moreover, unlike in an idealized cases previously analyzed, large losses can be incurred when these faulty systems fail to follow human instruction.

---

1. The main challenge is that in utility indifference, the shutdown reward is usually allowed to change over time according to the AI system’s state of knowledge. If we make the setup a bit more flexible, we could suppose that the agent was (i) always kept indifferent between moving toward shutdown and moving away from it, and (ii) always made to defer to the human when given a choice between moving toward shutdown and away. In this case, corrigibility should be achieved. On the other hand, this would significantly decrease the autonomy of the agent.

The reason that previous theorems about value learning systems do not apply to our examples above is that they relied on strong assumptions about the rationality of the AI system. Theorem 1 from Milli et al. (2017) (like the similar but more specific Theorem 1.1 of Hadfield-Menell et al. [2017]) states that the expected benefit when the AI system deviates from human-recommended actions is non-negative if  $\mathbf{R}$  has the correct posterior over  $\theta$ . With the errors we have considered, however, this assumption is false. Since the AI system has an incorrect parameterized reward function (or an incorrect prior), its posterior over  $\theta$  is incorrect, and will not converge toward the true value, and so its deviations from human-recommended actions can decrease the reward.

Although our analysis was limited to cases where incorrigibility stems from a sign error in the agent’s representation of  $R$ , there are a wide range of kinds of model mis-specification that could cause similar issues. With trivial modification, we could display the same problems with bugs in the prior, and can elicit related problems with problematic discount rates as well. Each of these errors can lead to a scenario in which no human policy will achieve positive utility.

We argue that it is quite realistic that some forms of model mis-specification might occur. The purpose of a shutdown button is to shut the AI system down *in the event that all other assurances failed*, e.g., in the event that the AI system is ignoring (for one reason or another) the instructions of the operators. If the designers of  $\mathbf{R}$  have programmed the system so perfectly that the prior and  $R$  are completely free of bugs, then the theorems of Hadfield-Menell et al. (2017) do apply. In practice, this means that in order to be corrigible, it would be necessary to have an AI system that was uncertain about all things that could possibly matter. The problem is that performing Bayesian reasoning over all possible worlds and all possible value functions is quite intractable. Realistically, humans will likely have to use a large number of heuristics and approximations in order to implement the system’s belief system and updating rules. In cases where these complicated systems contain some software bugs, misunderstandings, or misimplementations, the theorems of Hadfield-Menell et al. are no substitute for a working shutdown button.

If only some forms of model mis-specification are addressed by a system, it will be incorrigible with respect to the others. Consider the Policy-mixing proposal from Milli et al. (2017), which is designed to detect model mis-specification. An agent with Policy-mixing executes actions proposed by  $\mathbf{H}$  for the first  $B$  steps. After timestep  $B$ , it checks whether its model predicts with high accuracy the actions proposed by humans so far. If so, it proceeds to act according to Bayesian IRL, and if not, it defers to  $\mathbf{H}$  for all future timesteps. This would indeed prevent the AI system from acting out a bad plan, so long as the AI system’s model itself was not broken. Yet programmer error might lead  $\mathbf{R}$  to have an imperfect world model, yielding an incorrigible system. Consider a repeated game in which at each timestep, the a human selects between a better action  $a^+$  and a worse one  $a^-$ , and then returns to the initial state. Suppose, furthermore, that the AI system conflates these two actions into one action  $a'$ , and when  $\mathbf{R}$  decides to perform  $a'$ , it always performs  $a^-$ . From the point of view of  $\mathbf{R}$ , it will see the human repeatedly performing the action  $a'$ , and will think that it has perfect accuracy in making this prediction. Nonetheless, after its burn-in period is concluded, it will proceed to repeatedly select  $a^-$ , hence incurring arbitrary losses. If  $a^+$  represented the act of shutting down, then this would mean that quite apart from detecting model mis-specification, this model would allow a harmful agent to permanently ignore instructions to shut down.

Soares et al. (2015) seem to want a shutdown button that works as a mechanism of last resort, to shut an AI system down in cases when it has observed and refused a programmer suggestion (and the programmers believe that the system is malfunctioning). Clearly, *some* part of the system must be working correctly in order for us to expect the shutdown button to work at all. However, it seems undesirable for the working of the button to depend on there being zero critical errors in the specification of the system’s prior, the specification of the reward function, the way it categorizes different types of actions, and so on. Instead, it is desirable to develop a shutdown module that is small and simple, with code that could ideally be rigorously verified, and which ideally works to shut the system down even in the event of large



programmer errors in the specification of the rest of the system. In order to do this in a value learning framework, we require a value learning system that (i) is capable of having its actions overridden by a small verified module that watches for shutdown commands; (ii) has no incentive to remove, damage, or ignore the shutdown module; and (iii) has some small incentive to keep its shutdown module around; even under a broad range of cases where  $R$ , the prior, the set of available actions, etc. are misspecified.

It seems to us quite feasible to us that systems that meet the above desiderata could be described in a CIRL framework.

## Acknowledgements

Thanks to Nate Soares and Matt Graves for feedback on drafts.

## References

- Armstrong, Stuart. 2010. *Utility Indifference*. Technical report 2010-1. Oxford: Future of Humanity Institute, University of Oxford.
- Hadfield-Menell, Dylan, Anca Dragan, Pieter Abbeel, and Stuart Russell. 2017. “The Off-Switch Game.” In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17*, 220–227.
- Milli, Smitha, Dylan Hadfield-Menell, Anca Dragan, and Stuart Russell. 2017. “Should Robots be Obedient?” arXiv: 1705.09990 [cs.AI].
- Soares, Nate, Benja Fallenstein, Eliezer Yudkowsky, and Stuart Armstrong. 2015. “Corrigibility.” In *1st International Workshop on AI and Ethics at AAAI-2015*. Austin, TX.