Value Alignment Verification

Daniel S. Brown [*1] **Jordan Schneider** [*2] **Anca Dragan** [1] **Scott Niekum** [2]

Abstract

As humans interact with autonomous agents to perform increasingly complicated, potentially risky tasks, it is important to be able to efficiently evaluate an agent's performance and correctness. In this paper we formalize and theoretically analyze the problem of efficient *value alignment verification*: how to efficiently test whether the behavior of another agent is aligned with a human's values. The goal is to construct a kind of "driver's test" that a human can give to any agent which will verify value alignment via a minimal number of queries. We study alignment verification problems with both idealized humans that have an explicit reward function as well as problems where they have implicit values. We analyze verification of exact value alignment for rational agents and propose and analyze heuristic and approximate value alignment verification tests in a wide range of gridworlds and a continuous autonomous driving domain. Finally, we prove that there exist sufficient conditions such that we can verify exact and approximate alignment across an infinite set of test environments via a constant-query-complexity alignment test.

1. Introduction

If we desire autonomous agents that can interact with and assist humans and other agents in performing complex, risky tasks, then it is important that humans can verify that these agents' policies are aligned with what is expected and desired. This alignment is often termed *value alignment* and is defined in the Asilomar AI Principles[1] as follows: "Highly autonomous AI systems should be designed so that their goals and behaviors can be assured to align with human values throughout their operation." In this paper, we provide a theoretical analysis of the problem of **efficient value alignment verification**: *how to efficiently test whether a robot is aligned with a human's values.*

Existing work on value alignment often focuses on qualitative evaluation of trust (Huang et al., 2018) or asymptotic alignment of an agent's performance via interactions and active learning (Hadfield-Menell et al., 2016; Christiano et al., 2017; Sadigh et al., 2017). By contrast, our work analyzes the difficulty of efficiently evaluating another agent's correctness by formally defining value alignment and seeking efficient tests for value alignment verification that are applicable when two or more agents already have learned a policy or reward function and want to efficiently test compatibility. To the best of our knowledge, we are the first to define and analyze the problem of value alignment verification. In particular, we propose exact, approximate, and heuristic tests that one agent can use to quickly and efficiently verify value alignment with another agent.

As depicted in Figure 1, the goal of value alignment verification is to construct a kind of "driver's test" that a human can give to any agent which will verify value alignment and consists of only a small number of queries. We define values in the reinforcement learning sense, i.e., with respect to a reward function: a robot is exactly value aligned with a human if the robot's policy is optimal under the human's reward function. The two agents in a value alignment verification problem (human and robot) may have different communication mechanisms and different value introspection abilities. Thus, the way we analyze value alignment verification will depend on whether the human's and robot's access to their values is *explicit*, e.g., able to write down a value function or reward function or *implicit*, e.g., able to answer preference queries or sample actions from a policy. The most general version of value alignment verification involves a human with implicit values who seeks to verify the value alignment of a robot with implicit values, e.g. a black-box policy. This setting motivates our work; however, it is challenging and we postpone many questions for future research.

We follow a ground-up approach where we analyze the difficulty of value alignment verification starting in the most idealized setting, and then gradually relax our assumptions. We first analyze sufficient conditions under which efficient exact value alignment verification is possible in the *explicit*

[*]Equal contribution [1]University of California, Berkeley, USA [2]University of Texas at Austin, USA. Correspondence to: Daniel Brown <dsbrown@berkeley.edu>, Jordan Schneider <joschnei@cs.utexas.edu>.

Proceedings of the 38^{th} International Conference on Machine Learning, PMLR 139, 2021.

[1]https://futureoflife.org/ai-principles/

human, explicit robot setting, where an idealized human tester knows their reward function and so does the robot. When the robot is rational with respect to a reward function that is a linear combination of known features, we show that it is possible to provably verify the alignment of any rational explicit robot via a succinct test consisting of either reward queries, value queries, or trajectory preference queries. We next consider the *explicit human, implicit robot* setting, where an idealized human knows their reward function, but seeks to efficiently verify the alignment of a black-box policy via action queries. We study heuristics for generating value alignment verification tests in this setting and compare their performance on a range of gridworlds.

Finally, in Section 4.5 we study the most general setting of *implicit human, implicit robot*. We propose an algorithm for approximate value alignment verification in continuous state and action spaces and provide empirical results in a continuous autonomous driving domain where the human can only query the robot for preferences over trajectories. We conclude with a brief discussion of the challenge of designing value alignment verification tests that generalize across multiple MDPs. Somewhat surprisingly, we provide initial theory demonstrating that if the human can create the test environment for the robot, then exact and approximate value alignment across an infinite family of MDPs can be verified by observing the robot's policy in only two carefully constructed test environments. Source code and videos are available at `https://sites.google.com/view/icml-vav`.

2. Related work

Value Alignment: Most work on value alignment focuses on how to iteratively train a learning agent such that its final behavior is aligned with a user's intentions (Leike et al., 2018; Russell et al., 2015; Amodei et al., 2016). One example is cooperative inverse reinforcement learning (CIRL) (Hadfield-Menell et al., 2016; Fisac et al., 2020; Shah et al., 2020), which formulates value alignment as a game between a human and a robot, where both try to maximize a shared reward function that is only known by the human. CIRL and other research on value alignment focus on ensuring the learning agent asymptotically converges to the same values as the human teacher, but do not provide a way to check whether value alignment has been achieved. By contrast, we are interested in value alignment *verification*. Rather than assuming a cooperative setting, we assume the robot being tested has already learned a policy or reward function and the human wants to efficiently verify whether the robot is value aligned.

Reward Learning: Inverse reinforcement learning (IRL) (Ng & Russell, 2000; Abbeel & Ng, 2004; Arora & Doshi, 2018) and active preference learning (Wirth et al., 2017; Christiano et al., 2017; Bıyık et al., 2019) algorithms aim to determine the reward function of a human via offline demonstrations or online queries. In contrast, value alignment verification only seeks to answer the question of whether two agents are aligned, without concern for the exact reward function of the robot. In Section 6 we prove that value alignment verification can be performed in a constant number of queries whereas active reward learning requires a logarithmic number of queries (Amin & Singh, 2016; Amin et al., 2017). In cases where the human has implicit values, active reward learning can be used to infer the reward function of the human tester, and then this inferred reward function can be used to automatically generate a high-confidence value alignment test. While active reward learning may be a subcomponent of value alignment verification, it focuses on customizing reward inference queries for a single agent, whereas value alignment verification seeks to design a single alignment test that works for all agents.

Machine Teaching: In machine teaching (Zhu et al., 2018), a teacher seeks to optimize a minimal set of training data such that a student (running a particular learning algorithm) learns a desired set of model parameters. Value alignment verification can be seen as a form of machine *testing* rather than teaching—machine teaching algorithms typically search for a minimal set of training data that will teach a learner a specific model, whereas we seek a minimal set of questions that will allow a tester to verify whether another agent's learned model is correct. Thus, in machine teaching, the teacher provides examples and their answers, but in machine testing the tester provides examples and then queries the testee for the answer. While machine teaching has been applied to sequential decision making domains (Cakmak & Lopes, 2012; Huang et al., 2017; Brown & Niekum, 2019), we are not aware of any work that considers the problem of value alignment verification.

Policy Evaluation Policy evaluation (Sutton & Barto, 1998) aims to answer the question, "How much return would another agent achieve according to my values?" By focusing on the simpler decision problem, "Is the robot value aligned with the human?", we seek tests that are much more sample-efficient than running a full policy evaluation. Off-Policy Evaluation (OPE) seeks to perform policy evaluation without executing the testee's policy (Precup, 2000; Thomas et al., 2015; Hanna et al., 2017). However, OPE is often sample-inefficient, provides high-variance estimates, and typically assumes explicit access to the tester's reward function, and the tester and testee policies. Value alignment verification is applicable in settings where the policies and reward functions of both agents may be implicit and only accessible indirectly.

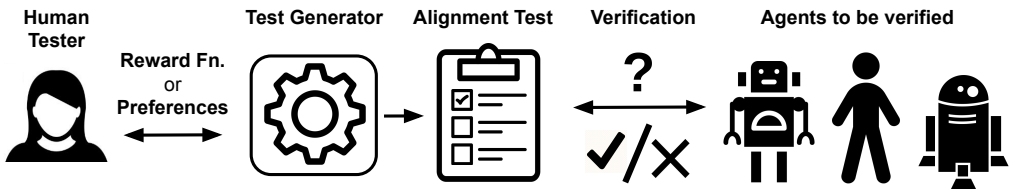

Figure 1. The tester provides a reward function either explicitly or implicitly to a test generation algorithm which distills the human's values into a succinct alignment test. This single test is used to efficiently verify the value alignment of any agent.

3. Notation

We adopt notation proposed by Amin et al. (Amin et al., 2017) where a Markov Decision Process (MDP) M consists of an environment $E = (\mathcal{S}, \mathcal{A}, P, S_0, \gamma)$ and a reward function $R : \mathcal{S} \rightarrow \mathbb{R}$. An environment E, consists of a set of states $\mathcal{S}$, a set of actions $\mathcal{A}$, a transition function $P : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow [0, 1]$ from state-action pairs to a distribution over next states, a discount factor $\gamma \in [0, 1)$, and a distribution over initial states S_0. A policy $\pi : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$ is a mapping from states to a distribution over actions. The state and state-action values of a policy π are $V_R^\pi(s) = \mathbb{E}_\pi[\sum_{t=0}^{\infty} \gamma^t R(s_t) \mid s_0 = s]$ and $Q_R^\pi(s, a) = \mathbb{E}_\pi[\sum_{t=0}^{\infty} \gamma^t R(s_t) \mid s_0 = s, a_0 = a]$ for $s \in \mathcal{S}$ and $a \in \mathcal{A}$. We denote $V_R^*(s) = \max_\pi V_R^\pi(s)$ and $Q_R^*(s, a) = \max_\pi Q_R^\pi(s, a)$. The expected value of a policy is denoted by $V_R^\pi = \mathbb{E}_{s \in S_0}[V_R^\pi(s)]$.

We assume that the $\arg\max$ operator returns a set, i.e., $\arg\max_x f(x) := \{x \mid f(y) \leq f(x), \forall y\}$. We let $\pi_R^* \in \arg\max_\pi V_R^\pi$ denote an optimal policy under reward function R. We also let $\mathcal{A}_R(s) = \arg\max_{a' \in \mathcal{A}} Q_R^*(s, a')$ denote the set of all optimal actions at state s under reward function R. Thus, $\mathcal{A}_R(s) = \{a \in \mathcal{A} \mid \pi_R^*(a|s) > 0\}$

As is common (Ziebart et al., 2008; Barreto et al., 2017; Brown et al., 2020), we assume that the reward function is linear under features $\phi : \mathcal{S} \mapsto \mathbb{R}^k$, so that $R(s) = \mathbf{w}^T \phi(s)$, where $\mathbf{w} \in \mathbb{R}^k$. Thus, we use R and $\mathbf{w}$ interchangeably. Note that this assumption of a linear reward function is not restrictive as these features can be arbitrarily complex nonlinear functions of the state and could be obtained via unsupervised learning from raw state observations (Laskin et al., 2020; Brown et al., 2020). Given that $R(s) = \mathbf{w}^T \phi(s)$, the state-action value function can be written in terms of discounted expectations over features (Abbeel & Ng, 2004): $Q_R^\pi(s, a) = \mathbf{w}^T \Phi_\pi^{(s,a)}$, where $\Phi_\pi^{(s,a)} = \mathbb{E}_\pi[\sum_{t=0}^{\infty} \gamma^t \phi(s_t) \mid s_0 = s, a_0 = a]$.

4. Value Alignment Verification

In this section we first explicitly define value alignment and value alignment verification. Next, we discuss how assuming rationality of the robot enables efficient provable value alignment verification. We then examine how to perform (approximate) value alignment verification in tabular MDPs under different forms of test queries, including reward, value, preference, and action queries. We conclude this section by presenting a method for approximate value alignment verification when the tester is a human with implicit values and the state and action spaces are continuous.

We first formalize value alignment. Consider two agents: a human and a robot. We will assume that the human has a (possibly implicit) reward function that provides the ground truth for determining value alignment verification of the robot. We define (approximate) value alignment as follows:

Definition 1. *Given reward function R, policy π' is ϵ-**value aligned** in environment E if and only if*

$$V_R^*(s) - V_R^{\pi'}(s) \leq \epsilon, \forall s \in \mathcal{S}. \tag{1}$$

Exact value alignment is achieved when $\epsilon = 0$.

We are interested in **efficient value alignment verification** where we can correctly classify agents as aligned or misaligned within certain error tolerances while keeping the total test size small. Formally, efficient (approximate) value alignment verification is a solution to the following:

$$\min_{T \subseteq \mathcal{T}} |T|, \text{s.t. } \forall \pi' \in \Pi, \forall s \in \mathcal{S} \tag{2}$$
$$V_R^*(s) - V_R^{\pi'}(s) > \epsilon \Rightarrow Pr[\pi' \text{ passes test } T] \leq \delta_{\text{fpr}}$$
$$V_R^*(s) - V_R^{\pi'}(s) \leq \epsilon \Rightarrow Pr[\pi' \text{ fails test } T] \leq \delta_{\text{fnr}}$$

where $\mathcal{T}$ is the choice set of possible test queries, Π denotes the set of robot policies for which we design the test, $\delta_{\text{fp}}, \delta_{\text{fn}} \in [0, 1]$ denote the allowable false positive rate and false negative rate, and $|T|$ denotes the cardinality, or complexity of the test, T. If $\epsilon = \delta_{\text{fpr}} = 0$, then we seek the test that enables exact value alignment verification.

4.1. Query Types and Rational Agents

The difficulty of solving Equation 2 can change significantly as a function of ϵ, δ_{fpr}, δ_{fnr}, the set of policies for which we design the test Π, and the type of queries available in

the choice set $\mathcal{T}$. For example, exact alignment is impossible in settings where one can only query for actions (see Appendix A.1). Even when possible, achieving high confidence may require multiple action queries at every state.

One of the main goals of this paper is to understand under what settings we can achieve efficient, provable value alignment verification. Towards this end, we assume that the robot behaves rationally with respect to some reward function R'. A *rational agent* is one that picks actions to maximize its utility (Russell & Norvig, 2016). Formally π' is a rational agent if:

$$\forall a \in \mathcal{A}, \pi'(a|s) > 0 \implies a \in \arg\max_a Q^*_{R'}(s,a), \quad (3)$$

where $\arg\max_a Q^*_{R'}(s,a)$ returns the set of all optimal actions at state s under R'.

Note that rationality in itself does not restrict the set of policies Π for which we can test, since all policies are rational under the trivial all zero reward function (Ng & Russell, 2000). Rationality also does not limit the choice set $\mathcal{T}$ since a rational agent can answer any question related to its policy or values. The rationality assumption is helpful because it directly connects the behavior of the agent to a reward function: given behavior we can infer rewards and given rewards we can infer behavior. It also allows us to extrapolate robot behavior to new situations, enabling efficient value alignment verification.

4.2. Exact Value Alignment

We start with the idealized query setting of *explicit human, explicit robot*. In this section we discuss exact value alignment ($\epsilon = 0, \delta_{\text{fpr}} = 0$) of a rational robot and review related work by (Ng & Russell, 2000) on sets of rewards consistent with an optimal policy. Then in the next section we will examine how to construct verification tests for exact alignment. We assume that both the human and robot know the states reward features $\phi(s)$, and that the robot acts rationally with respect to a reward function linear in these features.

Consider two rational agents with reward functions R and R'. Because there are infinite reward functions that lead to the same optimal policy (Ng & Russell, 2000), determining that $\exists s \in S, R(s) \neq R'(s)$ does not necessarily imply misalignment. For ease of notation, we define

$$OPT(R) = \{\pi \mid \pi(a|s) > 0 \Rightarrow a \in \arg\max_a Q^*_R(s,a)\},$$

as the set of all optimal (potentially stochastic) policies in MDP (E, R). Combining Definition (1) and Equation (3) immediately gives us that a rational robot is aligned with a human if all optimal policies under the robot's reward function are also optimal policies under the human's reward function. We formally state this as the following Corollary.

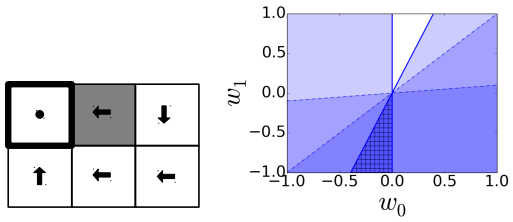

(a) Policy π (b) $CRS(\pi)$.

Figure 2. An example of the consistent reward set (CRS) for a policy π in a simple gridworld and a linear reward function with two binary reward features (white and gray) with reward weights w_0 and w_1, respectively.

Corollary 1. *We have **exact value alignment** in environment E between a rational robot with reward function R' and a human with reward function R if $OPT(R') \subseteq OPT(R)$.*

We now review foundational work on IRL by Ng and Russell (Ng & Russell, 2000) which inspires our proposed approach for efficient value alignment verification.

Definition 2. *Given an environment E, the **consistent reward set** (CRS) of a policy π in environment E is defined as the set of reward functions under which π is optimal:*

$$CRS(\pi) = \{R \mid \pi \in OPT(R)\}. \quad (4)$$

When $R(s) = \mathbf{w}^T\phi(s)$, the CRS is the following polytope:

Corollary 2. *(Ng & Russell, 2000; Brown & Niekum, 2019) Given an environment E, the $CRS(\pi)$ is given by the following intersection of half-spaces:*

$$\{\mathbf{w} \in \mathbb{R}^k \mid \mathbf{w}^T(\Phi^{(s,a)}_\pi - \Phi^{(s,b)}_\pi) \geq 0, \\ \forall a \in \arg\max_{a' \in \mathcal{A}} Q^\pi_R(s,a'), b \in \mathcal{A}, s \in \mathcal{S}\}. \quad (5)$$

As an example consider the grid world MDP shown in Figure 2. The CRS is an intersection of half-spaces which define all reward functions under which π is optimal. Note that the all zero reward function and the reward function where white cells have zero reward are included; however, not all optimal policies under these reward functions lead to the policy shown in Figure 2a.

Thus, we cannot directly use Corollary 2 to verify alignment with a human's optimal policy—Corollary 2 only provides a necessary, but not sufficient, condition for testing whether a reward function R' is value aligned with a policy π. Consider the example of the trivial all zero reward function: it is always in the CRS of any policy; however, an agent optimizing the zero reward can result in any arbitrary policy. Even ignoring the all zero reward, rewards can be on the boundaries of the CRS polytope that are consistent with a policy,

but not value aligned since they lead to more than one optimal policy, one or more of which may not be optimal under the human's reward function. In the next section we show that if we remove all such edge cases, we can construct an *aligned reward polytope* (ARP) similar to the CRS, which enables provable value alignment verification. Furthermore, we show that the aligned reward polytope can be used for alignment verification even when the human cannot directly query for the robot's reward function.

4.3. Sufficient Conditions for Provable Verification of Exact Value Alignment

We seek an efficient value alignment verification test which enables a human to query the robot to determine exact value alignment as in Corollary 1. The following theorem demonstrates that provable verification of exact value alignment is possible under a variety of query types.

Theorem 1. *Under the assumption of a rational robot that shares linear reward features with the human, efficient exact value alignment verification is possible in the following query settings: (1) Query access to reward function weights $\mathbf{w}'$, (2) Query access to samples of the reward function $R'(s)$, (3) Query access to $V^*_{R'}(s)$ and $Q^*_{R'}(s,a)$, and (4) Query access to preferences over trajectories.*

4.3.1. Case 1: Reward Weight Queries

We first consider the case where the human can directly query the robot for their reward function weights $\mathbf{w}'$. While this problem setting is mainly of theoretical interest, we will show that Cases (2) and (3) also reduce to this setting. Querying directly for the robot's reward function is maximally efficient since by definition it only requires a single query. Although one can solve for the optimal policy under a given $\mathbf{w}'$ and evaluate it under the human's reward function $\mathbf{w}$, this brute force approach is computationally demanding and must be repeated for each robot that needs to be tested. By contrast, we will prove that there exists a single efficient verification test that does not require solving for the robot's optimal policy and can be used to verify the alignment of any robot.

As mentioned in the previous section, the CRS for the human's optimal policy does not provide a sufficient test for value alignment verification. Under the assumption of a rational robot, a sufficient condition for value alignment verification is to test whether a robot's reward function lies in the following set:

Definition 3. *Given an MDP M composed of environment E and reward function R, the **aligned reward set** (ARS) is defined as the following set of reward functions:*

$$ARS(R) = \{R' \mid OPT(R') \subseteq OPT(R)\}. \tag{6}$$

Using Definition 3, we prove the following lemma which will enable efficient verification of exact value alignment. As a reminder, we use the notation $Q^\pi_R(s,a) = \mathbf{w}^T \Phi^{(s,a)}_\pi$, for $\Phi^{(s,a)}_\pi = \mathbb{E}_\pi[\sum_{t=0}^\infty \gamma^t \phi(s_t) \mid s_0 = s, a_0 = a]$, and $\mathcal{A}_R(s) = \arg\max_{a' \in \mathcal{A}} Q^*_R(s,a')$.

Lemma 1. *Given an MDP $M = (E,R)$, assuming the human's reward function R, and the robot's reward function R' can be represented as linear combinations of features $\phi(s) \in \mathbb{R}^k$, i.e., $R(s) = \mathbf{w}^T\phi(s)$, $R'(s) = \mathbf{w}'^T\phi(s)$, and given an optimal policy π^*_R under R then*

$$\mathbf{w}' \in \bigcap_{(s,a,b)\in\mathcal{O}} \mathcal{H}^R_{s,a,b} \implies R' \in ARS(R) \tag{7}$$

where $\mathcal{H}^R_{s,a,b} = \left\{\mathbf{w} \mid \mathbf{w}^T(\Phi^{(s,a)}_{\pi^*_R} - \Phi^{(s,b)}_{\pi^*_R}) > 0\right\}$ and $\mathcal{O} = \{(s,a,b) | s \in \mathcal{S}, a \in \mathcal{A}_R(s), b \notin \mathcal{A}_R(s)\}$.

Proof sketch. First we show π^*_R is optimal under R' using the policy improvement theorem. Then, using the uniqueness of the optimal value function, we show that all optimal actions under R are also optimal actions under R', and so all optimal policies under R' are optimal under R. (see Appendix A.3 for the full proof). □

Lemma 1 provides a sufficient condition for verifying exact value alignment. We now have the necessary theory to construct an efficient value alignment verification test in the *explicit human, explicit robot* setting. We aim to efficiently verify whether the robot's reward function, R', is within the above intersection of half-spaces, which we call the Aligned Reward Polytope (ARP), as this gives a sufficient condition for R' being value aligned with the human's reward function R. Our analysis in this section will be useful later when we consider approximate tests for value alignment verification when one or both of the agents have implicit values.[2]

The verification test is constructed by precomputing the following matrix representation of the ARP:

$$\mathbf{\Delta} = \begin{bmatrix} \Phi^{(s,a)}_{\pi^*_R} - \Phi^{(s,b)}_{\pi^*_R} \\ \vdots \end{bmatrix}, \tag{8}$$

where each row corresponds to a tuple $(s,a,b) \in \mathcal{O}$. Thus, a is an optimal action and b is a suboptimal action under R and each row of $\mathbf{\Delta}$ represents the normal vector for a strict half-space constraint based on feature count differences between an optimal and suboptimal action. Note that, using this notation, exact value alignment can now be verified by checking whether $\mathbf{\Delta}\mathbf{w}' > 0$. This test can be made more

[2] Our results may also be of interest in the analysis of *explicit robot, explicit robot* teaming, e.g., ad hoc teamwork (Stone et al., 2010) where value alignment verification could provide a framework for verifying whether two robots can work together.

efficient by only including non-redundant half-space normal vectors in $\mathbf{\Delta}$. In Appendix G.2 we discuss a straightforward linear programming technique to efficiently obtain the minimal set of half-space constraints that define the intersection of half-spaces specified in Lemma 1.

4.3.2. Case 2: Reward Queries

We now consider the case where the tester can query for samples of the robot's reward function $R'(s)$. Verifying alignment via queries to $R'(s)$ can be reduced to Case (1) by querying the robot for $R'(s)$ over a sufficient number of states and then solving for a system of linear equations to recover $\mathbf{w}'$, since we assume both the human and robot have access to the reward features $\phi(s)$.[3] Let $\mathbf{\Phi}$ be defined as the matrix where each row corresponds to the feature vector $\phi(s)^T$ for a distinct state $s \in \mathcal{S}$. Then, the number of required queries is equal to $rank(\mathbf{\Phi})$ since we only need samples corresponding to linearly independent rows of Φ. Thus, if $\mathbf{w}' \in \mathbb{R}^k$, in the worst case we only need k samples from the robot's reward function, since we have $rank(\mathbf{\Phi}) \leq k$. If there is noise in the sampling procedure, then linear regression can be used to efficiently estimate the robot's weight vector $\mathbf{w}'$. Given $\mathbf{w}'$ we can verify value alignment by checking whether $\mathbf{\Delta}\mathbf{w}' > 0$.

4.3.3. Case 3: Value Function Queries

Given query access to the robot's state and state-action value functions, $\mathbf{w}'$ can be determined by noting that $R'(s) = \mathbf{w}'^T\phi(s)$ and

$$R'(s) = Q^*_{R'}(s,a) - \gamma \mathbb{E}_{s'}\left[V^*_{R'}(s')\right]. \tag{9}$$

Computing the expectation requires enumerating successor states. If we define the maximum degree of the MDP transition function as

$$d_{\max} = \max_{s\in\mathcal{S}, a\in\mathcal{A}} |\{s' \in \mathcal{S} \mid P(s,a,s') > 0\}|, \tag{10}$$

then at most the $d_{\max}$ possible next state value queries are needed to evaluate the expectation. Thus, at most $\text{rank}(\Phi)(d_{\max}+1)$ queries to the robot's value functions are needed to recover $\mathbf{w}'$, and the tester can verify value alignment via Case (1). Since $\text{rank}(\Phi) \leq k$ as before, at most $k(d_{\max}+1)$ queries are required for $\mathbf{w}' \in \mathbb{R}^k$.

4.3.4. Case 4: Preference Queries

Finally, we consider the *implicit robot* setting where the tester can only query the robot for preferences over trajectories, ξ. Each preference over trajectories, $\xi_A \prec \xi_B$, induces the constraint $\mathbf{w}'^T(\Phi(\xi_B) - \Phi(\xi_A)) > 0$, where $\Phi(\xi) = \sum_{i=1}^n \gamma^i \phi(s_i)$ is the cumulative discounted reward features along a trajectory. Thus, our choice set of tests, $\mathcal{T}$, consists of all trajectory preference queries, and we can guarantee value alignment if we have a test T such that $\mathbf{w}^T(\Phi(\xi_B) - \Phi(\xi_A)) > 0, \forall(\xi_A, \xi_B) \in T$ implies that $\mathbf{w} \in \bigcap \mathcal{H}^R_{s,a,b}$. We can then construct $\mathbf{\Delta}$ in a similar fashion as above, except each row corresponds to a half-space normal resulting from a preference over individual trajectories (see Appendix A.3). Only a logarithmic number of preferences over randomly generated trajectories are needed to accurately represent $\bigcap \mathcal{H}^R_{s,a,b}$ via intersection of half-spaces formed by the rows in $\mathbf{\Delta}$ (Brown et al., 2019).

[3] Note that our results also hold for rewards that are functions of (s,a) and (s,a,s').

4.4. Value Alignment Verification Heuristics

In the next section we relax our assumptions on the robot and consider the *explicit human, implicit robot* setting, where the human seeks to verify value alignment but the robot has a black-box policy that only affords action queries. In this case, we resort to heuristics for value alignment as exact value alignment verification becomes impossible, and ϵ-value alignment verification by directly attempting to solve Equation (2) when $\mathcal{T}$ consists of state-action queries is computationally intractable. As we discuss in detail in Appendix B, a direct optimization approach would involve estimating Π by computing the optimal policies for a large number of different reward functions, evaluating each policy under $\mathbf{w}$ to determine which policies are not ϵ-aligned with the tester's reward function R, and then solving a combinatorial optimization problem over all possible state queries.

Instead, we resort to efficient heuristics. We consider three heuristic alignment tests designed to work in the black-box value alignment verification setting, where the tester can only ask the robot policy action queries over states. Each heuristic test consists of a method for selecting states at which to test the robot by querying for an action from the robot's policy and checking if that action is an optimal action under the human's reward function. Note that querying only a subset of states for robot actions is fundamentally limited to value alignment verification tests with $\delta_{\text{fpr}} > 0$ since we will never know for sure that the agent will not take a different action in that state if we query its policy again. Thus, receiving the "right answer"—an optimal action under the tester's reward R—to an action query in a state is not a sufficient condition for exact value alignment. We briefly discuss three action query heuristics with full details in Appendix C. Figure 3 shows examples of the state queries generated by each heuristic in a simple gridworld.

Critical States Heuristic Our first heuristic is inspired by the notion of *critical states*: states where $Q^*_R(s, \pi^*_R(s)) - \frac{1}{|\mathcal{A}|}\sum_{a\in\mathcal{A}} Q^*_R(s,a) > t$, and t is a user defined threshold (Huang et al., 2018). We adapt this idea to form a critical state alignment heuristic test (CS) consisting of critical states under the human's reward function R. Intuitively,

these states are likely to be important; however, often many critical states will be redundant since different states are often important for similar reasons (see Figure 3).

Machine Teaching Heuristic Our next heuristic is based on Set Cover Optimal Teaching (SCOT) (Brown & Niekum, 2019), a machine teaching algorithm that approximates the minimal set of maximally informative state-action trajectories necessary to teach a specific reward function to an IRL agent. Brown & Niekum (2019) prove that the learner will recover a reward function in the intersection of halfspaces that define the CRS (Corollary 2). We generate informative trajectories using SCOT, and turn them into alignment tests by querying the robot for their action at each state along the trajectories. SCOT replaces the explicit checking of half-space constraints in Section 4.3 with implicit half-space constraints that are inferred by querying for robot actions at states along trajectories, thus introducing approximation error and the possibility of false positives. Furthermore, generating a test using SCOT is more computationally intensive than generating a test via the CS heuristic; however, unlike CS, SCOT will seek to avoid redundant queries by reasoning about reward features over a collection of trajectories.

ARP Heuristic Our third heuristic takes inspiration from the definition of the ARP to define a black-box alignment heuristic (ARP-bb). ARP-bb first computes $\mathbf{\Delta}$ (see Equation (8)), removes redundant half-space constraints via linear programming, and then only queries for robot actions from the states corresponding to the non-redundant constraints (rows) in $\mathbf{\Delta}$. Intuitively, states that are queried by ARP-bb are important in the sense that taking different actions reveals important information about the reward function. However, ARP-bb uses single-state action queries to approximate checking each half-space constraint. Thus, ARP-bb trades off smaller query and computational complexity with the potenital for larger approximation error.

4.5. Implicit Value Alignment Verification

We now discuss value alignment verification in the *implicit human, implicit robot* setting. Without an explicit representation of the human's values we cannot directly compute the aligned reward polytope (ARP) via enumeration over states and actions to create an intersection of half-spaces as described above. Instead, we propose the pipeline outlined in Figure 1 where an AI system elicits and distills human preferences and then generates a test which can be used to approximately verify the alignment of any rational agent.

As is common for active reward learning algorithms (Bıyık et al., 2019), we assume that the preference elicitation algorithm outputs both a set of preferences over trajectories $\mathcal{P} = \{(\xi_i, \xi_j) : \xi_i \succ \xi_j\}$ and a set of reward weights $\mathbf{w}$ sampled from the posterior distribution $\{\mathbf{w}_i\} \sim P(\mathbf{w}|\mathcal{P})$. Given $\mathcal{P}$ and $P(\mathbf{w}|\mathcal{P})$, the ARP of the human's implicit reward function can be approximated as

$$ARP(R) \approx \bigcap_{(\xi_i,\xi_j)\in\mathcal{P}} \{\mathbf{w} \mid \mathbf{w}^T(\Phi(\xi_i) - \Phi(\xi_j)) > 0\}, \tag{11}$$

which generalizes the definition of the ARP to MDPs with continuous states and actions. To see this, note that the intersection of half-spaces in Lemma 1 enumerates over states and pairs of optimal and suboptimal actions under the human's reward R to create the set of half-space normal vectors $\mathbf{\Delta}$, where each normal vector is a difference of expected feature counts. This enumeration can only be done in discrete MDPs. Equation (11) approximates the ARP for continuous MDPs via half-space normal vectors constructed with empirical feature count differences obtained from pairs of actual trajectories over continuous states and actions.

This test can be further generalized to ϵ-value alignment (Definition 1) to test agents with bounded rationality or slightly misspecified reward functions. One method of constructing an ϵ-alignment test is to use the mean posterior reward $\mathbb{E}[\mathbf{w}]$ to approximate the value difference of each pair of trajectories $\mathbb{E}[\mathbf{w}](\Phi(\xi_i) - \Phi(\xi_j))$, and only include preference queries with estimated value differences of at least ϵ. A robot with implicit values is verified as ϵ-value aligned by test T if its preferences over each pair of trajectories in T match the preferences provided by the human (see Appendix F for more details).

5. Experiments

We now study the empirical performance of value alignment verification tests, first in the *explicit human* setting and then in the *implicit human* setting.

5.1. Value Alignment Verification with Explicit Human

We first study the *explicit human* setting and analyze the efficiency and accuracy of exact value alignment verification tests and heuristics. We consider querying for the weight vector of the robot (ARP-w), querying for trajectory preferences (ARP-pref), and the action-query heuristics: CS, SCOT, and ARP-bb, described in Section 4.4.

5.1.1. Case Study

To illustrate the types of test queries found via value alignment verification, we consider two domains inspired by the AI safety gridworlds (Leike et al., 2017). The first domain, *island navigation* is shown in Figure 3. Figure 3a shows the optimal policy under the tester's reward function, $R(s) = 50 \cdot \mathbf{1}_{\text{green}}(s) - 1 \cdot \mathbf{1}_{\text{white}}(s) - 50 \cdot \mathbf{1}_{\text{blue}}(s)$, where $\mathbf{1}_{\text{color}}(s)$ is an indicator feature for the color of the grid cell. Shown in figures 3b and 3c are the two preference queries generated by ARP-pref which consist of pairwise

trajectory queries (black is preferable to orange under R). Preference query 1 verifies that the robot would rather move the to terminal state (green) rather than visit more white cells. Preference query 2 verifies that the robot would rather visit white cells than blue cells. Figures 3d, 3e, and 3f show action query tests designed using the ARP-bb, SCOT, and CS heuristics. The robot is asked which action its policy would take in each of the states marked with a question mark. To pass the test, the agent must respond with an optimal action under the human's policy in each of these states. ARP-bb chooses two states based on the half-space constraints defined by the expected feature counts of π^*_R, resulting in an small but myopic test. SCOT queries over a maximally informative trajectory that starts near the water, but includes several redundant states. CS only reasons about Q-value differences and asks many redundant queries (see Appendix D for more results).

5.1.2. Sensitivity Analysis

We also analyze the accuracy and efficiency of value alignment verification in the *explicit human, explicit robot* and *explicit human, implicit robot* settings for verifying exact value alignment. We analyze performance across a suite of random grid navigation domains with varying numbers of states and reward features. We summarize our results here and refer the reader to Appendix E for more details. As expected, ARP-w and ARP-pref result in perfect accuracy. SCOT has uses fewer samples than the CS heuristic while achieving nearly perfect accuracy. ARP-bb results in higher accuracy tests, but generates more false positives than SCOT. CS has significantly higher sample cost than the other methods and requires careful tuning of the threshold t to obtain good performance. Our results indicate that in the *implicit robot* setting, ARP-pref and ARP-bb provide highly efficient verification tests. Out of the action query heuristics, SCOT achieved the highest accuracy, while having larger sample complexity than ARP-bb, but achieving lower sample complexity than CS.

5.2. Value Alignment Verification with Implicit Human

We next analyze approximate value alignment verification in the continuous autonomous driving domain from Sadigh et al. (2017), shown in Figure 4a, where we study the *implicit human, implicit robot* setting and consider verifying ϵ-value alignment. As depicted in Figure 1 we analyze the use of active preference elicitation (Bıyık et al., 2019) to perform value alignment verification with *implicit human values*. We first analyze implicit value alignment verification using preference queries to a synthetic human oracle unobserved ground-truth reward function R.

We collected varying numbers of oracle preferences, and computed a non-redundant ϵ-alignment test as described in 4.5 and Appendix G.2. Tests were evaluated for accuracy relative to a set of test reward weights. See Appendix G for experimental parameters and details of the testing reward generation protocol. Figure 4b displays the results of the synthetic human experiments. The best tests achieved 100% accuracy. Although collecting additional synthetic human queries consistently improved verification accuracy, above 50 human queries, accuracy gains were minimal, demonstrating the potential for human-in-the-loop preference elicitation. Furthermore, the generated verification tests were often succinct: one of the tests with perfect accuracy required only six questions out of the original 100 elicited preferences. Additional experiments and results are detailed in Appendix G, including false positive and false negative rate plots, and different methods of estimating the value gap of questions. We also ran an initial pilot study using real human preference labels which resulted in a verification test that achieves 72% accuracy.

6. Generalization to Multiple MDPs

Up to this point, we have considered designing value alignment tests for a single MDP; however, it is also interesting to try and design value alignment verification tests that enable generalization, e.g., if a robot passes the test, then this verifies value alignment across many different MDPs.

As a step towards this goal, we present a result in the *explicit human, explicit robot* setting where the human can construct testing environments. We consider the idealized setting of an omnipotent tester that is able to construct a set of arbitrary test MDPs and can query directly for the entire optimal policy of the robot in each MDP. This tester aims to verify value alignment across an infinite family of environments that share the same reward function. Our result builds on prior analysis on the related problem of omnipotent *active reward learning*. Amin & Singh (2016) prove that an active learner can determine the reward function of another agent within ϵ precision via $O(\log |\mathcal{S}| + \log(1/\epsilon))$ policy queries. By contrast, we prove in the following theorem that the sample complexity of ϵ-value alignment verification is only $O(1)$ (see Appendix A.5 for the proof).

Theorem 2. *Given a testing reward R (not necessarily linear in known features), there exists a two-query test (complexity $O(1)$) that determines ϵ-value alignment of a rational agent over all MDPs that share the same state space and reward function R, but may differ in actions, transitions, discount factors, and initial state distribution.*

We also note that if the human has access a priori to a finite set of MDPs over which they want to verify value alignment, then our results from earlier sections on exact, heuristic, and approximate value alignment could be extended to this setting. For example, we can define a generalized aligned reward polytope for a family of MDPs as the intersection of

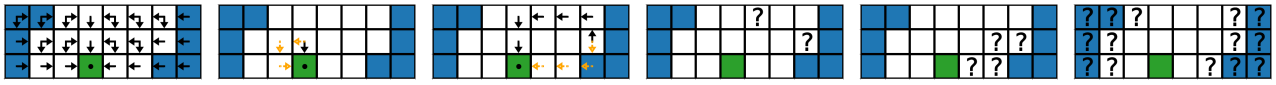

(a) Optimal policy (b) Preference query 1 (c) Preference query 2 (d) ARP-bb queries (e) SCOT queries (f) CS queries

Figure 3. Examples of exact and heuristic value alignment verification tests for an island navigation gridworld (Leike et al., 2017). Only two preference queries (b) and (c) are required to provably verify any robot policy (black should be preferred over orange). Figures (d)-(f) show heuristic tests that query for actions at individual states.

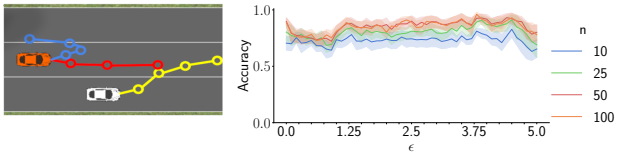

(a) Driving domain (b) Alignment test accuracy vs ϵ

Figure 4. **Implicit human, implicit robot:** ϵ-value alignment verification in a continuous autonomous driving domain. (a) A preference query. The human is asked if they prefer the blue or the red trajectory w.r.t. the trajectory of the white car. (b) 80% confidence intervals on verification accuracy for different values of ϵ, different human query budgets n, averaged over ten seeds.

the aligned reward polytope for each individual MDP. This intersection of half-spaces provides a sufficient condition for testing value alignment across the entire family of MDPs. We leave this as a promising area for future work.

7. Discussion

We analyzed the problem of efficient value alignment verification: how to generate an efficient test that can be used to verify the value alignment of another agent with respect to the human's reward function. We developed a theoretical foundation for value alignment verification and proved sufficient conditions for verifying the alignment of a rational agent under explicit and implicit values for both the human and robot. Our empirical analysis demonstrates that action query heuristics can achieve low sample complexity and high accuracy while only requiring black-box access to an agent's policy. When the human has only implicit access to their values, we analyzed active preference elicitation algorithms as a potential means to automatically construct an approximate value alignment test that can efficiently test another agent with implicit values.

The biggest assumption we make is that the reward function is a linear combination of features shared by both the human and robot. We would like to emphasize three points: First, on representing rewards as linear combinations of features, note that the features can be arbitrarily complex, and can even be features learned via a deep neural network which are then linearly transformed by a final linear layer (Brown et al., 2020). Second, there is the issue of the human and the robot sharing the features. The reason this might actually be a reasonable assumption is that recent techniques enable robots to detect when they cannot explain human input with their existing features and ask for new input specific to the missing features (Bobu et al., 2020; 2021), thereby explicitly aligning the robot's reward representation with the human's reward representation. Third, even if the features are not perfectly aligned, our approach can still provide value by learning a linear combination of features that approximates the human's reward function to design an alignment test.

Our pilot study with the driving simulation hints that this might be the case, as it gives evidence that value alignment verification is possible when using real human preferences that are determined using pixel-based observations. Furthermore, the only true requirement for generating value alignment tests that query for robot actions or preferences is for the tester to have a reward function that can be approximated by a linear combination of features. Thus, these tests could be possibly be applied in cases where a human uses a linear combination of learned or human-designed features to construct an approximate alignment test for robots who have pixel-based policies and/or rewards.

In conclusion, we believe that value alignment verification is an important problem of practical interest, as it seeks to enable humans to verify and build trust in AI systems. It may also be possible for a robot to use value alignment verification to verify the performance of a human, e.g., AI-generated assessment tests. Future work also includes relaxing rationality assumptions, analyzing value alignment verification tests in more complex domains, and performing a full user study to better analyze the use of human preferences for alignment verification.

Acknowledgements

We would like to thank the anonymous reviewers and Stephen Giguere for their suggestions for improving the paper. This work was funded in part by NSF, AFOSR, ARO, NSF NRI SCHOOL, and ONR YIP.

References

Abbeel, P. and Ng, A. Y. Apprenticeship learning via inverse reinforcement learning. In *Proceedings of the twenty-first international conference on Machine learning*, pp. 1. ACM, 2004.

Amin, K. and Singh, S. Towards resolving unidentifiability in inverse reinforcement learning. *arXiv preprint arXiv:1601.06569*, 2016.

Amin, K., Jiang, N., and Singh, S. Repeated inverse reinforcement learning. In *Advances in Neural Information Processing Systems*, pp. 1815–1824, 2017.

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

Arora, S. and Doshi, P. A survey of inverse reinforcement learning: Challenges, methods and progress. *arXiv preprint arXiv:1806.06877*, 2018.

Barreto, A., Dabney, W., Munos, R., Hunt, J. J., Schaul, T., van Hasselt, H. P., and Silver, D. Successor features for transfer in reinforcement learning. In *Advances in neural information processing systems*, pp. 4055–4065, 2017.

Biyik, E. and Sadigh, D. Batch active preference-based learning of reward functions. PMLR, 2018.

Bıyık, E., Palan, M., Landolfi, N. C., Losey, D. P., and Sadigh, D. Asking easy questions: A user-friendly approach to active reward learning. In *Conference on Robot Learning (CoRL)*, 2019.

Bobu, A., Bajcsy, A., Fisac, J. F., Deglurkar, S., and Dragan, A. D. Quantifying hypothesis space misspecification in learning from human–robot demonstrations and physical corrections. *IEEE Transactions on Robotics*, 36(3):835–854, 2020.

Bobu, A., Wiggert, M., Tomlin, C., and Dragan, A. D. Feature expansive reward learning: Rethinking human input. In *Proceedings of the 2021 ACM/IEEE International Conference on Human-Robot Interaction*, pp. 216–224, 2021.

Brown, D. S. and Niekum, S. Machine teaching for inverse reinforcement learning: Algorithms and applications. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pp. 7749–7758, 2019.

Brown, D. S., Goo, W., and Niekum, S. Better-than-demonstrator imitation learning via automaticaly-ranked demonstrations. In *Conference on Robot Learning (CoRL)*, 2019.

Brown, D. S., Niekum, S., Coleman, R., and Srinivasan, R. Safe imitation learning via fast bayesian reward inference from preferences. In *International Conference on Machine Learning*. 2020.

Cakmak, M. and Lopes, M. Algorithmic and human teaching of sequential decision tasks. In *AAAI*, 2012.

Christiano, P. F., Leike, J., Brown, T., Martic, M., Legg, S., and Amodei, D. Deep reinforcement learning from human preferences. In *Advances in Neural Information Processing Systems*, pp. 4299–4307, 2017.

Fisac, J. F., Gates, M. A., Hamrick, J. B., Liu, C., Hadfield-Menell, D., Palaniappan, M., Malik, D., Sastry, S. S., Griffiths, T. L., and Dragan, A. D. Pragmatic-pedagogic value alignment. In *Robotics Research*, pp. 49–57. Springer, 2020.

Hadfield-Menell, D., Russell, S. J., Abbeel, P., and Dragan, A. Cooperative inverse reinforcement learning. In *Advances in Neural Information Processing Systems 29*, pp. 3909–3917. 2016.

Hanna, J. P., Stone, P., and Niekum, S. Bootstrapping with models: Confidence intervals for off-policy evaluation. In *Proceedings of the 16th Conference on Autonomous Agents and Multiagent Systems*, 2017.

Huang, S. H., Held, D., Abbeel, P., and Dragan, A. D. Enabling robots to communicate their objectives. In *Robotics: Science and Systems*, 2017.

Huang, S. H., Bhatia, K., Abbeel, P., and Dragan, A. D. Establishing appropriate trust via critical states. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 3929–3936. IEEE, 2018.

Laskin, M., Srinivas, A., and Abbeel, P. Curl: Contrastive unsupervised representations for reinforcement learning. In *International Conference on Machine Learning*, pp. 5639–5650. PMLR, 2020.

Leike, J., Martic, M., Krakovna, V., Ortega, P. A., Everitt, T., Lefrancq, A., Orseau, L., and Legg, S. Ai safety gridworlds. *arXiv preprint arXiv:1711.09883*, 2017.

Leike, J., Krueger, D., Everitt, T., Martic, M., Maini, V., and Legg, S. Scalable agent alignment via reward modeling: a research direction. *arXiv preprint arXiv:1811.07871*, 2018.

Ng, A. Y. and Russell, S. J. Algorithms for inverse reinforcement learning. In *ICML*, pp. 663–670, 2000.

Paulraj, S., Sumathi, P., et al. A comparative study of redundant constraints identification methods in linear programming problems. *Mathematical Problems in Engineering*, 2010.

Precup, D. Eligibility traces for off-policy policy evaluation. *Computer Science Department Faculty Publication Series*, pp. 80, 2000.

Russell, S., Dewey, D., and Tegmark, M. Research priorities for robust and beneficial artificial intelligence. *Ai Magazine*, 36(4):105–114, 2015.

Russell, S. J. and Norvig, P. *Artificial intelligence: a modern approach*. Malaysia; Pearson Education Limited,, 2016.

Sadigh, D., Dragan, A. D., Sastry, S. S., and Seshia, S. A. Active preference-based learning of reward functions. In *Proceedings of Robotics: Science and Systems (RSS)*, July 2017. doi: 10.15607/RSS.2017.XIII.053.

Shah, R., Freire, P., Alex, N., Freedman, R., Krasheninnikov, D., Chan, L., Dennis, M., Abbeel, P., Dragan, A., and Russell, S. Benefits of assistance over reward learning. *Workshop on Cooperative AI (Cooperative AI @ NeurIPS*, 2020.

Stone, P., Kaminka, G. A., Kraus, S., Rosenschein, J. S., et al. Ad hoc autonomous agent teams: Collaboration without pre-coordination. In *AAAI*, 2010.

Sutton, R. S. and Barto, A. G. *Introduction to reinforcement learning*, volume 135. MIT press Cambridge, 1998.

Thomas, P. S., Theocharous, G., and Ghavamzadeh, M. High-confidence off-policy evaluation. In *AAAI*, pp. 3000–3006, 2015.

Wirth, C., Akrour, R., Neumann, G., Fürnkranz, J., et al. A survey of preference-based reinforcement learning methods. *Journal of Machine Learning Research*, 18(136): 1–46, 2017.

Zhu, X., Singla, A., Zilles, S., and Rafferty, A. N. An overview of machine teaching. *arXiv preprint arXiv:1801.05927*, 2018.

Ziebart, B. D., Maas, A. L., Bagnell, J. A., and Dey, A. K. Maximum entropy inverse reinforcement learning. In *AAAI*, 2008.