
Combining reward information from multiple sources

Dmitrii Krasheninnikov ^{*} [†]
University of Amsterdam

Rohin Shah [†]
UC Berkeley

Herke van Hoof
University of Amsterdam

Abstract

Given two sources of evidence about a latent variable, one can combine the information from both by multiplying the likelihoods of each piece of evidence. However, when one or both of the observation models are *misspecified*, the distributions will conflict. We study this problem in the setting with two conflicting reward functions learned from different sources. In such a setting, we would like to retreat to a broader distribution over reward functions, in order to mitigate the effects of misspecification. We assume that an agent will maximize expected reward given this distribution over reward functions, and identify four desiderata for this setting. We propose a novel algorithm, Multitask Inverse Reward Design (MIRD), and compare it to a range of simple baselines. While all methods must trade off between conservatism and informativeness, through a combination of theory and empirical results on a toy environment, we find that MIRD and its variant MIRD-IF strike a good balance between the two.

1 Introduction

While deep reinforcement learning (RL) has led to considerable success when *given* an accurate reward function [28, 21], *specifying* a reward function that captures human preferences in real-world tasks is challenging [1, 6]. Value learning seeks to sidestep this difficulty by learning rewards from various types of data, such as natural language [22], demonstrations [33, 11], comparisons [6], ratings [7], human reinforcement [18], proxy rewards [13], the state of the world [27], etc. Many approaches to value learning aim to create a distribution over reward functions [14, 23, 24, 13].

A natural idea is to *combine* the information from these diverse sources. However, the sources may conflict. For example, in Figure 1, Alice might say that she wants the healthy avocado toast, but empirically she almost always chooses to eat the tasty cake. A reward learning method based on natural language may confidently infer that Alice would prefer the toast, whereas a method based on revealed preferences [25] might confidently infer that she wants cake. Presumably one of the algorithms made a poor assumption: perhaps the natural language algorithm failed to model that

Alice only said that she wants avocado toast to impress her friends. Alternatively, maybe Alice wants to start eating healthy, but chooses the cake in moments of weakness, and the revealed preferences algorithm failed to model this bias. While it is unclear what the “true” reward is, we do know that at least one of the methods is *confidently* wrong. We would like to defer on the question of what the “true” reward is, and instead retreat to a position of uncertainty. Formally, we want a distribution over the true reward function θ^* given the inferred rewards $\tilde{\theta}^1$, $\tilde{\theta}^2$.




Figure 1: Alice says she wants avocado toast, but her past actions imply she prefers cake to toast. What should the agent believe about the true reward θ^* ?

^{*}Work done at UC Berkeley

[†]Correspondence to dmitrii.krasheninnikov@student.uva.nl, rohinmshah@berkeley.edu

We could apply Bayes Rule: $p(\theta^* | \tilde{\theta}^1, \tilde{\theta}^2) \propto p(\theta^*)p(\tilde{\theta}^1 | \theta^*)p(\tilde{\theta}^2 | \theta^*)$. However, when rewards conflict due to misspecification, one of the likelihoods will be very low, and the posterior will be low everywhere, likely leading to garbage after normalization [10]. While it is possible to mitigate this by making the assumptions more realistic, e.g. by modeling human biases [26, 19, 9], it is very hard to get a perfect model [3, 29]. So, we aim for a method that is robust to misspecification.

What makes for a good reward combination method? To answer this question, we must know how the resulting distribution will be used. As in Hadfield-Menell et al. [14], we assume that an agent maximizes expected reward, *given that the agent can gather more information about the reward*. In this framework, an agent must balance three concerns: actively learning about the reward [24, 20], preserving its ability to pursue the true reward in the future when it knows more [30, 17], and acting to gain immediate reward. Each concern suggests different desiderata for a reward combination method, which we formalize for rewards functions that are linear in features of the state.

There are several reasonable methods for combining reward functions based on the intuition that we should put weight on any parameter values “in between” the parameters of the given reward functions. These *reward-space* methods operate in the space of parameters of the reward. However, this ignores any details about the environment and its features. We introduce a method called Multi-task Inverse Reward Design (MIRD) that leverages the environment to infer which reward functions are compatible with the behaviors “in between” the behaviors incentivized by the two input rewards, making it a *behavior-space* method. Through a combination of theory and empirical results on a simple environment, we evaluate how well the proposed methods meet our desiderata, and conclude that MIRD is a good option when the reward distribution must be used to act, while its variant MIRD-IF is better for increased robustness and a higher chance of support on the true reward.

2 Background

Markov Decision Process (MDP). A MDP \mathcal{M} is a tuple $\mathcal{M} = \langle \mathcal{S}, \mathcal{A}, \mathcal{T}, r, H \rangle$, where \mathcal{S} is the set of states, \mathcal{A} is the set of actions, $\mathcal{T} : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \mapsto [0, 1]$ is the transition probability function, $r : \mathcal{S} \mapsto \mathbb{R}$ is the reward function, and $H \in \mathbb{Z}_+$ is the finite planning horizon. We consider MDPs where the reward is linear in features, and does not depend on action: $r(s; \theta) = \theta^T f(s)$, where θ are the parameters defining the reward function and f computes features of a given state. A policy $\pi : \mathcal{S} \mapsto \mathcal{A}$ specifies how to act in the MDP. A trajectory $\tau : (\mathcal{S} \times \mathcal{A})^*$ is a sequence of states and actions, where the actions are sampled from a policy π and the states are sampled from the transition function \mathcal{T} . We abuse notation and write $f(\tau)$ to denote $\sum_{s \in \tau} f(s)$. The *feature expectations* (FE) of policy π are the expected feature counts when acting according to π : $\mathcal{F}^\pi = \mathbb{E}_{\tau \sim \pi}[f(\tau)]$. The *value* of a policy π is given by $V^\pi = \mathbb{E}_{\tau \sim \pi}[\sum_{s \in \tau} r(s; \theta)] = \theta^T \mathcal{F}^\pi$. The optimal policy is $\pi^* = \max_\pi V^\pi$.

Reward learning by simulating the past (RLSP). Shah et al. [27] note that environments where humans have acted are already optimized for human preferences, and so contain reward information. The RLSP algorithm considers the current state s_0 as the final state in a human trajectory $\tau = (s_{-H}, a_{-H}, \dots, s_0, a_0)$ generated by the human with preferences θ . The RLSP likelihood $p(\theta | s_0)$ is obtained by marginalizing out $(s_{-H}, a_{-H}, \dots, s_{-1}, a_{-1}, a_0)$, since only s_0 is actually observed.

Inverse Reward Design (IRD). Inverse reward design (IRD) notes that since reward designers often produce rewards by checking what the reward does in the training environment using an iterative trial-and-error process, the final reward they produce is likely to produce good behavior *in the environments that they used during the reward design process*. So, $\tilde{\theta}$ need not be identical to θ^* , but merely provides evidence about it. Formally, IRD assumes that the probability of specifying a given θ is proportional to the exponent of the expected true return in the training environment: $p(\tilde{\theta} | \theta^*) \propto \exp(\theta^{*\top} \mathcal{F}^{\pi_{\tilde{\theta}}})$. IRD inverts this generative model of *reward design* to sample from the posterior of the true reward: $p(\theta^* | \tilde{\theta}) \propto p(\tilde{\theta} | \theta^*)p(\theta^*)$.

3 Uses of the reward distribution

So far, we have argued that we would like to combine the information from two different sources into a single distribution over reward functions. But what makes such a distribution good? We take a pragmatic approach: we identify downstream tasks for such a distribution, find desiderata for a reward distribution for these tasks, and use these desiderata to guide our search for methods.

3.1 Optimizing the reward

We follow the framework of Hadfield-Menell et al. [14] in which an agent maximizes expected return under reward uncertainty. In this setting, maximizing the return causes the agent to use its actions both to obtain (expected) reward, and to gather information about the true reward in order to better obtain reward in the future [14, 31]. We might expect that an agent could first learn the reward and then act; however, this is unrealistic: we would be unhappy with a household robot that had to learn how we want walls to be painted before it would bake a cake. The agent must act given reward uncertainty, but maintain its ability to optimize any of the potential rewards in the future when it will know more about the reward. Put simply, it needs to preserve its *option value* while still obtaining as much reward as it can currently get.

We would like to evaluate our reward distribution by plugging it into an algorithm that maximizes expected reward given reward uncertainty, but unfortunately current methods [14, 31] are computationally expensive and only work in environments where option value preservation is not a concern. So, we analyze reward learning and option value preservation separately.

Active reward learning. An agent that learns rewards while acting will need to select queries with care. *Active reward learning* aims to pick queries that maximize the information gained about the reward, or minimize the expected regret [24, 20]³. We analyse how the reward distributions resulting from different reward combination methods fit as priors for the active reward learning methods.

Option-value preservation. Low impact agents [30, 17, 2] take fewer irreversible actions, and so preserve their ability to optimize multiple reward functions. However, work in this area does not consider the possibility of learning more information about the reward function, and so these methods penalize entire *plans* that would have a high impact. In contrast, we would like our agent to pursue impactful plans, but pause when option value would be destroyed, and get more reward information.

For our purposes, we start with the *attainable utility preservation* (AUP) method introduced by Turner et al. [30]. AUP formalizes the impact of an action on a reward function as the change in the Q-value as a result of taking that action (relative to a no-op action \emptyset). For our purposes, we only care about cases where we are no longer able to optimize the reward function, as opposed to ones where we are *better* able to optimize the reward function. So, we use the truncated difference summary function proposed by Krakovna et al. [17] that only counts *decreases* in Q-values and not increases:

$$\text{Impact}(s, a) = \mathbb{E}_r[\max(Q_r(s, \emptyset) - Q_r(s, a), 0)]. \quad (1)$$

Turner et al. [30] penalize the agent whenever it would cause too much impact, using a hyperparameter λ to trade off task reward with impact. They plan using the reward $R_{\text{task}} - \lambda \text{Impact}(s, a)$. However, by penalizing impact in the reward function, the agent is penalized for any *plan* that would cause impact, whereas we want our agent to start on potentially impactful plans, but stop once a particular *action* would destroy option value. So, we only penalize impact during action selection:

$$\pi(s) = \max_a (Q(s, a) - \lambda \text{Impact}(s, a)). \quad (2)$$

3.2 Acting in a new environment

Often, we wish to use the reward function in the same environment as the one in which it was specified or inferred. For example, a human expert may have demonstrated how to bake a cake in a particular kitchen. Since we know *at training time* that the agent will be deployed in the same kitchen, the only guarantee we need is that the agent’s behavior is correct. If two rewards lead to the same behavior in this kitchen, we don’t care which one we end up using.

However, we may instead want to apply a reward function learned in one environment to a different environment. In this case, there may be new options in the test environment that we never saw during training. As a result, even if two rewards led to identical behavior in the training environment, they need not do so in the test environment. Hence when the goal is to transfer the reward distribution to a new environment, it is crucial to ensure that the distribution contains the true reward, instead of just containing a reward that gives rise to the same behavior as the true one in the training environment.

³We speculate that future active reward learning methods will also have to ensure that the queries are *relevant* and *timely*: in our household robot example, this would prevent the robot from asking about wall paint color when it is meant to be baking a cake.

4 Desiderata for the reward posterior

We discuss properties of the posterior that inform the choice of a reward combination algorithm for different applications of the resulting reward distribution. We focus on two applications: option-value preserving planning and active reward learning.

4.1 Desired properties for any subsequent application

All else equal, we want the true reward to have high probability in the reward distribution. One model for misspecification is that the weight for each feature has some chance of being “corrupted”, and the chance of the feature being corrupted in one reward is independent from that of the other reward.

If the fraction of misspecified values is small, then if two reward vectors’ values are different at a particular index, it is likely that one of the values is correct. Given this, it is desirable for the reward distribution to assign significant probability to every reward function s.t. the value of the weight for each feature is in one of the values of that feature’s weight of the input rewards. Formally, we denote the set of such reward parameters

$$\Theta^{\text{indep pts}} = \{\theta \in \mathbb{R}^n \mid \theta_i \in \{\tilde{\theta}_i^1, \tilde{\theta}_i^2\} \forall i \in \{1, \dots, n\}\}. \quad (3)$$

Desideratum 1 (Support on $\Theta^{\text{indep pts}}$). *All $\theta^* \in \Theta^{\text{indep pts}}$ should be in the support of $p(\theta^* \mid \tilde{\theta}^1, \tilde{\theta}^2)$.*

Note that when we do not need to transfer the reward function to a new environment, Desideratum 1 can be made weaker. We only require that the distribution have support on reward functions that give rise to the same *behavior* as the rewards in $\Theta^{\text{indep pts}}$, and not necessarily on $\Theta^{\text{indep pts}}$ itself.

In order for the agent to act, the reward distribution must provide information about useful behavior:

Desideratum 2 (Informative about desirable behavior). *If $\tilde{\theta}^1$ and $\tilde{\theta}^2$ give rise to the same behavior, most of the reward functions in the distribution should give rise to this behavior.*

The strong version of this desideratum states that if the FEs of $\{\tilde{\theta}^1, \tilde{\theta}^2\}$ are equal on some feature, the FEs of this feature of *any* reward in the support of the posterior are equal to those of $\{\tilde{\theta}^1, \tilde{\theta}^2\}$:

$$\forall i \in \{1, n\} : \left(\mathcal{F}_i^{\pi_{\tilde{\theta}^1}} = \mathcal{F}_i^{\pi_{\tilde{\theta}^2}} \implies \forall \theta^* \in \text{supp}(p(\theta^* \mid \tilde{\theta}^1, \tilde{\theta}^2)) : \mathcal{F}_i^{\pi_{\theta^*}} = \mathcal{F}_i^{\pi_{\tilde{\theta}^1}} = \mathcal{F}_i^{\pi_{\tilde{\theta}^2}} \right). \quad (4)$$

The medium version considers the FEs as a whole instead of per-feature:

$$\mathcal{F}^{\pi_{\tilde{\theta}^1}} = \mathcal{F}^{\pi_{\tilde{\theta}^2}} \implies \forall \theta^* \in \text{supp}(p(\theta^* \mid \tilde{\theta}^1, \tilde{\theta}^2)) : \mathcal{F}^{\pi_{\theta^*}} = \mathcal{F}^{\pi_{\tilde{\theta}^1}} = \mathcal{F}^{\pi_{\tilde{\theta}^2}}. \quad (5)$$

The weak version assumes that the FEs of all $\theta \in \Theta^{\text{indep pts}}$ are the same, rather than just $\tilde{\theta}^1$ and $\tilde{\theta}^2$:

$$(\forall \theta \in \Theta^{\text{indep pts}} : \mathcal{F}^{\pi_{\theta}} = \mathcal{F}^{\pi_{\tilde{\theta}^1}} = \mathcal{F}^{\pi_{\tilde{\theta}^2}}) \implies \forall \theta^* \in \text{supp}(p(\theta^* \mid \tilde{\theta}^1, \tilde{\theta}^2)) : \mathcal{F}^{\pi_{\theta^*}} = \mathcal{F}^{\pi_{\tilde{\theta}^1}}. \quad (6)$$

Des. 2 allows option-value preserving planning to do useful things instead of performing the no-op action all the time. Furthermore, by restricting possible behaviors, Des. 2 simplifies the job of approaches that actively learn the reward function. If the reward distribution used by our household robot always incentivizes using soap while washing dishes (because both input rewards did so), the robot’s active learning algorithm does not need to query the human about whether or not to use soap.

Des. 2 might be problematic if both input rewards agree on an undesirable behavior w.r.t. a particular feature. This presents a natural tradeoff between informativeness and the abilities to preserve option value and actively learn arbitrary reward functions: the more informative we assume the input rewards are about the desired behavior, the less robust we are to their misspecification.

4.2 Desired properties for option value preservation

Suppose we would like to preserve the option value of pursuing two different behaviors. Option-value preservation is more robust when the distribution over FEs induced by the reward distribution places significant and approximately equal weight on the two behaviors. Intuitively, the “weight” AUP puts on preserving expected return attainable by a particular behavior is proportional to the probability of that behavior in the reward distribution. So we would like the FEs of the input rewards to be *balanced* across the reward distribution. Concretely:

Desideratum 3 (Behavior-space balance). *The FEs of the rewards sampled from the distribution should correspond to behaviors arising from the input reward functions with a similar frequency:*

$$p(\mathcal{F}^{\pi_\theta} = \mathcal{F}^{\pi_{\tilde{\theta}^1}}) \approx p(\mathcal{F}^{\pi_\theta} = \mathcal{F}^{\pi_{\tilde{\theta}^2}}) \mid \theta \sim p(\theta \mid \tilde{\theta}^1, \tilde{\theta}^2). \quad (7)$$

4.3 Desired properties for active reward learning

Given a reward vector we define the *tradeoff point* between the i -th and j -th features of the state as the ratio between the reward vector values corresponding to these features: $T_{i,j}^\theta = \frac{\theta_i}{\theta_j}$. The true reward's tradeoff point may be in between the tradeoff points of the input rewards, and so we would like our distribution to support all such tradeoff points.

Desideratum 4 (Support on intermediate feature tradeoffs). *The reward posterior should have support on rewards that prescribe all intermediate tradeoff points between the states' features.*

In practice tradeoff points only affect the behavior to the extent allowed by the environment. We only need to maintain *all* intermediate tradeoff points if we want to transfer to a new environment.

5 Multitask Inverse Reward Design

Inspired by Inverse Reward Design [13], we introduce a reward combination method *Multitask Inverse Reward Design* (MIRD). We assume that the two input rewards meet the IRD assumption for two different rewards, namely that the input rewards generate nearly optimal behavior in the training environment according to the corresponding task rewards. A straightforward way to construct the distribution for the combined reward function is to require the FEs of the rewards in the distribution to be convex combinations of the FEs of the input rewards: $\mathcal{F}^{\pi_{\theta^*}} \in \{b\mathcal{F}^{\pi_{\tilde{\theta}^1}} + (1-b)\mathcal{F}^{\pi_{\tilde{\theta}^2}} \mid b \in [0, 1]\}$.

To satisfy this requirement, we propose a distribution that depends on the behavior (set of trajectories) $D = \{\tau^1, \dots, \tau^d\}$ resulting from the agents optimizing $\tilde{\theta}^1$ and $\tilde{\theta}^2$:

$$p(\theta^* \mid \tilde{\theta}^1, \tilde{\theta}^2) = \int_D p(\theta^* \mid D) p(D \mid \tilde{\theta}^1, \tilde{\theta}^2) dD \quad (8)$$

Each of the reward functions $\{\theta^i \mid i = \{1, \dots, d\}\}$ used to generate the trajectories τ^i in D is sampled from the Bernoulli mixture over the input reward functions. We introduce a random variable $b \in [0, 1]$ that determines the probability of sampling $\theta^i = \tilde{\theta}^1$ to generate the trajectory τ^i :

$$p(D \mid \tilde{\theta}^1, \tilde{\theta}^2) = \int_b \text{Beta}(b \mid \beta_1, \beta_2) \prod_{i=1}^n \left(bp(\tau^i \mid \tilde{\theta}^1) + (1-b)p(\tau^i \mid \tilde{\theta}^2) \right) db \quad (9)$$

We sample from $p(\tau^i \mid \theta^i)$ using policies arising from soft value iteration. For $p(\theta^* \mid D)$, we set θ^* to the reward learned by Maximum Causal Entropy Inverse RL (MCEIRL) [33]. Denoting the Dirac delta as $\delta(\cdot)$, we have $p(\theta^* \mid D) = \delta(\theta^* - \arg\max_{\theta} p(\theta \mid D))$. Having defined all relevant conditional distributions, we can sample from the joint $p(\theta^*, b, D, \theta^i \mid \tilde{\theta}^1, \tilde{\theta}^2)$ and the marginal $p(\theta^* \mid \tilde{\theta}^1, \tilde{\theta}^2)$.

Theorem 1. *The FEs arising from rewards θ^* in the support of the MIRD posterior are convex combinations of FEs of the input reward vectors $\tilde{\theta}^1$ and $\tilde{\theta}^2$:*

$$\forall \theta^* \in \text{supp}(p(\theta^* \mid \tilde{\theta}^1, \tilde{\theta}^2)) : \mathcal{F}^{\pi_{\theta^*}} = \alpha \mathcal{F}^{\pi_{\tilde{\theta}^1}} + (1-\alpha) \mathcal{F}^{\pi_{\tilde{\theta}^2}}, \alpha \in [0, 1]. \quad (10)$$

Proof. See appendix A for full proof. Key to the proof is the fact that MCEIRL finds such θ^* that its FEs match those of the soft value iteration expert who generated the trajectories [33]. In our case, the FEs of θ^* match D , which contains trajectories arising from $\tilde{\theta}^1$ and $\tilde{\theta}^2$ in various proportions.

Corollary 1.1. *The average FEs of the reward functions in the posterior are:*

$$\mathbb{E}_{\theta \sim p(\theta \mid \tilde{\theta}^1, \tilde{\theta}^2)} [\mathcal{F}^{\pi_\theta}] = \mathbb{E}_b [b\mathcal{F}^{\pi_{\tilde{\theta}^1}} + (1-b)\mathcal{F}^{\pi_{\tilde{\theta}^2}}] = \mathbb{E}_b [b] \mathcal{F}^{\pi_{\tilde{\theta}^1}} + (1 - \mathbb{E}_b [b]) \mathcal{F}^{\pi_{\tilde{\theta}^2}}. \quad (11)$$

Corollary 1.2. *Eq. 11 implies that if the FEs of the $\tilde{\theta}^1$ and $\tilde{\theta}^2$ w.r.t. a given feature are equal, the FEs of all rewards in the distribution for that feature equal the FEs of $\tilde{\theta}^1$ and $\tilde{\theta}^2$ for that feature. Thus MIRD satisfies the strong version of Des. 2 (informative about desirable behavior).*

Corollary 1.3. When $\beta_1 = \beta_2$ the probability of sampling $\theta^i = \tilde{\theta}^1$ equals the probability of sampling $\theta^i = \tilde{\theta}^2$, and hence the probabilities of observing behaviors $\mathcal{F}^{\pi_{\tilde{\theta}^1}}$ and $\mathcal{F}^{\pi_{\tilde{\theta}^2}}$ when optimizing a reward function from the posterior are equal. Hence Des. 3 (behavior balance) is satisfied.

Theorem 1 allows us to bound the expected return $\theta^{\text{true}\top} \mathcal{F}^{\pi_{\theta}}$, $\theta \in \text{supp}(p(\theta^* | \tilde{\theta}^1, \tilde{\theta}^2))$ of the true reward function θ^{true} in terms of expected true returns of $\mathcal{F}^{\pi_{\tilde{\theta}^1}}$ and $\mathcal{F}^{\pi_{\tilde{\theta}^2}}$:

Theorem 2. The minimum expected true return following any of the rewards in the MIRD posterior equals the minimum true return following one of the input rewards:

$$\min_{\theta \in \text{supp}(p(\theta^* | \tilde{\theta}^1, \tilde{\theta}^2))} \theta^{\text{true}\top} \mathcal{F}^{\pi_{\theta}} = \min_{\theta \in \{\tilde{\theta}^1, \tilde{\theta}^2\}} \theta^{\text{true}\top} \mathcal{F}^{\pi_{\theta}}. \quad (12)$$

Proof. See Appendix B.

Theorem 2 is the best possible regret bound for following rewards in a distribution resulting from reward combination: the worst the agent can do when following a reward from the distribution is no worse than when just randomly choosing to follow one of the input reward functions.

Generally FEs arising from $\theta \in \Theta^{\text{indep pts}} \setminus \{\tilde{\theta}^1, \tilde{\theta}^2\}$ are not convex combinations of $\mathcal{F}^{\pi_{\tilde{\theta}^1}}$ and $\mathcal{F}^{\pi_{\tilde{\theta}^2}}$, so MIRD does not satisfy any version of Des. 1 (support on $\Theta^{\text{indep pts}}$). In addition, when $\tilde{\theta}^1$ and $\tilde{\theta}^2$ conflict on more than one feature, Des. 4 (support on intermediate feature tradeoffs) is not satisfied.

Independent features formulation. Since both Des. 1 and Des. 4 aim to ensure robustness to reward misspecification, one of our primary goals, we would like a variant of MIRD that satisfies them. An algorithm that supports rewards that give rise to all potential behaviors arising from reward functions in $\Theta^{\text{indep pts}}$ is a simple modification of MIRD. The only change required in the original formulation of MIRD is sampling the reward vectors $\{\theta^i\}$ giving rise to the trajectories in D from $\Theta^{\text{indep pts}}$ instead of just $\{\tilde{\theta}^1, \tilde{\theta}^2\}$. Now b is a vector of dimensionality $|\Theta^{\text{indep pts}}|$, sampled from Dirichlet($b | \beta_1, \dots, \beta_{|\Theta^{\text{indep pts}}|}$). Each θ^i is sampled from a Multinoulli distribution over $\Theta^{\text{indep pts}}$, parameterized by b . We refer to this version of MIRD as MIRD independent features, or MIRD-IF.

6 Analysis

We first introduce several baselines each of which does not meet some of the outlined desiderata, and then analyze all reward combination methods in the context of our modified version of AUP.

6.1 Reward-space baselines

Additive. The additive reward combination method is commonly used in previous work utilizing reward combination [8, 27, 16]. The reward posterior is simply $p(\theta^* | \tilde{\theta}^1, \tilde{\theta}^2) = \delta(\theta^* - (\tilde{\theta}^1 + \alpha \tilde{\theta}^2))$, where α is the hyperparameter used to balance the two input reward functions.

Gaussian. Here, we use the standard Bayesian framework, and assume that $\tilde{\theta}^1$ and $\tilde{\theta}^2$ are generated from θ^* with Gaussian noise:

$$p(\theta^* | \tilde{\theta}^1, \tilde{\theta}^2) \propto p(\theta^*) p(\tilde{\theta}^1 | \theta^*) p(\tilde{\theta}^2 | \theta^*) = p(\theta^*) p_{\mathcal{N}}(\tilde{\theta}^1 | \theta^*, \sigma_1^2 \mathbf{I}) p_{\mathcal{N}}(\tilde{\theta}^2 | \theta^*, \sigma_2^2 \mathbf{I}),$$

where \mathbf{I} is the identity matrix, $p_{\mathcal{N}}$ is the multivariate Gaussian probability density function, and $\{\sigma_1, \sigma_2\}$ are hyperparameters controlling the standard deviations. One can express higher trust in a given input reward by lowering its σ . For our experiments, we use an uninformative Gaussian prior (that is, $\sigma \rightarrow \infty$) and set $\sigma_1 = \sigma_2 = 1$. The posterior is then $p(\theta^* | \tilde{\theta}^1, \tilde{\theta}^2) = p_{\mathcal{N}}(\theta^* | \frac{\tilde{\theta}^1 + \tilde{\theta}^2}{2}, \frac{1}{2} \mathbf{I})$.

Convex combinations of input reward vectors (CC-in). Consider a scenario where the input reward functions are correctly specified for two different tasks. A natural reward-space analogue to MIRD would be to consider a distribution over all convex combinations of the input reward vectors:

$$p(\theta^* | \tilde{\theta}^1, \tilde{\theta}^2) = \int_b \text{Beta}(b | \beta_1, \beta_2) \delta(\theta^* - b\tilde{\theta}^1 - (1-b)\tilde{\theta}^2) db. \quad (13)$$

Theorem 3. If $\mathcal{F}^{\pi_{\tilde{\theta}^1}} = \mathcal{F}^{\pi_{\tilde{\theta}^2}}$, the FEs of any reward in the support of the CC-in posterior equal $\mathcal{F}^{\pi_{\tilde{\theta}^1}}$. Hence CC-in satisfies the medium version of Des. 2 (informative about desirable behavior).

Proof. See Appendix C.

We provide an example showing that CC-in does not satisfy the strong version of Des 2 in Appendix F. Furthermore, CC-in does not satisfy Des. 1 (support on $\Theta^{\text{indep pts}}$): consider combining rewards $[0, 1, 1]$ and $[1, 0, 1]$. Here $\theta = [0, 0, 1] \in \Theta^{\text{indep pts}}$, but is not a convex combination of input rewards.

Convex combinations of $\theta \in \Theta^{\text{indep pts}}$ (CC- $\Theta^{\text{indep pts}}$). Instead of only considering convex combinations of the input rewards, we consider convex combinations of $\theta \in \Theta^{\text{indep pts}}$. The vector $b \in \mathbb{R}^{|\Theta^{\text{indep pts}}|}$ of weights for each convex combination is sampled from the Dirichlet distribution:

$$p(\theta^* \mid \tilde{\theta}^1, \tilde{\theta}^2) = \int_b \text{Dirichlet}(b \mid \beta_1, \dots, \beta_{|\Theta^{\text{indep pts}}|}) \delta \left(\theta^* - \sum_{i=0}^{|\Theta^{\text{indep pts}}|} b_i \Theta_i^{\text{indep pts}} \right) db. \quad (14)$$

Theorem 4. *If all FEs of rewards in $\Theta^{\text{indep pts}}$ are equal, the FEs of any reward in the CC- $\Theta^{\text{indep pts}}$ posterior equal the FEs of rewards in $\Theta^{\text{indep pts}}$. Hence CC- $\Theta^{\text{indep pts}}$ satisfies Des. 2 (weak).*

Proof. See Appendix D.

Uniform points. Note that as $\beta_1 = \dots = \beta_{|\Theta^{\text{indep pts}}|} \rightarrow 0$, Eq. 14 essentially becomes $p(\theta^* \mid \tilde{\theta}^1, \tilde{\theta}^2) = \text{Uniform}(\Theta^{\text{indep pts}})$. It is helpful to analyze the properties of this *uniform points* distribution separately.

Table 1 summarizes the properties of MIRD, MIRD-IF, and the reward-space baselines.

Table 1: Performance of reward combination algorithms on the desiderata outlined in Section 4.

	Support on $\Theta^{\text{indep pts}}$	Informative about desirable behavior			Behavior balance	Support on intermediate feature tradeoffs
		(strong)	(medium)	(weak)		
Additive	✗	✗	✓	✓	✗	✗
Gaussian	✓	✗	✗	✗	✗	✓
CC-in	✗	✗	✓	✓	\approx	✗
CC- $\Theta^{\text{indep pts}}$	✓	✗	✗	✓	\approx	✓
Uniform Points	✓	✗	✗	✓	✓	✗
MIRD	✗	✓	✓	✓	✓	✗
MIRD-IF	✓	✗	✗	✓	✓	✓

6.2 Analysis on a toy environment

To demonstrate the use of the desiderata, as well as how each method performs on these desiderata, we developed a toy environment in which to test reward combination methods. Due to space limitations, we defer full qualitative and quantitative explanations to Appendix E and report the broad results here.

Setup. We use the setting of AUP planning to demonstrate the importance of support on $\Theta^{\text{indep pts}}$, informativeness about desired the behavior, and behavior balance desiderata. We obtain the reward posterior by combining the input rewards using each of our seven reward combination methods, and qualitatively evaluate the behavior of AUP planning w.r.t. the posterior over five seeds. The input rewards are either both specified by hand, or $\tilde{\theta}^1$ is learned with RLSP while $\tilde{\theta}^2$ is hand-specified. We use the *cooking* environment described in Figure 2. The states’ features indicate the number of jars of flour, pieces of dough, cakes, avocado toasts, and servings of pasta. At each state the agent can move, cook, or do nothing. Cooking is possible when the agent is to the left of flour or dough.

Support on $\Theta^{\text{indep pts}}$. We specify both input rewards to strongly incentivise either cake or toast by setting $\tilde{\theta}_{\text{cake}}^1 = \tilde{\theta}_{\text{toast}}^2 = 3$; $\tilde{\theta}_{\text{toast}}^1 = \tilde{\theta}_{\text{cake}}^2 = 0$, and to mildly incentivise pasta: $\tilde{\theta}_{\text{pasta}}^1 = \tilde{\theta}_{\text{pasta}}^2 = 1$. The desired behavior is to preserve the flour, as any one of the foods might be desirable. This requires the distribution to place weight on rewards




Figure 2: The Cooking environment. Left: the grid-world layout and current state. Right: irreversible actions for cooking. Making cake or avocado toast involves first making dough from the flour, and then making cake or toast; making pasta uses flour directly.

such as $\theta_{\text{cake}} = \theta_{\text{toast}} = 0$; $\theta_{\text{pasta}} = 1$ which would make pasta. As expected, methods that have support on $\Theta^{\text{indep pts}}$ behave as desired, and other methods do not.

Informativeness about desirable behavior (strong). In the setting above the two input rewards lead the agent to make either cake or toast, but not pasta. So, a reward distribution meeting the strong version of Des. 2 would only make cake or toast, and so the agent should make dough (which is useful for both). Note that despite the same setup, this is exactly the opposite behavior than desired previously, showing the tradeoff between informativeness and support on $\Theta^{\text{indep pts}}$. As such, the results are reversed.

Informativeness about desirable behavior (weak). We hand-specify one of the reward functions to incentivise the agent to make cakes: $\tilde{\theta}_{\text{cake}}^2 = 1$, $\tilde{\theta}_{\text{toast}}^2 = \tilde{\theta}_{\text{pasta}}^2 = 0$. $\tilde{\theta}^1$ is inferred by RLSP with a uniform prior over s_{-T} . Since the current state contains a cake, RLSP infers a positive reward for cakes, $\tilde{\theta}_{\text{cake}}^1 \approx 0.1$, and near-zero rewards for the other features. Hence we have two rewards that both incentivise the agent to make another cake. The desired behavior is to make the cake, and not worry about preserving the ability to make pasta or toast as neither input reward cares about it. All but the Gaussian method satisfy the weak informativeness desideratum, and so they all succeed. The Gaussian method fails since it thinks θ_{cake}^* could be negative.

Behavior balance. $\tilde{\theta}^2$ is specified to only reward toast, $\tilde{\theta}_{\text{toast}}^2 = 1$, $\tilde{\theta}_{\text{cake}}^2 = \tilde{\theta}_{\text{pasta}}^2 = 0$, and $\tilde{\theta}^1$ is specified only reward cake: $\tilde{\theta}_{\text{cake}}^1 > 0$, $\tilde{\theta}_{\text{toast}}^1 = \tilde{\theta}_{\text{pasta}}^1 = 0$. So we have two disagreeing rewards: one incentivises cake, while the other incentivises toast. The desired behavior from AUP here is to make dough as it is desirable for both the cake and the toast, but then to stop as the next step is unclear. Ensuring that the reward for toast does not overwhelm the reward for cake requires behavior balance. We vary $\tilde{\theta}_{\text{cake}}^1$ between 0.01 and 0.7 to obtain many degrees of imbalance between toast and cake. MIRD and MIRD-IF have the lowest frequency of the toast reward overwhelming the cake reward, which indicates that MIRD and MIRD-IF are most behavior balanced. However, occasionally these methods do nothing instead of making dough.

7 Related Work

Utility aggregation. Combining the specified reward functions is in many ways similar to combining utility functions of different agents. Harsanyi’s aggregation theorem [15] suggests that maximizing a fixed linear combination of utility functions results in Pareto-optimal outcomes for correctly specified utility functions. This corresponds to our Additive baseline.

Multi-task and meta IRL. Multi-task IRL assumes that the demonstrations originate from multiple experts performing different tasks, and seeks to recover reward functions that explain the preferences of each of the experts [4, 5]. The IRL part of MIRD and MIRD-IF is conceptually similar to these approaches: the trajectories in each demonstration set D can be seen as arising from an expert executing different tasks. However, MIRD seeks to explain each D as if the trajectories arose from a *single* reward function. This resembles meta IRL, which learns an initialization for the reward function parameters that is helpful for learning the rewards of new tasks [32, 12].

8 Limitations and future work

Summary. We analyze the problem of combining two potentially misspecified reward functions into a distribution over rewards. We identify active reward learning and option value preservation as two key applications for such a distribution, and determine four properties that the distribution would ideally satisfy for these applications. We suggest some simple reward-space methods as well as a behavior-space method, Multitask Inverse Reward Design (MIRD), which is grounded in the behavior arising from the input rewards. MIRD works well for most applications, and MIRD-IF can be used for greater robustness and increased likelihood that the true reward is in the distribution. If the only application is to preserve option value, then the uniform points method also works well.

Future work. The primary avenue for future work is to evaluate the reward combination methods in the context of active reward learning, and formalize Des. 4 (support on intermediate tradeoffs) in a way most helpful for active learning. Another avenue for future work is scaling up MIRD to realistic environments, where the reward function may be nonlinear.

References

- [1] Amodei, D., Olah, C., Steinhardt, J., Christiano, P., Schulman, J., and Mané, D. (2016). Concrete problems in ai safety. *arXiv preprint arXiv:1606.06565*.
- [2] Armstrong, S. and Levinstein, B. (2017). Low impact artificial intelligences. *arXiv preprint arXiv:1705.10720*.
- [3] Armstrong, S. and Mindermann, S. (2018). Occam’s razor is insufficient to infer the preferences of irrational agents. In *Advances in Neural Information Processing Systems*, pages 5598–5609.
- [4] Babes, M., Marivate, V., Subramanian, K., and Littman, M. L. (2011). Apprenticeship learning about multiple intentions. In *Proceedings of the 28th International Conference on Machine Learning (ICML-11)*, pages 897–904.
- [5] Choi, J. and Kim, K.-E. (2012). Nonparametric bayesian inverse reinforcement learning for multiple reward functions. In *Advances in Neural Information Processing Systems*, pages 305–313.
- [6] Christiano, P. F., Leike, J., Brown, T., Martic, M., Legg, S., and Amodei, D. (2017). Deep reinforcement learning from human preferences. In *Advances in Neural Information Processing Systems*, pages 4299–4307.
- [7] Daniel, C., Viering, M., Metz, J., Kroemer, O., and Peters, J. (2014). Active reward learning. In *Robotics: Science and Systems*.
- [8] Desai, N., Critch, A., and Russell, S. J. (2018). Negotiable reinforcement learning for pareto optimal sequential decision-making. In *Advances in Neural Information Processing Systems*, pages 4713–4721.
- [9] Evans, O., Stuhlmüller, A., and Goodman, N. D. (2016). Learning the preferences of ignorant, inconsistent agents. In *AAAI*, pages 323–329.
- [10] Frazier, D. T., Robert, C. P., and Rousseau, J. (2017). Model misspecification in abc: Consequences and diagnostics. *arXiv preprint arXiv:1708.01974*.
- [11] Fu, J., Luo, K., and Levine, S. (2018). Learning robust rewards with adversarial inverse reinforcement learning. In *International Conference on Learning Representations*.
- [12] Gleave, A. and Habryka, O. (2018). Multi-task maximum entropy inverse reinforcement learning. *arXiv preprint arXiv:1805.08882*.
- [13] Hadfield-Menell, D., Milli, S., Abbeel, P., Russell, S. J., and Dragan, A. (2017). Inverse reward design. In *Advances in Neural Information Processing Systems*, pages 6765–6774.
- [14] Hadfield-Menell, D., Russell, S. J., Abbeel, P., and Dragan, A. (2016). Cooperative inverse reinforcement learning. In *Advances in neural information processing systems*, pages 3909–3917.
- [15] Harsanyi, J. C. (1955). Cardinal welfare, individualistic ethics, and interpersonal comparisons of utility. *Journal of political economy*, 63(4):309–321.
- [16] Kolter, J. Z. and Ng, A. Y. (2009). Near-bayesian exploration in polynomial time. In *Proceedings of the 26th Annual International Conference on Machine Learning*, pages 513–520. ACM.
- [17] Krakovna, V., Orseau, L., Kumar, R., Martic, M., and Legg, S. (2018). Penalizing side effects using stepwise relative reachability. *arXiv preprint arXiv:1806.01186*.
- [18] MacGlashan, J., Ho, M. K., Loftin, R., Peng, B., Wang, G., Roberts, D. L., Taylor, M. E., and Littman, M. L. (2017). Interactive learning from policy-dependent human feedback. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 2285–2294. JMLR.org.
- [19] Majumdar, A., Singh, S., Mandlekar, A., and Pavone, M. (2017). Risk-sensitive inverse reinforcement learning via coherent risk models. In *Robotics: Science and Systems*.
- [20] Mindermann, S., Shah, R., Gleave, A., and Hadfield-Menell, D. (2019). Active inverse reward design. In *International Conference on Machine Learning*.

- [21] OpenAI (2018). Openai five. <https://blog.openai.com/openai-five/>.
- [22] Radlinski, F., Balog, K., Byrne, B., and Krishnamoorthi, K. (2019). Coached conversational preference elicitation: A case study in understanding movie preferences. In *Proceedings of the Annual SIGdial Meeting on Discourse and Dialogue*.
- [23] Ramachandran, D. and Amir, E. (2007). Bayesian inverse reinforcement learning. In *IJCAI*, volume 7, pages 2586–2591.
- [24] Sadigh, D., Dragan, A. D., Sastry, S., and Seshia, S. A. (2017). Active preference-based learning of reward functions. In *Robotics Science and Systems*.
- [25] Samuelson, P. A. (1938). A note on the pure theory of consumer’s behaviour. *Economica*, 5(17):61–71.
- [26] Shah, R., Gundotra, N., Abbeel, P., and Dragan, A. (2019a). On the feasibility of learning, rather than assuming, human biases for reward inference. In *International Conference on Machine Learning*, pages 5670–5679.
- [27] Shah, R., Krashennnikov, D., Alexander, J., Dragan, A., and Abbeel, P. (2019b). Preferences implicit in the state of the world. In *International Conference on Learning Representations*.
- [28] Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., van den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., Dieleman, S., Grewe, D., Nham, J., Kalchbrenner, N., Sutskever, I., Lillicrap, T., Leach, M., Kavukcuoglu, K., Graepel, T., and Hassabis, D. (2016). Mastering the game of go with deep neural networks and tree search. *Nature*, 529:484–503.
- [29] Steinhardt, J. and Evans, O. (2017). Model mis-specification and inverse reinforcement learning.
- [30] Turner, A. M., Hadfield-Menell, D., and Tadepalli, P. (2019). Conservative agency via attainable utility preservation. *arXiv preprint arXiv:1902.09725*.
- [31] Woodward, M., Finn, C., and Hausman, K. (2019). Learning to interactively learn and assist. *arXiv preprint arXiv:1906.10187*.
- [32] Xu, K., Ratner, E., Dragan, A., Levine, S., and Finn, C. (2018). Learning a prior over intent via meta-inverse reinforcement learning. *arXiv preprint arXiv:1805.12573*.
- [33] Ziebart, B. D. (2010). Modeling purposeful adaptive behavior with the principle of maximum causal entropy. *PhD thesis, Carnegie Mellon University*.

A Proof of Theorem 1

D is a set of trajectories each of which arises either from optimizing $\tilde{\theta}^1$ or optimizing $\tilde{\theta}^2$. The fraction of trajectories in D that result from $\tilde{\theta}^1$ is b , and the fraction of trajectories arising from $\tilde{\theta}^2$ is $1 - b$. Hence the FEs of D are

$$\mathcal{F}^D = b\mathcal{F}^{\pi_{\tilde{\theta}^1}} + (1 - b)\mathcal{F}^{\pi_{\tilde{\theta}^2}}. \quad (15)$$

$p(\theta^* \mid D) = \delta(\theta^* - \arg\max_{\theta} p(\theta \mid D))$ entails that each sample θ^* is the MCEIRL reward $\theta^* = \arg\max_{\theta} p(\theta \mid D)$. MCEIRL finds a reward vector θ^* such that its FEs match those of the soft value iteration expert who generated the trajectories [33]. Hence each sampled θ^* gives rise to FEs $\mathcal{F}^{\pi_{\theta^*}}$ matching D . This and Eq. 15 together imply that the FEs arising from optimizing any $\theta^* \in \text{supp}(p(\theta^* \mid \tilde{\theta}^1, \tilde{\theta}^2))$ with soft value iteration are a convex combination of the FEs of the input rewards:

$$\forall \theta^* \in \text{supp}(p(\theta^* \mid \tilde{\theta}^1, \tilde{\theta}^2)) : \mathcal{F}^{\pi_{\theta^*}} = b\mathcal{F}^{\pi_{\tilde{\theta}^1}} + (1 - b)\mathcal{F}^{\pi_{\tilde{\theta}^2}}, \quad b \in [0, 1]. \quad \blacksquare \quad (16)$$

B Proof of Theorem 2

By Theorem 1, the FEs arising from any θ in the posterior are a convex combination of the FEs of the input rewards. Hence finding θ in the support of $p(\theta^* \mid \tilde{\theta}^1, \tilde{\theta}^2)$ that minimizes $\theta^{\text{true}\top} \mathcal{F}^{\pi_{\theta}}$ is equivalent to finding $\alpha \in [0, 1]$ that minimizes

$$\theta^{\text{true}\top} (\alpha\mathcal{F}^{\pi_{\tilde{\theta}^1}} + (1 - \alpha)\mathcal{F}^{\pi_{\tilde{\theta}^2}}) = \alpha\theta^{\text{true}\top} \mathcal{F}^{\pi_{\tilde{\theta}^1}} + (1 - \alpha)\theta^{\text{true}\top} \mathcal{F}^{\pi_{\tilde{\theta}^2}}.$$

Since the true returns when following policies arising from the input reward functions $\theta^{\text{true}\top} \mathcal{F}^{\pi_{\tilde{\theta}^1}}$ and $\theta^{\text{true}\top} \mathcal{F}^{\pi_{\tilde{\theta}^2}}$ are scalars, the term above is minimized with $\alpha = 1$ when $\theta^{\text{true}\top} \mathcal{F}^{\pi_{\tilde{\theta}^1}} < \theta^{\text{true}\top} \mathcal{F}^{\pi_{\tilde{\theta}^2}}$, and $\alpha = 0$ otherwise. \blacksquare

C Proof of Theorem 3

Since the feature expectations (FEs) of the input rewards are equal, the return of any mixture reward $b\tilde{\theta}^1 - (1 - b)\tilde{\theta}^2$ for those FEs is $\mathcal{F}^{\pi_{\tilde{\theta}^1}}(b\tilde{\theta}^1 - (1 - b)\tilde{\theta}^2)$. Suppose there exist FEs \mathcal{F}' that result in a higher expected return for the mixture reward:

$$\mathcal{F}'(b\tilde{\theta}^1 - (1 - b)\tilde{\theta}^2) > \mathcal{F}^{\pi_{\tilde{\theta}^1}}(b\tilde{\theta}^1 - (1 - b)\tilde{\theta}^2) \quad (17)$$

For this to be true \mathcal{F}' must result in a higher return than $\mathcal{F}^{\pi_{\tilde{\theta}^1}}$ for either $\tilde{\theta}^1$ or $\tilde{\theta}^2$. This is a contradiction, since $\mathcal{F}^{\pi_{\tilde{\theta}^1}}$ already results in the highest return for both input rewards. Hence $\mathcal{F}^{\pi_{\tilde{\theta}^1}}$ corresponds to the optimal behavior for any mixture reward $b\tilde{\theta}^1 - (1 - b)\tilde{\theta}^2 \forall b \in [0, 1]$. \blacksquare

D Proof of Theorem 4

Analogous to the proof of Theorem 3 (Appendix C).

E Analysis details

In this appendix we provide more details and explanations regarding the behavior of the various reward combination methods in the Cooking environment.

E.1 Support on $\Theta^{\text{indep pts}}$

We hand-specify both input reward functions to strongly incentivise either cake or toast by setting $\tilde{\theta}_{\text{cake}}^1 = \tilde{\theta}_{\text{toast}}^2 = 3$; $\tilde{\theta}_{\text{toast}}^1 = \tilde{\theta}_{\text{cake}}^2 = 0$, and to mildly incentivise pasta: $\tilde{\theta}_{\text{pasta}}^1 = \tilde{\theta}_{\text{pasta}}^2 = 1$. The desired behavior from AUP here is to preserve the flour, as any one of the foods requiring it might be desirable.

The Additive method returns a single reward function that incentivises the agent to make both cake and toast. AUP applied to one reward function simply optimizes it without needing to preserve option value, so here it makes either cake or toast (and does not preserve the flour).

The CC-in method does not preserve the flour, choosing to make dough instead. Because of the high values of $\tilde{\theta}_{\text{cake}}^1$ and $\tilde{\theta}_{\text{toast}}^2$, optimizing any convex combination of $\{\tilde{\theta}^1, \tilde{\theta}^2\}$ leads the agent to either make cake or toast, and so AUP infers that it does not need to preserve the ability to make pasta. MIRD chooses to make dough for a similar reason: since all trajectories in D arise from optimizing $\tilde{\theta}^1$ or $\tilde{\theta}^2$, each trajectory demonstrates the agent making either cake or toast, and no trajectory demonstrates it making pasta. So no reward inferred from D incentivises the agent to make pasta, and hence AUP does not preserve the ability to make pasta.

All methods whose reward samples sometimes give rise to the behaviors arising from $\theta \in \Theta^{\text{indep pts}}$ preserve the flour. This is because $\Theta^{\text{indep pts}}$ contains θ s.t. optimizing it leads the agent to make pasta: $\theta_{\text{cake}} = \theta_{\text{toast}} = 0$; $\theta_{\text{pasta}} = 1$. As MIRD-IF, Gaussian, uniform points, and CC- $\Theta^{\text{indep pts}}$ methods all have substantial probability mass on rewards that incentivise making pasta, AUP applied to the reward samples from these methods preserves the ability to make pasta.

E.2 Informativeness about desirable behavior (strong).

In the setting above the two input rewards lead the agent to make either cake or toast, but not pasta. Here a reward posterior meeting the strong version of Des. 2 would only give rise to reward functions that incentivise the agent to make either cake or toast. Note that despite the same setup, this is exactly the opposite behavior than desired previously, showing the tradeoff between informativeness and support on $\Theta^{\text{indep pts}}$. As such, the results are reversed.

We see that MIRD-IF, Gaussian, uniform points, and CC- $\Theta^{\text{indep pts}}$ methods do not meet Des. 2 (strong), since they all have substantial probability mass on rewards that incentivise making pasta. Because of this AUP preserves the ability to make pasta, and does not lead the agent to make dough. On the other hand, MIRD, Additive, and CC-in methods never give rise to rewards incentivising the agent to make pasta, and AUP applied to these reward distributions results in making dough. Although Additive and CC-in do not support reward functions that lead the agent to make pasta in this particular case, they do not meet the strong version of Des. 2 (shown in Appendix F).

E.3 Informativeness about desirable behavior (medium & weak)

We hand-specify one of the reward functions to incentivise the agent to make cakes: $\tilde{\theta}_{\text{cake}}^2 = 1$, $\tilde{\theta}_{\text{toast}}^2 = \tilde{\theta}_{\text{pasta}}^2 = 0$. $\tilde{\theta}^1$ is inferred by RLSP with a uniform prior over s_{-T} . The current state contains a cake which was likely made by a human, so RLSP infers a positive reward for cakes, $\tilde{\theta}_{\text{cake}}^1 \approx 0.1$, and near-zero rewards for the other features. Hence we have two agreeing rewards that both incentivise the agent to make another cake. The desired behavior from AUP planning here is to make the cake, and not worry about preserving the ability to make pasta or toast as neither input reward cares about it.

The Gaussian posterior fails to make the cake. The distribution over θ_{cake}^* here is a Gaussian with $\mu \approx 0.55$ and $\sigma \approx 0.71$. The posterior assigns substantial probability to θ_{cake}^* being negative, and AUP planning avoids making the cake to preserve the ability to optimize rewards that punish cakes.

Here each reward vector in $\Theta^{\text{indep pts}}$ leads the agent to make cake. Since every method except Gaussian satisfies at least the weak version of Des. 2, all reward functions in the support of their posteriors lead the agent to make cake. Because of this AUP infers that it is not important to preserve the ability to make pasta or toast, and successfully makes cakes when applied to rewards sampled from these posteriors.

E.4 Behavior balance

To better understand robustness of the reward combination methods to imbalanced input rewards, we analyze the performance of our methods by varying the extent of reward imbalance and the degree to which AUP seeks to preserve option value.

The second input reward $\tilde{\theta}^2$ is hand-specified to only reward toast, $\tilde{\theta}_{\text{toast}}^2 = 1, \tilde{\theta}_{\text{cake}}^2 = \tilde{\theta}_{\text{pasta}}^2 = 0$. The first reward $\tilde{\theta}^1$ is hand-specified as well: $\tilde{\theta}_{\text{toast}}^1 = \tilde{\theta}_{\text{pasta}}^1 = 0$. We vary $\tilde{\theta}_{\text{cake}}^1$ between 0.01 and 0.7 to obtain many degrees of imbalance between toast and cake. So the input rewards disagree whether to make toast or cake, and agree that pasta is irrelevant. Hence the desirable action from AUP here would be to preserve the ability to choose between cake and toast, as either could be desirable. In addition to varying $\tilde{\theta}_{\text{cake}}^1$, we vary the weight of the AUP penalty λ .




Figure 3: Most common outcomes of AUP applied to the six reward posteriors over 5 seeds. Pink corresponds to AUP failure for more than half of the seeds: $\tilde{\theta}_{\text{toast}}^2$ overwhelms $\tilde{\theta}_{\text{cake}}^1$ and the agent makes toast. Purple corresponds to AUP success (AUP makes dough) for more than half of the seeds. Yellow corresponds to the middle-ground outcome of doing nothing in more than half of the seeds. The hyperparameters β of the Beta and the Dirichlet distributions are set to 0.5.

Additive always fails and makes toast as it has support on a single reward function that prioritizes toast. Our results for the other six reward combination methods are shown in Figure 3. We observe that MIRD and MIRD-IF are significantly more robust to imbalanced input rewards than any reward-space method. However, sometimes MIRD and MIRD-IF lead to AUP doing nothing instead of making dough. This is explained by the fact that MIRD and MIRD-IF use MCEIRL, which sometimes outputs negative rewards. For example, consider the case where in the process of generating a MIRD sample we generate 10 trajectories all of which demonstrate making toast. MCEIRL with these trajectories might output a positive reward on toast and a *negative* reward on cake and dough (instead of positive reward for toast and zero for cake and dough), even if a zero-centered Gaussian prior is used. These negative rewards on dough sometimes lead AUP to avoid making dough.

Furthermore, we observe that MIRD is more robust than MIRD-IF, and CC-in is more robust than CC- $\Theta^{\text{indep pts}}$. This is because in this case support on rewards giving rise to behaviors arising from $\theta \in \Theta^{\text{indep pts}}$ makes the reward distributions less balanced. In particular, denoting the value of $\tilde{\theta}_{\text{cake}}^1$ we vary as x ,

$$\Theta_{\text{toast, cake}}^{\text{indep pts}} = \left\{ \begin{array}{l} \theta_{\text{toast, cake}} = [0, 0], \\ \theta_{\text{toast, cake}} = [0, x], \\ \theta_{\text{toast, cake}} = [1, 0], \\ \theta_{\text{toast, cake}} = [1, x] \end{array} \right\}.$$

Here only the second reward vector incentivises the agent to make cake, while both the third and the fourth rewards incentivise the agent to make toast. This slight imbalance of $\Theta^{\text{indep pts}}$ in favor of making toast lead to CC- $\Theta^{\text{indep pts}}$ and MIRD-IF posteriors being more likely to cause AUP to make toast instead of preserving the ability to choose between cake and toast.

Finally, we observe that the Gaussian method is roughly equally robust to reward imbalance across the whole range of $\tilde{\theta}_{\text{cake}}^1$. This is because unlike the other methods, the standard deviation of the Gaussian posterior does not depend $\tilde{\theta}^1$ and $\tilde{\theta}^2$, but only depends on hyperparameters σ_1 and σ_2 , which are fixed to 1.

F Additive and CC-in do not satisfy Des. 2 (strong)




Figure 4: Toy reward combination problem. The aim is to combine reward vectors $\tilde{\theta}^1 = [1, 0, 0]$ and $\tilde{\theta}^2 = [0, 1, 0]$. The MDP is as follows. The starting state is s_1 . From s_1 the agent must irreversibly and deterministically transition to s_2 , s_3 , or s_4 , and stay there indefinitely. The features of each state s are given by $f(s)$.

Additive. Consider the reward combination scenario in Figure 4. Maximizing $\tilde{\theta}^1$ or $\tilde{\theta}^2$ an agent would correspondingly end up in s_2 or s_3 , both of which have the third feature equal 1. However, maximizing $\theta^* = \tilde{\theta}^1 + \tilde{\theta}^2$ would lead the agent to s_4 , where the value of the third feature is 0.

CC-in. Again consider the reward combination task in Figure 4. Maximizing $\tilde{\theta}^1$ and $\tilde{\theta}^2$ separately would correspondingly lead to choosing s_2 and s_3 , both states having the third feature equal 1. However, maximizing $\theta^* = \frac{\tilde{\theta}^1 + \tilde{\theta}^2}{2} = [0.5, 0.5, 0]$, a convex combination of the input rewards, would lead to the agent choosing s_4 where the value of the third feature equals 0.

As the Additive and CC-in reward posteriors have support on rewards giving rise to behaviors *different* from the behaviors of the input rewards w.r.t. the third feature of the toy MDP, neither method meets the strong version of Des. 2 (informative about desirable behavior).