

Active Preference Learning using Maximum Regret

Nils Wilde, Dana Kulić, and Stephen L. Smith

Abstract—We study active preference learning as a framework for intuitively specifying the behaviour of autonomous robots. A user chooses the preferred behaviour from a set of alternatives, from which the robot learns the user’s preferences, modeled as a parameterized cost function. Previous approaches present users with alternatives that minimize the uncertainty over the parameters of the cost function. However, different parameters might lead to the same optimal behaviour; as a consequence the solution space is more structured than the parameter space. We exploit this by proposing a query selection that greedily reduces the maximum error ratio over the solution space. In simulations we demonstrate that the proposed approach outperforms other state of the art techniques in both learning efficiency and ease of queries for the user. Finally, we show that evaluating the learning based on the similarities of solutions instead of the similarities of weights allows for better predictions for different scenarios.

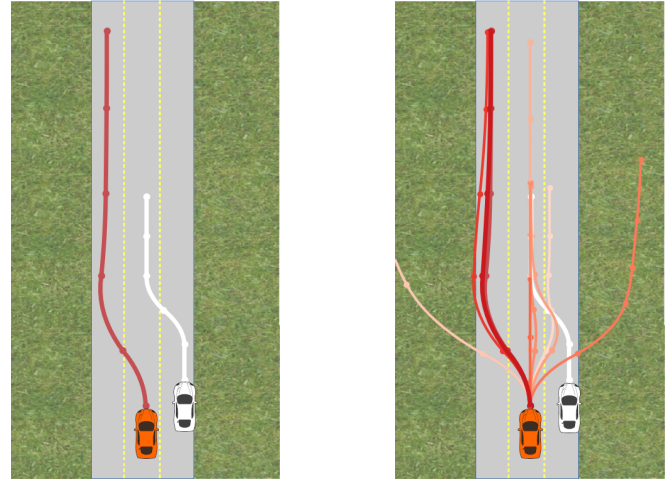
I. INTRODUCTION

Recently, research in human robot interaction (HRI) has focused on the design of frameworks that enable inexperienced users to efficiently deploy robots [1]–[6]. Autonomous mobile robots for instance are capable of navigating with little to no human guidance; however, user input is required to ensure their behaviour meets the user’s expectations. For example, in industrial facilities, a robot might need to be instructed about the context and established workflows or safety regulations [7] or an autonomous car should learn which driving style a passenger would find comfortable [8], [9]. Users who are not experts in robotics find it challenging to specify robot behaviour that meets their preferences [3].

Active preference learning offers a methodology for a robot to learn user preferences through interaction [1]–[3], [10]–[12]. Users are presented with a sequence of alternative behaviours to a specific robotic task and choose their preferred alternative. Figure 1 shows an example of learning user preferences for an autonomous vehicle where alternative behaviours are presented on an interface. Usually, the user is assumed to make their choice based on an internal, hidden cost function. The objective is to learn this cost function such that a robot can optimize its behaviour accordingly. Often, the user cost function is modelled as a weighted sum of predefined features [1]. Hence, learning the cost function is reduced to learning the weights. The key questions in this methodology are (1) how to select a set of possible solutions that are presented to the user such that the

This research is partially supported by the Natural Sciences and Engineering Research Council of Canada (NSERC).

N. Wilde and S. L. Smith are with the Department of Electrical and Computer Engineering, University of Waterloo, D. Kulić is with the University of Waterloo and Monash University. (dana.kulic@monash.edu; nwilde@uwaterloo.ca; stephen.smith@uwaterloo.ca)



(a) Optimal behaviour.

(b) Learned behaviour.

Fig. 1: Behaviour of an autonomous car (red) in the presence of another vehicle (white). In (a) we show the optimal behaviour for some user. In (b) we show alternative paths presented during active preference learning. Darker shades of red indicate behaviour that was presented later. The figure was created using code from [2].

cost function can be learned from few queries to the user, and (2) can the user choose reliably between these solutions.

In this work we propose a new approach for selecting solutions in active preference learning. In contrast to the work of [1], [2], [13] our approach does not focus on reducing the uncertainty of the belief over the weights; instead, we consider the set of all possible solutions to the task. Different weights in the user cost function might correspond to similar or even equal optimal solutions; in optimization problems this is known as sensitivity [14]. Thus, even if the estimated weights do not equal the true user weights, the corresponding solution might be the same. Therefore, we propose a new measure for active preference learning: The regret of the learned path. The concept of regret is known in robust shortest path problems [15], [16]. Consider two sets of weights for a user cost function, one that is optimal for a user and one that was estimated through active preference learning. The regret of the estimate captures the suboptimality of the solution found using the estimated weights, i.e., the ratio of the cost of the estimated solution *evaluated* by the optimal weights and the cost of the optimal solution, evaluated by the optimal weights.

We use the notion of regret to select alternatives to show to

the user. From a set of solutions that are considered equally good for the user given the feedback obtained so far, we choose the pair of solutions (P, Q) that, if P is the optimum, the ratio of costs is maximised. As the user either rejects P or Q we remove the most sub-optimal alternative from our solution space. In each iteration, our proposed approach optimizes over the set of all solutions that are consistent with the user feedback obtained so far. It then presents the user with the pair of solutions where the regret is maximized.

Following the motivation for regret, we evaluate the results of active preference learning based on the learned solution, instead of the learned weights. Therefore, we use the relative error in the cost of paths as a metric. This mirrors how an actual user would evaluate a robot's behaviour: Users are not interested in what weights are used by a robot's motion planner; one of the main motivations for active preference learning is that users find it challenging to express weights for cost or reward functions. Instead users judge a robot's behaviour by how similar it is to what they deem as optimal.

A. Related Work

The concept of learning a hidden cost or reward function from a user is widely used in various human-robot interaction frameworks, such as learning from demonstrations (LfD) [4], [17], learning from corrections [18], [19] and learning from preferences [1], [3], [12], [13], [17].

Closely related to our work, the authors of [1], [13] and [2] investigate how active preference learning can be used to shape a robot's behaviour. Thereby a general robot control problem with continuous states and actions is considered. The user cost function is modelled as weighted sum of features. They show that the robot is able to learn user preferred behaviours from few iterations using active preference learning. In [1] and [13], Dragan and colleagues investigate a measure for selecting a new pair of possible solutions to be shown to the user based on the posterior belief over the set of all weights. In detail, new solutions are selected such that the integral over the unnormalized posterior, called volume, is minimized in expectation. This approach is revised in [2], where a failure case for the volume removal is demonstrated. As an alternative measure, the authors propose the information entropy of the posterior belief over the weights. We show that both of the above approaches disregard the sensitivity of the underlying motion planning problem: Learning about weights of a cost function can be inefficient, as different weights can lead to the same optimal behaviour. In our previous work [12] we discretized the weight space into *equivalence regions*, i.e., sets of weights where the optimal solution of the planning problem is the same.

Another concern during active preference learning is to present alternatives to the user that are easy for them to differentiate which leads to a lower error rate. The authors of [20] investigate strategies for active learning that consider the flow of the queries to reduce the mental effort of the user and thus decrease the user's error rate. Similarly, [2] optimizes for queries that are easy to answer. In our work, we present an active query strategy that features these properties

intrinsically: By maximizing the regret of the presented paths, we automatically choose paths that are different with respect to the user cost function and thus are expected to be easily distinguishable for the user.

B. Contributions

We contribute to the ongoing research in active preference learning as a framework for specifying complex robot behaviours. We propose a measure for evaluating the solution found by preference learning based on the robot's learned behaviour instead of the learned weights in the cost function. Further, we propose a new active query selection guided by the maximum error ratio between solutions. Thereby, users are presented with the pair of solutions that has the maximum error ratio among all paths in the feasible solution space. We demonstrate the performance of our approach by comparing it to a competing state of the art technique and show that our proposed method learns the desired behaviour more efficiently. Moreover, the queries the user is presented with are easier to answer and thus lead to more reliable user feedback. Finally, we demonstrate how our measure based on solutions gives better predictions about the behaviour of the robot in different scenarios that were not part of the learning.

II. PROBLEM STATEMENT

Preliminaries: Let \mathcal{X} be the state space of a robot and the environment it is acting in and x_0 some start state. Further, we have an action space \mathcal{A} where each $a \in \mathcal{A}$ potentially only affect parts of the state, i.e., there might be static or dynamic obstacles unaffected by the robot's actions.

Further let P be a path of finite length starting at x_0 . A path is evaluated by a column vector of predefined features $\phi(P) = [\phi_1(P), \phi_2(P), \dots, \phi_d(P)]^T$. Together with a row vector of weights \mathbf{w} we define the cost of a path as

$$c(P, \mathbf{w}) = \phi(P)\mathbf{w}. \quad (1)$$

Given some weight \mathbf{w}' let the *optimal path* be $P' = \arg \min_P c(P, \mathbf{w}')$. The optimal cost for a weight is

$$c^*(\mathbf{w}') = c(P', \mathbf{w}'). \quad (2)$$

For any other weight \mathbf{w} , we call $c(P', \mathbf{w})$ the cost of P' evaluated by \mathbf{w} .

Problem Formulation: We consider a robot's state and action space $(\mathcal{X}, \mathcal{A})$ and some start state x_0 . Further, let a vector of weights \mathbf{w}^{user} , describing a user's preference for the robot's behaviour and the corresponding optimal path P^{user} . Each weight w_i^{user} has a lower and upper bound l_i and u_i . However, \mathbf{w}^{user} itself is hidden. We can learn about \mathbf{w}^{user} by presenting the user with pairs of paths (P, Q) over k iterations. The objective is to find an estimated path P^k that reflects the user preferences \mathbf{w}^{user} , i.e., is as similar to P^{user} as possible. To evaluate the result of learning \mathbf{w}^{user} the authors of [1] propose the *alignment metric*, i.e., the cosine of the angle between the learned weight vector \mathbf{w} and \mathbf{w}^{user} .

We adapt this metric and transform it to a normalized error between 0 and 1, which we call the *weight error*:

$$\text{Err}_{\text{weight}}(\mathbf{w}, \mathbf{w}^{\text{user}}) = \frac{1}{2} \left(1 - \frac{\mathbf{w} \cdot \mathbf{w}^{\text{user}}}{\|\mathbf{w}\|_2 \|\mathbf{w}^{\text{user}}\|_2} \right). \quad (3)$$

The alignment metric was also used in [2], [13]. However, this metric has two potential shortcomings: 1) It does not consider the sensitivity of the optimization problem that finds an optimal path for a given weight vector. Thus, an error in $\text{Err}_{\text{weight}}$ might actually not result in a different optimal path. Moreover, even if the learned weight has a relatively small error, the corresponding path might be suboptimal to the user. 2) The weight error is not suitable as a test error (i.e., to test whether the learned user preferences generalize well to new task instances not encountered during learning) since it does not consider the robot's resulting behaviour: $\text{Err}_{\text{weight}}(\mathbf{w}, \mathbf{w}^{\text{user}})$ is equal for all training and test instances. Hence, the weight error gives no insight into how well the estimated preferences translate into different scenarios, unless $\text{Err}_{\text{weight}}(\mathbf{w}, \mathbf{w}^{\text{user}}) = 0$, i.e., the optimal weights are found. Therefore, we choose a different metric for evaluating the learned behaviour: Instead of the learned weight \mathbf{w}^k we consider the learned path P^k . We compare the cost of P^k , evaluated by the user's true cost \mathbf{w}^{user} to the optimal cost path of \mathbf{w}^{user} :

$$\text{Err}_{\text{Path}}(P^k, \mathbf{w}^{\text{user}}) = \frac{c(P^k, \mathbf{w}^{\text{user}})}{c^*(\mathbf{w}^{\text{user}})} - 1. \quad (4)$$

This error was proposed in [21] and we refer to it as the *path error*. A similar error was used in [22] for finding risk-aware policies in inverse reinforcement learning. Based on this metric we can now formally pose the learning problem.

Problem 1. Given $(\mathcal{X}, \mathcal{A})$ and x_0 , and a user with hidden weights \mathbf{w}^{user} who can be queried over k iterations about their preference between two paths P and Q , find a weight \mathbf{w}^k with the corresponding optimal path P^k starting at x_0 that minimizes $\text{Err}_{\text{Path}}(P, \mathbf{w}^{\text{user}})$.

III. ACTIVE PREFERENCE LEARNING

We introduce the user model and learning framework of our active preference learning approach and then discuss several approaches for selecting new solutions in each iteration.

A. User Model

To learn about \mathbf{w}^{user} and thus find P' , we can iteratively present the user with a pair of paths (P, Q) and they return the one they prefer:

$$\begin{aligned} c(P, \mathbf{w}^{\text{user}}) \leq c(Q, \mathbf{w}^{\text{user}}) &\implies \text{the user returns } P, \\ c(P, \mathbf{w}^{\text{user}}) > c(Q, \mathbf{w}^{\text{user}}) &\implies \text{the user returns } Q. \end{aligned} \quad (5)$$

However, a user might not always follow this model exactly. For instance, they might consider features that are not in the model or they are uncertain in their decision when P and Q are relatively similar. Thus, we extend equation (5) to a probabilistic model, similar to our previous work in [12]. Let

$I^{P,Q}$ be a binary random variable where $I^{P,Q} = 1$ if the user prefers path P over Q , and -1 otherwise. Then we have

$$\begin{aligned} \mathbb{P}(I^{P,Q} = 1 | c(P, \mathbf{w}^{\text{user}}) \leq c(Q, \mathbf{w}^{\text{user}})) &= p, \\ \mathbb{P}(I^{P,Q} = -1 | c(P, \mathbf{w}^{\text{user}}) \leq c(Q, \mathbf{w}^{\text{user}})) &= 1 - p, \end{aligned} \quad (6)$$

where $1/2 < p \leq 1$. If $p = 1$ we recover the deterministic case from equation (5). In this very simple model the user's choice does not depend on how similar P and Q are. In the simulations we simulate the user with to the more complex model in [2], which poses the user's error rate as a function of the similarity between alternatives, and show that equation (6) nonetheless allows us to achieve strong performance.

B. Learning Framework

Over multiple iterations, equation (5) yields a collection of inequalities of the form $(\phi^P - \phi^Q)\mathbf{w} \leq 0$. We write the feedback obtained after k iterations as a sequence $U^k = \{(P^1, Q^1), (P^2, Q^2), \dots, (P^k, Q^k)\}$. Without loss of generality, we assume that for any pair (P, Q) in U^k the path P was preferred over the path Q . We then summarize the left-hand-sides $(\phi^P - \phi^Q)$ for all k iterations using a matrix \mathbf{A}^k . Based on the sequence U^k we can compute an estimate \mathbf{w}^k of \mathbf{w}^{user} by taking the expectation.

Deterministic case: In the deterministic case, i.e., $p = 1$, the estimate \mathbf{w}^k must satisfy $\mathbf{A}^k \mathbf{w}^k \leq 0$ to be consistent with the user feedback obtained thus far. The set of all such weights constitutes the *feasible set* $\mathcal{F} = \{\mathbf{w} \in \mathbb{R}^d | l_i \leq w_i \leq u_i, \mathbf{A}^k \mathbf{w} \leq 0\}$.

C. Active Query Selection

In active preference learning we can choose a pair of paths (P, Q) to present to the user in each iteration k . Throughout this work we only consider paths P and Q that are optimal for some weights \mathbf{w}^P and \mathbf{w}^Q . Given the user feedback obtained until iteration k , a new pair (P^k, Q^k) is found by maximizing some measure $f(\mathbf{w}^P, \mathbf{w}^Q, U^{k-1})$ describing the expected learning effect from showing (P^k, Q^k) to the user. Recently, several measures have been introduced: Removing the *Volume*, i.e., minimizing the integral of the unnormalized posterior over the weights [1], [13], maximizing the information entropy over the weights [2] and removing *equivalence regions*, i.e., sets of weights where for each weight has the same optimal path [12].

Parameter space and solution space: The first two approaches maximize information about the parameter space, i.e., the weights \mathbf{w} , instead of the solution space, i.e., the set of all possible paths P . Despite its motivation based on inverse reinforcement learning, this has a major drawback: The difference in the parameters does not map linearly to the difference in the features of corresponding optimal solutions. Given some \mathbf{w}^P and \mathbf{w}^Q , we can compute optimal paths P and Q with features ϕ^P and ϕ^Q , respectively. Then $\|\mathbf{w}^P - \mathbf{w}^Q\| \propto \|\phi^P - \phi^Q\|$ does *not* necessarily hold. Thus, learning efficiently about \mathbf{w} does not guarantee efficient or effective learning about paths. Moreover, learning

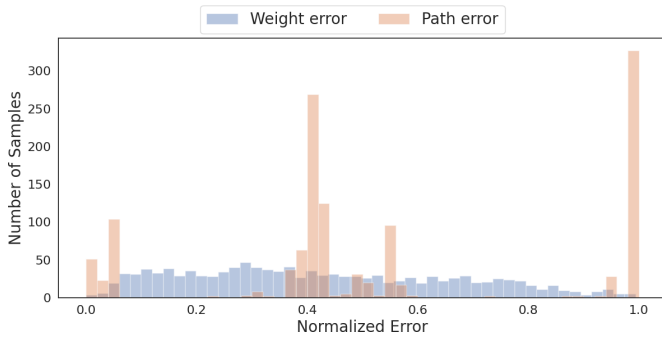


Fig. 2: Example of the sensitivity of a continuous motion planning problem. We show the histogram of the *normalized* weight error $\text{Err}_{\text{Weight}}$ and the *normalized* path error Err_{Path} for 3000 uniformly sampled random weights.

about w might allow for disregarding a large number of weights. However, the corresponding optimal paths might be very similar and thus the learning step is potentially less informative in the solution space.

Example 1. We consider the autonomous driving example from [2] which is posed in a continuous state and action space, illustrated in Figure 1. In Figure 2 we compare the weight error $\text{Err}_{\text{Weight}}(w, w^{\text{user}})$ and the path error $\text{Err}_{\text{Path}}(P, w^{\text{user}})$ of 3000 uniformly random samples. While the weight error is distributed uniformly, the path error distribution takes nearly a discrete form, despite the continuous action space. This illustrates how different weights do not necessarily lead to different solutions, making the solution space more structured than the parameter space.

In our previous work [12] we proposed a query selection based on a discretization of the weight space: Sets of weights that have the same optimal path are labeled as *equivalence regions*. The objective then is to maximally reduce the posterior belief over equivalence regions, i.e., to reject as many equivalence regions as possible. A drawback of this approach is that there exists cases where any query only allows for updating the belief of few equivalence regions, resulting in slow convergence. Because of these limitations of the existing approaches we study a new measure $f(w^P, w^Q, U^k)$ based on the solution space.

IV. MIN-MAX REGRET LEARNING

We propose a new measure $f(w^P, w^Q, U^k)$ called the maximum regret, which we seek to minimize.

Definition 1 (Regret of weights). Given a weight w^P with its corresponding optimal path P and some weight w^Q , the regret of P under w^Q is

$$r(w^P, w^Q) = \frac{c(P, w^Q)}{c^*(w^Q)}. \quad (7)$$

Regret expresses how sub-optimal a path P is when evaluated by some weights w^Q . In active learning, this can be interpreted as follows: *If P is the final estimate, but Q is the optimal solution, how large is the ratio between the*

cost of P , evaluated by w^Q , and the optimal cost? We now formulate an approach for selecting which alternatives to show to the user by using regret.

A. Deterministic Regret

When assuming a deterministic user, we need to assure that $w^P \in \mathcal{F}^k$ and $w^Q \in \mathcal{F}^k$, such that the presented paths reflect the user feedback obtained so far. Given w^P we pose the *Maximum Regret under Constraints Problem (MRuC)* as

$$\begin{aligned} \max_{w^Q} r(w^P, w^Q) \\ \text{s.t. } A^k w^Q \leq 0, \\ l_i \leq w_i^Q \leq u_i. \end{aligned} \quad (8)$$

The objective can be written in the form $\max_{w^Q, \phi^Q} \phi^Q w^Q / c^*(w^P)$. This is a bi-linear program, which are a generalization of quadratic programs. Unfortunately, in our case the objective function is non-convex; generally, such problems are hard to solve.

Symmetric Regret: In equation (8) we have defined the maximum regret problem when one path is given. While presenting users with a new pair of paths (P, Q) , we want to find paths where the regret of w^P under w^Q is maximized *and vice versa*. Thus, we rewrite the objective in (8) to $r(w^P, w^Q) + r(w^Q, w^P)$, which we call the *symmetric regret*. The maximum symmetric regret of a feasible set \mathcal{F}^k can be found with the following bi-linear program:

$$\begin{aligned} \max_{w^P, w^Q} r(w^P, w^Q) + r(w^Q, w^P) \\ \text{s.t. } A^k w^P \leq 0, \\ A^k w^Q \leq 0, \\ l_i \leq w_i^P \leq u_i, \quad l_i \leq w_i^Q \leq u_i. \end{aligned} \quad (9)$$

Similar to equation (8) this is a non-convex optimization problem. In the evaluation we solve this problem by sampling a set of weights and pre-computing the corresponding optimal paths, following the approach in [1].

B. Probabilistic Regret

We now formulate regret with consideration of the user's uncertainty when choosing among paths. Taking a Bayesian perspective we treat w^{user} as a random vector. This allows us to express a posterior belief over w^{user} given an observation $I^{P,Q}$. Let $c^P = c(P, w^{\text{user}})$ and $c^Q = c(Q, w^{\text{user}})$, respectively. Further, we assume a uniform prior over w^{user} . For any estimate w where $(\phi^P - \phi^Q)w \leq 0$ we have

$$\mathbb{P}(w^{\text{user}} = w | I^{P,Q}) \propto \mathbb{P}(I^{P,Q} | c^P \leq c^Q). \quad (10)$$

Let $\mathbb{P}(w | I^{P,Q})$ denote $\mathbb{P}(w^{\text{user}} = w | I^{P,Q} = 1)$. We calculate the posterior given a sequence of user feedback $U^k = \{(P^1, Q^1), (P^2, Q^2), \dots, (P^{k-1}, Q^{k-1})\}$ as

$$\mathbb{P}(w | U^k) \propto \prod_{(P,Q) \in U^k} \mathbb{P}(w | I^{P,Q}). \quad (11)$$

We formulate the symmetric regret in the probabilistic case by weighting the regret by the posterior of w^P and w^Q :

Algorithm 1: Maximum Regret Learning

Input: $(\mathcal{X}, \mathcal{A}), x_0, K$
Output: P^{best}

- 1 Initialize $U^0 = \emptyset$
- 2 Sample a set of weights Ω
- 3 **for** $k = 1$ to K **do**
- 4 $\mathbf{w}^P, \mathbf{w}^Q \leftarrow \text{max_regret}((\mathcal{X}, \mathcal{A}), x_0, U^{k-1}, \Omega)$
- 5 $P, Q \leftarrow \text{opt_path}(\mathbf{w}^P), \text{opt_path}(\mathbf{w}^Q)$
- 6 $I^k \leftarrow \text{user_feedback}(P, Q)$
- 7 **if** $I^k = -1$ **then**
- 8 $U^k = U^{k-1} \cup (P, Q)$
- 9 **else**
- 10 $U^k = U^{k-1} \cup (Q, P)$
- 11 **return** $\text{opt_path}(\mathbb{E}_{\mathbf{w}}[\mathbf{w}|U^k])$

$$\begin{aligned} \mathcal{R}^k(\mathbf{w}^P, \mathbf{w}^Q|U^k) &= \mathbb{P}(\mathbf{w}^P|U^k)\mathbb{P}(\mathbf{w}^Q|U^k) (r(\mathbf{w}^P, \mathbf{w}^Q) + r(\mathbf{w}^Q, \mathbf{w}^P)). \end{aligned} \quad (12)$$

That is, we discount the *symmetric regret* such that we only consider pairs (P, Q) where both $\mathbf{w}^{\text{user}} = \mathbf{w}^P$ and $\mathbf{w}^{\text{user}} = \mathbf{w}^Q$ are likely given the user feedback U^k .

Finally, we adapt the problem of finding the maximum symmetric regret from equation (9) to the probabilistic case. As we cannot formulate a feasible set \mathcal{F} for a probabilistic user, we consider a finite set Ω where each $\mathbf{w} \in \Omega$ is uniformly randomly sampled from the set $\{\mathbf{w} \in \mathbb{R}^d | l_i \leq w_i \leq u_i\}$. We then take the maximum over all $\mathbf{w} \in \Omega$ to compute the *probabilistic maximum regret*

$$\mathcal{R}_{\max}^k(U^k) = \max_{\mathbf{w}^P, \mathbf{w}^Q \in \Omega} [\mathcal{R}^k(\mathbf{w}^P, \mathbf{w}^Q|U^k)]. \quad (13)$$

In min-max regret learning, we choose the pair of paths (P, Q) that is the maximizer of equation (13).

C. Preference Learning with Probabilistic Maximum Regret

Our proposed solution for active preference learning using *probabilistic maximum regret* is summarized in Algorithm 1. In each iteration we find the pair $(\mathbf{w}^P, \mathbf{w}^Q)$ that maximizes the probabilistic symmetric regret as in equation (13) over a set of samples Ω (line 4). We then obtain user feedback $I_k = 1$ if the user prefers path P and $I_k = -1$ otherwise (line 7) and add the feedback to a sequence (line 6-10). After K iterations, we return the path that is optimal for the expected weight, given the observed user feedback (line 11). Using the maximum regret in the query selection is a greedy approach to minimize the maximum error. Given the current belief over the weights, we choose the pair P, Q with the maximum error ratio, discounted by the likelihoods of \mathbf{w}^P and \mathbf{w}^Q .

V. EVALUATION

We evaluate the proposed approach using the simulation environment from [2], allowing us to compare our approach to theirs in the same experimental setup. To label the

approaches let Entropy denote the maximum entropy learning from [2] and Regret our maximum regret learning.

Learning experiments: First, we consider one of the experiments in [2]: The autonomous driving scenario (*Driver*) where an autonomous car moves on a three lane road in the presence of a human-driven vehicle as shown in Figure 1. Paths are described by four features: Heading relative to the road, staying in the lane, vehicle speed, and the distance to the other car. Every feature is averaged over the entire path. Furthermore, we introduce the *Extended Driver* experiment with additional features to create a more complex scenario. In addition to the above features we add the distance travelled along the road, the summed lateral movement, summed and maximum lateral and angular acceleration, the minimal speed, and the minimum distance to the other vehicle. We choose the Driver example because the entropy approach from [2] showed strong results and this scenario was already previously investigated in [1]. The extension aims to show how the learning techniques behave in higher dimensions.

Additionally we consider a third experiment adapted from [3], [12]: An autonomous mobile robot navigates between given start and goal locations in a known environment. However, there are n areas in the environment that a user marked as desired or undesired for robot traffic. Each such area is a soft constraint, i.e., there is a penalty or reward associated with it, which can be expressed by a weight. By defining features for all areas describing if a robot trajectory passes through yields a cost function of the form $c(P) = \phi(P)\mathbf{w}$. Here ϕ is an $n + 1$ dimensional vector, the first n entries are the features describing the length of path P in area i for all $i = 1, 2, \dots, n$. The $n + 1$ -th feature is the time it takes the robot to execute path P . The robot is unaware of the value of each penalty and reward, i.e., the weights \mathbf{w} are not given to the robot, yielding an instance of Problem 1. We will refer to this experiment as *Mobile*. The instance of the problem used for evaluation consists of 18 areas; thus, the dimensionality of the feature and weight space is 19.

Optimal paths: Given a weight \mathbf{w} we need to find the corresponding optimal path P in order to evaluate the path error (and to compute regret in Algorithm 1). In [2] no motion planner is given; in the experiments we rely on the generic non-linear optimizer L-BFGS [23] for the *Driver* and *Extended Driver* experiments. However, depending on the problem, this solver can return suboptimal solutions. To mitigate this effect, we pre-sample paths which are used as a look-up-table. Given a path that was found using L-BFGS, we iterate over all pre-sampled paths; if a sampled path yields a better cost for the given weight, we use that path instead. In both experiments, the entropy approach uses the implementation provided by [2] where queries are chosen from 500,000 pre-sampled pairs of random paths. Since regret requires optimal paths, the regret approach uses a set of 200 weights with their corresponding optimal paths,

yielding 40,000 possible pairs of paths¹. Finally, in these experiments the behaviour is actually captured by a reward and not a cost. Thus, an optimal path is found by minimizing the negative cost and we change the definition of regret to $r(\mathbf{w}^P, \mathbf{w}^Q) = 1 - c(P, \mathbf{w}^Q)/c^*(\mathbf{w}^Q)$.

In the *Mobile* experiment the robot moves using a state lattice planner [24]; given a weight \mathbf{w} we can always find an optimal path in polynomial time. We varied the problem setup by choosing three different start and goal locations for the robot to navigate between. For each start goal pair we need to pre-sample paths individually. The discrete state space led to a significantly smaller set of pre-samples, varying between 5 and 32. However, using randomly generated paths as in [2] for Entropy led to very poor performance. Therefore, we slightly modified the Entropy approach for this experiment, such that the same pre-samples were used as for the Regret approach.

Simulated users: We simulate user feedback using the probabilistic user model from [2]. Given two paths, the user’s uncertainty depends on how similar the paths are with respect to the cost function evaluated for \mathbf{w}^{user} :

$$\mathbb{P}(I = 1 | (P, Q), \mathbf{w}^{\text{user}}) = \frac{e^{c(P, \mathbf{w}^{\text{user}})}}{e^{c(P, \mathbf{w}^{\text{user}})} + e^{c(Q, \mathbf{w}^{\text{user}})}}. \quad (14)$$

The probabilistic regret is computed using pre-sampled weights as described in Algorithm 1, with an uncertainty of $p = 0.85$ in equation (6). Similar to [2] for each experiment we sample a user preference \mathbf{w}^{user} uniformly randomly from the unit circle, i.e., $\|\mathbf{w}^{\text{user}}\|_2 = 1$. We notice that this can include *irrational* user behaviour: A negative weight on heading for instance would encourage the autonomous car to *not* follow the road.

A. Learning error

In Figure 3 we compare Entropy to Regret on both metrics over 10 iterations for the two experiments, each repeated 200 times. In the boxplots the center line shows the median and the green triangle shows the mean.

In the driver example, Entropy overall achieves a smaller weight error and smaller deviations from the mean, reproducing the results from [2]. In the path space we observe that Entropy achieves a slightly better result in the last two iterations. However, between iteration 2 and 8, the Regret approach performs better, i.e., learns more quickly. Overall, both approaches perform equally well.

For the *Extended Driver* example in Figure 3b, both approaches make limited progress on the weight metric and exhibit large deviations. For the path error we observe that Entropy performs better initially, but makes little progress after iteration 6. The final median lies at ≈ 0.05 at iteration 10, but the highest quartile still reaches up to 0.2. The Regret approach achieves a lower mean and median error in iteration 4 and subsequently improves further. At iteration 8

¹Using sampled optimal paths for the entropy approach did not lead to different results in the experiments, therefore we show the results using the original implementation.

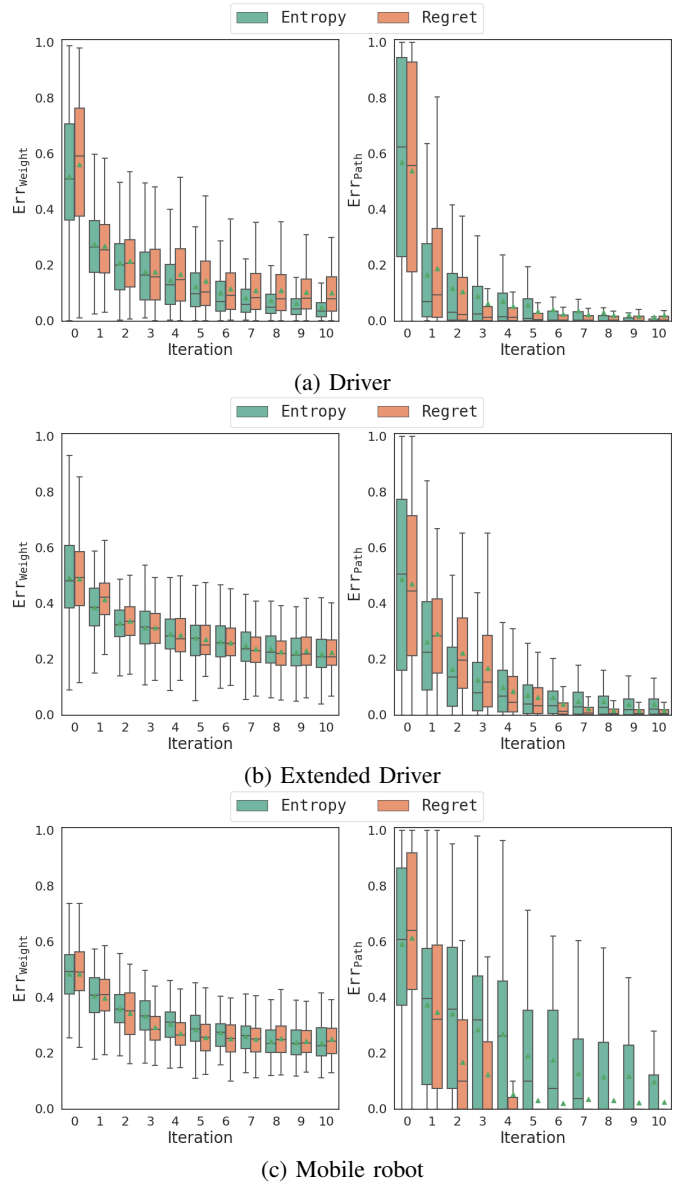


Fig. 3: Comparison of active preference learning with maximizing entropy and minimizing regret.

the mean and median are close to 0, the box plot also shows that three quarters of all trials are very close to convergence.

Figure 3c illustrates the result for the *Mobile* experiment. Here, the weight error shows no difference between the two approaches; both perform equally poorly and inconsistently. At the same time the path error shows a large difference. Regret achieves convergence for nearly all trials after just 5 iterations (some outliers cause the mean value to still be at ≈ 0.05). At the same time the performance of Entropy is inferior: Even though the median error becomes 0 in iteration 8, the mean value is still ≈ 0.15 with large deviations.

In conclusion, Regret achieves an equally good result as Entropy on the path error for the *Driver* experiment, despite having a larger weight error. That is, while the weights found by Entropy are more similar to \mathbf{w}^{user} based on the

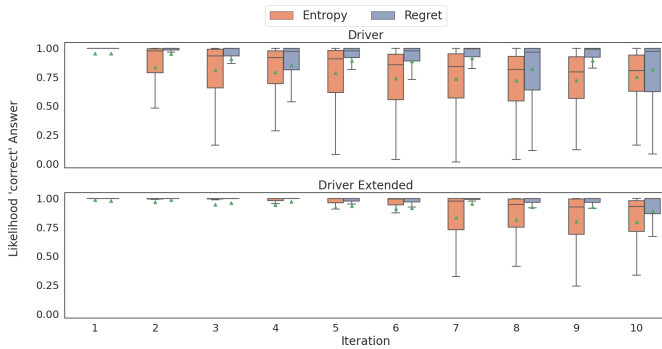


Fig. 4: The likelihood that the simulated user gives the ‘correct’ answer, i.e., the probability in equation (14).

alignment metric, the resulting behaviour of Regret and Entropy are equally good. Moreover, the *Extended Driver* and *Mobile* experiments have shown that the performance of Entropy deteriorates for higher dimensions, i.e., larger sets of features. In contrast, Regret still achieves a very strong performance on the path error.

B. Easiness of queries

A major contribution of [2] is the design of queries that are *easy* for the user to answer, i.e., the probability that the user choice is inconsistent with the assumed cost function from equation (14) is low. In maximum regret learning we do not directly consider the user’s uncertainty when choosing a new pair of paths. However, as the paths maximizing the probabilistic symmetric regret have a large difference on cost, our approach implicitly selects paths P, Q that potentially are easy for a user to answer. To compare the easiness of the queries presented to the user, we consider the probability that the user would choose the path with lower cost, evaluated by w^{user} (14). In Figure 4 we compare the probability of correct user answers for Entropy and Regret.

In the *Driver* experiment, we recorded a mean probability of 0.84 for Entropy, which is slightly worse than reported for the strict queries in [2], where correct answers occurred in 87% of cases. Nonetheless, with Regret the simulated answers had a mean probability of being correct of 0.94, outperforming Entropy. In Figure 4, we observe that both approaches achieve very high probabilities for correct user answers in the first iteration, i.e., ask an *easy* question. Afterwards, the probabilities get smaller: The median of Entropy decreases to 0.9 in iteration 3 and the deviations increase significantly. The Regret approach maintains higher median values for all iterations. Interestingly, we observe cyclic decreases of the mean (and increases for the deviations) in iterations 4, 8 and 10. According to the user model the presented paths were very similar, indicating that the learning might have been close to convergence. This aligns with the small errors of the expected weight reported in Figure 3a.

In the *Extended Driver* experiment the user behaviour is much more accurate for both approaches with a mean of 0.92 for Entropy and 0.96 for Regret. This indicates that the sampled paths differ more in cost. From iteration

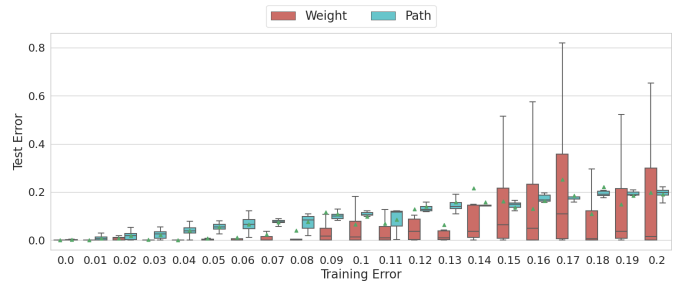


Fig. 5: Relationship between training errors measured by the path and weight metric to test errors in the path metric.

7 onwards, Entropy starts to show larger deviations, i.e., questions become more difficult to answer, implying that the presented paths are very similar. Together with the very small decrease in path error the we observed in Experiment 1 (Figure 3b) at the same iteration, this leads to the conjecture that Entropy is converging to a local optimum. Finally, the *Mobile* experiment did not show any difference between the two approaches, both achieving a very high accuracy of 0.99.

Overall, these results strongly support our claim that maximizing regret implicitly creates queries that are easy for the user to answer.

C. Generalization of the error

Finally, we investigate how the two error metrics generalize to different scenarios, independent of whether the error is a result of learning with Entropy or Regret. That is, we investigate how useful each error metric is for predicting the robot’s performance when deployed in a new instance of the problem not encountered during learning. For the *Driver* experiment we use the setup from Figure 1 as a training case and construct five test cases by changing the initial state of the human driven vehicle (white). The weight error is scenario independent, it directly describes how similar the estimated weight is to w^{user} . Thus, the weight error is the same in training and test cases and cannot be used as test error, as this would contain no additional information about performance on the test case. Hence, we use the path error as the test error. Further, we notice that if the weight error is zero, i.e., the weights have been learned perfectly, then the path error is zero in all scenarios. However, as shown in Figure 3 and in [1], [2] the weights typically do not converge to the true user weight within few iterations. Given some weight the path errors are fixed values in every test scenario. We are interested in how well the weight error and the path error of the training predict the path error of the test scenario.

We generate 40 different random user weights w^{user} and then generate 200 estimates of each of these weights. For every estimate we find the optimal path and compute the path and weight error which are used as training errors for the estimate. In Figure 5 we show how these training errors relate to the test error. We compare the path and weight error as a measure of generalisation performance (i.e., how well the weight and path errors predict the test case performance).

We observe that the path error translates linearly between

training and test scenarios: Given a weight with a certain path error in the training scenario, the weight yields paths in the test scenarios that have a similar path error, on average. The relationship between weight error and test error is more complex. For a weight error of 0 – 0.01 during training, we observe a test error of 0, i.e., if the weights are very close to the optimum, the optimal solution is found in every scenario. However, for larger training errors the test error shows large deviations, implying that a low weight error in training is not a robust measure of how good the resulting behaviour is in test cases. The observation is supported by a strong Pearson correlation of $p = 0.92$ between training and test error for the path error, but a much weaker correlation of $p = 0.28$ for the weight error. This lends support to the claim that the path error is better suited for making predictions of the performance in scenarios that were not part of the training.

We conducted a the same experiment for the *Extended Driver* scenario. The correlation of the path error is weaker, but with $p = 0.80$ still substantially stronger than for the weight error where we observed $p = 0.27$. In the *Mobile* scenario, training and test instances are defined by multiple start-goal pairs. However, we observed no correlation for both path error and weight error. The features in this scenario are local, i.e., describe if the robot visits a certain part of the environment. Learning about one task gains insufficient information to always find a good path for a different task.

In summary, we observe that the path error is more suitable than the weight error for reliable predictions of the test performance in scenarios with global features. However, higher dimensions can weaken the reliability, and local features may not allow for any predictions.

VI. DISCUSSION

In this paper we investigated a new technique for generating queries in active preference learning for robot tasks. We have shown that competing state of the art techniques have shortcomings as they focus on the weight space only. As an alternative, we introduced the regret of the cost of paths as a heuristic for the query selection, which allows to greedily minimize the maximum error. Further, we studied an error function that captures the cost ratio between the behaviour of estimated preferences and the optimal behaviour, instead of the similarity of weights. In simulations we demonstrated that using regret in the query selection leads to faster convergence than entropy while the queries are even easier for the user to answer. Moreover, we have shown that the path error allows for better predictions for other scenarios.

For future work special cases such as discrete action spaces in the form of lattice planners should be investigated. This would give further insight into the computational hardness of finding the maximum regret and potentially allow for solution strategies that do not require pre-sampling weights and paths. Richer user feedback such as an *equal preference* option could be of interest, promising results for this approach were presented in [2], [13]. Finally, regret based preference learning should be investigated in a user study to show the practicality of this approach.

REFERENCES

- [1] D. Sadigh, A. D. Dragan, S. Sastry, and S. A. Seshia, "Active preference-based learning of reward functions," in *RSS*, 2017.
- [2] E. Bıyık, M. Palan, N. C. Landolfi, D. P. Losey, and D. Sadigh, "Asking easy questions: A user-friendly approach to active reward learning," 2019.
- [3] N. Wilde, A. Blidaru, S. L. Smith, and D. Kulić, "Improving user specifications for robot behavior through active preference learning: Framework and evaluation," *IJRR*, vol. 39, no. 6, pp. 651–667, 2020.
- [4] P. Abbeel and A. Y. Ng, "Apprenticeship learning via inverse reinforcement learning," in *Proceedings of the twenty-first international conference on Machine learning*. ACM, 2004, p. 1.
- [5] A. Jain, S. Sharma, T. Joachims, and A. Saxena, "Learning preferences for manipulation tasks from online coactive feedback," *IJRR*, vol. 34, no. 10, pp. 1296–1313, 2015.
- [6] B. Akgun, M. Cakmak, J. W. Yoo, and A. L. Thomaz, "Trajectories and keyframes for kinesthetic teaching: A human-robot interaction perspective," in *ACM/IEEE international conference on Human-Robot Interaction*. ACM, 2012, pp. 391–398.
- [7] M. C. Gombolay, R. J. Wilcox, and J. A. Shah, "Fast scheduling of robot teams performing tasks with temporospatial constraints," *IEEE Transactions on Robotics*, vol. 34, no. 1, pp. 220–239, 2018.
- [8] D. S. González, O. Erkent, V. Romero-Cano, J. Dibangoye, and C. Laugier, "Modeling driver behavior from demonstrations in dynamic environments using spatiotemporal lattices," in *2018 IEEE ICRA*. IEEE, 2018, pp. 1–7.
- [9] T. Gu, J. Atwood, C. Dong, J. M. Dolan, and J.-W. Lee, "Tunable and stable real-time trajectory planning for urban autonomous driving," in *2015 IEEE/RSJ IROS*. IEEE, 2015, pp. 250–256.
- [10] C. Daniel, M. Viering, J. Metz, O. Kroemer, and J. Peters, "Active Reward Learning," *RSS*, vol. 10, no. July, 2014.
- [11] P. F. Christiano, J. Leike, T. Brown, M. Martic, S. Legg, and D. Amodei, "Deep reinforcement learning from human preferences," in *NIPS*, 2017, pp. 4299–4307.
- [12] N. Wilde, D. Kulić, and S. L. Smith, "Bayesian active learning for collaborative task specification using equivalence regions," *IEEE RA-L*, vol. 4, no. 2, pp. 1691–1698, April 2019.
- [13] C. Basu, M. Singhal, and A. D. Dragan, "Learning from richer human guidance: Augmenting comparison-based learning with feature queries," in *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction*, ser. HRI '18. New York, NY, USA: ACM, 2018, pp. 132–140.
- [14] D. Bertsimas and J. N. Tsitsiklis, *Introduction to linear optimization*. Athena Scientific Belmont, MA, 1997, vol. 6.
- [15] R. Montemanni and L. M. Gambardella, "An exact algorithm for the robust shortest path problem with interval data," *Computers & Operations Research*, vol. 31, no. 10, pp. 1667–1680, 2004.
- [16] A. Kasperski and P. Zielinski, "An approximation algorithm for interval data minmax regret combinatorial optimization problems," *Inf. Process. Lett.*, vol. 97, no. 5, pp. 177–180, 2006.
- [17] M. Palan, N. C. Landolfi, G. Shevchuk, and D. Sadigh, "Learning reward functions by integrating human demonstrations and preferences," *arXiv preprint arXiv:1906.08928*, 2019.
- [18] D. P. Losey and M. K. O'Malley, "Including uncertainty when learning from human corrections," in *Conference on Robot Learning*, 2018, pp. 123–132.
- [19] J. Y. Zhang and A. D. Dragan, "Learning from extrapolated corrections," in *IEEE ICRA*, 2019, pp. 7034–7040.
- [20] M. Racca, A. Oulasvirta, and V. Kyrki, "Teacher-aware active robot learning," in *2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 2019, pp. 335–343.
- [21] N. Wilde, D. Kulić, and S. L. Smith, "Learning user preferences in robot motion planning through interaction," in *IEEE ICRA*, May 2018, pp. 619–626.
- [22] D. S. Brown, Y. Cui, and S. Niekum, "Risk-aware active inverse reinforcement learning," in *Conference on Robot Learning*, 2018, pp. 362–372.
- [23] G. Andrew and J. Gao, "Scalable training of l_1 -regularized log-linear models," in *Proceedings of the 24th international conference on Machine learning*, 2007, pp. 33–40.
- [24] M. Pivtoraiko, R. A. Knepper, and A. Kelly, "Differentially constrained mobile robot motion planning in state lattices," *Journal of Field Robotics*, vol. 26, no. 3, pp. 308–333, 2009.