UNIVERSITY OF AMSTERDAM

UvA-DARE (Digital Academic Repository)

Optimizing Interactive Systems via Data-Driven Objectives

Li, Z.; Kiseleva, Y.; Agarwal, A.; de Rijke, M.; White, R.W.

DOI
10.48550/arXiv.2006.12999

Publication date
2020

Document Version
Submitted manuscript

License
Unspecified

Link to publication

Citation for published version (APA):
Li, Z., Kiseleva, Y., Agarwal, A., de Rijke, M., & White, R. W. (2020). *Optimizing Interactive Systems via Data-Driven Objectives*. (v1 ed.) ArXiv.
https://doi.org/10.48550/arXiv.2006.12999

General rights
It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations
If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: https://uba.uva.nl/en/contact, or a letter to: Library of the University of Amsterdam, Secretariat, P.O. Box 19185, 1000 GD Amsterdam, The Netherlands. You will be contacted as soon as possible.

UvA-DARE is a service provided by the library of the University of Amsterdam (https://dare.uva.nl)

Download date:07 Oct 2026

arXiv:2006.12999v1 [cs.AI] 19 Jun 2020

Optimizing Interactive Systems via Data-Driven Objectives

Ziming Li Z.LI@UVA.NL
University of Amsterdam
The Netherlands

Julia Kiseleva JULIA.KISELEVA@MICROSOFT.COM
Microsoft Research
Redmond, USA

Alekh Agarwal ALEKHA@MICROSOFT.COM
Microsoft Research
Redmond, USA

Maarten de Rijke M.DERIJKE@UVA.NL
University of Amsterdam
The Netherlands

Ryen W. White RYENW@MICROSOFT.COM
Microsoft Research
Redmond, USA

Abstract

Effective optimization is essential for real world interactive systems to provide a satisfactory user experience in response to changing user behavior. However, it is often challenging to find an objective to optimize for interactive systems (e.g., policy learning in task-oriented dialog systems). Generally, such objectives are manually crafted and rarely capture complex user needs in an accurate manner. We propose an approach that infers the objective directly from observed user interactions. These inferences can be made regardless of prior knowledge and across different types of user behavior. We introduce Interactive System Optimizer (ISO), a novel algorithm that uses these inferred objectives for optimization. Our main contribution is a new general principled approach to optimizing interactive systems using data-driven objectives. We demonstrate the high effectiveness of ISO over several simulations.

Keywords: interactive systems, reinforcement learning, reward learning

1. Introduction

Interactive systems (White, 2016) play an important role in assisting people in a wide range of tasks. For instance, if users are seeking information, interactive systems can assist them in the form of web search engines (Williams et al., 2016b; Borisov et al., 2016; Dehghani et al., 2017; Williams et al., 2016a), dialog systems (Li et al., 2016, 2019; Dhingra et al., 2016; Williams et al., 2017; Peng et al., 2018), digital assistants (Kiseleva et al., 2016b,a; Kiseleva and de Rijke, 2017; ter Hoeve et al., 2020), recommender systems (Schnabel et al., 2019; Sepliarskaia et al., 2018), or virtual reality (Argelaguet et al., 2016). The described instances of interactive systems can be considered as examples of machine learning applications where the goal is to assist users in real world day-to-day tasks. These systems are characterized by repeated interactions with humans which follow the request-response schema, where the user takes an action, followed by a response from the interactive system. Such interactions can continue for several iterations until the user decides to stop, e.g., when they are either

satisfied or frustrated with their experience. Interaction with the system produces traces/trajectories of user interactions. Importantly, an interactive system and its users always have a shared goal: for users to have the best experience in the premise of successfully completing the user's task.

Thus, both a system and its users are expected to behave accordingly, e.g., a searcher issues a query that he expects will lead him to the desired results and the interactive system provides the search results that are most helpful to him. However, despite their shared goal, only the user can observe their own experience, leaving interactive systems unable to directly optimize their own behavior. Given as an example, in most cases, users will not leave explicit rates about their experience after interacting with digital assistants and this brings difficulties in optimizing the system.

Understanding user objectives and acting accordingly has been shown to be a difficult task, even for humans (Perner and Lang, 1999). However, studies in behavioral economics provide supporting evidence that users intend to maximize expected utility or minimize expected cost and effort (Varian, 1999; Lovett, 2006; Blume et al., 2008). Following this line of research, in this paper, we assume that the general population of users has a shared goal that is achieved through interactions, so user behaviors are aligned with their preferences but rational noises are allowed. We call such behavior *approximately rational*. Obviously, the exact user utility function is inherently complex, but we can approximate it via some meaningful *objective function*, recovered directly from observed traces of user interactions. A similar principle has been successfully employed in robotics (Jeon et al., 2020; Reddy et al., 2019), and understanding user behavior on web (Azzopardi, 2014; Kosinski et al., 2013; Wei et al., 2017). Hence, knowing the *approximate* user objective function can help us to improve the flow of interactive systems.

Currently, optimizing interactive systems relies on explicit assumptions about users' objectives in terms of their needs and frustrations (Li et al., 2017b). Commonly, an objective function is manually designed for a particular task to reflect the quality of an interactive system, e.g., in terms of user satisfaction (Kelly, 2009, 2015), user effort (Yilmaz et al., 2014) or other domain-specific metrics, such as relevance judgements in information retrieval (Järvelin and Kekäläinen, 2002; Saracevic, 1975; Saracevic et al., 1988; Drutsa et al., 2015; Dupret and Lalmas, 2013), user feedbacks (e.g. click, order, skip) in recommender systems (Zhao et al., 2018; Zheng et al., 2018; Chen et al., 2019).

The drawbacks of this approach are that a handcrafted objective function is heavily based on domain knowledge, that it is expensive to maintain, and that it does not generalize over different tasks, e.g., clicks on search results, gestures for mobile digital assistants (Kiseleva et al., 2016a; Williams et al., 2016b), the cross-entropy between generated replies and predefined answers (Li et al., 2016; Cui et al., 2019). Consequently, manually crafted objective functions rarely correspond to the actual user experience. Therefore, even an interactive system that maximizes an objective function is not expected to provide an optimal experience as long as that objective function is hand-crafted. Moreover, it is impossible to design such functions when there is a lack of domain knowledge. Also, we have witnessed how the badly designed objective functions can lead to wrong results. For example, Liu et al. (2016) validated that applying evaluation metrics (e.g., BLEU score (Papineni et al., 2002)) in the machine-translation field to dialog systems is problematic because there is significant diversity in the space of valid responses to a given context.

Given an objective function, optimization can be done following the Reinforcement Learning (RL) paradigm (Sutton and Barto, 2018), which is successfully applied to physically constrained environments (Silver et al., 2016; Levine et al., 2016a,b; Finn et al., 2017). The majority of previous work in the area of interactive systems does this by considering the interactive system as the agent and the underlying stochastic environment induced by the user (Hofmann et al., 2013a; Li et al.,

2016; Peng et al., 2017; Lipton et al., 2018; Su et al., 2018) where the system policies are optimized by interacting with real users or user simulators. However, this setup does not allow us to apply the principle, outlined earlier, that it is a user (not the interactive system) who is getting reward by interacting with the system while maximizing their utility. Recently, Leike et al. (2018) showed how agent alignment, cast in an RL framework, can be applied for optimizing general purpose interactive systems via reward modeling. Jeon et al. (2020) and Reddy et al. (2019) demonstrate how this approach can be applied in the robotics domain. However, this setup requires a quantity of user feedback that may not always be available in practice (Fox et al., 2005; Joachims et al., 2005), which leave us with unlabeled user trajectories.

In this paper, we assume that users continue their interactions with the system if their goals are fulfilled[1], so they are getting rewards after each action. we propose a general perspective on how to improve interactive systems by simultaneously (1) inferring an objective function directly from data, namely unlabeled trajectories of user interactions with the system, and (2) iteratively, and step by step, optimizing the system for this data-driven objective. Since users have difficulties in comprehending dramatic changes in an interactive system (Mitchell and Shneiderman, 1989; White et al., 2002; Obendorf et al., 2007; Teevan, 2008), changes should be made gradually so as to let users adapt to a newly optimized interactive system. The proposed setup is schematically outlined in Figure 1. It embodies a principled approach by concurrently inferring data-driven objectives from user interactions and optimizing the interactive system accordingly. Thus, our approach does not depend on any domain knowledge.

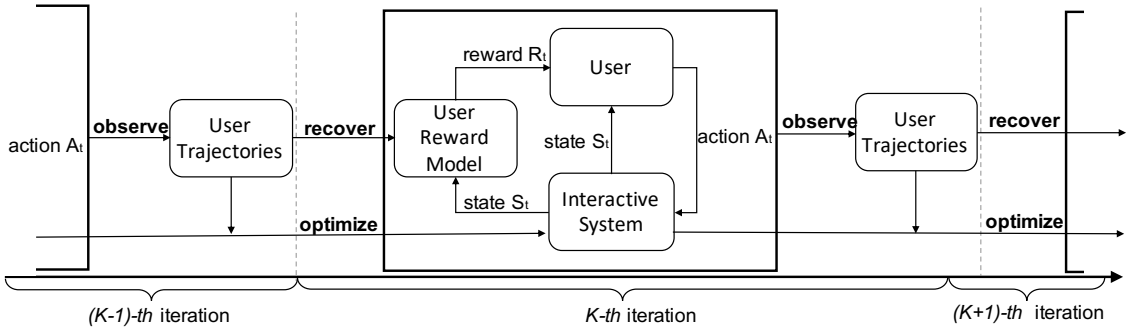

Figure 1: Schematic illustration of the proposed setup of iterative gradual optimization of the flow of the interactive system: the user reward model is *recovered* from the logs collected while users are interacting with a interactive system; Interactive System Optimizer (ISO) is used to *optimize* the interactive system at each iteration.

Below, we start by outlining relevant research areas (Section 2). Then we make the following contributions:

- A new way of modeling user-system interactions, which is depicted as k^{th} iteration in Figure 1 (System 3).
- A novel optimization setup to infer data-driven objectives that accurately reflect the users' needs solely from interaction, without using any domain knowledge to handcraft an optimizing goal, which is partially reflected by the arrow 'recover' in Figure 1 (Section 4).

1. or at least partially fulfilled

- A novel algorithm, ISO, that optimizes an interactive system through data-driven objectives, which is depicted as the arrow labeled 'optimize' in Figure 1 (Section 5).
- To validate the success of the proposed method, we apply it to two different simulated interactive systems. We show how the proposed optimizer can improve the system performance in the designed setups. We also show that by inferring user reward functions, we can optimize the interactive system without real users in the loop and real users are only involved while collecting user-system interaction trajectories (Section 6).

2. Related Work

Relevant work for this paper comes in two broad strands: how to optimize interactive systems (Section 2.1) and what reward signal can be used for optimization (Section 2.2).

2.1 Optimizing Interactive Systems

The flow of interactive systems (White, 2016) can be improved by direct and indirect optimization. Direct optimization aims to maximize user satisfaction directly; in contrast, indirect optimization solves a related problem while hoping that its solution also maximizes user satisfaction (Dehghani et al., 2017). Direct optimization can be performed using supervised learning or RL (Mohri et al., 2012). In RL, an agent learns to alter its behavior through trial-and-error interactions with its environment (Sutton and Barto, 1998). The goal of the agent is to learn a policy that maximizes the expected return. RL algorithms have successfully been applied to areas ranging from traditional games to robotics (Mnih et al., 2015; Silver et al., 2016; Levine et al., 2016a,b; Duan et al., 2016; Wang et al., 2016; Zhu et al., 2017; Schulman et al., 2017; Haarnoja et al., 2018; Vinyals et al., 2019; Akkaya et al., 2019; Hafner et al., 2019; Schrittwieser et al., 2019).

Many applications of RL to optimizing interactive systems come from such fields as Information Retrieval (IR), recommender systems, and dialogue systems. General assumption of users trying to maximize their utility proposed in Reddy et al. (2019); Jeon et al. (2020) holds for interactive systems as well Azzopardi (2014). (Hofmann et al., 2011, 2013b) apply RL to optimize IR systems; they use RL for online learning to rank and use interleaving to infer user preferences (Hofmann et al., 2013a). Shani et al. (2005) describe an early MDP-based recommender system and report on its live deployment. Li et al. (2016) apply RL to optimize dialogue systems; in particular, they optimize handcrafted reward signals such as ease of answering, information flow, and semantic coherence. A number of RL methods, including Q-learning (Peng et al., 2017; Lipton et al., 2018; Li et al., 2017a; Su et al., 2018) and policy gradient methods (Dhingra et al., 2016; Williams et al., 2017; Takanobu et al., 2019), have been applied to optimize dialogue policies by interacting with real users or user simulators. With the help of RL, the dialogue agent is able to explore contexts that may not exist in previously observed data. A key component in RL is the quality of the reward signal used to update the agent policy. Most existing RL-based methods require access to a reward signal from user feedback or a predefined reward.

However it still remains non-trivial to apply RL paradigm towards scalable real world machine learning tasks (Leike et al., 2018) due to the lack of general approach of recovering data-driven objectives, which we discuss next.

2.2 Rewards for Interactive Systems

When applying RL to the problem of optimizing interactive systems, we need to have rewards for at least some state-action pairs. Previous work typically handcrafts those, using, e.g., normalized discounted cumulative gain (nDCG) (Odijk et al., 2015) or clicks (Kutlu et al., 2018; Zhao et al., 2018) before the optimization or the evaluation of the algorithm. Instead of handcrafting rewards, we recover them from observed interactions between the user and the interactive system using Inverse Reinforcement Learning (IRL). The main motivation behind IRL is that designing an appropriate reward function for most RL problems is non-trivial; this includes animal and human behavior (Abbeel and Ng, 2004), where the reward function is generally assumed to be fixed and can only be ascertained through empirical investigation. Thus inferring the reward function from historical behavior generated by an agent's policy can be an effective approach. Another motivation comes from imitation learning, where the aim is to teach an agent to behave like an *expert* agent. Instead of directly learning the agent's policy, other work first recovers the expert's reward function and then uses it to generate a policy that maximizes the expected accrued reward (Ng and Russell, 2000). Since the inception of IRL (Russell, 1998), several IRL algorithms have been proposed, including maximum margin approaches (Abbeel and Ng, 2004; Ratliff et al., 2009), and probabilistic approaches such as (Ziebart et al., 2008) and (Boularias et al., 2011). In the last few years, a number of adversarial IRL methods (Finn et al., 2016a,b; Ho and Ermon, 2016; Fu et al., 2017; Qureshi et al., 2019; Seyed Ghasemipour et al., 2019) have been proposed because of its ability to adapt training samples to improve learning efficiency. One more aspect aspect IRL methods are differ is availability of feedback or score for the user trances. Christiano et al. (2017); Leike et al. (2018) suggest setup where system can learn from user feedback which is not always available in practice. In our paper, we tackle the case with no explicit feedback.

Regarding of the applications of IRL, Ziebart et al. (2012) use IRL for predicting the desired target of a partial pointing motion in graphical user interfaces. Monfort et al. (2015) use IRL to predict human motion when interacting with the environment. IRL has also been applied to dialogues to extract the reward function and model the user (Pietquin, 2013; Takanobu et al., 2019; Li et al., 2020, 2019). IRL is used to model user behavior in order to make predictions about it. But we use IRL as a way to recover the rewards from user behavior instead of handcrafting them and optimize an interactive system using these recovered rewards. Lowe et al. (2017) learns a function to evaluate dialogue responses. However, the authors stop at evaluation and do not actually optimize the interactive system.

Recent work (Leike et al., 2018; Zhang and Dragan, 2019; Jeon et al., 2020) demonstrate impressive results and outline new research direction while modeling user-system interaction using the agent alignment problem (Sutton and Barto, 2018). In contrast to our work, the reward modeling heavily relies on a user feedback loop, which is mostly not available in the internet based interactive systems (Kiseleva et al., 2014, 2015).

The key difference between our work and previous studies is that we first use recovered rewards from observed user interactions to reflect user needs and define interactive system objectives. Subsequently the interactive system can be optimized according to the defined data-driven objectives so as to improve the user experience. We regard the interactions between a user and an interactive system as an agent interacting with an changeable environment, where the transition distribution of the environment can be updated. Treating an interactive system as a changeable and programmable environment is novel and reasonable because we have complete control on the behaviors of interactive

systems since we are the system designers. Lowe et al. (2017) learned a function to evaluate dialogue responses but does not actually optimize the interactive system. Leike et al. (2018) formulated the optimization problem in a complex multi-agent setup because their environment is physical and non programmable. Besides, the reward modeling by (Leike et al., 2018) heavily relies on user feedback loop, which is mostly not available in the interactive systems.

3. Modeling User-System Interactions

In this section, we first introduce our assumptions about collaborations between a user and an interactive system (Section 3.1), and then we explain how we model these interactions (Section 3.2).

3.1 Assumptions

Our goal is to design an interactive system that can successfully assist users with completing some real world tasks. We have formulated a set of assumptions to formalize user-system interactions, which are schematically depicted in Figure 1. They can be roughly separated into two groups: assumptions about the system design (S) and assumptions regarding to user goals and behavior (U):

Assumption 1 **S:** system's goal is to accommodate better user experience, namely maximize utility a user gets from the system by minimizing their efforts;

Assumption 2 **S:** system setup allows us to iteratively and gradually improve the system in a sequential manner to accommodate better user experience, given that at the beginning the system provides 'non-zero' utility for users but can be significantly improved;

Assumption 3 **S:** a system designer has the ability to transform an interactive system, but it has some obligatory steps a user needs to take to complete their task due to system design constraints;

Assumption 4 **U:** users have incentives to continue their iterations with a interactive system if they are getting some value from it;

Assumption 5 **U:** users of a interactive system have *approximately* homogeneous behaviour, namely users have a shared notion of utility that can be approximated by some objective function [2];

Assumption 6 **U:** users try to maximize their utility while interacting with a system;

Assumption 7 **U:** users are not required to provide feedback about their experience. However, user actions can be considered as *implicit signals* reflecting their satisfaction/frustration with an interactive system.

3.2 Modeling Interactions

While employing Reinforcement Learning (RL) formalism our Assumption 5 and Assumption 6 can be reformulated as follows: the user is seen the *optimal* agent who interacts with the environment, an interactive system, with the goal of maximizing their expected rewards.

2. Terms 'users' and 'a user' are used interchangeably.

As a *running example* we can consider a user who is interacting with a search engine. The process of user-system interaction is modeled using a finite Markov Decision Process (MDP) (S, A, T, r, γ), in the following way:[3]

- S is a set of states that represent responses from the interactive system to the user. S is finite as there are limited predefined number of responses that interactive system can return.
- A is a finite set of actions that the user can perform on the system to move between states. In case of search engine, a user can run a query, click on the returned results, reformulate a query etc.
- T is a transition distribution and $T(s, a, s')$ is the probability of transitioning from state s to state s' under action a at time t:

$$T(s' \mid s, a) = \mathbb{P}(S_{t+1} = s' \mid S_t = s, A_t = a). \tag{1}$$

  For search engines, being at the start page (which is s) a user is making an action a, e.g. running a query, and the engine redirects him to a result page (which is s').

- $r(s, a, s')$ is the expected immediate reward after transitioning from s to s' by taking action a.

  We compute the expected rewards for (state, action, next state) triples as:

$$r(s, a, s') = \mathbb{E}[R_t \mid S_t = s, A_t = a, S_{t+1} = s'], \tag{2}$$

  where R_t is reward at time t. In case of search engine, a user is getting a reward for finding a desired information. However, the rewards are not observed in practise (Assumption 7). For simplicity in exposition, we write rewards as $r(s)$ rather than $r(s, a, s')$ in our setting; the extension is trivial (Ng and Russell, 2000).

- $\gamma \in (0, 1]$ is a discount factor.

We write $\mathcal{P}$ to denote the set of interactive systems, i.e., triples of the form (S, A, T). Following Assumption 3, system designers have control over the sets S, A, and the transition distribution, T, and T can be changed to optimize an interactive system.

The *user behavior strategy* for accomplishing their tasks is represented by a policy, which is a mapping, $\pi \in \Pi$, from states, $s \in S$, and actions, $a \in A$, to $\pi(a|s)$, which is the probability of performing action $A_t = a$ by the user when in state $S_t = s$. The observed history of interactions between the user and the interactive system, H[4], is represented as a set of trajectories, $\{\zeta_i\}_{i=1}^n$, drawn from a distribution Z, which is brought about by T, π, and D_0, where D_0 is the initial distribution of states. following Assumption 5, which proposes homogeneity in user behavior, simplifies the problem, i.e. as if one user generated H. A *trajectory* is a sequence of state-action pairs, where a user does not provide explicit feedback (Assumption 7):

$$\zeta_i = S_0, A_0, S_1, A_1, \ldots, S_t, A_t, \ldots. \tag{3}$$

To conclude, we suppose that the user is an optimal agent who is trying to maximize its reward under the system dynamics it faces and that the system wants to improve the user experience over

3. We follow the notation proposed in (Sutton and Barto, 2018)
4. H can be referred further as logs of user interactions, or log, or user trajectories/traces.

time by creating progressively easier MDPs to solve for the user. However, a interactive system cannot transition from all initial to goal states in one step due to design constraints. For example, if a user is searching for holiday destinations, the system cannot redirect him to the final stage of booking a hotel because he needs to go through a necessary step, e.g., providing payment details.

To summarize, we have described the basic principles of modeling interactions between users and an interactive system. Next, we detail how to define data-driven objectives that are used to optimize an interactive system.

4. Defining Data-driven Objectives

In this section, we first present our approach to convert user needs to data-driven objectives of an interactive system (Section 4.1), and then we explain how these objectives can be estimated (Section 4.2).

4.1 Defining Interactive System Objectives

We define the *quality* of an interactive system as the expected state value under an optimal user policy. The value of a state S_0 under a policy π is given as (Sutton and Barto, 2018):

$$V^{\pi}(S_0) = \mathbb{E}_{\pi}\left[\sum_{t=0}^{\infty} \gamma^t R_{t+1}\right], \tag{4}$$

where the expectation $\mathbb{E}_{\pi}[\cdot]$ is taken with respect to sequences of states $S_0, S_1, \ldots, S_t, \ldots$ drawn from the policy π and transition distribution T. We use V_T^{π} to denote the value of a policy π under the current transition distribution T, and hide the initial states S_0 for simplification.

In the proposed setting, the user goal is to find the best policy π^* such that V_T^{π} is maximized. $V_*(T)$ defines the maximum possible value of V_T^{π} under transition distribution T as follows:

$$V_*(T) = \max_{\pi \in \Pi} V_T^{\pi}, \tag{5}$$

where Π is the set of possible user policies. We formulate the problem of finding the optimal interactive system's transition distribution, denoted T^*, in the following terms:

$$T^* = \arg\max_{T \in T} V_*(T). \tag{6}$$

Therefore, Eq. 6 represents the objective function, mentioned in Assumption 6, which is derived from user trajectories (Eq. 3) directly. After finding T^* through solving the proposed optimization problem, the system designer has the ability to transform the current system to a new one, which should deliver a better user experience as it reflects user needs better. This process is illustrated in Figure 1 by the arrow marked *optimize* between two consecutive iterations.

With the transition distribution T, the interactive system will respond with the next state s' given the current state s and the user action a. In real life, it is not guaranteed that the tuple (s, a, s') exists. For example, in task-oriented dialog systems, the system first needs to collect essential information for booking a hotel (e.g., hotel name, room type) step by step. In some cases, the system also needs to recommend potential hotels and asks the user to make a choice. After successfully collected all information, the system can guide the user to a payment page. Obviously, it is not possible

to deliver a payment state to the user when the information contained in the current state is not complete. Therefore, inherent constraints exist in interactive systems and this makes finding the optimal interactive system's transition distribution a meaningful and interesting task. Otherwise, the system can always deliver the most valuable state to the user in one step at any states.

To estimate the data-driven objectives interactive system presented in Eq. 6, we first need to recover R_t, which we will discuss next.

4.2 Recovering User Rewards

Assumption 4 suggests that continued user interactions with the system indicate a certain level of user satisfaction, which can be reflected by experienced rewards. In contrast with $\zeta_i \in H$ presented in Eq. 3, the complete history of interactions, $\hat{H}$, consists of trajectories $\hat{\zeta}_i \sim \hat{Z}$, which include the user reward R_t:

$$\hat{\zeta}_i = S_0, A_0, R_1, S_1, A_1, R_2 \ldots, R_t, S_t, A_t, \ldots. \tag{7}$$

The problem is that the true user reward function is hidden from a interactive system and inherently difficult due to the complexity of the real world surrounding users. Our goal is using the collected incomplete user trajectories, H, shown in Eq. 3, to find a way to approximate true user rewards. To address this challenge we apply Inverse Reinforcement Learning (IRL) methods[5], which are proposed to recover the rewards of different states, $r(s)$, for $\zeta_i \in H$. Our assumption about the form of user reward function is: given state feature functions $\phi : S_t \to \mathbb{R}^k$ that describe S_t as a k-dimensional feature vector, the true reward function $r(s)$ is a linear combination of the state features $\phi(s)$, which can be given as $r(s) = \theta^T \phi(s)$. To uncover the reward weights θ, we employ the following approaches.

Maximum Entropy Inverse Reinforcement Learning (MaxEnt-IRL): The core idea of MaxEnt-IRL Ziebart et al. (2008) is that trajectories with equivalent rewards have equal probability to be selected and trajectories with higher rewards are exponentially more preferred, which can be formulated as:

$$\mathbb{P}(\zeta_i \mid \theta) = \frac{1}{Z(\theta)} \exp(\theta^T \phi(\zeta_i)) = \frac{1}{Z(\theta)} \exp(\sum_{t=0}^{|\zeta_i|-1} \theta^T \phi(S_t)), \tag{8}$$

where $Z(\theta)$ is the partition function. MaxEnt-IRL maximizes the likelihood of the observed data under the maximum entropy (exponential family) distribution.

Adversarial Inverse Reinforcement Learning (AIRL): Based on MaxEnt-IRL, (Finn et al., 2016b) combine sample-based MaxEnt-IRL with forward reinforcement learning to estimate the partition function Z, where:

$$L(\theta) = -\,\mathbb{E}_{\zeta_i \sim p}\, r_\theta(\zeta_i) + \log\left(\mathbb{E}_{\zeta_j \sim q}\left[\frac{\exp(r_\theta(\zeta_j))}{q(\zeta_j)}\right]\right). \tag{9}$$

Here, $r_\theta(\zeta_i)$ is the reward of trajectory ζ_i, p represents the distribution of demonstrated samples, while q is the background distribution for estimating the partition function $\int \exp(r_\theta(\zeta))d\zeta$. Due to

5. IRL methods are described in greater details in Section 2.2

high variance from operating over entire trajectories, Fu et al. (2017) extend the algorithm to single state-action pairs and the proposed method, AIRL, which is a practical and scalable IRL algorithm based on an adversarial reward learning formulation. We use AIRL to recover the reward function for complex interactive systems since AIRL can estimate non-linear reward functions.

Distance Minimization Inverse Reinforcement Learning (DM-IRL): For completeness, we also employ DM-IRL (El Asri et al., 2013; Burchfiel et al., 2016), which deal with scored trajectories, to have a case of perfectly recovered reward weight θ for comparison. DM-IRL directly attempts to regress the user's actual reward function that explains the given score. DM-IRL uses discounted accrued features to represent the trajectory: $\psi(\zeta_i) = \sum_{t=0}^{|\zeta_i|-1} \gamma^t \phi(S_t)$, where γ is the discount factor. The score of a trajectory ζ_i is $\text{score}_{\zeta_i} = \theta^T \psi(\zeta_i)$. Since the exact score for each trajectory is supplied, the recovered rewards with DM-IRL are exactly the ground truth of reward functions, which can be regarded as oracle rewards.
This process is depicted in Figure 1 by the arrow marked *recover*. Once we have recovered the reward function $r(s)$, we can proceed to the optimization objectives presented in Eq. 6.

5. Optimizing Interactive Systems

We start by explaining how to maximize the quality of an interactive system for a user behaving according to a fixed stationary policy π:

$$T_\pi^* = \arg\max_{T \in T} V_T^\pi. \tag{10}$$

To solve this problem, we first build an MDP as proposed above, where the user is the agent and the system is the environment Following Assumption 2, the system can be optimized to improve the user experience which we characterized by the quality of the interactive system. This problem is *equivalent* to finding the optimal policy in a reformulated MDP$^+(S^+, A^+, T^+, r^+, \gamma^+)$, where the agent is an interactive system and the stochastic environment is a user. It should be noted that the roles of agent and environment in the reformulated MDP$^+$ are exactly the opposite of the roles in the original MDP. We also convert the state space and action space correspondingly. We rely on the first MDP for inferring the user reward functions, while we rely on the second one, MDP$^+$, for updating interactive systems. In MDP$^+$, the state S_t^+ is represented by a concatenation of the state S_t the user is in and the action A_t the user takes at time step t from the original MDP; the action A_t^+ is the original state S_{t+1}. The interactive system observes the current state S_t^+ and picks an action A_t^+ under the interactive system policy $\pi^+(A_t^+|S_t^+)$. Then the user returns the next state S_{t+1}^+ according to the transition distribution $T^+(S_{t+1}^+|S_t^+, A_t^+)$ which is inferred from the policy model $\pi(A_{t+1}|S_{t+1})$. Therefore, finding the optimal transition T_π^* from Eq. 10 is equivalent to finding the optimal policy π_*^+ in the reformulated MDP$^+$ as follows:

$$\pi_*^+ = \arg\max_{\pi^+ \in \Pi^+} V_{T^+}^{\pi^+}, \tag{11}$$

which can be done using an appropriate RL method such as Q-learning or Policy Gradient. D_0^+ is the initial distribution of states in MDP$^+$. After we have demonstrated how to optimize the interactive system for a given stationary policy, we return to the original problem of optimizing the interactive system for an optimal policy π_*.

Algorithm 1 Interactive System Optimizer (ISO)

1: **Input:** Original system (S, A, T), r, γ, D_0.
2: Construct original MDP(S, A, T, r, γ)
3: $\pi_*(a|s) = RL(S, A, T, r, \gamma)$ *// finding the current user policy*
4: Construct system MDP$^+(S^+, A^+, T^+, r^+, \gamma^+)$: *// reformulate the original MDP by switching the roles of agent and environment*
   - $S_t^+ = S_t \oplus A_t$ *// build the new state space by concatenation*
   - $A_t^+ = S_{t+1}$ *// build the new action space*
   - $T^+(S_{t+1}^+|S_t^+, A_t^+) = \pi_*(A_{t+1}|S_{t+1})$ *// build the transitions in* MDP^+
   - $r(S_t^+)^+ = r(S_t)$ *// convert the reward function*
   - $\gamma^+ = \gamma$ *// both MDPs share the same discount factor*
5: $D_0^+ \sim (S_0 \sim D_0, A_0 \sim \pi_*(a|S_0))$ *// sample initial states in* MDP^+
6: $\pi^+(A_t^+|S_t^+) = T(S_{t+1}|S_t, A_t)$ *// find the optimal transition distribution in the original MDP is formulated as fining the optimal policy in a reformulated* MDP^+
7: $\pi_*^+(a^+|s^+) = RL(S^+, A^+, T^+, r^+, \gamma^+)$ *// optimize the system policy in* MDP^+
8: $T^*(S_{t+1}|S_t, A_t) = \pi_*^+(A_t^+|S_t^+)$ *// replace the transition distribution in original MDP with the newly updated system policy*
9: **Output:** Optimized system (S, A, T^*)

To summarize, we propose a formal procedure for optimizing interactive systems, called ISO, presented in Algorithm 1, with the following steps:

Line 1 We assume that we have an estimate of the reward function $r(s)$ using one of the IRL methods described in Section 4.2. So we have as input: the original system (S, A, T), the reward function r, the discount factor γ, and the initial distribution of states D_0.

Line 2 ISO formulates the original system as MDP(S, A, T, r, γ).

Line 3 ISO uses an appropriate RL algorithm to find the current user policy $\pi_*(a|s)$ given the reward function r.

Line 4 ISO transforms the original MDP(S, A, T, r, γ) into the new one MDP$^+(S^+, A^+, T^+, r^+, \gamma^+)$, where the roles of the agent and environment are switched. In our setting, S_t^+ has the same reward value as S_t. The discount factor γ^+ remains the same.

Line 5 ISO transforms D_0 to D_0^+ to match the distribution of first state-action pairs

Line 6 The equivalence $\pi^+(A_t^+|S_t^+) = T(S_{t+1}|A_t, S_t)$ means that finding the optimal π_*^+ according to Eq. 11 is equivalent to finding the optimal T_π^* according to Eq. 10. Therefore, the transition distribution can be regarded as a policy network or a policy table from MDP's perspective depending on the policy learning method.

Line 7 We can use an appropriate RL algorithm to find $\pi_*^+(A_t^+|S_t^+)$.

Line 8 ISO extracts $T^*(S_{t+1}|S_t, A_t)$ from the optimal system policy $\pi_*^+(A_t^+|S_t^+)$. The extraction process is trivial: $T^*(S_{t+1}|S_t, A_t) = \pi^+(A_t^+|S_t^+)$. Then, ISO terminates by returning the *optimized* interactive system.

Line 9 ISO outputs the optimized interactive system (S, A, T^*).

Once ISO has delivered the optimized system (S, A, T^*), we expose it to users so they can interact with it as illustrated in Figure 1. We assume that users adjust their policy to T^*. After enough iterations the user policy will converge to the optimal one. Iterations between optimizing the interactive system for the current policy and updating the user policy for the current interactive system continue until both converge.

In summary, we have presented the Interactive System Optimizer (ISO). It optimizes an interactive system using data-driven objectives. It works by transforming the original MDP, solving it and using its solution to yield the optimal transition distribution in the original MDP.

6. Experiments and Results

In this section, we apply our proposed method, ISO, to two different simulated interactive setups. In the first setup, the interactive system operates in a tabular-based world with finite states and actions (Section 6.1). The second one has a more realistic setup, where the agent, the environment, and the reward function are all represented by neural separate networks (Section 6.2). Each proposed experimental setup is described with regard to the three following components: the design of the interactive system, modeling user behavior, and suitable evaluation process. For both experimental setups, our results demonstrate that ISO can improve significantly improve the system performance in the designed setups. We conclude this section openly discussing a list of foreseeing limitations (Section 6.3).

6.1 Optimizing Interactive Systems in a Tabular-based World

6.1.1 Experimental Setup

Designing an Interactive System We simulate an arbitrary interactive system where we need a finite set of states S, a finite set of actions A and a transition distribution T. Features of a state $\phi(s)$ are fixed. For our experimental setup, we simulate the interactive system where $|S| = 64$ and $|A| = 4$. We work with a complex environment where a user can transition between any two states if these two states are connected. The connections between two states are predefined and fixed, but the transition distribution is changeable. In another word, for the same system in different runs, the connectivity graph of this systems is fixed and will not be changed once it is sampled at the very beginning. This setup corresponds to the inherent constraints between state transitions in real interactive systems Section 4.1. We use a hyper-parameter, the connection factor cf, to define the number of possible next states after the user has taken one specific action at the current state. For an initial interactive system, D_0 is randomly sampled as well as T. At each iteration ISO delivers T^*, which substitutes the initial T obtained at the previous iteration. The *optimized* interactive system is used for the next iteration until the process converges.

Modeling User Behavior To model user behavior we require a true reward function $r_{real}(s)$, and an optimal user policy π^*_{user}. We utilize a linear reward function $r_{real}(s)$ by randomly assigning 25% of the states reward 1, while all others receive 0. As we use one-hot features for each state, $r_{real}(s)$ is guaranteed to be linear.

We use a soft value iteration method (Ziebart, 2010) to obtain the optimal user policy π^*_{user}.

The quality of the recovered reward functions is influenced by how trajectories are created, which in turn can affect the performance of ISO as it relies on $r_{real}(s)$ to optimize the transition distribution T behind the interactive system with reinforcement learning.

We experiment with the following types of user trajectories:

- *Optimal:* Users know how to behave optimally in an interactive system to satisfy their needs. To simulate the user interactions H, we use π^*_{user} trained with the real reward function $r_{real}(s)$.

- *SubOptimal:* Not all users know the system well, which means that the demonstrated behavior is a mixture of optimal and random. We propose two different methods to simulate suboptimal behavior. The degree of optimality of user behavior is controlled by either of two following factors: (1) the proportion of random behavior (this is called 'wandering' behavior in (White et al., 2005)); or (2) the user action noise, which are collectively called the noise factor (NF) $\in [0.0, 1.0]$.

  Mix of Behaviors (MB): The log of user interactions H is a mix of trajectories generated by the optimal policy and the adversarial policy[6].

  Noise in Behavior (NB): In this case, the trajectories in H are generated from the optimal policy but we add noise to the user actions to get suboptimal behavior [7].

The generated history of user interactions H represents the case of trajectories without a score which will be fed to MaxEnt-IRL. In terms of DM-IRL, interaction history should be given along with scores for each trajectory – $\hat{H}$. To generate the required dataset $\hat{H}$, we calculate the score using the true reward function $r_{real}(s)$. $\hat{H}$ is the input to DM-IRL.

At each iteration, we sample the following datasets reflecting different types of history of user interactions: $\hat{H}$, $H_{Optimal}$, $H_{SubOptimal-0.2-MB}$, $H_{SubOptimal-0.6-MB}$, $H_{SubOptimal-0.2-NB}$, $H_{SubOptimal-0.6-NB}$ each of size $15,000$ and $|\zeta_i| \in [30, 40]$.

Evaluation Process To evaluate the performance of ISO, we report the *expected state value* under optimal policy Eq. 5 for an *initial* interactive system (S, A, T_{init}) and an *optimized* one (S, A, T_{opt}), which we derive after around 100 *iterations* (one *iteration* means we sequentially recover the reward function and run Algorithm 1 once).

A higher expected state value means users are more satisfied while interacting with the interactive system. We initialize 40 different initial interactive systems by randomly sampling reward functions and transition distribution, and report the overall performance over these 40 systems.

The true reward functions and the connectivity graphs of these sampled systems are fixed in the whole optimizing process. We use the recovered reward function with DM-IRL as the oracle reward for in this setup.

6.1.2 Results and Discussion

6. To model suboptimal user behavior we use two user policies: (1) an optimal user policy π^*_{user}; and (2) an adversarial policy $(1 - \pi^*_{user})$, which means we choose the action that has the lowest likelihood according to π^*_{user}. We include an adversarial policy instead of a random one because it is the hardest case as users behave opposite of what we expect. E.g., $NF = 0.2$ means that 20% of the trajectories are generated with the adversarial policy.
7. E.g., $NF = 0.2$ means the probability is 20% that the user will not choose the action with the highest probability in the optimal policy.

Improving Interactive Systems with ISO Figure 2,3,4 show how the quality of the interactive system increases with each iteration of ISO in terms of different connection factors. The final relative improvements after optimization can be found in Table 1. We use IRL-lableled to reprsent the system optimized with the recovered reward function by method DM-IRL. As expected, when the user gives feedback about the quality of the trajectories (IRL-labelled), the task is simpler and ISO manages to get high improvements with the oracle rewards. However, the picture changes when we hide the scores from the trajectories. Without scores, ISO relies on the optimality of user behavior to recover the reward function. As the optimality decreases, so does the behavior of ISO, and the performance decays. With the oracle rewards from DM-IRL, ISO converges quite fast – as we can see in Figure 3 and Figure 4, after 20 iterations the expected state value begins to plateau. Most improvements happen in the first several iterations. Thus, ISO works with accurately labeled trajectories, but usually obtaining high-quality scores is intractable and expensive in a real interactive system because the real rewards are invisible. We report it as the oracle performance in our experiment.

With respect to trajectories without scores, ISO is able to improve the initial expected state value.

In Figure 2, the influence of the noise factor and types of trajectories (MB or NB) is clear. However, in Figure 4 where there are fewer connections between two states, only the convergence speeds of different curves are different but they all converge to the same state value eventually. ISO manages to optimize the interactive system even though the user trajectories are quite noisy.

More remarkable, the convergence speed and final converged values are different depending on the connection factors. As we can see, it is more difficult to get high performance when there are more connections between different states in the predefined systems. More connections mean that more possible trajectories could be taken and it is intractable for MaxEnt-IRL to learn a reward function from this kind of situation. By contrast, in Figure 4, each state-action pair can only have two possible next states and the final average state value is much higher than the system in Figure 2.

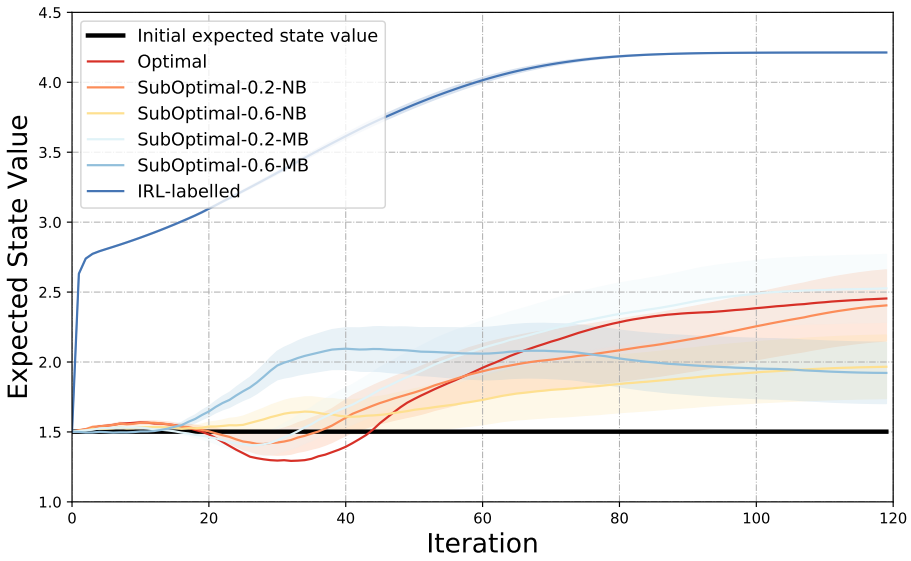

Figure 2: Performance of ISO over 40 randomly sampled systems when connection_factor=32. The error bounds denote the standard error of the mean (±SEM). The x-axis is the number of iterations of ISO and the y-axis is the expected state value.

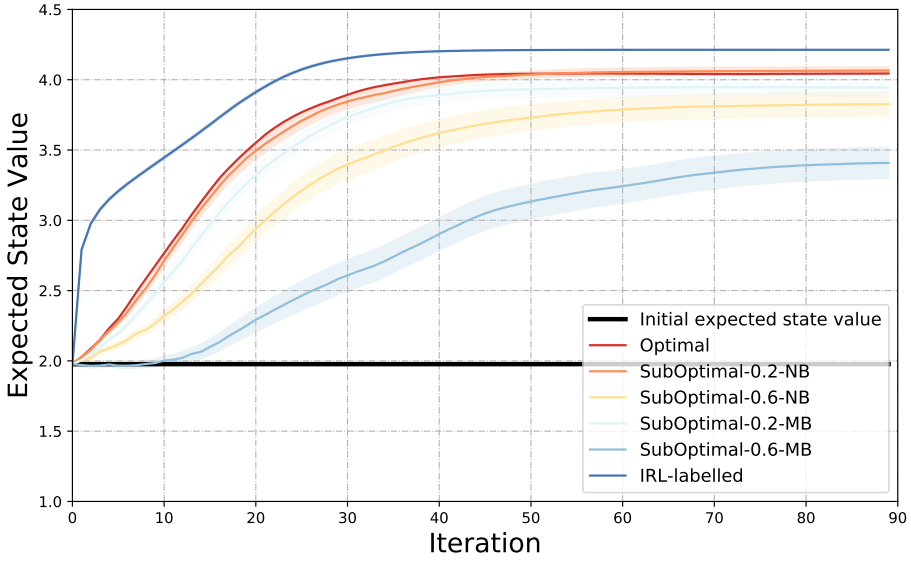

Figure 3: Performance of ISO over 40 randomly sampled systems when connection_factor=8. The error bounds denote the standard error of the mean (±SEM). The x-axis is the number of iterations of ISO and the y-axis is the expected state value.

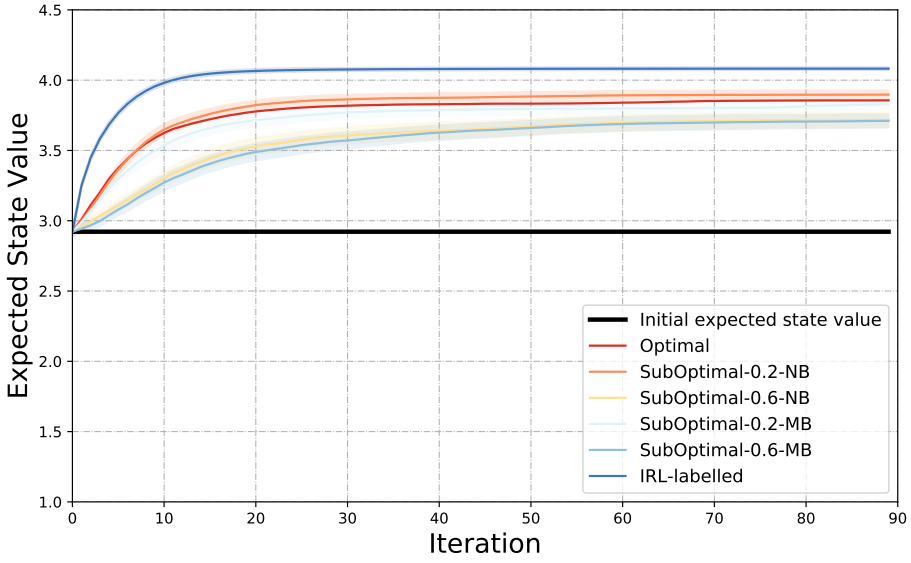

Figure 4: Performance of ISO over 40 randomly sampled systems when connection_factor=2. The error bounds denote the standard error of the mean (±SEM). The x-axis is the number of iterations of ISO and the y-axis is the expected state value.

Impact of ISO Components The performance of ISO depends on its two components: (1) RL methods used to optimize the user policy π_{user} for the original MDP and system policy π^{+}_{sys} for the reformulated MDP^{+}; and (2) IRL methods – to estimate the true reward function $r_{real}(s)$. The dependence on RL methods is obvious – the end result will only be as good as the quality of the final optimization, so an appropriate method should be used. The performance of ISO can be influenced by the quality of the recovered reward functions, $r(s)$.

CF \ BT	(a) IRL-labelled			(b) Optimal (NF=0.0)			(c) SubOptimal-0.2-MB (NF=0.2)		
	Initial	Optimized	Impr.	Initial	Optimized	Impr.	Initial	Optimized	Impr.
32	1.50	4.21	281%*	1.50	2.45	164%*	1.50	2.52	169%*
8	1.98	4.21	213%*	1.98	4.04	205%*	1.98	3.94	200%*
2	2.92	4.08	140%*	2.92	3.86	132%*	2.92	3.83	131%*

CF \ BT	(d) SubOptimal-0.6-MB (NF=0.6)			(e) SubOptimal-0.2-NB (NF=0.2)			(f) SubOptimal-0.6-NB (NF=0.6)		
	Initial	Optimized	Impr.	Initial	Optimized	Impr.	Initial	Optimized	Impr.
32	1.50	1.92	128%*	1.50	2.40	160%*	1.50	1.96	131%*
8	1.98	3.41	173%*	1.98	4.06	206%*	1.98	3.83	194%*
2	2.92	3.71	127%*	2.92	3.90	134%*	2.92	3.72	128%*

Table 1: The performance of ISO, measured as relative improvement (Impr.) in expected state value over the Initial interactive system of the Optimized version (after 120 and 90 iterations) for different types of user behaviors: (a) IRL-labelled, (b) Optimal, (c) SubOptimal-0.2-MB, (d) SubOptimal-0.6-MB, (e) SubOptimal-0.2-NB, (f) SubOptimal-0.6-NB. Only IRL-labelled has access to the trajectory labels. * indicates statistically significant changes ($p < 0.01$) using a paired t-test over the initial expected state value and the optimized expected state value.

For the case of labeled trajectories, the values of $r(s)$ recovered by DM-IRL are identical to the ground truth $r_{real}(s)$ since a regression model is used and we have the exact score for each user trajectory.

For the case of trajectories without scores, the quality of the recovered reward function is worse than DM-IRL. MaxEnt-IRL can only give a general overview of $r_{real}(s)$ if the user trajectories are optimal. If there are not enough constraints on the connections between states, with each iteration of running ISO, the shape of the sampled trajectories becomes more similar, which means that most trajectories pass by the same states and the diversity of trajectories decreases. We found that this makes it even more difficult to recover $r_{real}(s)$ and the MaxEnt-IRL quality deteriorates with the number of iterations, which results in lower performance in Figure 2.

Hence, improving the performance of IRL methods is likely to significantly boost the performance of ISO and more advanced IRL methods could be adopted according to the real task.

6.2 Optimizing Interactive Systems in a Network-based World

6.2.1 Experimental Setup

Designing an Interactive System with Neural Networks In this setup, we first present a simulated framework used for optimizing the interactive system (S, A, T) with ISO. Based on the two-step optimization setup in Section 5 we designed two separate optimizing modules respectively. Figure 5 shows the architecture of the optimizing module for the original MDP(S, A, T, r, γ), while Figure 6 describes the optimizing module for the reformulated MDP$^+(S^+, A^+, T^+, r^+, \gamma^+)$ respectively. As described in Section 5, we use MDP(S, A, T, r, γ) for reward learning and the reformulated MDP$^+(S^+, A^+, T^+, r^+, \gamma^+)$ for system optimization.

In the proposed setup, we have continuous state space S and discrete action space A for the original MDP(S, A, T, r, γ), where the dimension of S is $S_{dim} = 50$ and action number is $|A| = 10$.

The user policy π_{user}, the system policy π_{sys} and the reward function $r(s)$ are represented with multi-layer perceptrons separately. Following Algorithm 1, the transition distribution $T(S_{t+1}|S_t, A_t)$ is exactly the system policy π_{sys} which is fixed in this step. We assume the state distribution follows multivariate Gaussian distribution with a diagonal covariance matrix and the system policy π^+_{sys} will produce the corresponding mean and variance. Since state space S is continuous, the output of π_{sys} will be a sampled continuous state s_{t+1} at next step $t+1$ given s_t and a_t. Here we use Proximal Policy Optimization (PPO) (Schulman et al., 2017), a policy gradient based method, to optimize user policy π_{user}. With respect to the reformulated MDP$^+(S^+, A^+, T^+, r^+, \gamma^+)$, the state s_t and action a_t from the original MDP will be concatenated to form the new state s^+_t following Algorithm 1. The action a^+ is continuous and the transition distribution $T^+(S^+_{t+1}|S^+_t, A^+_t)$ is exactly the user policy π_{user} in the original MDP. Different from π_{sys} in the original MDP, π^+_{sys} will be updated with PPO and it will be used to replace π_{sys} in the original MDP after optimization finished. $r^+(s)$ is the learned reward function in the first optimizing step.

With respect to the optimizing module for MDP(S, A, T, r, γ) shown in Figure 5, the user policy is wrapped up with a PPO agent and the reward function is loaded to the reward agent. To estimate the user reward function, we utilize Adversarial Inverse Reinforcement Learning (AIRL) in the reward learning step. The user policy agent and the reward agent make up of the AIRL agent. As an adversarial learning method, the AIRL agent needs user traces generated by real user to update the reward function. In this setup, we use the user agent $\pi^*_{user_real}$ trained with the true reward function $r_{real}(s)$ to produce necessary user-system interaction traces, which will be stored in the *Expert Behavior* area. The environment *Environment-1* for AIRL training and behavior generation mainly consists of the system policy π_{sys} to deliver the next state s_{t+1} given state s_t and action a_t according to $T(S_{t+1}|S_t, A_t)$. The system policy π_{sys} will keep fixed in MDP(S, A, T, r, γ). It should be noted that there are two different user reward functions in Figure 5. The reward function $r_{airl}(s)$ in AIRL agent is updated during AIRL training while the reward function $r_{real}(s)$ in the expert agent is the true reward function. The AIRL agent and system policy has no access to the true reward function $r_{real}(s)$ and we use $r_{airl}(s)$ to approximate $r_{real}(s)$, which is also the motivation of AIRL.

The optimizing module for the reformulated MDP$^+(S^+, A^+, T^+, r^+, \gamma^+)$ shown in Figure 6 is responsible of updating the system policy $\pi_{sys}$with the recovered reward function $r_{airl}(s)$. Just like other reinforcement learning setups, it mainly has three components: environment, PPO agent, and the reward function. The system policy π_{sys} is wrapped up with a PPO agent and the reward agent is the function $r_{airl}(s)$ learned in MDP(S, A, T, r, γ). Given state s^+_t and action a^+_t, the step function of the environment *Environment-2* will return the next state s^+_{t+1} according to $T^+(S^+_{t+1}|S^+_t, A^+_t)$ in Line 4 of Algorithm 1, where the user policy π_{user} is involved.

Modeling User Behavior Given the current system policy π_{sys} and the real user reward function $r_{real}(s)$, we optimize the user policy π_{user} by running PPO method. The optimized user policy will be saved as the oracle user policy $\pi^*_{user_real}$. Then by making the user policy $\pi^*_{user_real}$ interact with the system policy π_{sys}, we can collect a bunch of interaction trajectories ($20K$ in our experiments) and all these behavior data will be loaded to the expert behavior bucket. The maximum length of collected trajectories is 40. The stored user interaction traces will be used to learn the user reward function $r_{airl}(s)$ (we use AIRL method in this setup).

Evaluation Process To evaluate the performance of ISO in the proposed framework, we report the *Average Return* of m sampled trajectories ($m = 1000$ in our setup) under optimal policy $\pi^*_{user_real}$ under the real reward $r_{real}(s)$ for an *initial* interactive system and an *optimized* one, which we derive

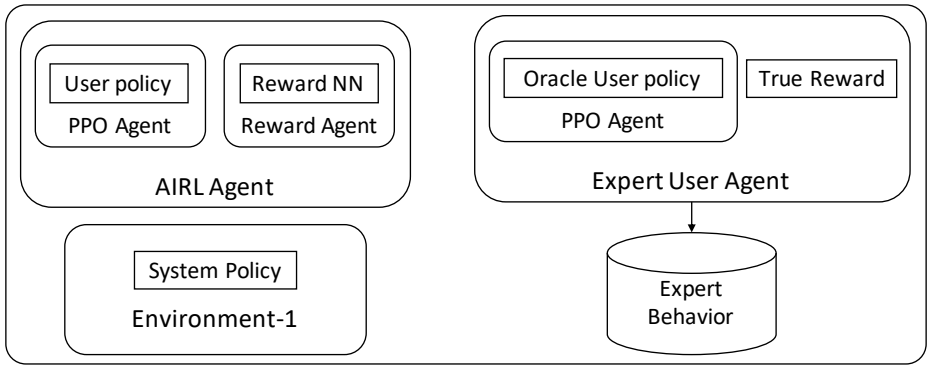

Figure 5: The architecture of the optimizing module in the original MDP(S, A, T, r, γ), which is responsible of generating user behavior and recovering user reward functions.

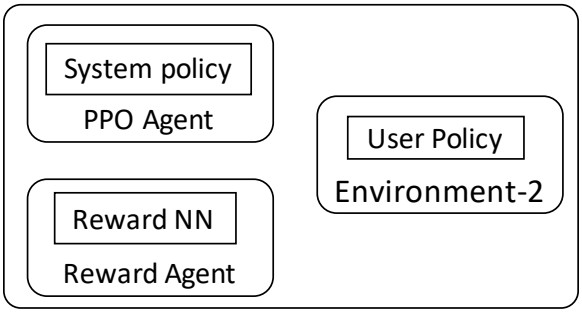

Figure 6: The architecture of the optimizing module in the reformulated MDP$^+(S^+, A^+, T^+, r^+, \gamma^+)$, responsible for optimizing the system agent.

after 3 *iterations* (one *iteration* means we sequentially recover the reward function and run Algorithm 1 once).

A higher average return means users are more satisfied while interacting with the interactive system. We initialize 5 different initial interactive systems by randomly sampling the system policy π_{sys}, and report the overall performance over these 5 systems. Besides, we want to avoid the situations that the optimized system has totally different behaviors compared to the initial system because the dramatic change may hurt users' experience. To make sure the systems before and after optimized follow similar behaviors, we introduce a regularization term, the KL-Divergence $\lambda * D_{KL}(T_{opt} \| T_{init})$, to control the distance between these two system policies. This term can also be regarded as the inherent constraints between state transitions, just like the "connection factor" in Section 6.1.1. The hyperparameter λ is applied to control the affect of the term. Due to the training complexity of network-based simulations, we run 5 times for each parameter setup rather than 40 times in the tabular world.

The Ground Truth of User Reward Functions With respect to the true reward function r_{real}, we have two different setups: *a handcrafted reward function* and *a randomly initialized reward function*. For the handcrafted one, we use $r_{real}(s) = \frac{1}{s_{dim}}\phi(s) * \phi(s)$ as the reward for the given state s. In terms of the sampled reward function, we initialize the parameters of the reward network with uniform distributions, and this makes recovering the reward function more difficult because there are

no patterns in the sampled reward function. In the real world, users always have their preferences and the reward function in users' minds is not likely to be random. The true reward function $r_{real}(s)$ is fixed in the whole optimizing process.

6.2.2 Results and Discussion

In this section, we first discuss the results of the experiments with a manually designed real reward function. Then, we move to the discussion of the experimental results with randomly initialized reward function.

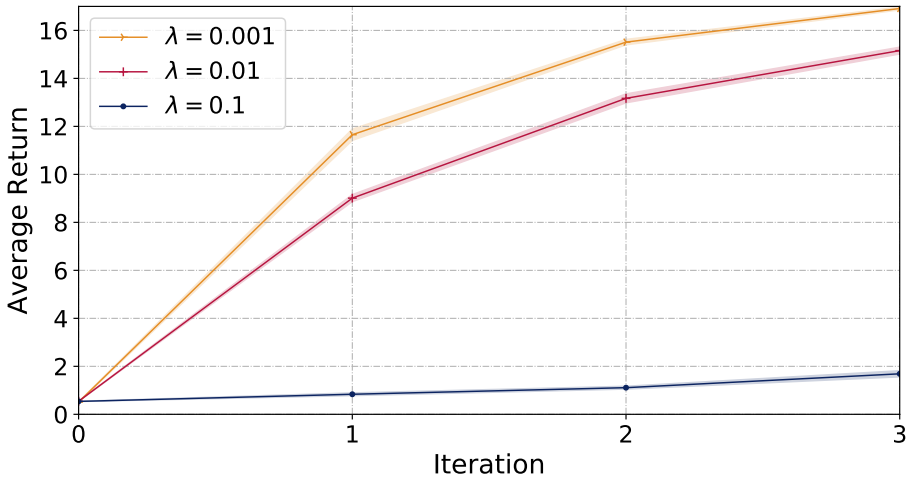

Figure 7: The state performance during optimization with oracle reward function and oracle user policy. The real reward function is manually designed. The error bounds denote the standard error of the mean (±SEM).

Manually Designed Real Reward Function To verify if the proposed two-MDP framework works or not, we first skip the reward learning step and use the oracle reward function $r_{real}(s)$ as the "learned" reward function with collected user behaviors. With respect to the user policy π_{user} used to interact with the system agent in the second optimizing module, we use the oracle user policy $\pi^*_{user_real}$ trained with true reward function $r_{real}(s)$. Other modules keep the same and we obtain the performance in Figure 7. An interactive sytem at iteration 0 is the initial system and not optimized yet. As we can see, with a looser restriction (i.e., a smaller λ value) on the distance between the optimized system and the initial system, we can achieve higher performance with respect to the average trajectory returns. After we bring back the reward learning step and use the learned reward function $r_{airl}(s)$ to optimize the system policy, we have the results shown in Figure 8. The system can still achieve higher performance by running Algorithm 1. If we compare the results between systems $\lambda = 0.001$ in Figure 7 and Figure 8, we can find that the system trained with oracle reward $r_{real}(s)$ can hit higher returns after two iterations. The finding still holds with respect to the systems $\lambda = 0.01$ in both setups. However, this is not the case when we set $\lambda = 0.1$. We suspect this is because the large regularization term $D_{KL}(T_{opt} \| T_{init})$ has brought too many uncontrollable factors into the optimization step, which may disturb the training.

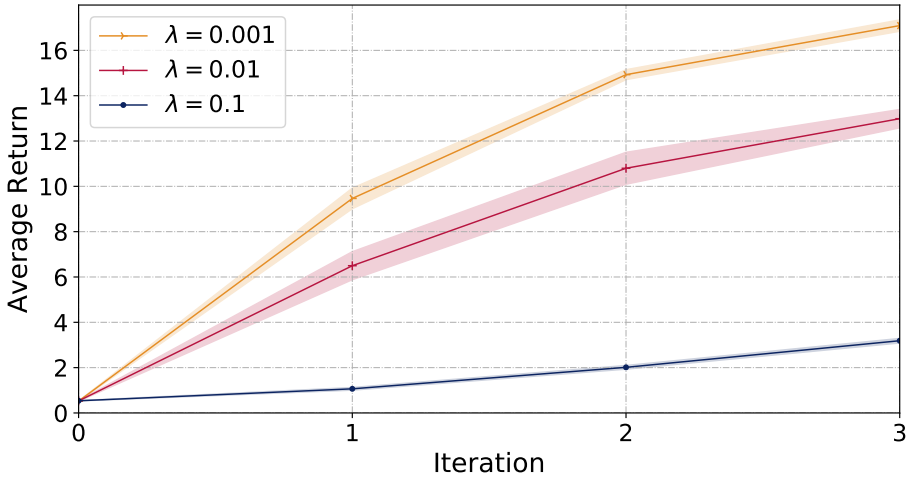

Figure 8: The state performance during optimization with recovered reward and oracle user policy. The real reward function is manually designed. The error bounds denote the standard error of the mean (±SEM).

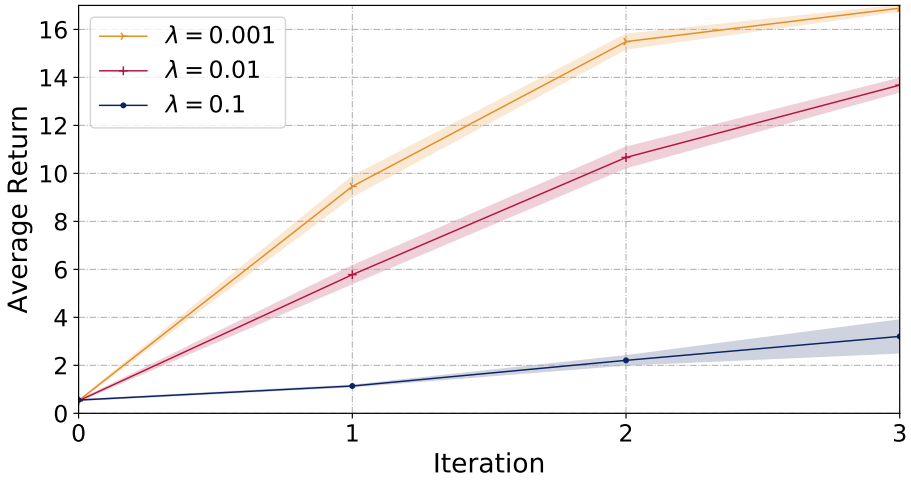

Figure 9: The state performance during optimization with recovered reward function and recovered user policy. The real reward function is manually designed. The error bounds denote the standard error of the mean (±SEM).

As mentioned in Section 6.2.1, the user policy π_{user} is essential while optimizing the system π_{sys}. In Algorithm 1, the user policy π_{user} plays the role of the transition distribution T^+ in the environment *Environment-2*. In addition to the two reward function setups above, we need to conduct an experiment with the user policy π_{user} trained with recovered reward function $r_{airl}(s)$ for system optimization in the reformulated MDP$^+$. Since we use AIRL to learn the reward function in this framework, we have the estimated user policy $\pi^*_{user_airl}$ which is rebuilt during the adversarial training process. We replace $\pi^*_{user_real}$ in the environment *Environment-2* with this rebuilt policy

$\pi^*_{user_airl}$. In terms of the reward function r^+ in MDP$^+(S^+, A^+, T^+, r^+, \gamma^+)$, we use the reward function $r_{airl}(s)$. By running Algorithm 1, we have the results in Figure 9. It is clear that the rebuilt policy $\pi^*_{user_airl}$ can still help with improving the system performance. This is meaningful because by using adversarial training we can rebuilt the user policy and user reward function simultaneously. The accuracy of the estimated user policy will definitely benefit from a high quality estimation of the user reward function. The only moment that real users are involved happens when we are collecting user-system interaction trajectories. This perfectly matches the scenarios in real life, where we first collect interaction histories from users and then infer the user preferences (r_{airl}) and user behavior patterns (π_{user_airl}) according to the collected data. In the next step, the system policy π_{sys} will be optimized based on user preferences and user behavior patterns. In the end, the newly updated system (S, A, T^*) will be presented to users to improve their user experience. If necessary, new interaction trajectories will be collected and another optimization turn can start again.

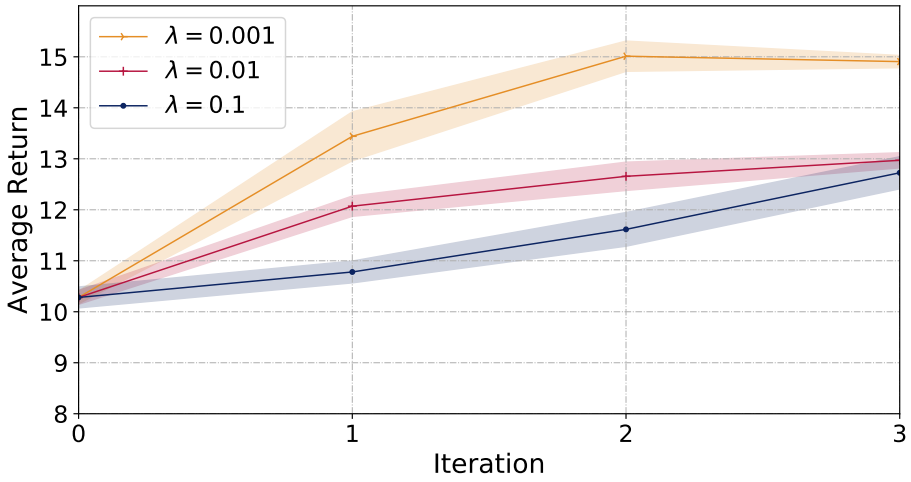

Figure 10: The state performance during optimization with oracle reward function and oracle user policy. The real reward function is randomly initialized. The error bounds denote the standard error of the mean (±SEM).

Randomly Initialized Reward Function In this section, we show how the interactive optimizer performs when the reward function $r_{real}(s)$ is randomly initialized. In Figure 10, with the real reward function $r_{real}(s)$, the system can still achieve relatively large improvements in terms of average return. All curves have higher starting points is because the randomly initialized system policy has advantage to hit higher reward for a random reward function and this will not hold when the reward function has special pattern like Section 6.2.2. We also find that looser restrictions on the distance between the optimized system and the initial system can bring larger performance improvements, as we observed in Section 6.2.2.

With respect to Figure 11 and Figure 12, the improvements still exist but are not so significant compared to those with the handcrafted reward function in Figure 8 and Figure 9. The potential reason is that it is hard to recover a high quality reward function given user behaviors generated by a random reward function. Especially before the first iteration, the system still performs randomly (the initial system is randomly initialized) and it is difficult to collect useful interaction traces for

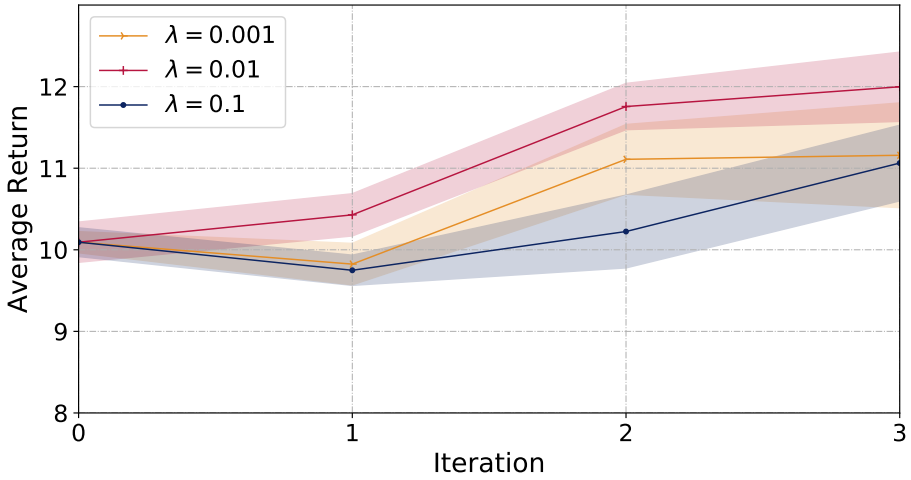

Figure 11: The state performance during optimization with recovered reward and oracle user policy. The real reward function is randomly initialized. The error bounds denote the standard error of the mean (±SEM).

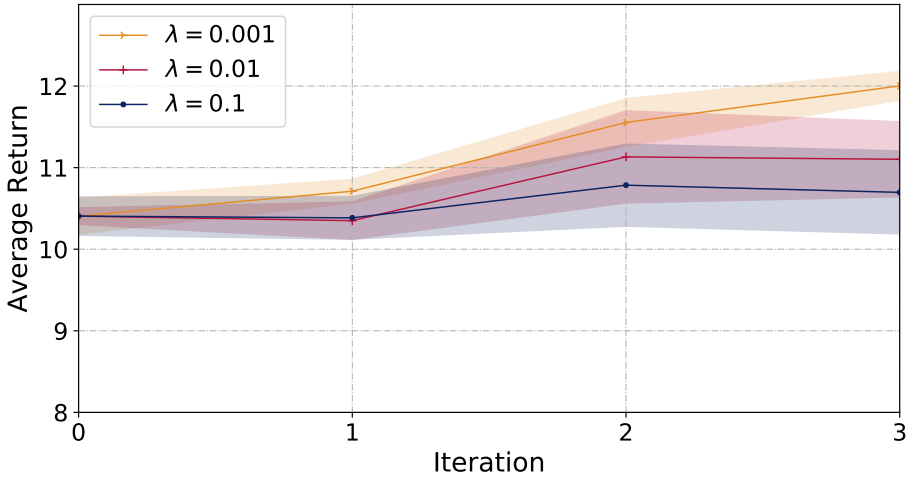

Figure 12: The state performance during optimization with recovered reward function and recovered user policy. The real reward function is randomly initialized. The error bounds denote the standard error of the mean (±SEM).

reward learning, and this is also the reason why, in Figure 11 and Figure 12, the average returns of some curves even drop after the first iteration. However, in the real world, users always have their preferences and the reward function in users' minds is unlikely to be random. Besides, the initial system will not behave randomly because in most cases a real interactive system (e.g., search engine, digital assistant) will be tested offline first and will not be deployed before it can achieve reasonable performance. These will alleviate the reward learning stress at some degree. We have this random

reward function here simply to validate how well the method could perform with most uncontrollable behaviors.

6.3 Limitations (as Future Directions)

First, to recover a reliable reward function, a large number of high-quality user interaction traces are essential, which can come with a great cost in real life (but not impossible). Furthermore, an interactive system usually serves different users, which can lead to the violation of Assumption 5 about the homogeneity of user behavior. Therefore, we would need to work on the personalizing of the recovered reward functions. A possible way to address this limitation in the future is to incorporate the user features into the state space, but this still needs to be explored.

Second, as shown in Section 6.1.2, the final performance of the optimized system highly relies on the quality of the recovered reward function. With respect to the more advanced extension of Maximum Entropy Inverse Reinforcement Learning (MaxEnt-IRL), which is Adversarial Inverse Reinforcement Learning (AIRL), the adversarial training process is intractable for complex real behaviors. The two limitations above will boil down to the quality of recovered reward functions, given limited user traces in real life.

Third, after we have inferred the reward function, we will update the system in a reformulated MDP setup, where we switch the roles between the agent and environment. The potential problem that can arise is that the action space for the new MDP could be extremely large and this may present challenges for the scalability of the Reinforcement Learning (RL) process.

Finally, we validate our method in two simulated experimental setups. Despite the fact that we try to design our setups as close as possible to the real-word scenarios, there is a potential gap between the designed systems and real-world applications. But the positive verification of our method in simulated setup helps us to understand better the pros and cons of the proposed approach and help us with planning the experiments with the real-world scenarios in the near future.

To summarize, we have proposed two experimental setups to test the proposed framework: tabular and neural. In both cases, the results demonstrate significant improvements of interactive systems when applying our method. We conclude this section acknowledging a number of possible limitations, some of which can be considered as future directions.

7. Conclusions and Future Work

We have recognized that previous work on interactive systems has relied on the assumption that the handcrafted objective functions can accurately reflect users' preferences and intentions while interacting with interaction systems. As a result, interactive systems have been optimized for manually designed objectives that do not always align with the true user preferences and cannot be generalized across different domains. To overcome this discrepancy, we have proposed a novel two-step framework to optimize interactive systems, which first infer the user reward model given collected user interaction traces and then update the system with the inferred reward functions via a novel algorithm: the Interactive System Optimizer (ISO).

Firstly, we modeled user-system interactions using MDP, where the agent is the user, and the stochastic environment is the interactive system. User satisfaction is modeled via rewards received from interactions, and the user interaction history is represented by a set of trajectories. We followed the previously justified assumption that user incentive to interact with the system if they are

rewarded. Treating an interactive system as a changeable and programmable environment is novel and reasonable because we have complete control of the interactive systems since we are the system designers.

Secondly, we formalized an optimization problem to infer the user needs from the observed user-system interactions, in the form of a data-driven objective. Importantly, our method works without any domain knowledge, and is thus even applicable when prior knowledge is absent.

Thirdly, we proposed a novel, Interactive System Optimizer (ISO), that iterates between optimizing the interactive system for the current inferred objective; and letting the user adapt to the new system behavior. This process repeats until both the user and system policies converge. Our experimental results show that ISO robustly improves the user satisfaction.

The newly proposed approach to optimize an interactive system based on data-driven objectives is novel, many promising directions for future work are possible. For instance, while ISO performs well for users with a single goal, this approach could be extended to settings with multiple goals. Similarly, extensions considering more personalized goals could benefit the overall user experience. Finally, investigating the scalability and real world applicability of ISO could open many research possibilities.

References

P. Abbeel and A. Y. Ng. Apprenticeship learning via inverse reinforcement learning. In *ICML*, pages 1–8. ACM, 2004.

I. Akkaya, M. Andrychowicz, M. Chociej, M. Litwin, B. McGrew, A. Petron, A. Paino, M. Plappert, G. Powell, R. Ribas, et al. Solving rubik's cube with a robot hand. *arXiv preprint arXiv:1910.07113*, 2019.

F. Argelaguet, L. Hoyet, M. Trico, and A. Lécuyer. The role of interaction in virtual embodiment: Effects of the virtual hand representation. In *VR*, pages 3–10. IEEE, 2016.

L. Azzopardi. Modelling interaction with economic models of search. In *SIGIR*, pages 3–12. ACM, 2014.

L. E. Blume, S. Durlauf, and L. E. Blume. *The new Palgrave dictionary of economics*. Palgrave Macmillan Manchester, 2008.

A. Borisov, I. Markov, M. de Rijke, and P. Serdyukov. A neural click model for web search. In *WWW*, pages 531–541, 2016.

A. Boularias, J. Kober, and J. Peters. Relative entropy inverse reinforcement learning. In *AISTATS*, pages 182–189, 2011.

B. Burchfiel, C. Tomasi, and R. Parr. Distance minimization for reward learning from scored trajectories. In *AAAI*, pages 3330–3336. AAAI Press, 2016.

M. Chen, A. Beutel, P. Covington, S. Jain, F. Belletti, and E. H. Chi. Top-k off-policy correction for a reinforce recommender system. In *WSDM*, pages 456–464. ACM, 2019.

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

C. Cui, W. Wang, X. Song, M. Huang, X.-S. Xu, and L. Nie. User attention-guided multimodal dialog systems. In *SIGIR*, pages 445–454. ACM, 2019.

M. Dehghani, H. Zamani, A. Severyn, J. Kamps, and W. B. Croft. Neural ranking models with weak supervision. In *SIGIR*, pages 65–74, 2017.

B. Dhingra, L. Li, X. Li, J. Gao, Y.-N. Chen, F. Ahmed, and L. Deng. Towards end-to-end reinforcement learning of dialogue agents for information access. *arXiv preprint arXiv:1609.00777*, 2016.

A. Drutsa, G. Gusev, and P. Serdyukov. Engagement periodicity in search engine usage: Analysis and its application to search quality evaluation. In *WSDM*, pages 27–36, 2015.

Y. Duan, J. Schulman, X. Chen, P. L. Bartlett, I. Sutskever, and P. Abbeel. RL^2: Fast reinforcement learning via slow reinforcement learning. *arXiv preprint arXiv:1611.02779*, 2016.

G. Dupret and M. Lalmas. Absence time and user engagement: evaluating ranking functions. In *WSDM*, pages 173–182, 2013.

L. El Asri, R. Laroche, and O. Pietquin. Reward shaping for statistical optimisation of dialogue management. In *SLSP*, pages 93–101. Springer, 2013.

C. Finn, P. Christiano, P. Abbeel, and S. Levine. A connection between generative adversarial networks, inverse reinforcement learning, and energy-based models. *arXiv preprint arXiv:1611.03852*, 2016a.

C. Finn, S. Levine, and P. Abbeel. Guided cost learning: Deep inverse optimal control via policy optimization. In *ICML*, pages 49–58, 2016b.

C. Finn, T. Yu, T. Zhang, P. Abbeel, and S. Levine. One-shot visual imitation learning via meta-learning. In *CoRL*, pages 357–368, 2017.

S. Fox, K. Karnawat, M. Mydland, S. T. Dumais, and T. White. Evaluating implicit measures to improve web search. *ACM Transactions on Information Systems*, 23(2):147–168, 2005.

J. Fu, K. Luo, and S. Levine. Learning robust rewards with adversarial inverse reinforcement learning. *arXiv preprint arXiv:1710.11248*, 2017.

T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. *arXiv preprint arXiv:1801.01290*, 2018.

D. Hafner, T. Lillicrap, J. Ba, and M. Norouzi. Dream to control: Learning behaviors by latent imagination. *arXiv preprint arXiv:1912.01603*, 2019.

J. Ho and S. Ermon. Generative adversarial imitation learning. In *Advances in neural information processing systems*, pages 4565–4573, 2016.

K. Hofmann, S. Whiteson, and M. de Rijke. Balancing exploration and exploitation in learning to rank online. In *ECIR*, pages 251–263. Springer, 2011.

K. Hofmann, A. Schuth, S. Whiteson, and M. de Rijke. Reusing historical interaction data for faster online learning to rank for IR. In *WSDM*, pages 183–192. ACM, 2013a.

K. Hofmann, S. Whiteson, and M. de Rijke. Balancing exploration and exploitation in listwise and pairwise online learning to rank for information retrieval. *Information Retrieval Journal*, 16(1): 63–90, 2013b.

K. Järvelin and J. Kekäläinen. Cumulated gain-based evaluation of ir techniques. *ACM Transactions on Information Systems*, 20(4):422–446, 2002.

H. J. Jeon, S. Milli, and A. D. Dragan. Reward-rational (implicit) choice: A unifying formalism for reward learning. *arXiv preprint arXiv:2002.04833*, 2020.

T. Joachims, L. Granka, B. Pan, H. Hembrooke, and G. Gay. Accurately interpreting clickthrough data as implicit feedback. In *SIGIR*, pages 154–161, 2005.

D. Kelly. Methods for evaluating interactive information retrieval systems with users. *Foundations and Trends in Information Retrieval*, 3(1–2):1–224, 2009.

D. Kelly. When effort exceeds expectations: A theory of search task difficulty. In *ECIR Supporting Complex Search Task Workshop '15*, 2015.

J. Kiseleva and M. de Rijke. Evaluating personal assistants on mobile devices. *arXiv preprint arXiv:1706.04524*, 2017.

J. Kiseleva, E. Crestan, R. Brigo, and R. Dittel. Modelling and detecting changes in user satisfaction. In *CIKM*, pages 1449–1458, 2014.

J. Kiseleva, J. Kamps, V. Nikulin, and N. Makarov. Behavioral dynamics from the serp's perspective: What are failed serps and how to fix them? In *Submission of SIGIR*, 2015.

J. Kiseleva, K. Williams, A. H. Awadallah, I. Zitouni, A. Crook, and T. Anastasakos. Predicting user satisfaction with intelligent assistants. In *SIGIR*, pages 45–54. ACM, 2016a.

J. Kiseleva, K. Williams, J. Jiang, A. H. Awadallah, I. Zitouni, A. Crook, and T. Anastasakos. Understanding user satisfaction with intelligent assistants. In *CHIIR*, pages 121–130, 2016b.

M. Kosinski, D. Stillwell, and T. Graepe. Private traits and attributes are predictable from digital records of human behavior. *PNAS*, 110:5802–5805, 2013.

M. Kutlu, V. Khetan, and M. Lease. Correlation and prediction of evaluation metrics in information retrieval. *arXiv preprint arXiv:1802.00323*, 2018.

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

S. Levine, C. Finn, T. Darrell, and P. Abbeel. End-to-end training of deep visuomotor policies. *Journal of Machine Learning Research*, 17(39):1–40, 2016a.

S. Levine, P. Pastor, A. Krizhevsky, J. Ibarz, and D. Quillen. Learning hand-eye coordination for robotic grasping with deep learning and large-scale data collection. *The International Journal of Robotics Research*, pages 173–184, 2016b.

J. Li, W. Monroe, A. Ritter, D. Jurafsky, M. Galley, and J. Gao. Deep reinforcement learning for dialogue generation. In *EMNLP*, pages 1192–1202, 2016.

X. Li, Y.-N. Chen, L. Li, J. Gao, and A. Celikyilmaz. End-to-end task-completion neural dialogue systems. *arXiv preprint arXiv:1703.01008*, 2017a.

Z. Li, J. Kiseleva, M. de Rijke, and A. Grotov. Towards learning reward functions from user interactions. In *ICTIR*, pages 941–944. ACM, 2017b.

Z. Li, J. Kiseleva, and M. de Rijke. Dialogue generation: From imitation learning to inverse reinforcement learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 6722–6729, 2019.

Z. Li, S. Lee, B. Peng, J. Li, S. Shayandeh, and J. Gao. Guided dialog policy learning without adversarial learning in the loop. *arXiv preprint arXiv:2004.03267*, 2020.

Z. Lipton, X. Li, J. Gao, L. Li, F. Ahmed, and L. Deng. Bbq-networks: Efficient exploration in deep reinforcement learning for task-oriented dialogue systems. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.

C.-W. Liu, R. Lowe, I. Serban, M. Noseworthy, L. Charlin, and J. Pineau. How NOT to evaluate your dialogue system: An empirical study of unsupervised evaluation metrics for dialogue response generation. In *EMNLP*, 2016.

F. Lovett. Rational choice theory and explanation. *Rationality and Society*, 18(2):237–272, 2006.

R. Lowe, M. Noseworthy, I. V. Serban, N. Angelard-Gontier, Y. Bengio, and J. Pineau. Towards an automatic turing test: Learning to evaluate dialogue responses. In *ACL*, pages 1116–1126, 2017.

J. Mitchell and B. Shneiderman. Dynamic versus static menus: an exploratory comparison. *ACM SigCHI Bulletin*, 20(4):33–37, 1989.

V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529, 2015.

M. Mohri, A. Rostamizadeh, and A. Talwalkar. *Foundations of Machine Learning*. MIT Press, 2012.

M. Monfort, A. Liu, and B. Ziebart. Intent prediction and trajectory forecasting via predictive inverse linear-quadratic regulation. In *AAAI*, pages 3672–3678. AAAI Press, 2015.

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

H. Obendorf, H. Weinreich, E. Herder, and M. Mayer. Web page revisitation revisited: implications of a long-term click-stream study of browser usage. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 597–606, 2007.

D. Odijk, E. Meij, I. Sijaranamual, and M. de Rijke. Dynamic query modeling for related content finding. In *SIGIR*, pages 33–42. ACM, 2015.

K. Papineni, S. Roukos, T. Ward, and W.-J. Zhu. Bleu: a method for automatic evaluation of machine translation. In *ACL*, pages 311–318. ACL, 2002.

B. Peng, X. Li, L. Li, J. Gao, A. Celikyilmaz, S. Lee, and K.-F. Wong. Composite task-completion dialogue policy learning via hierarchical deep reinforcement learning. *arXiv preprint arXiv:1704.03084*, 2017.

B. Peng, X. Li, J. Gao, J. Liu, Y.-N. Chen, and K.-F. Wong. Adversarial advantage actor-critic model for task-completion dialogue policy learning. In *ICASSP*, pages 6149–6153. IEEE, 2018.

J. Perner and B. Lang. Development of theory of mind and executive control. *Trends in Cognitive Sciences*, 3(9):337–344, 1999.

O. Pietquin. Inverse reinforcement learning for interactive systems. In *Workshop on Machine Learning for Interactive Systems*, pages 71–75. ACM, 2013.

A. H. Qureshi, B. Boots, and M. C. Yip. Adversarial imitation via variational inverse reinforcement learning. In *ICLR*, 2019.

N. D. Ratliff, D. Silver, and J. A. Bagnell. Learning to search: Functional gradient techniques for imitation learning. *Autonomous Robots*, 27(1):25–53, 2009.

S. Reddy, A. D. Dragan, S. Levine, S. Legg, and J. Leike. Learning human objectives by evaluating hypothetical behavior. *arXiv preprint arXiv:1912.05652*, 2019.

S. Russell. Learning agents for uncertain environments. In *COLT*, pages 101–103. ACM, 1998.

T. Saracevic. Relevance: A review of and a framework for the thinking on the notion in information science. *Journal of the American Society for Information Science and Technology*, 26:321–343, 1975.

T. Saracevic, P. B. Kantor, A. Y. Chamis, and D. Trivison. A study of information seeking and retrieving. I. background and methodology. II. users, questions and effectiveness. III. searchers, searches, overlap. *Journal of the American Society for Information Science*, 39:161–176; 177–196; 197–216, 1988.

T. Schnabel, P. N. Bennett, and T. Joachims. Shaping feedback data in recommender systems with interventions based on information foraging theory. In *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining*, pages 546–554, 2019.

J. Schrittwieser, I. Antonoglou, T. Hubert, K. Simonyan, L. Sifre, S. Schmitt, A. Guez, E. Lockhart, D. Hassabis, T. Graepel, et al. Mastering atari, go, chess and shogi by planning with a learned model. *arXiv preprint arXiv:1911.08265*, 2019.

J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.

A. Sepliarskaia, J. Kiseleva, F. Radlinski, and M. de Rijke. Preference elicitation as an optimization problem. In *Proceedings of the 12th ACM Conference on Recommender Systems*, pages 172–180, 2018.

S. K. Seyed Ghasemipour, S. S. Gu, and R. Zemel. Smile: Scalable meta inverse reinforcement learning through context-conditional policies. In *Advances in Neural Information Processing Systems 32*, pages 7881–7891. 2019.

G. Shani, D. Heckerman, and R. I. Brafman. An MDP-based recommender system. *Journal of Machine Learning Research*, 6(Sep):1265–1295, 2005.

D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. Van Den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot, et al. Mastering the game of go with deep neural networks and tree search. *Nature*, 529(7587):484–489, 2016.

S.-Y. Su, X. Li, J. Gao, J. Liu, and Y.-N. Chen. Discriminative deep dyna-q: Robust planning for dialogue policy learning. In *EMNLP*, 2018.

R. S. Sutton and A. G. Barto. *Reinforcement learning: An introduction*. MIT press Cambridge, 1998.

R. S. Sutton and A. G. Barto. *Reinforcement learning: An introduction*. MIT press, 2018.

R. Takanobu, H. Zhu, and M. Huang. Guided dialog policy learning: Reward estimation for multi-domain task-oriented dialog. *arXiv preprint arXiv:1908.10719*, 2019.

J. Teevan. How people recall, recognize, and reuse search results. *ACM Transactions on Information Systems (TOIS)*, 26(4):1–27, 2008.

M. ter Hoeve, R. Sim, E. Nouri, A. Fourney, M. de Rijke, and R. W. White. Conversations with documents: An exploration of document-centered assistance. In *CHIIR*, pages 43–52. ACM, 2020.

H. R. Varian. Economics and search. In *ACM SIGIR Forum*, volume 33, pages 1–5. ACM New York, NY, USA, 1999.

O. Vinyals, I. Babuschkin, W. M. Czarnecki, M. Mathieu, A. Dudzik, J. Chung, D. H. Choi, R. Powell, T. Ewalds, P. Georgiev, et al. Grandmaster level in starcraft ii using multi-agent reinforcement learning. *Nature*, 575(7782):350–354, 2019.

J. X. Wang, Z. Kurth-Nelson, D. Tirumala, H. Soyer, J. Z. Leibo, R. Munos, C. Blundell, D. Kumaran, and M. Botvinick. Learning to reinforcement learn. *arXiv preprint arXiv:1611.05763*, 2016.

H. Wei, F. Zhang, N. J. Yuan, C. Cao, H. Fu, X. Xie, Y. Rui, and W.-Y. Ma. Beyond the words: Predicting user personality from heterogeneous information. In *WSDM*, pages 305–314. ACM, 2017.

R. W. White. *Interactions with search systems*. Cambridge University Press, 2016.

R. W. White, I. Ruthven, and J. M. Jose. Finding relevant documents using top ranking sentences: an evaluation of two alternative schemes. In *SIGIR*, pages 57–64, 2002.

R. W. White, I. Ruthven, J. M. Jose, and C. van Rijsbergen. Evaluating implicit feedback models using searcher simulations. *ACM Transactions on Information Systems (TOIS)*, 23(3):325–361, 2005.

J. D. Williams, K. Asadi, and G. Zweig. Hybrid code networks: practical and efficient end-to-end dialog control with supervised and reinforcement learning. *arXiv preprint arXiv:1702.03274*, 2017.

K. Williams, J. Kiseleva, A. Crook, I. Zitouni, A. H. Awadallah, and M. Khabsa. Is this your final answer? evaluating the effect of answers on good abandonment in mobile search. In *SIGIR*, pages 889–892, 2016a.

K. Williams, J. Kiseleva, A. C. Crook, I. Zitouni, A. H. Awadallah, and M. Khabsa. Detecting good abandonment in mobile search. In *WWW*, pages 495–505, 2016b.

E. Yilmaz, M. Verma, N. Craswell, F. Radlinski, and P. Bailey. Relevance and effort: An analysis of document utility. In *CIKM*, pages 91–100, 2014.

J. Y. Zhang and A. D. Dragan. Learning from extrapolated corrections. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 7034–7040. IEEE, 2019.

X. Zhao, L. Zhang, Z. Ding, L. Xia, J. Tang, and D. Yin. Recommendations with negative feedback via pairwise deep reinforcement learning. In *KDD*, pages 1040–1048, 2018.

G. Zheng, F. Zhang, Z. Zheng, Y. Xiang, N. J. Yuan, X. Xie, and Z. Li. Drn: A deep reinforcement learning framework for news recommendation. In *WWW*, pages 167–176, 2018.

Y. Zhu, R. Mottaghi, E. Kolve, J. J. Lim, A. Gupta, L. Fei-Fei, and A. Farhadi. Target-driven visual navigation in indoor scenes using deep reinforcement learning. In *ICRA*, pages 3357–3364. IEEE, 2017.

B. Ziebart, A. Dey, and J. A. Bagnell. Probabilistic pointing target prediction via inverse optimal control. In *IUI*, pages 1–10. ACM, 2012.

B. D. Ziebart. *Modeling Purposeful Adaptive Behavior with the Principle of Maximum Causal Entropy*. PhD thesis, Carnegie Mellon University, 2010.

B. D. Ziebart, A. L. Maas, J. A. Bagnell, and A. K. Dey. Maximum entropy inverse reinforcement learning. In *AAAI*, pages 1433–1438. AAAI Press, 2008.