

# WFA-IRL: Inverse Reinforcement Learning of Autonomous Behaviors Encoded as Weighted Finite Automata

Tianyu Wang

Nikolay Atanasov

**Abstract**—This paper presents a method for learning logical task specifications and cost functions from demonstrations. Constructing specifications by hand is challenging for complex objectives and constraints in autonomous systems. Instead, we consider demonstrated task executions, whose logic structure and transition costs need to be inferred by an autonomous agent. We employ a spectral learning approach to extract a weighted finite automaton (WFA), approximating the unknown task logic. Thereafter, we define a product between the WFA for high-level task guidance and a labeled Markov decision process for low-level control. An inverse reinforcement learning (IRL) problem is considered to learn a cost function by backpropagating the loss between agent and expert behaviors through the planning algorithm. Our proposed model, termed WFA-IRL, is capable of generalizing the execution of the inferred task specification in a suite of MiniGrid environments.

## I. INTRODUCTION

Autonomous systems are expected to achieve reliable performance in increasingly complex environments with increasingly complex objectives. Yet, it is often challenging to design a mathematical formulation that captures all safety and liveness requirements across various operational conditions. Minimizing a misspecified cost function may lead to undesirable performance, regardless of the quality of the optimization algorithm. However, a domain expert is often able to demonstrate desirable or undesirable behavior that implicitly captures the task specifications. As a simple illustration, consider the navigation task in Fig. 1, requiring a door to be unlocked before reaching a goal state. Instead of encoding the task requirements as a cost function, an expert may provide several demonstrations of navigating to the goal, some of which require picking up the key whenever the door is locked. A reinforcement learning agent should infer the underlying logic sequence of the demonstrated task in order to learn the desired behavior.

Inverse reinforcement learning (IRL) [1]–[3] focuses on inferring the latent costs of expert demonstrations. Early works assume that the cost is linear in a set of state features and minimize the feature expectation difference between learned policy and demonstrations. [3] use dynamic programming to find a maximum entropy (MaxEnt) policy which maximizes the likelihood of the demonstrated actions. Later works [4], [5] introduce Gaussian process or deep neural networks to learn nonlinear cost functions. [6] builds a connection between MaxEnt IRL and adversarial learning

We gratefully acknowledge support from ONR SAI N00014-18-1-2828. The authors are with the Department of Electrical and Computer Engineering, University of California San Diego, La Jolla, CA 92093, USA {tw161, natanasov}@eng.ucsd.edu.




Fig. 1: (Left) An example trajectory in the MiniGrid environment [9], where an agent has to pick up a key, open the door, and navigate to the goal. (Right) The trajectory can be decomposed into three segments, identified by hidden states  $\alpha_t$ . Transitions between the high-level states are triggered by events, such as picking up key ( $\sigma_0$ ), or opening door ( $\sigma_1$ ).

and solves continuous control problems. Most IRL models, however, consider general cost formulations that do not explicitly capture sequencing and compositional requirements of the demonstrated task [7], [8]. Compared to a general cost formulation, this paper shows the logical structure of a complex task can be inferred from demonstrations. Exploiting the underlying task logic in planning ensures that the learned agent behavior mode matches the demonstrations.

Hierarchical reinforcement learning and options framework [10]–[13] are formulations that learn task decomposition and temporal abstraction. Options are high-level macro-actions consisting of primitive actions. [14] introduces a multi-level hierarchical model to discover options from demonstrations where option boundaries are inferred for trajectory segmentation. [15] uses an unsupervised encoder-decoder model to predict subtask segmentation and categorical latent encoding. [16] uses graph recurrent neural networks with relational features between objects for high-level planning and low-level primitive dynamics prediction.

Formal methods have been applied in robotics to prove and guarantee different behavioral properties such as safety and correctness [17]–[20]. For example, linear temporal logic (LTL) [21] is used to specify safety and liveness objectives with temporal ordering constraints in control and reinforcement learning problems [20], [22]–[26]. Specification mining of LTL formulas can learn finite state automata from execution traces [27], [28]. LTL formulas can also be inferred from Bayesian inference [29] or from graph connectivity of directed acyclic graphs over atomic propositions [30]. In this work, we consider weighted finite automata (WFA) in which the transitions carry weights. Whereas classical automata determine whether a word is accepted or rejected, WFA can

compute quantitative values as a function of the weighted transitions from the execution of words [31]. WFA offer the expressive power to model quantitative properties, such as resources, time or cost, of the demonstrated behavior. Under certain assumptions of the semiring on which the WFA is defined, it can be shown that WFA is expressively equivalent to weighted monadic second-order (MSO) logic [31], [32].

We introduce an IRL model that learns to infer high-level task specifications and low-level control costs to imitate demonstrated behavior. Given a set of demonstrations, we use a spectral method to learn a WFA which encodes the task logic structure. The agent’s interaction with the environment is modeled as a product between the learned WFA and a labeled Markov decision process (L-MDP). We propose a planning algorithm to search over the product space for a policy that satisfies task requirements encoded by the WFA. Since the true transition cost is not directly observable, we differentiate the error between the agent’s policy and the demonstrated controls through the planning algorithm using a subgradient method introduced in [2], [33]. We demonstrate that our WFA-IRL method correctly classifies accepting and rejecting sequences and learns a cost function that generalizes the demonstrated behavior to new settings in several MiniGrid environments [9]. In summary, our **contribution** is to recognize that the logic structure of a demonstrated task can be learned as a weighted finite automaton and, in turn, can be integrated with differentiable task planning to learn generalizable behavior from demonstrations.

## II. PRELIMINARIES

### A. Agent and environment models

The agent’s interaction with the environment is modeled as an L-MDP [34].

**Definition 1.** A labeled Markov decision process is a tuple  $\{\mathcal{X}, \mathcal{U}, \mathbf{x}_0, f, c, \mathcal{AP}, \ell\}$ , where  $\mathcal{X}, \mathcal{U}$  are finite sets of states and controls,  $\mathbf{x}_0 \in \mathcal{X}$  is an initial state,  $f : \mathcal{X} \times \mathcal{U} \rightarrow \mathcal{X}$  is a deterministic transition function, and  $c : \mathcal{X} \times \mathcal{U} \rightarrow \mathbb{R}_{\geq 0}$  assigns a non-negative cost when control  $\mathbf{u} \in \mathcal{U}$  is applied at state  $\mathbf{x} \in \mathcal{X}$ . A finite set of atomic propositions  $\mathcal{AP}$  provides logic statements that must be true or false (e.g., “the agent is 1 meter away from the closest obstacle” or “the agent possesses a key”). A labeling function  $\ell : \mathcal{X} \times \mathcal{U} \rightarrow 2^{\mathcal{AP}}$  assigns a set of atomic propositions that evaluate true for a given state transition.

We assume that the state  $\mathbf{x}$  is fully observable and captures both endogenous variables for the agent, such as position and orientation, and exogenous variables, such as an environment containing objects of interest as illustrated in Fig. 1. The transition function  $f(\mathbf{x}, \mathbf{u})$  specifies the change of state  $\mathbf{x}$  when control  $\mathbf{u}$  is executed, and  $c(\mathbf{x}, \mathbf{u})$  assigns a non-negative cost to this transition. The alphabet of the L-MDP is the set of labels  $\Sigma = 2^{\mathcal{AP}}$  that can be assigned to the transitions. The labeling function  $\sigma = \ell(\mathbf{x}, \mathbf{u})$  provides the atomic propositions  $\sigma \in \Sigma$  which are satisfied during the transition  $f(\mathbf{x}, \mathbf{u})$ . The set of words on  $\Sigma$  is denoted by  $\Sigma^*$  and consists of all strings  $\sigma_{0:T} = \sigma_0 \dots \sigma_T$  for  $\sigma_t \in \Sigma$  and

$T \in \mathbb{N}$ . We assume that the transition  $f$  and labeling  $\ell$  are known. However, the cost function  $c$  is unknown and needs to be inferred from expert demonstrations.

### B. Expert model

The agent needs to execute a task, whose success is evaluated based on the word  $\sigma_{0:T} \in \Sigma^*$  resulting from the agent’s actions. We model the quality of the task execution by a function  $h : \Sigma^* \rightarrow \mathbb{R}$ . An execution  $\sigma_{0:T}$  is deemed successful if  $h(\sigma_{0:T}) \geq \xi$  for a known performance threshold  $\xi$ , and unsuccessful otherwise. As argued in the introduction, defining the function  $h$  explicitly is challenging in many applications. Instead, we consider a training set  $\mathcal{D} = \{(\mathbf{x}_{0:T_n}^n, \mathbf{u}_{0:T_n}^n, s^n)\}_{n=1}^N$  of  $N$  demonstrations of the same task in different environment configurations provided by an expert. Each demonstration  $n$  contains the controls  $\mathbf{u}_{0:T_n}^n = \mathbf{u}_0^n \dots \mathbf{u}_{T_n}^n$  executed by the expert, the resulting agent-environment states  $\mathbf{x}_{0:T_n}^n = \mathbf{x}_0^n \dots \mathbf{x}_{T_n}^n$ , and the success level  $s^n \in \mathbb{R}$  of the execution, measured by  $h(\sigma_{0:T_n}^n)$ , where  $\sigma_t^n = \ell(\mathbf{x}_t^n, \mathbf{u}_t^n)$  is the label encountered by the expert at time  $t$ . We assume that the expert knows the *true* task  $h$  and the *true* cost  $c$  and can solve a finite-horizon first-exit deterministic optimal control problem [35] over the L-MDP:

$$Q^*(\mathbf{x}, \mathbf{u}) := \min_{T, \mathbf{u}_{1:T}} \sum_{t=0}^T c(\mathbf{x}_t, \mathbf{u}_t) \quad (1)$$

s.t.  $\mathbf{x}_{t+1} = f(\mathbf{x}_t, \mathbf{u}_t)$ ,  $\mathbf{x}_0 = \mathbf{x}$ ,  $\mathbf{u}_0 = \mathbf{u}$ ,  
 $\sigma_t = \ell(\mathbf{x}_t, \mathbf{u}_t)$ ,  $h(\sigma_{0:T}) \geq \xi$ ,

where  $Q^*(\mathbf{x}, \mathbf{u})$  is the optimal value function. Since (1) is a deterministic optimal control problem, there exists an open-loop control sequence which is optimal, i.e., achieves the same cost as an optimal closed-loop policy function [35, Chapter 6]. However, we consider experts that do not necessarily choose strictly rational controls. Instead, we model the expert behavior using a stochastic Boltzmann policy over the optimal values  $\pi^*(\mathbf{u}|\mathbf{x}) \propto \exp\left(-\frac{1}{\eta} Q^*(\mathbf{x}, \mathbf{u})\right)$ , where  $\eta \in (0, \infty)$  is a temperature parameter representing a continuous spectrum of rationality. For example,  $\eta \rightarrow 0$  means that the expert takes strictly optimal controls while  $\eta \rightarrow \infty$  means random controls are selected. The Boltzmann expert model was previously introduced and studied in [33], [36], [37]. It provides an exponential preference for controls that incur low long-term costs. This expert model also allows efficient policy search, as shown in Sec. IV-B, and computation of the policy gradient with respect to the cost needed to optimize the cost parameters, as shown in Sec. IV-C.

## III. PROBLEM STATEMENT

The agent needs to infer the unknown task model  $h$  and unknown cost function  $c$  from the expert demonstrations  $\mathcal{D} = \{(\mathbf{x}_{0:T_n}^n, \mathbf{u}_{0:T_n}^n, s^n)\}_{n=1}^N$ .

**Problem 1.** Given the demonstrations  $\mathcal{D}$  and labeling  $\sigma_t^n = \ell(\mathbf{x}_t^n, \mathbf{u}_t^n)$ , optimize the parameters  $\psi$  of an approximation  $h_\psi$  of the unknown task function  $h$  to minimize the mean




Fig. 2: Inferring the hidden state progression  $\alpha_t$  from events  $\sigma_t$  can be achieved by an RNN with initial hidden state  $\alpha_0$ , hidden state transition  $\alpha_{t+1} = g_1(\sigma_t, \alpha_t, \mathbf{W})$  and output  $\hat{s} = h_\psi(\sigma_{0:T}) = g_2(\alpha_{T+1}, \beta)$ , where  $g_1, g_2$  are nonlinear functions. The weights  $\psi = (\alpha_0, \mathbf{W}, \beta)$  can be learned via the loss  $\mathcal{L}(\hat{s}, s)$  in (2) between RNN outputs  $\hat{s}$  and demonstration scores  $s$ .

squared error:

$$\min_{\psi} \mathcal{L}_h(\psi) := \frac{1}{N} \sum_{n=1}^N (h_\psi(\sigma_{0:T_n}^n) - s^n)^2. \quad (2)$$

Similarly, the agent needs to obtain an approximation  $c_\theta$  with parameters  $\theta$  of the unknown cost function  $c$ . This allows the agent to obtain a control policy:

$$\pi_\theta(\mathbf{u}|\mathbf{x}) \propto \exp\left(-\frac{1}{\eta} Q_\theta(\mathbf{x}, \mathbf{u})\right), \quad (3)$$

approximating the expert model using a value function  $Q_\theta$  computed according to (1) with  $c$  and  $h$  replaced by  $c_\theta$  and  $h_\psi$ , respectively.

**Problem 2.** Given the demonstrations  $\mathcal{D}$ , optimize the parameters  $\theta$  of an approximation  $c_\theta$  of the unknown cost function  $c$  such that the log-likelihood of the demonstrated controls  $\mathbf{u}_t^n$  is maximized under the agent policy in (3):

$$\min_{\theta} \mathcal{L}_c(\theta) := -\sum_{n=1}^N \mathbb{1}_{\{s^n \geq \xi\}} \sum_{t=0}^{T_n} \log \pi_\theta(\mathbf{u}_t^n | \mathbf{x}_t^n), \quad (4)$$

where  $\mathbb{1}$  is an indicator function and  $\xi$  is the known task satisfaction threshold.

#### IV. TECHNICAL APPROACH

We first discuss how to learn a task model  $h_\psi$  from demonstrations  $\mathcal{D}$  in Sec. IV-A. Next, in Sec. IV-B, we learn a cost model  $c_\theta$  by solving the optimal control problem in (1) to obtain an agent policy  $\pi_\theta$ . Finally, in Sec. IV-C, we show how to backpropagate the policy loss  $\mathcal{L}_c(\theta)$  in (4) through the optimal control problem to update the cost parameters  $\theta$ .

##### A. Spectral learning of task specifications

Fitting a single cost neural network  $c_\theta$  that is capable of generalizing to various environment configurations and tasks is difficult when state and control spaces are large and the task horizon is long. An alternative is to consider the cost function and its corresponding policy only for small segments of the task, associated with different subtasks. This idea is based on the observation that task specifications often have a compositional logic structure. For example, the demonstrated trajectory in the DoorKey environment in

Fig. 1 can be decomposed into three segments, each denoted by a high-level state  $\alpha$ . The transitions between the high-level states are triggered by events like  $\sigma_0$ : a key is picked up, and  $\sigma_1$ : a door is opened. Note that there is no direct transition between  $\alpha_1$  and  $\alpha_3$  because the door cannot be opened without possessing a key. Such high-level state abstraction and transitions are commonly learned via recurrent neural network (RNN) or memory architectures [38], [39]. For example, to solve Problem 1, we can use an RNN  $h_\psi$  in Fig. 2 with initial hidden state  $\alpha_0$ , hidden state transition  $\alpha_{t+1} = g_1(\sigma_t, \alpha_t, \mathbf{W})$  and output function  $h_\psi(\sigma_{0:T}) = g_2(\alpha_{T+1}, \beta)$ , where  $g_1, g_2$  are nonlinear functions and  $\psi = (\alpha_0, \mathbf{W}, \beta)$  are learnable weights. Instead of an RNN model, in this work, we propose to use a weighted finite automaton (WFA) [40] to represent  $h_\psi$ . A WFA is less expressive than an RNN [41] but can be trained more effectively from small demonstration dataset. Moreover, a WFA generalizes deterministic and nondeterministic finite automata, which are commonly used to model logic task specifications for autonomous agents [20], [22]–[24]. Hence, a WFA model is sufficiently expressive to represent a complex task and allows one to focus on a temporal abstraction without reliance on the low-level system dynamics.

**Definition 2.** A weighted finite automaton (WFA) with  $m$  states is a tuple  $\psi = \{\alpha_0, \beta, \{\mathbf{W}_\sigma\}_{\sigma \in \Sigma}\}$  where  $\alpha_0, \beta \in \mathbb{R}^m$  are initial and final weight vectors and  $\mathbf{W}_\sigma \in \mathbb{R}^{m \times m}$  are transition matrices associated with each symbol  $\sigma \in \Sigma$ . A WFA  $\psi$  represents a function  $h_\psi: \Sigma^* \rightarrow \mathbb{R}$  by  $h_\psi(\sigma_{0:T}) = \alpha_0^\top \mathbf{W}_{\sigma_0} \mathbf{W}_{\sigma_1} \dots \mathbf{W}_{\sigma_T} \beta$ .

A WFA represents the task progress for a given word  $\sigma_{0:t}$  via  $h_\psi(\sigma_{0:t}) = \alpha_0^\top \mathbf{W}_{\sigma_0} \mathbf{W}_{\sigma_1} \dots \mathbf{W}_{\sigma_t} \beta$ , where the high-level state at time  $t+1$  is  $\alpha_{t+1} = (\alpha_0^\top \mathbf{W}_{\sigma_0} \mathbf{W}_{\sigma_1} \dots \mathbf{W}_{\sigma_t})^\top$ . When the WFA is learned correctly, its prediction for an expert word should approximate the expert score  $s$ , i.e.,  $h_\psi(\sigma_{0:T}) \approx s$ . This can be used to guide a task planning algorithm by providing a task satisfaction criterion. A trajectory with corresponding word  $\sigma_{0:T}$  is identified as successful if the WFA prediction passes the known performance threshold introduced in Sec. II-B, i.e.,  $h_\psi(\sigma_{0:T}) = \alpha_{T+1}^\top \beta \geq \xi$ .

Our approach to learn a minimal WFA is based on the spectral learning method developed by [40]. The spectral method makes use of a Hankel matrix  $\mathbf{H}_h \in \mathbb{R}^{\Sigma^* \times \Sigma^*}$  associated with the function  $h: \Sigma^* \rightarrow \mathbb{R}$ , which is a bi-infinite matrix with entries  $\mathbf{H}_h(u, v) = h(uv)$  for  $u, v \in \Sigma^*$ . We assume the class of functions  $h$  that can be represented by a WFA are rational power series functions and their associated Hankel matrix  $\mathbf{H}_h$  has finite rank [42], [43]. It can be shown that under certain assumptions WFA are expressively equivalent to monadic second-order logic. The quantitative property of WFA allows us to model the performance score  $s$  of the demonstrated trajectories.

**Assumption 1.** The Hankel matrix  $\mathbf{H}_h$  associated with the true task specification  $h$  has finite rank.

In practice, only finite sub-blocks of the Hankel matrix, constructed from the expert demonstrations  $\mathcal{D}$ , can be con-

sidered. Given a basis  $\mathcal{B} = (\mathcal{P}, \mathcal{S})$  where  $\mathcal{P}, \mathcal{S} \subset \Sigma^*$  are finite sets of prefixes and suffixes respectively, define  $\mathbf{H}_{\mathcal{B}}$  and  $\{\mathbf{H}_{\sigma}\}_{\sigma \in \Sigma}$  as the finite sub-blocks of  $\mathbf{H}_h$  such that  $\mathbf{H}_{\mathcal{B}}(u, v) = h(uv)$ ,  $\mathbf{H}_{\sigma}(u, v) = h(u\sigma v)$ ,  $\forall u \in \mathcal{P}, v \in \mathcal{S}$ . The foundation of the spectral learning method is summarized in the following theorem.

**Theorem 1** ([40]). *Given a basis  $\mathcal{B} = (\mathcal{P}, \mathcal{S})$  such that the empty string  $\lambda \in \mathcal{P} \cap \mathcal{S}$  and  $\text{rank}(\mathbf{H}_h) = \text{rank}(\mathbf{H}_{\mathcal{B}})$ , for any rank  $m$  factorization  $\mathbf{H}_{\mathcal{B}} = \mathbf{P}\mathbf{S}$  where  $\mathbf{P} \in \mathbb{R}^{|\mathcal{P}| \times m}$  and  $\mathbf{S} \in \mathbb{R}^{m \times |\mathcal{S}|}$ , the WFA  $\{\alpha_0, \beta, \{\mathbf{W}_{\sigma}\}\}$  is a minimal WFA representing  $h$ , where  $\alpha_0^\top = \mathbf{P}(\lambda, \cdot)$  is the row vector of  $\mathbf{P}$  corresponding to prefix  $\lambda$ ,  $\beta = \mathbf{S}(\cdot, \lambda)$  is the column vector of  $\mathbf{S}$  corresponding to suffix  $\lambda$ , and  $\mathbf{W}_{\sigma} = \mathbf{P}^\dagger \mathbf{H}_{\sigma} \mathbf{S}^\dagger$ ,  $\forall \sigma \in \Sigma$ .*

A basis can be chosen empirically from demonstrations  $\mathcal{D}$ . For example, we can choose a basis that includes all prefixes and suffixes that appear in the words  $\{\sigma_{0:T_n}^n\}$  or one with desired cardinality for the most frequent prefixes and suffixes. Given a basis, the Hankel blocks  $\mathbf{H}_{\mathcal{B}}, \{\mathbf{H}_{\sigma}\}$  are constructed from  $\mathcal{D}$ . For example, given a word and its score  $(\sigma_{0:T}, s)$ , we set the entries  $\mathbf{H}_{\mathcal{B}}(\lambda, \sigma_{0:T}), \mathbf{H}_{\mathcal{B}}(\sigma_0, \sigma_{1:T}), \dots, \mathbf{H}_{\mathcal{B}}(\sigma_{0:T}, \lambda)$ , and  $\mathbf{H}_{\sigma}(\sigma_{0:t-1}, \sigma_{t+1:T})$ , where  $\sigma = \sigma_t$  with value  $s$ . To find a low rank factorization of  $\mathbf{H}_{\mathcal{B}}$ , we use truncated singular value decomposition,  $\mathbf{H}_{\mathcal{B}} = \mathbf{U}_m \mathbf{\Lambda}_m \mathbf{V}_m^\top$  where  $\mathbf{\Lambda}_m$  is a diagonal matrix of the  $m$  largest singular values and  $\mathbf{U}_m, \mathbf{V}_m$  are the corresponding column vectors, and set  $\mathbf{P} = \mathbf{U}_m$  and  $\mathbf{S} = \mathbf{\Lambda}_m \mathbf{V}_m^\top$ . Finally, the vectors and matrices  $\psi = \{\alpha_0, \beta, \{\mathbf{W}_{\sigma}\}\}$  of the WFA can be obtained from  $\mathbf{P}, \mathbf{S}, \{\mathbf{H}_{\sigma}\}$  using Theorem 1.

### B. Planning in a product WFA-MDP system

Given a learned WFA representation  $h_\psi$  and an initial cost estimate  $c_\theta$ , we propose a planning algorithm to solve the deterministic optimal control problem in (1) and obtain a control policy  $\pi_\theta(\mathbf{u}|\mathbf{x})$  as in (3). To determine the termination condition for the problem in (1), we define the product of the WFA, modeling the task, and the L-MDP, modeling the agent-environment interactions.

**Definition 3.** Given an L-MDP  $\{\mathcal{X}, \mathcal{U}, \mathbf{x}_0, f, c, \mathcal{AP}, \ell\}$  and a WFA  $\{\alpha_0, \beta, \{\mathbf{W}_{\sigma}\}\}$ , a product WFA-MDP model is a tuple  $\{\mathcal{S}, \mathcal{U}, \mathbf{s}_0, T, \mathcal{S}_F, c, \mathcal{AP}, \ell\}$  where  $\mathcal{S} = \mathcal{X} \times \mathbb{R}^m$  is the product state space,  $\mathbf{s}_0 = (\mathbf{x}_0, \alpha_0)$  is the initial state, and  $\mathcal{S}_F = \{(\mathbf{x}, \alpha) \in \mathcal{S} \mid \alpha^\top \beta \geq \xi\}$  are the final states. The function  $T : \mathcal{S} \times \mathcal{U} \rightarrow \mathcal{S}$  is a deterministic transition function such that  $T((\mathbf{x}_t, \alpha_t), \mathbf{u}_t) = (\mathbf{x}_{t+1}, \alpha_{t+1})$  where  $\mathbf{x}_{t+1} = f(\mathbf{x}_t, \mathbf{u}_t)$ , emitting symbol  $\sigma_t = \ell(\mathbf{x}_t, \mathbf{u}_t)$  and causing transition  $\alpha_{t+1} = \mathbf{W}_{\sigma_t}^\top \alpha_t$ .

To obtain the agent policy in (3) for any state  $\mathbf{x}_t \in \mathcal{X}$  and control  $\mathbf{u}_t \in \mathcal{U}$ , our goal is to compute the optimal cost-to-go values for the WFA-MDP model:

$$\begin{aligned} Q_\theta(\mathbf{s}_t, \mathbf{u}_t) &= c_\theta(\mathbf{x}_t, \mathbf{u}_t) + V_\theta(T(\mathbf{s}_t, \mathbf{u}_t)) \\ &= c_\theta(\mathbf{x}_t, \mathbf{u}_t) + \min_{T, \mathbf{u}_{t+1:T}} \sum_{k=t+1}^T c_\theta(\mathbf{x}_k, \mathbf{u}_k) \end{aligned} \quad (5)$$

where  $\mathbf{s}_{t+1} = T(\mathbf{s}_t, \mathbf{u}_t)$  and  $\alpha_{T+1}^\top \beta \geq \xi$ . We have rewritten the terminal state condition as  $\alpha_{T+1}^\top \beta \geq \xi$ , where we keep track of the task hidden state  $\alpha_t$  using the WFA-MDP transition function  $T$ . Our key observation is that (5) is a deterministic shortest path problem and  $V_\theta(T(\mathbf{s}_t, \mathbf{u}_t))$  can be obtained via any shortest path algorithm, such as Dijkstra [44], A\* [45] or RRT\* [46]. When we use a shortest path algorithm to update the cost-to-go values of successor states  $\mathbf{s}_{t+1} = T(\mathbf{s}_t, \mathbf{u}_t)$ , we concurrently compute the corresponding WFA state  $\alpha_{t+1} = \mathbf{W}_{\sigma_t}^\top \alpha_t$  where  $\sigma_t = \ell(\mathbf{x}_t, \mathbf{u}_t)$ . A goal state  $\mathbf{s}_{T+1}$  is reached when its WFA state  $\alpha_{T+1}$  satisfies  $\alpha_{T+1}^\top \beta \geq \xi$ . The agent policy  $\pi_\theta$  in (3) with respect to the current cost estimate  $c_\theta$  can be obtained from the cost-to-go values  $Q_\theta$  in (5) computed by the shortest path algorithm.

### C. Optimizing cost parameters

We discuss how to differentiate the loss function  $\mathcal{L}_c(\theta)$  in (4) with respect to  $\theta$  through the deterministic shortest path problem defined by the product WFA-MDP model. [33] introduce a sub-gradient descent approach to differentiate the log likelihood of expert demonstrations from the Boltzman policy in (3) through the optimal cost-to-go values in (5). The cost parameters can be updated by (stochastic) subgradient descent at each iteration  $k$  with learning rate  $\gamma^{(k)}$ ,  $\theta^{(k+1)} = \theta^{(k)} - \gamma^{(k)} \nabla \mathcal{L}_c(\theta^{(k)})$ . Intuitively, the subgradient descent makes the trajectory starting with a demonstrated control more likely, while those with other controls less likely. The analytic subgradient computation is presented below.

**Proposition 1.** [33, Proposition 1] *Consider an expert transition  $(\mathbf{x}_t, \mathbf{u}_t)$ . Define  $\tau(\mathbf{x}_t, \mathbf{u})$  as the optimal path starting from state  $\mathbf{x}_t$  and any control  $\mathbf{u} \in \mathcal{U}$  that achieves  $Q_\theta(\mathbf{x}_t, \mathbf{u})$  in (5) under cost estimate  $c_\theta$ . A subgradient of the agent policy (3) evaluated at expert transition  $(\mathbf{x}_t, \mathbf{u}_t)$  with respect to cost parameters  $\theta$  can be obtained via the chain rule as:*

$$\begin{aligned} \frac{\partial \log \pi_\theta(\mathbf{u}_t \mid \mathbf{x}_t)}{\partial \theta} &= \sum_{\mathbf{u}' \in \mathcal{U}} \frac{d \log \pi_\theta(\mathbf{u}_t \mid \mathbf{x}_t)}{d Q_\theta(\mathbf{x}_t, \mathbf{u}')} \frac{\partial Q_\theta(\mathbf{x}_t, \mathbf{u}')}{\partial \theta} \\ &= \sum_{\mathbf{u}' \in \mathcal{U}} \frac{1}{\eta} (\mathbb{1}_{\{\mathbf{u}' = \mathbf{u}_t\}} - \pi_\theta(\mathbf{u}_t \mid \mathbf{x}_t)) \\ &\quad \times \sum_{(\mathbf{x}, \mathbf{u}) \in \tau(\mathbf{x}_t, \mathbf{u}')} \frac{\partial Q_\theta(\mathbf{x}_t, \mathbf{u}')}{\partial c_\theta(\mathbf{x}, \mathbf{u})} \frac{\partial c_\theta(\mathbf{x}, \mathbf{u})}{\partial \theta} \end{aligned} \quad (6)$$

Substituting (6) in the gradient of  $\mathcal{L}_c(\theta)$  in (4), Proposition 1 provides an explicit subgradient computation to allow backpropagation with respect to  $\theta$  through the value function  $Q_\theta$  of the deterministic shortest path problem in (5). The subgradient only affects the cost parameters through the optimal trajectories  $\tau(\mathbf{x}_t, \mathbf{u}')$ ,  $\forall \mathbf{u}' \in \mathcal{U}$  for expert transitions  $(\mathbf{x}_t, \mathbf{u}_t)$  which can be retrieved from any optimal planning algorithm applied in Sec IV-B. Thereafter, the cost parameters can be optimized depending on the specific form of  $\frac{\partial c_\theta(\mathbf{x}, \mathbf{u})}{\partial \theta}$ .

Our complete approach WFA-IRL is summarized in Fig. 3. We first solve Problem 1 to find a WFA  $h_\psi$  which models the demonstrated task. The learned WFA provides termination conditions for a deterministic shortest path problem in the



Fig. 3: WFA-IRL architecture for joint learning of a task specification  $h_\psi$  and cost function  $c_\theta$ . Given demonstrations  $\mathcal{D}$  and a labeling function  $\ell$ , we learn the unknown task specification with a weighted finite automaton. We construct a product WFA-MDP space from the learned WFA  $\psi = \{\alpha_0, \beta, \{\mathbf{W}_\sigma\}_{\sigma \in \Sigma}\}$  to solve a deterministic shortest path problem with cost estimate  $c_\theta$ . The agent policy  $\pi_\theta$  is compared with the demonstrated controls to backpropagate the loss  $\mathcal{L}_c(\theta)$  with respect to  $\theta$ .




Fig. 4: MiniGrid environments [9] of our experiments. In T1 (*MiniGrid-MultiRoom-N4-S5-v0*) the agent must pick up *Key* to unlock *Door* and reach *Goal* in the other room. In T2 (*MiniGrid-MultiRoom-N4-S5-v0*) it has to open a series of *Doors* to reach *Goal* in the last room. In T3 (*MiniGrid-BlockedUnlockPickup-v0*) it has to move away a blocking *Ball*, unlock *Door* with *Key* and pick up *Box*. The state  $\mathbf{x}_t$  includes the grid image  $\mathbf{m}_t \in \{\text{Wall}, \text{Key}, \text{Door}, \text{Box}, \text{Ball}, \text{Empty}\}^{H \times W}$ , the agent position  $\mathbf{p}_t \in \{1, \dots, H\} \times \{1, \dots, W\}$ , direction  $d_t \in \{\text{Up}, \text{Left}, \text{Down}, \text{Right}\}$ , and the object carried  $o_t \in \{\text{Key}, \text{Ball}, \text{Box}, \text{Empty}\}$ . The control space  $\mathcal{U}$  is defined as turn left/right, move forward, pick up/drop/toggle an object.

product WFA-MDP. Cost parameters are optimized by back-propagating the loss in (4) through the planning algorithm.

## V. EVALUATION

We consider three MiniGrid tasks shown in Fig. 4 whose atomic propositions are shown in Table I. The task specifications can be expressed in terms of these propositions, e.g., one possible trajectory that fulfills task T1 is to evaluate the propositions  $p_1, p_2, p_3$  as true sequentially.

TABLE I: Atomic propositions used in each task.

	T1	T2	T3
$p_1$	<i>Key</i> is picked up	<i>Door</i> 1 is open	<i>Ball</i> is 2 steps away from <i>Door</i>
$p_2$	<i>Door</i> is open	<i>Door</i> 2 is open	<i>Key</i> is picked up
$p_3$	Agent reaches <i>Goal</i>	<i>Door</i> 3 is open	<i>Door</i> is open
$p_4$	—	Agent reaches <i>Goal</i>	<i>Box</i> is picked up

### A. Demonstrations

An expert trajectory is collected by iteratively rolling out the controls sampled from the expert policy  $\pi^*$  at each state  $\mathbf{x}$ , where  $Q^*(\mathbf{x}, \mathbf{u})$  in (1) is computed via the Dijkstra’s algorithm with cost of 1 for any feasible transition. For each task, we consider two sets of expert demonstrations  $\mathcal{D}_1$  and  $\mathcal{D}_2$ , each with 32 trajectories collected from expert policies with temperatures  $\eta \in \{0, 0.5\}$ . The expert trajectories in  $\mathcal{D}_1$  and  $\mathcal{D}_2$  are strictly optimal and suboptimal, respectively, and are labeled with score  $s = 1$ . In each set we also add 128 failed trajectories with score  $s = 0$  from a random exploration

TABLE II: Results on MiniGrid environment tasks. In each entry, Green / Orange are results trained from demonstrations  $\mathcal{D}_1$  with expert policy temperature  $\eta = 0$  and  $\mathcal{D}_2$  with  $\eta = 0.5$ , respectively. Top: Best Scikit-SpLearn hyperparameters that solve Problem 1 for each task. Bottom: Mean episode returns (or negative cumulative true cost, higher is better) are reported across 64 randomly generated test environments.

	T1	T2	T3
rank	5/9	6/9	7/11
rows	4/5	5/5	6/6
cols	4/5	5/6	6/7

	T1	T2	T3
BC	0.364/0.253	0.338/0.307	0.284/0.192
GAIL	0.483/0.429	0.274/0.185	0.342/0.257
WFA-IRL(ours)	0.797/0.708	0.776/0.642	0.733/0.602
WFA-IRL w/o WFA	0.683/0.514	0.652/0.488	0.566/0.390
Expert	0.798/0.718	0.776/0.668	0.734/0.639
Optimal	0.798	0.776	0.734
Random	0.000	0.000	0.000

policy (effectively setting expert policy temperature  $\eta \rightarrow \infty$ ). The full demonstration set is used in each case to learn a WFA representation  $h_\psi$  of the task via the spectral method in Sec. IV-A while only the successful trajectories are used to learn the cost function  $c_\theta$ , as in Sec. IV-B and IV-C.

### B. Our method and baselines

Our method WFA-IRL uses a neural network architecture to represent the cost function. For a detailed description, please refer to the extended version of our paper [47]. We use the spectral learning algorithm in the Scikit-SpLearn toolbox [48] to learn the parameters  $\psi$  of the WFA from the expert demonstrations  $\mathcal{D}$ . In the implementation, we first compress the demonstration words  $\sigma_{0:T_n}^n$  where consecutive identical symbols are removed. This greatly reduces the complexity of learning the WFA while keeping the symbol sequences unchanged. The hyperparameters for the spectral learning method are the rank of the automaton ( $m$  in Theorem 1) and the sizes (rows and cols) of the prefix and suffix basis ( $\mathcal{B} = (\mathcal{P}, \mathcal{S})$  in Theorem 1), which determine the size of the Hankel matrix estimated empirically from demonstrations. The complexity is  $O(n \times \text{rank} \times \text{rows} \times \text{cols})$  since we iterate through the ranks to find a minimal WFA and count the prefix and suffix frequencies for a given word. The spectral method can learn almost perfectly with near zero loss in (2) for all tasks and the best hyperparameter configurations are shown in Table II. We observe that larger WFA capacity is required to learn from suboptimal trajectories and thus more diverse words from  $\mathcal{D}_2$ .

As baselines, we first ablate the WFA component in our method to understand its effects. Instead of planning in the product WFA-MDP space and checking the WFA termination condition in (5), the agent without WFA component simply plans in the original MDP and checks whether a goal state is achieved. Additionally, we compare our method with standard imitation learning and inverse reinforcement learning algorithms, including behavioral cloning (BC) [49] and GAIL [50]<sup>1</sup>. The value/policy functions in these baselines

<sup>1</sup>The implementations are adapted from the imitation learning library [51]

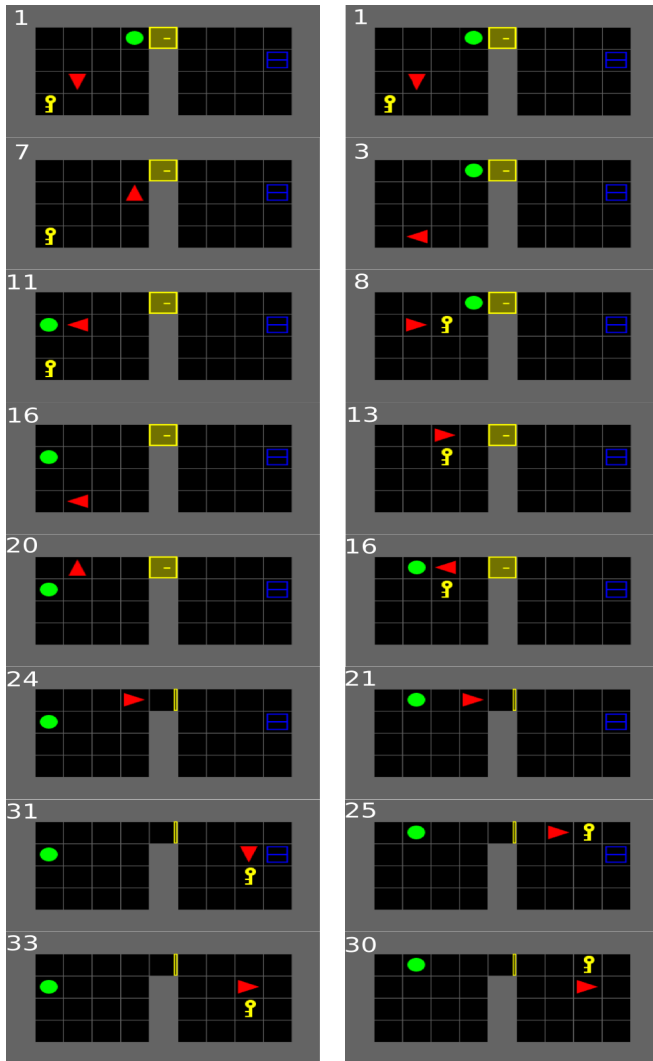


Fig. 5: Agent trajectory (left column) trained with  $\mathcal{D}_1$  and expert trajectory (right column) in task T3 during testing. The agent WFA only learns words that appear in demonstration  $\mathcal{D}_1$ , which always moves *Ball* away from *Door* first before picking up *Key*. In testing it fails to recognize a lower cost trajectory of a different word sequence, where *Key* is carried closer to *Door* before moving away *Ball*.

follow our cost neural network architecture to fit the state-control input format and to compare fairly in representation power across methods. This includes the policy network in BC, the discriminator in GAIL, the policy and value networks in PPO [52], used as generator in GAIL. Only the size of the last fully-connected layer is modified depending on whether it is action or value prediction. GAIL is known to achieve stable performance in fixed horizon environments while the MiniGrid environments terminate as soon as the agent fulfills the tasks. We fix this issue by adding a virtual absorbing state as suggested in [53] when training GAIL.

### C. Results

We report the average performance of each method in Table II by testing on 64 new environment configurations generated randomly for each task. First, we observe that our method can achieve almost perfect performance when trained




Fig. 6: Visualization of policy probabilities of each method trained on  $\mathcal{D}_2$  at a critical state in T2. Our method shows a stronger preference towards controls (toggle door) that can make task progress.

on  $\mathcal{D}_1$ . This is expected since the learned WFA strictly chooses planned trajectories whose words would match the optimal behavior. Interestingly, the learned WFA could make the agent suboptimal if the optimal word in testing is not seen in training, as shown in Fig. 5. Next, our method matches the expert performance well using either  $\mathcal{D}_1$  or  $\mathcal{D}_2$  and outperforms BC and GAIL (even without WFA). This demonstrates that using planning to solve tasks that encode logical structures performs better than a reactive policy employed by BC. Moreover, the performance gap between our method and ours without WFA shows that learning logic specifications explicitly with a WFA can further improve the policy. On the other hand, we find the performance of GAIL is limited as PPO cannot easily generate successful samples similar to the demonstrations (notice that the random policy never succeeds) to improve the cost discriminator and, in turn, the generator itself. We visualize the agent policy in Fig. 6 and observe that our method has a stronger bias on controls that follow the learned logical sequences.

## VI. CONCLUSION

We present WFA-IRL which solves tasks with high-level reasoning and outperforms prior imitation learning and IRL methods that do not exploit logical structures from demonstrations. We show that cost functions learned via solving deterministic shortest path problems in the product WFA-MDP can generalize well in unseen environments and across demonstrations of different optimality levels.

## REFERENCES

[1] A. Y. Ng and S. Russell, “Algorithms for inverse reinforcement learning,” in *International Conference on Machine Learning*, 2000.

- [2] N. D. Ratliff, J. A. Bagnell, and M. A. Zinkevich, "Maximum margin planning," in *International Conference on Machine Learning*, 2006.
- [3] B. D. Ziebart, A. Maas, J. Bagnell, and A. K. Dey, "Maximum entropy inverse reinforcement learning," in *AAAI Conference on Artificial Intelligence*, 2008.
- [4] S. Levine, Z. Popovic, and V. Koltun, "Nonlinear inverse reinforcement learning with gaussian processes," *Advances in Neural Information Processing Systems*, vol. 24, pp. 19–27, 2011.
- [5] M. Wulfmeier, D. Z. Wang, and I. Posner, "Watch this: Scalable cost-function learning for path planning in urban environments," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2016.
- [6] C. Finn, S. Levine, and P. Abbeel, "Guided cost learning: Deep inverse optimal control via policy optimization," in *International Conference on Machine Learning*, 2016.
- [7] M. Vazquez-Chanlatte, S. Jha, A. Tiwari, M. K. Ho, and S. A. Seshia, "Learning task specifications from demonstrations," *Advances in Neural Information Processing Systems*, 2017.
- [8] S. Krishnan, A. Garg, R. Liaw, B. Thananjeyan, L. Miller, F. T. Pokorny, and K. Goldberg, "SWIRL: A sequential windowed inverse reinforcement learning algorithm for robot tasks with delayed rewards," *The International Journal of Robotics Research*, vol. 38, no. 2-3, pp. 126–145, 2019.
- [9] M. Chevalier-Boisvert, L. Willems, and S. Pal, "Minimalistic grid-world environment for openai gym," <https://github.com/maximecb/gym-minigrid>, 2018.
- [10] R. S. Sutton, D. Precup, and S. Singh, "Between mdps and semi-mdps: A framework for temporal abstraction in reinforcement learning," *Artificial intelligence*, vol. 112, no. 1-2, pp. 181–211, 1999.
- [11] T. D. Kulkarni, K. Narasimhan, A. Saedi, and J. Tenenbaum, "Hierarchical deep reinforcement learning: Integrating temporal abstraction and intrinsic motivation," in *Advances in Neural Information Processing Systems*, 2016.
- [12] P.-L. Bacon, J. Harb, and D. Precup, "The option-critic architecture," in *AAAI Conference on Artificial Intelligence*, 2017.
- [13] M. Riemer, M. Liu, and G. Tesauro, "Learning abstract options," 2018.
- [14] R. Fox, S. Krishnan, I. Stoica, and K. Goldberg, "Multi-level discovery of deep options," *arXiv preprint arXiv:1703.08294*, 2017.
- [15] T. Kipf, Y. Li, H. Dai, V. Zambaldi, A. Sanchez-Gonzalez, E. Grefenstette, P. Kohli, and P. Battaglia, "CompILE: Compositional imitation learning and execution," in *International Conference on Machine Learning*, 2019.
- [16] F. Xie, A. Chowdhury, M. Kaluza, L. Zhao, L. L. Wong, and R. Yu, "Deep imitation learning for bimanual robotic manipulation," in *Advances in Neural Information Processing Systems*, 2020.
- [17] E. Plaku and S. Karaman, "Motion planning with temporal-logic specifications: Progress and challenges," *AI communications*, pp. 151–162, 2016.
- [18] M. Luckcuck, M. Farrell, L. Dennis, C. Dixon, and M. Fisher, "Formal specification and verification of autonomous robotic systems," *ACM Computing Surveys (CSUR)*, 2019.
- [19] M. Farrell, M. Luckcuck, and M. Fisher, "Robotics and integrated formal methods: Necessity meets opportunity," in *International Conference on Integrated Formal Methods*, 2018.
- [20] G. E. Fainekos, A. Girard, H. Kress-Gazit, and G. J. Pappas, "Temporal logic motion planning for dynamic robots," *Automatica*, vol. 45, no. 2, pp. 343–352, 2009.
- [21] C. Baier and J.-P. Katoen, *Principles of Model Checking*. MIT press, 2008.
- [22] H. Kress-Gazit, G. E. Fainekos, and G. J. Pappas, "Where's waldo? sensor-based temporal logic motion planning," in *IEEE International Conference on Robotics and Automation*, 2007.
- [23] —, "Temporal-logic-based reactive mission and motion planning," *IEEE Transactions on Robotics*, vol. 25, no. 6, pp. 1370–1381, 2009.
- [24] G. E. Fainekos, H. Kress-Gazit, and G. J. Pappas, "Hybrid controllers for path planning: A temporal logic approach," in *IEEE Conference on Decision and Control*, 2005.
- [25] A. Bhatia, L. E. Kavraki, and M. Y. Vardi, "Sampling-based motion planning with temporal goals," in *IEEE International Conference on Robotics and Automation*, 2010.
- [26] J. Fu, N. Atanasov, U. Topcu, and G. J. Pappas, "Optimal temporal logic planning in probabilistic semantic maps," in *IEEE International Conference on Robotics and Automation*, 2016.
- [27] C. Lemieux, D. Park, and I. Beschastnikh, "General ltl specification mining (t)," in *2015 30th IEEE/ACM International Conference on Automated Software Engineering (ASE)*, 2015.
- [28] H. J. Kang and D. Lo, "Adversarial specification mining," *ACM Transactions on Software Engineering and Methodology (TOSEM)*, 2021.
- [29] A. J. Shah, P. Kamath, S. Li, and J. A. Shah, "Bayesian inference of temporal task specifications from demonstrations," 2018.
- [30] G. Chou, N. Ozay, and D. Berenson, "Explaining multi-stage tasks by learning temporal logic formulas from suboptimal demonstrations," in *Robotics: Science and Systems (RSS)*, 2020.
- [31] M. Droste, W. Kuich, and H. Vogler, *Handbook of Weighted Automata*. Springer Science & Business Media, 2009.
- [32] M. Droste and P. Gastin, "Weighted automata and weighted logics," in *International Colloquium on Automata, Languages, and Programming*. Springer, 2005, pp. 513–525.
- [33] T. Wang, V. Dhiman, and N. Atanasov, "Learning navigation costs from demonstration in partially observable environments," in *IEEE International Conference on Robotics and Automation*, 2020.
- [34] X. C. Ding, S. L. Smith, C. Belta, and D. Rus, "Mdp optimal control under temporal logic constraints," in *IEEE Conference on Decision and Control and European Control Conference*, 2011.
- [35] D. Bertsekas, *Dynamic Programming and Optimal Control*. Athena Scientific, 1995.
- [36] G. Neu and C. Szepesvári, "Apprenticeship learning using inverse reinforcement learning and gradient methods," in *Conference on Uncertainty in Artificial Intelligence*, 2007.
- [37] D. Ramachandran and E. Amir, "Bayesian inverse reinforcement learning," in *International Joint Conference on Artificial Intelligence*, 2007.
- [38] M. Hausknecht and P. Stone, "Deep recurrent q-learning for partially observable mdps," in *AAAI Fall Symposium Series*, 2015.
- [39] P. Mirowski, R. Pascanu, F. Viola, H. Soyer, A. J. Ballard, A. Banino, M. Denil, R. Goroshin, L. Sifre, K. Kavukcuoglu, et al., "Learning to navigate in complex environments," *International Conference on Learning Representations*, 2017.
- [40] B. Balle and M. Mohri, "Spectral learning of general weighted automata via constrained matrix completion," in *Advances in Neural Information Processing Systems*, 2012.
- [41] G. Rabusseau, T. Li, and D. Precup, "Connecting weighted automata and recurrent neural networks through spectral learning," in *International Conference on Artificial Intelligence and Statistics*, 2019.
- [42] J. Berstel and C. Reutenauer, *Rational Series and Their Languages*. Springer-Verlag, 1988, vol. 12.
- [43] A. Salomaa and M. Soittola, *Automata-theoretic Aspects of Formal Power Series*. Springer Science & Business Media, 2012.
- [44] E. W. Dijkstra, "A note on two problems in connexion with graphs," *Numerische Mathematik*, vol. 1, no. 1, pp. 269–271, 1959.
- [45] M. Likhachev, G. Gordon, and S. Thrun, "ARA\*: Anytime a\* with provable bounds on sub-optimality," in *Advances in Neural Information Processing Systems*, 2004.
- [46] S. Karaman and E. Frazzoli, "Sampling-based algorithms for optimal motion planning," *The International Journal of Robotics Research*, vol. 30, no. 7, pp. 846–894, 2011.
- [47] T. Wang and N. Atanasov, "WFA-IRL: Inverse reinforcement learning of autonomous behaviors encoded as weighted finite automata," *arXiv preprint arXiv:2103.05895*, 2021.
- [48] D. Arrivault, D. Benielli, F. Denis, and R. Eyraud, "Scikit-splearn: a toolbox for the spectral learning of weighted automata compatible with scikit-learn," in *Conférence francophone sur l'Apprentissage Automatique*, 2017.
- [49] S. Ross and D. Bagnell, "Efficient reductions for imitation learning," in *International Conference on Artificial Intelligence and Statistics*, 2010, pp. 661–668.
- [50] J. Ho and S. Ermon, "Generative adversarial imitation learning," *Advances in Neural Information Processing Systems*, 2016.
- [51] S. Wang, S. Toyer, A. Gleave, and S. Emmons, "The imitation library for imitation learning and inverse reinforcement learning," <https://github.com/HumanCompatibleAI/imitation>, 2020.
- [52] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," *arXiv preprint arXiv:1707.06347*, 2017.
- [53] I. Kostrikov, K. K. Agrawal, D. Dwibedi, S. Levine, and J. Tompson, "Discriminator-actor-critic: Addressing sample inefficiency and reward bias in adversarial imitation learning," *International Conference on Learning Representations*, 2019.