

# Inverse Reinforcement Learning via Nonparametric Spatio-Temporal Subgoal Modeling

**Adrian Šošić**

**Abdelhak M. Zoubir**

*Signal Processing Group*

*Technische Universität Darmstadt*

*64283 Darmstadt, Germany*

ADRIAN.SOSIC@SPG.TU-DARMSTADT.DE

ZOUBIR@SPG.TU-DARMSTADT.DE

**Elmar Rueckert**

*Institute for Robotics and Cognitive Systems*

*University of Lübeck*

*23538 Lübeck, Germany*

RUECKERT@ROB.UNI-LUEBECK.DE

**Jan Peters**

*Autonomous Systems Labs*

*Technische Universität Darmstadt*

*64289 Darmstadt, Germany*

MAIL@JAN-PETERS.NET

**Heinz Koepl**

*Bioinspired Communication Systems*

*Technische Universität Darmstadt*

*64283 Darmstadt, Germany*

HEINZ.KOEPPL@BCS.TU-DARMSTADT.DE

**Editor:** George Konidaris

## Abstract

Advances in the field of inverse reinforcement learning (IRL) have led to sophisticated inference frameworks that relax the original modeling assumption of observing an agent behavior that reflects only a single intention. Instead of learning a global behavioral model, recent IRL methods divide the demonstration data into parts, to account for the fact that different trajectories may correspond to different intentions, e.g., because they were generated by different domain experts. In this work, we go one step further: using the intuitive concept of *subgoals*, we build upon the premise that even a single trajectory can be explained more efficiently *locally* within a certain context than globally, enabling a more compact representation of the observed behavior. Based on this assumption, we build an implicit intentional model of the agent's goals to forecast its behavior in unobserved situations. The result is an integrated Bayesian prediction framework that significantly outperforms existing IRL solutions and provides smooth policy estimates consistent with the expert's plan. Most notably, our framework naturally handles situations where the intentions of the agent change over time and classical IRL algorithms fail. In addition, due to its probabilistic nature, the model can be straightforwardly applied in active learning scenarios to guide the demonstration process of the expert.

**Keywords:** Learning from Demonstration, Inverse Reinforcement Learning, Bayesian Nonparametric Modeling, Subgoal Inference, Graphical Models, Gibbs Sampling

## 1. Introduction

Inverse reinforcement learning (IRL) refers to the problem of inferring the intention of an agent, called *the expert*, from observed behavior. Under the Markov decision process (MDP) formalism (Sutton and Barto, 1998), that intention is encoded in the form of a reward function, which provides the agent with instantaneous feedback for each situation encountered during the decision-making process. Classical IRL methods (Ng and Russell, 2000; Abbeel and Ng, 2004; Ziebart et al., 2008; Ramachandran and Amir, 2007; Levine et al., 2011) assume there exists a single *global* reward model that explains the entire set of demonstrations provided by the expert. In order to relax this rather restrictive modeling assumption, recent IRL methods allow that the agent’s intention can change over time (Nguyen et al., 2015), or they presume that the demonstration data set is inherently composed of several parts (Dimitrakakis and Rothkopf, 2011), where different trajectories reflect the intentions of different domain experts.

In this work, we go a step further and start from the premise that—even in the case of a single expert or trajectory—the demonstrated behavior can be explained more efficiently *locally* (i.e., within a certain context) than by a global reward model. As an illustrative example, we may consider the task shown in Figure 1a, where the expert approaches a set of intermediate target positions before finally heading toward a global goal state. Similarly, in Figure 1b, the agent eventually returns to its initial position, from where the cyclic process repeats. Despite the simplicity of these tasks, the encoding of such behaviors in a global intention model requires a reward structure that comprises a comparably large number of redundant state-action-based rewards. Alternative modeling strategies rely on task-dependent expansions of the agent’s state representation, e.g., to memorize the last visited goal (Krishnan et al., 2016), or they resort to more general decision-making frameworks like semi-MDPs/options (Bradtke and Duff, 1994; Sutton et al., 1999) in order to achieve the necessary level of task abstraction.

In this paper, we present a substantially simpler modeling framework that requires only minimal adaptations to the standard MDP formalism but comes with a hypothesis space of behavioral models that is sufficiently large to cover a broad class of expert policies. The key insight that motivates our approach is that many tasks, like those in Figure 1, can be decomposed into smaller subtasks that require considerably less modeling effort. The resulting low-level task descriptions can then be used as building blocks to synthesize arbitrarily complex behavioral strategies through a suitable sequencing of subtasks. This offers the possibility to learn comparably simple task representations using the intuitive concept of *subgoals*, which is achieved by efficiently encoding the expert behavior using task-adapted partitionings of the system state space/the expert data.

The proposed framework builds upon the method of Bayesian nonparametric inverse reinforcement learning (BNIRL, Michini and How, 2012), which can be used to build a subgoal representation of a task based on demonstration data—however, without learning the underlying subgoal relationships or providing a policy model that can generalize the strategy of the demonstrator. In order to address this limitation, we generalize the BNIRL model using insights from our previous works on nonparametric subgoal modeling (Šošić et al., 2018a) and policy recognition (Šošić et al., 2018b), building a compact intentional model of the expert’s behavior that explicitly describes the local dependencies between the




Figure 1: Two simple behavior examples that motivate the subgoal principle. The setting is based on the grid world dynamics described in Section 5.1. In both cases, a task description based on a global reward function is inefficient as it requires many state-action-based rewards to explain the observed trajectory structures. However, the data can be described efficiently through subgoal-based encodings. Both scenarios are analyzed in detail in Section 5.

demonstrations and the underlying subgoal structure. The result is an integrated Bayesian prediction framework that exploits the spatio-temporal context of the demonstrations and is capable of producing smooth policy estimates that are consistent with the expert’s plan. Furthermore, capturing the full posterior information of the data set enables us to apply the proposed approach in an active learning setting, where the data acquisition process is controlled by the posterior predictive distribution of our model.

In our experimental study, we compare the proposed approach with common baseline methods on a variety of benchmark tasks and real-world scenarios. The results reveal that our approach performs significantly better than the original BNIRL model and alternative IRL solutions on all considered tasks. Interestingly enough, our algorithm outperforms the baselines even when the expert’s true reward structure is dense and the underlying subgoal assumption is violated.

### 1.1 Related Work

The idea of decomposing complex behavior into smaller parts has been around for long and researchers have approached the problem in many different ways. While the overall field of methods is too large to be covered here, most existing approaches can be clearly categorized according to certain criteria. Often, two approaches differ in their exact problem formulation, i.e., we can distinguish between *active* methods, where the learning algorithm can interact freely with the environment (e.g., hierarchical reinforcement learning, Botvinick, 2012; Al-Emran, 2015), and *passive* methods, where the behavioral model is trained solely through observation (learning from demonstration, Argall et al., 2009). Furthermore, we

can discriminate between methods that build an explicit intentional model of the underlying task (IRL and option-based models, Choi and Kim, 2012; Sutton et al., 1999), and such that work directly on the control/trajectory level (skill learning, movement primitives, Konidaris et al., 2012; Schaal et al., 2005). The latter distinction is sometimes also referred to as *intentional/subintentional* approaches (Albrecht and Stone, 2017; Panella and Gmytrasiewicz, 2017). In order to give a concise summary of the work that is most relevant to ours, we restrict ourselves to passive approaches, with a focus on intentional methods, the field of which is considerably smaller. For an overview of active approaches, we refer to existing literature, e.g., the work by Daniel et al. (2016a).

First, there is the class of methods that pursue a decomposition of the observed behavior on the global level, using trajectory-based IRL approaches. For example, Dimitrakakis and Rothkopf (2011) proposed a hierarchical prior over reward functions to account for the fact that different trajectories in a data set could reflect different behavioral intentions, e.g., because they were generated by different domain experts. Similarly, Babeş-Vroman et al. (2011) follow an expectation-maximization-based clustering approach to group individual trajectories according to their underlying reward functions. Choi and Kim (2012) generalized this idea by proposing a nonparametric Bayesian model in which the number of intentions is a priori unbounded.

While the above methods consider the expert data at a global scale, our work is concerned with the problem of *subgoal modeling*, which is often conducted in the form of option-based reasoning (Sutton et al., 1999). For instance, Tamassia et al. (2015) proposed a clustering approach based on state distances to find a minimal set of options that can explain the expert behavior. While the method provides a simple alternative to handcrafting options, it does not allow any probabilistic treatment of the data and involves many ad-hoc design choices. Going in the same direction, Daniel et al. (2016a) presented a more principled, probabilistic option framework based on expectation-maximization. Not only is the framework capable of inferring sub-policies automatically, it can be also used in a reinforcement learning context for intra-option learning. However, the resulting behavioral model is based on point estimates of the policy parameters, and the number of sub-policies needs to be specified manually. The latter problem was solved by Krishnan et al. (2016), who proposed a hierarchical nonparametric IRL framework to learn a sequential representation of the demonstrated task, based on a set of transition regions that are defined through local changes in linearity of the observed behavior. However, in contrast to the work by Daniel et al. (2016a), inference is not performed jointly but in several isolated stages where, again, each stage only propagates a point estimate of the associated model parameters. Moreover, the temporal relationship of the demonstration data, used to identify the local linearity changes, is considered only in an ad-hoc fashion with the help of a windowing function.

Another general class of models, which explicitly addresses this issue, employs a hidden Markov model (HMM) structure to establish a temporal relationship between the demonstrations. For instance, the work presented by Nguyen et al. (2015) can be regarded as a generalization of the model by Babeş-Vroman et al. (2011), which extends the expectation-maximization framework by imposing a Markov structure on the reward model. Similarly, Niekum et al. (2012) use an extended HMM to segment the demonstrations into vector autoregressive models, in order to learn a suitable set of movement primitives. However, the learning of those primitives is done in a post-processing step, meaning that the quality of

the final representation crucially depends on the success of the initial segmentation stage. In contrast, the method by Rueckert et al. (2013) automatically learns the position and timing of subgoals in the form of via-points, but the number of via-points is assumed to be known and the system objective gets finally encoded in form of a global cost function. Recently, Lioutikov et al. (2017) presented a related approach based on probabilistic movement primitives that jointly solves the segmentation and learning step for an unknown number of primitives, using an expectation-maximization framework. Yet, the model operates purely on the trajectory level and cannot reveal the latent intentions of the demonstrator. Another variant of the approach by Niekum et al. (2012) that explicitly addresses this problem was proposed by Surana and Srivastava (2014). In their paper, the authors propose to replace the HMM emission model with an MDP model, in order to infer a policy model from the segmented trajectories instead of recognizing changes in the dynamics. The model was later extended by Ranchod et al. (2015), who augmented the HMM representation with a beta process model to facilitate skill sharing across trajectories. While the resulting model formulation is highly flexible, its major drawback is that inference becomes computationally expensive as it involves multiple IRL iterations per Gibbs step.

In contrast to the HMM-based solutions, which by their sequential nature focus on the *temporal* relationship of subtasks, the approach presented in this paper establishes a more general correlation structure between demonstrations by employing *non-exchangeable prior distributions* over subgoal assignments, i.e., without committing to purely temporal factorizations of subgoals. This results in a compact model representation (e.g., it avoids the need of estimating latent subgoal transition probabilities required in an HMM structure) and adds the flexibility to capture both, the temporal and the spatial dependencies between subtasks.

## 1.2 Paper Outline

The organization of the paper is as follows: in Section 2, we briefly revisit the BNIRL model and discuss its limitations, which forms the basis for our work. Section 3 then introduces a new intentional subgoal framework, which addresses the shortcomings of BNIRL discussed in Section 2. In Section 4, we derive a sampling-based inference scheme for our model and explain how the new framework can be used for subgoal extraction and action prediction. Experimental results on both synthetic and real-world data are presented in Section 5 before we finally conclude our work in Section 6.

## 2. Bayesian Nonparametric Inverse Reinforcement Learning

The purpose of this section is to recapitulate the principle of Bayesian nonparametric inverse reinforcement learning. After briefly discussing all building blocks of the model, we focus on the limitations of the framework, which motivates the need for an extended model formulation and finally leads to a new inference approach, presented afterwards in Section 3.

### 2.1 Revisiting the BNIRL Framework

Following the common IRL paradigm (Ng and Russell, 2000; Zhifei and Joo, 2012), the goal of BNIRL is to infer the intentions of an agent based on demonstration data. Starting from a standard MDP model, the problem is formalized on a finite state space  $\mathcal{S}$ , assuming a

time-invariant state transition model  $T : \mathcal{S} \times \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$ , where  $\mathcal{A}$  is a finite set of actions available to the agent at each state. For notational convenience, we represent the states in  $\mathcal{S}$  by the integer values  $\{1, \dots, |\mathcal{S}|\}$ , where  $|\mathcal{S}|$  denotes the cardinality of the state space.

In BNIRL, it is assumed that we can observe a number of expert demonstrations provided in the form of state-action pairs,  $\mathcal{D} := \{(s_d, a_d)\}_{d=1}^D$ , where each pair  $(s_d, a_d) \in \mathcal{S} \times \mathcal{A}$  consists of a state  $s_d$  visited by the agent and the corresponding action  $a_d$  taken. Herein,  $D$  denotes the size of the demonstration set. Throughout the rest of this paper, we will use the shorthand notations  $\mathbf{s} := \{s_d\}_{d=1}^D$  and  $\mathbf{a} := \{a_d\}_{d=1}^D$  to access the collections of expert states and actions individually. Note that the BNIRL model makes no assumptions about the temporal ordering of the demonstrations, i.e., each state-action pair is considered to have arisen from a specific but arbitrary time instant of the agent’s decision-making process. We will come back to this point later in Sections 2.2 and 3.3.

In contrast to the classical MDP formalism and most other IRL frameworks, BNIRL does *not* presuppose that the observed expert behavior necessarily originates from a single underlying reward function. Instead, it introduces the concept of *subgoals* (and corresponding *subgoal assignments*) with the underlying assumption that, at each decision instant, the expert selects a particular subgoal to plan the next action. Each subgoal is herein represented by a certain reward function defined on the system state space; in the simplest case, it corresponds to a single reward mass placed at a particular goal state in  $\mathcal{S}$ , which we identify with a reward function  $R_g : \mathcal{S} \rightarrow \{0, C\}$  of the form

$$R_g(s) := \begin{cases} C & \text{if } g = s, \\ 0 & \text{otherwise,} \end{cases} \quad (1)$$

where  $g \in \{1, \dots, |\mathcal{S}|\}$  indicates the subgoal location and  $C \in (0, \infty)$  is some positive constant (compare Şimşek et al., 2005; Stolle and Precup, 2002; Tamassia et al., 2015).

Although in principle it is legitimate to associate each subgoal with an arbitrary reward structure to encode more complex forms of goal-oriented behavior (see, for example, Rancho et al., 2015), the restriction to the reward function class in Equation (1) is sufficient in the sense that the same behavioral complexity can be synthesized through a combination of subgoals. This is made possible by the nonparametric nature of BNIRL, i.e., because the number of possible subgoals is assumed to be unbounded. The use of the reward model in Equation (1) has the advantage, however, that posterior inference about the expert’s subgoals becomes computationally tractable, as will be explained in Section 4.5. In the following, we therefore focus on the above reward model and summarize the infinite collection of subgoals in the multiset  $\mathcal{G} := \{g_k\}_{k=1}^\infty \in \times_{k=1}^\infty \mathcal{S}$ , where we adopt the assumption that  $p(\mathcal{G} | \mathbf{s}) = \prod_{k=1}^\infty p_g(g_k | \mathbf{s})$ .<sup>1</sup>

The subgoal assignment in BNIRL is achieved using a set of indicator variables  $\tilde{\mathbf{z}} := \{\tilde{z}_d \in \mathbb{N}\}_{d=1}^D$ , which annotate each demonstration pair  $(s_d, a_d)$  with its unique subgoal index. The prior distribution  $p(\tilde{\mathbf{z}})$  is modeled by a Chinese restaurant process (CRP, Aldous, 1985),

---

1. Notice that the subgoal prior distribution in the original BNIRL formulation does not take the state variable  $\mathbf{s}$  as an argument. Nonetheless, the authors of BNIRL suggest to restrict the support of the distribution to the set of visited states, which indeed implies a conditioning on  $\mathbf{s}$ .

which assigns the event that indicator  $\tilde{z}_d$  points to the  $j$ th subgoal the prior probability

$$p(\tilde{z}_d = j \mid \tilde{\mathbf{z}}_{\setminus d}) \propto \begin{cases} n_j & \text{if } j \in \{1, \dots, K\}, \\ \alpha & \text{if } j = K + 1, \end{cases}$$

where  $\tilde{\mathbf{z}}_{\setminus d} := \{\tilde{z}_d\} \setminus \tilde{z}_d$  is a shorthand notation for the collection of all indicator variables except  $\tilde{z}_d$ . Further,  $n_j$  denotes the number of assignments to the  $j$ th subgoal in  $\tilde{\mathbf{z}}_{\setminus d}$ ,  $K$  represents the number of distinct entries in  $\tilde{\mathbf{z}}_{\setminus d}$ , and  $\alpha \in [0, \infty)$  is a parameter controlling the diversity of assignments.

Having targeted a particular subgoal  $g_{\tilde{z}_d}$  while being at some state  $s_d$ , the expert is assumed to choose the next action  $a_d$  according to a softmax decision rule,  $\pi : \mathcal{A} \times \mathcal{S} \times \mathcal{S} \rightarrow [0, 1]$ , which weighs the expected returns of all actions against one another,

$$\pi(a_d \mid s_d, g_{\tilde{z}_d}) := \frac{\exp\{\beta Q^*(s_d, a_d \mid g_{\tilde{z}_d})\}}{\sum_{a \in \mathcal{A}} \exp\{\beta Q^*(s_d, a \mid g_{\tilde{z}_d})\}}. \quad (2)$$

Herein,  $Q^*(s, a \mid g)$  denotes the state-action value (or *Q-value*, Sutton and Barto, 1998) of action  $a$  at state  $s$  under an optimal policy for the subgoal reward function  $R_g$ ,

$$Q^*(s, a \mid g) := \max_{\bar{\pi}} \mathbb{E} \left[ \sum_{n=0}^{\infty} \gamma^n R_g(s_{t=n}) \mid s_{t=0} = s, a_{t=0} = a, \bar{\pi} \right], \quad (3)$$

where the expectation is with respect to the stochastic state-action sequence induced by the fixed policy  $\bar{\pi} : \mathcal{S} \rightarrow \mathcal{A}$ , with initial action  $a$  executed at the starting state  $s$ . The explicit notation  $s_{t=n}$  and  $a_{t=n}$  is used to disambiguate the temporal index of the decision-making process from the demonstration index of the state-action pairs  $\{(s_d, a_d)\}$ .

The softmax policy  $\pi$  models the expert’s (in-)ability to maximize the future expected return in view of the targeted subgoal, while the coefficient  $\beta \in [0, \infty)$  is used to express the expert’s level of confidence in the optimal action. Combined with the subgoal prior distribution  $p_g$  and the partitioning model  $p(\tilde{\mathbf{z}})$ , we obtain the joint distribution of all demonstrated actions  $\mathbf{a}$ , subgoals  $\mathcal{G}$ , and subgoal assignments  $\tilde{\mathbf{z}}$  as

$$p(\mathbf{a}, \tilde{\mathbf{z}}, \mathcal{G} \mid \mathbf{s}) = p(\tilde{\mathbf{z}}) \prod_{k=1}^{\infty} p_g(g_k \mid \mathbf{s}) \prod_{d=1}^D \pi(a_d \mid s_d, g_{\tilde{z}_d}). \quad (4)$$

The structure of this distribution is visualized in form of a Bayesian network in Figure 2a. It is worth emphasizing that  $\pi$ —although referred to as the likelihood model for the state-action *pairs* in the original BNIRL paper—is really just a model for the actions *conditional on the states*. In contrast to what is stated in the original paper, the distribution in Equation (4) therefore takes the form of a *conditional* distribution (i.e., conditional on  $\mathbf{s}$ ), which does not provide any generative model for the state variables.

Posterior inference in BNIRL refers to the (approximate) computation of the conditional distribution  $p(\tilde{\mathbf{z}}, \mathcal{G} \mid \mathcal{D})$ , which allows to identify potential subgoal locations and the corresponding subgoal assignments based on the available demonstration data. For further details, the reader is referred to the original paper (Michini and How, 2012).

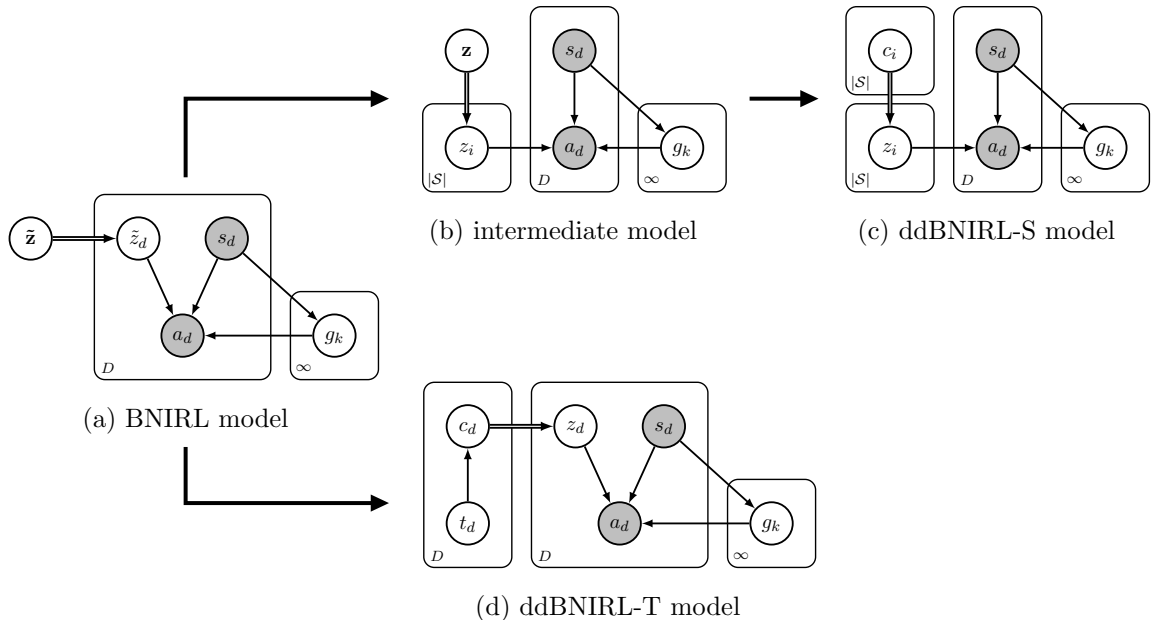


Figure 2: Relationships between all discussed subgoal models, illustrated in the form of Bayesian networks. Shaded nodes represent observed variables; deterministic dependencies are highlighted using double strokes.

## 2.2 Limitations of BNIRL

Subgoal-based inference is a well-motivated approach to IRL and the BNIRL framework has shown promising results in a variety of real-world scenarios. Yet, the model formulation by Michini and How (2012) comes with a number of significant conceptual limitations, which we explain in detail in the following paragraphs.

### LIMITATION 1: SUBGOAL EXCHANGEABILITY AND POSTERIOR PREDICTIVE POLICY

The central limitation of BNIRL is that the framework is restricted to pure subgoal extraction and does *not* inherently provide a reasonable mechanism to generalize the expert behavior based on the inferred subgoals. The reason lies in the particular design of the framework, which, at its heart, treats the subgoal assignments  $\tilde{\mathbf{z}}$  as *exchangeable random variables* (Aldous, 1985). By implication, the induced partitioning model  $p(\tilde{\mathbf{z}})$  is agnostic about the covariate information contained in the data set and the resulting behavioral model is unable to propagate the expert knowledge to new situations.

To illustrate the problem, let us investigate the predictive action distribution that arises from the original BNIRL formulation. For simplicity and without loss of generality, we may assume that we have perfectly inferred all subgoals  $\mathcal{G}$  and corresponding subgoal assignments  $\tilde{\mathbf{z}}$  from the demonstration set  $\mathcal{D}$ . Denoting by  $a^* \in \mathcal{A}$  the predicted action at some



new state  $s^* \in \mathcal{S}$ , the BNIRL model yields

$$\begin{aligned}
 p(a^* | s^*, \mathcal{D}, \tilde{\mathbf{z}}, \mathcal{G}) &= \sum_{\tilde{z}^* \in \mathbb{N}} p(a^*, \tilde{z}^* | s^*, \mathcal{D}, \tilde{\mathbf{z}}, \mathcal{G}) \\
 &= \sum_{\tilde{z}^* \in \mathbb{N}} p(a^* | \tilde{z}^*, s^*, \mathcal{D}, \tilde{\mathbf{z}}, \mathcal{G}) p(\tilde{z}^* | s^*, \mathcal{D}, \tilde{\mathbf{z}}, \mathcal{G}) \\
 &\stackrel{(\star)}{=} \sum_{\tilde{z}^* \in \mathbb{N}} p(a^* | \tilde{z}^*, s^*, g_{\tilde{z}^*}) p(\tilde{z}^* | \tilde{\mathbf{z}}), \tag{5}
 \end{aligned}$$

where  $\tilde{z}^* \in \mathbb{N}$  is the latent subgoal index belonging to  $s^*$ . Note that  $p(a^* | \tilde{z}^*, s^*, g_{\tilde{z}^*})$  can either represent the softmax decision rule  $\pi(a^* | s^*, g_{\tilde{z}^*})$  from Equation (2) or an optimal (deterministic) policy for subgoal  $g_{\tilde{z}^*}$ , depending on whether we aspire to describe the *noisy expert behavior* at  $s^*$  or want to determine an *optimal action* according to the inferred reward model. The last equality in Equation (5), indicated by  $(\star)$ , follows from the conditional independence properties implied by Equation (4), which can be easily verified using d-separation (Koller and Friedman, 2009) on the graphical model in Figure 2a.

As Equation (5) reveals, the predictive model is characterized by the posterior distribution  $p(\tilde{z}^* | s^*, \mathcal{D}, \tilde{\mathbf{z}}, \mathcal{G})$  of the latent subgoal assignment  $\tilde{z}^*$  of state  $s^*$  — the intuition being that, in order to generalize the expert’s plan to a new situation, we need to take into account the gathered information about what would be a likely subgoal targeted by the expert at  $s^*$ . However, in BNIRL, the distribution  $p(\tilde{z}^* | s^*, \mathcal{D}, \tilde{\mathbf{z}}, \mathcal{G})$  is modeled *without consideration of the query state  $s^*$ , or any other observed variable*. By conditional independence (Equation 4), the distribution effectively reduces  $(\star)$  to the CRP prior  $p(\tilde{z}^* | \tilde{\mathbf{z}})$ , which, due to its intrinsic exchangeability property, only considers the subgoal frequencies of the readily inferred assignments  $\tilde{\mathbf{z}}$ . Clearly, a subgoal assignment mechanism based solely on frequency information is of little use when it comes to predicting the expert behavior as it will inevitably ignore the structural information contained in the demonstration set and always return the same subgoal probabilities at all query states, regardless of the agent’s actual situation. By contrast, a reasonable assignment mechanism should inherently take into account the context of the agent’s current state  $s^*$  when deciding about the next action.

While the authors of BNIRL discuss the action selection problem in their paper and propose an assignment strategy for new states based on action marginalization, their approach does not provide a satisfactory solution to the problem because the alleged conditioning on the query state (see Equation 19 in the original paper, Michini and How, 2012) has no effect on the involved subgoal indicator variable, as shown by Equation (5) above. The only way to remedy the problem without modifying the model is to use an external post-processing scheme like the waypoint method, discussed in the next section.

## LIMITATION 2: SPATIAL AND TEMPORAL CONTEXT

The waypoint method, described at full length in a follow-up paper by Michini et al. (2015), is a post-processing routine to convert the subgoals identified through BNIRL into a valid option model (Sutton et al., 1999). The obtained model reconstructs the high-level plan of the demonstrator by sequencing the inferred subgoals in a way that complies with the spatio-temporal relationships of the expert’s decisions as observed during the demonstration phase. To this end, the required initiation and termination sets of the option-policies are




Figure 3: A diagram to illustrate the implications of the exchangeability assumption in BNIRL. Similar to a bag-of-words model (Blei et al., 2003; Yang et al., 2007), the BNIRL partitioning mechanism ignores the spatio-temporal context of the data, which makes it difficult to discriminate demonstration noise from a real change of the agent’s intentions. Note that the diagram illustrates the partitioning process in a simplified way as it only shows the effect of the prior  $p(\tilde{\mathbf{z}})$  but neglects the impact of the likelihood model  $\pi$ . While the latter *does* indeed consider the state context of the actions, it cannot account for spatial or temporal patterns in the data as it processes all state-action pairs separately.

constructed by considering the state distances to the identified subgoals as well as their temporal ordering prescribed by the expert.

When combined with BNIRL, this method allows to synthesize a behavioral model that mimics the observed expert behavior. However, the strategy comes with a number of significant drawbacks:

- (i) Using the waypoint method, the spatio-temporal relationships between the individual demonstrations are explored only in a post-hoc fashion and are largely ignored during the actual inference procedure (the state information enters via the likelihood model  $\pi$  but is not considered by the partitioning model  $p(\tilde{\mathbf{z}})$ , as explained in Limitation 1). This lack of context-awareness makes the inference mechanism overly prone to demonstration noise (see Figure 3 and results in Section 5).
- (ii) Measuring proximities to subgoals in order to determine the right visitation order requires some form of distance metric defined on the state space. If the system states correspond to physical locations, constructing such a metric is usually straightforward. However, in the general case where states encode arbitrary abstract information (see example in Section 5.2), it can become difficult to design that metric by hand. Unfortunately, the BNIRL framework does not provide any solution to this problem.
- (iii) The waypoint method cannot be applied to multiple unaligned trajectories (e.g., obtained from different experts) or in cases where the data set does not carry any temporal information. This situation occurs, for instance, when the expert data is provided as separate state-action pairs with unknown timestamps and not given in form of coherent trajectories (see again example in Section 5.2).

- (iv) Assigning a particular visitation order to the inferred subgoals is meaningful only if the expert eventually reaches those subgoals during the demonstration phase (or if, at least, the subgoals lie “close” to the visited states in terms of the aforementioned distance metric). Finding subgoals with such properties can be guaranteed by constraining the support of the subgoal prior distribution  $p_g$  to states that are near to the expert data (see footnote on page 6) but this reduces the flexibility of the model and potentially disables compact encodings of the task (Figure 4).

### LIMITATION 3: INCONSISTENCY UNDER TIME-INVARIANCE

Reasoning about the intentions of an agent, there are two basic types of behavior one may encounter:

- either the agent follows a static strategy to optimize a fixed objective (as assumed in the standard MDP formalism, Sutton and Barto, 1998), or
- the intentions of the agent change over time.

The latter is clearly the more general case but also poses a more difficult inference problem in that it requires us both, to identify the intentions of the agent and to understand their temporal relationship. The static scenario, in contrast, implies that there exists an optimal policy for the task in form of a simple state-to-action mapping  $\pi : \mathcal{S} \rightarrow \mathcal{A}$  (Puterman, 1994), which from the very beginning imprints a specific structure on the inference problem.

The BNIRL model generally falls into the second category since it freely allocates its subgoals per *decision instant* and not per state, allowing a flexible change of the agent’s objective. Yet, it is important to understand that the model does not actually distinguish between the two described scenarios. As explained in Limitation 2, the temporal aspect of the data is not explicitly modeled by the BNIRL framework, even though the waypoint method subsequently tries to capture the overall chronological order of events. As a consequence, the model is not tailored to either of the two scenarios: on the one hand, it ignores the valuable temporal context that is needed in the time-varying case to reliably discriminate demonstration noise from a real change of the agent’s intention. On the other hand, the model is agnostic about the predefined time-invariant nature of the optimal policy in the static scenario. This lack of structure not only makes the inference problem harder than necessary in both cases; it also allows the model to learn inconsistent data representations in the static case since the same state can be potentially assigned to more than one subgoal, violating the above-mentioned state-to-action rule (Figure 5).

### LIMITATION 4: SUBGOAL LIKELIHOOD MODEL

Apart from the discussed limitations of the BNIRL partitioning model, it turns out there are two problematic issues concerning the softmax likelihood model in Equation (2). On the following pages, we demonstrate that the specific form of the model encodes a number of properties that are indeed contradictory to our intuitive understanding of subgoals. While these properties are less critical for the final prediction of the expert behavior, it turns out they drastically affect the *localization* of subgoals. Since the cause of these effects is somewhat hidden in the model equation, we defer the detailed explanation to Section 3.1.

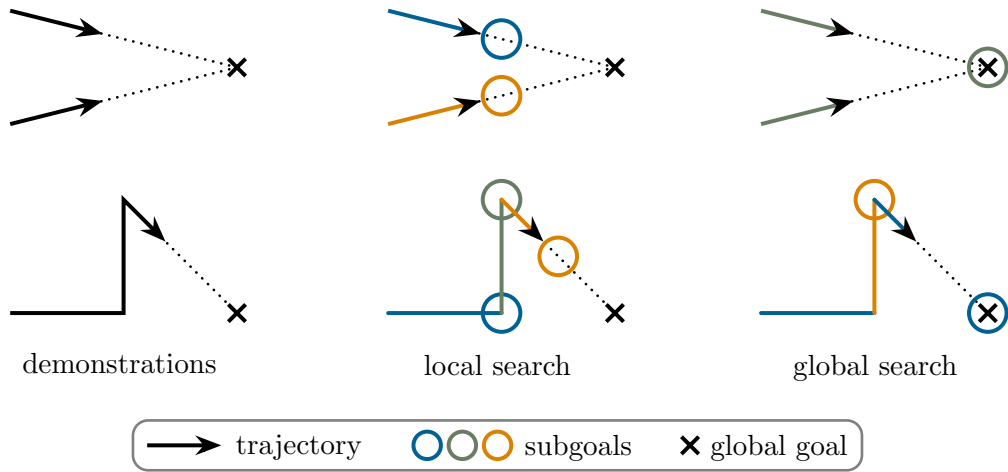


Figure 4: Difference between local (constrained) and global (unconstrained) subgoal search. The top and the bottom row depict two different sets of demonstration data (solid lines), together with potential goal/subgoal locations (crosses/circles) that explain the observed behavior. Color indicates the corresponding subgoal assignment of each trajectory segment. **Top:** two trajectories approaching the same goal. **Bottom:** the agent is heading toward a global goal, gets temporarily distracted, and then follows up on its original plan. **Left:** observed trajectories. **Center:** example partitioning under the assumption that the expert reached all subgoals during the demonstration. **Right:** example partitioning without restriction on the subgoal locations, yielding a more compact encoding of the task.




Figure 5: Schematic comparison of the two basic behavior types, illustrated using two different agent trajectories. Color indicates the temporal progress. (a) Time-varying intentions may cause the agent to perform a different action when revisiting a state (dotted circle). (b) By contrast, time-invariant intentions imply a simple state-to-action policy: the agent has no incentive to perform a different action at an already visited state since — by definition — the underlying objective has remained unchanged. Diverging actions, as observed at the crossing point in the left subfigure, can therefore only be explained as a result of suboptimal behavior.

## LIMITATION 5: STATE-ACTION DEMONSTRATIONS

Lastly, a minor problem of the original BNIRL framework is that the inference algorithm expects the demonstration data to be provided in the form of state-action pairs, which requires full access to the expert’s action record. This assumption is restrictive from a practical point of view as it confines the application of the model to settings with laboratory-like conditions that allow a complete monitoring of the expert. For this reason, it is important to note that an estimate of the expert’s action sequence can be recovered through BNIRL with the help of an additional sampling stage (omitted in the original paper), provided that we know the successor state reached by the expert after each decision. For the marginalized inference scheme described in this paper, we present the corresponding sampling stage in Section 4.4.

### 3. Nonparametric Spatio-Temporal Subgoal Modeling

In this section, we introduce a redesigned inference framework, which, in analogy to BNIRL, we refer to as *distance-dependent Bayesian nonparametric IRL* (ddBNIRL). We derive the model by making a series of modifications to the original BNIRL framework that address the previously described shortcomings on the conceptual level. Rethinking each part of the original framework, we begin with a discussion of the commonly used softmax action selection strategy (Equation 2) in the context of subgoal inference, which finally leads to a redesign of the subgoal likelihood model (Limitation 4). Next, we focus on the subgoal allocation mechanism itself and introduce two closely related model formulations, each targeting one of the basic behavior types described in Figure 5, thereby addressing Limitations 1, 2 and 3. For the time-invariant case, we begin with an intermediate model that introduces a subtle yet important structural modification to the BNIRL framework. In a second step, we generalize that new model to account for the spatial structure of the control problem, which finally allows us to extrapolate the expert behavior to unseen situations. As part of this generalization, we present a new state space metric that arises naturally in the context of subgoal inference (see Limitation 2, second point). Lastly, we tackle the time-varying case and present a variant of the model that explicitly considers the temporal aspect of the subgoal problem. A solution to Limitation 5 is discussed later in Section 4.

In contrast to BNIRL, both presented models can be used likewise for *subgoal extraction* and *action prediction*. Moreover, sticking with the Bayesian methodology, the presented approach provides complete posterior information at all levels.

#### 3.1 The Subgoal Likelihood Model

Like many other approaches found in the (I)RL literature, BNIRL exploits a softmax weighting (Equation 2) to transform the Q-values of an optimal policy into a valid subgoal likelihood model. The softmax action rule has its origin in RL where it is known as the Boltzmann exploration strategy (Cesa-Bianchi et al., 2017; Sutton and Barto, 1998), which is commonly applied to cope with the *exploration-exploitation dilemma* (Ghavamzadeh et al., 2015). In recent years, however, it has also become the de facto standard for describing the (imperfect) decision-making strategy of an observed demonstrator (see, for example, Dimitrakakis and Rothkopf, 2011; Ramachandran and Amir, 2007; Rothkopf and Dimitrakakis, 2011; Choi and Kim, 2012; Neu and Szepesvári, 2007; Babeş-Vroman et al., 2011).

In the following paragraphs, we focus on the implications of this model on the subgoal extraction problem and show that it contradicts our intuitive understanding of what characteristics a reasonable subgoal model should have. In particular, we argue that the subgoal posterior distribution arising from the BNIRL softmax model is of limited use for inferring the latent intention of the agent, due to subgoal artifacts caused by the system dynamics that cannot be reconciled with the evidence provided by the demonstrations. Based on these insights, we propose an alternative transformation scheme that is more consistent with the subgoal principle.

### 3.1.1 SCALE OF THE REWARD FUNCTION

The first implication of the softmax likelihood model concerns the choice of the uncertainty coefficient  $\beta$ . To explain the problem, we consider the thought experiment of an agent located at some state  $s$  targeting a particular subgoal  $g$ . The likelihood  $\pi(a | s, g)$  in Equation (2) quantifies the probability that the agent decides for a specific action  $a$ , based on the corresponding state-action values  $Q(\cdot, s | g)$ . Since those values are linear in the underlying reward function  $R_g$  (Equation 3), the softmax likelihood model implies that the expert’s ability to maximize the long-term reward, reflected by the spread of the probability mass in  $\pi(\cdot | s, g)$ , rises with the magnitude  $C$  of the assumed subgoal reward (more concentrated probability mass signifies a higher confidence in the action choice). In other words, assuming a higher goal reward virtually increases our level of confidence in the expert, even though the difficulty of the underlying task and the optimal policy remain unchanged. Nonetheless, the BNIRL model requires us to readjust the uncertainty coefficient  $\beta$  in order to keep both models consistent. However, as the model provides no reference level for the expert’s uncertainty across different scenarios, the choice of  $\beta$  becomes nontrivial. Yet, the parameter has a significant impact on the granularity of the learned subgoal model as it trades off purposeful goal-oriented behavior against random decisions.

Note that the described effect is not specific to the subgoal reward model in Equation (1) but is really a consequence of the softmax transformation in Equation (2). In fact, the same problem occurs when the model is applied in a regular MDP environment with arbitrary reward function, for example, when the agent is provided an additional constant reward at all states. Clearly, such a constant reward provides no further information about the underlying task and should hence not affect the agent’s belief about the optimal choice of actions (compare discussion on constant reward functions and transformations of rewards, Ng and Russell, 2000; Ng et al., 1999). Based on these two observations, our intuition tells us that we seek for a rationality model that is *invariant to affine transformations of the reward signal*, meaning that any two reward functions  $R : \mathcal{S} \rightarrow \mathbb{R}$  and  $\bar{R} := xR + y$  with  $x \in (0, \infty), y \in \mathbb{R}$ , should give rise to the same intentional representation. As we shall see in Section 3.1.3, this can be achieved by modeling the behavior of an agent based on the *relative advantages* of actions rather than on their absolute expected returns.

### 3.1.2 IMPACT OF THE TRANSITION DYNAMICS

The second implication of the softmax likelihood model is less immediate and inherently tied to the dynamics of the system. To explain the problem, we consider a scenario where we have a precise idea about the potential goals of the expert. For our example, we adopt

the grid world dynamics described in Section 5.1 and consider a simple upward-directed trajectory of state-action pairs, which we aspire to explain using a single (sub-)goal. The complete setting is depicted in Figure 6.

Intuitively, the shown demonstration set should lead to goals that are located in the upper region of the state space and concentrated around the vertical center line. Moreover, as we move away from that center line, we expect to observe a smooth decrease in the subgoal likelihood, while the rate of the decay should reflect our assumed level of confidence in the expert. As it turns out, the induced BNIRL subgoal posterior distribution, shown in the top row for different values of  $\beta$ , contradicts this intuition. In particular, we observe that the model yields unreasonably high posterior values at the upper border states and corners of the state space, which, according to our intuitive understanding of the problem, cannot be justified by the given demonstration set.

To pin down the cause of this effect, we recall from Equation (2) that the likelihood of an action grows with the corresponding Q-value. Hence, we need to ask what causes the Q-values of the demonstrated actions to be large when the subgoal is assumed to be located at one of the upper corner/border states of the space. Using Bellman’s principle, we can express the optimal Q-function for any subgoal  $g$  as

$$\begin{aligned} Q^*(s, a | g) &= R_g(s) + \gamma \mathbb{E}_T [V^*(s' | g) | s, a] \\ &= R_g(s) + \gamma \mathbb{E}_T [\mathbb{E}_{\rho^{\pi_g}} [R_g(s'') | s'] | s, a] \\ &= R_g(s) + \gamma \mathbb{E}_T [C \rho^{\pi_g}(g | s') | s, a], \end{aligned} \tag{6}$$

where  $V^*(s | g) := \max_{a \in \mathcal{A}} Q^*(s, a | g)$ ,  $\pi_g(s) := \arg \max_{a \in \mathcal{A}} Q^*(s, a | g)$  is the optimal policy for subgoal  $g$ , and  $C$  is the subgoal reward from Equation (1). Lastly,  $\rho^{\pi_g}(s' | s) := \sum_{t=0}^{\infty} \gamma^t p_t(s' | s, \pi_g)$  denotes the (improper) discounted state distribution generated by executing policy  $\pi_g$  from the considered initial state  $s$ , where  $p_t(s' | s, \pi_g)$  refers to the probability of reaching state  $s'$  from state  $s$  under policy  $\pi_g$  after exactly  $t$  steps, which is defined implicitly via the transition model  $T$ .

The outer expectation in Equation (6) accounts for the stochastic transition to the successor state  $s'$ , while the inner expectation evaluates the expected cumulative reward over all states  $s''$  that are reachable from  $s'$ . It is important to note that — by the construction of the Q-function — only the first move of the agent to state  $s'$  depends on the choice of action  $a$  whereas all remaining moves (i.e., the argument of the expectation in the last line) are purely determined by the system dynamics and the subgoal policy  $\pi_g$ . Focusing on that inner part, we conclude that, *regardless of the chosen action  $a$* , the Q-values will be large whenever the assumed subgoal induces a high state visitation frequency  $\rho^{\pi_g}$  at its own location  $g$ . The latter is fulfilled if

- (i) the chance of reaching the goal in a small number of steps is high so that the effect of discounting is small and/or
- (ii) the controlled transition dynamics  $T(s' | s, \pi_g(s))$  that are induced by the subgoal lead to a high chance of hitting the goal frequently.

Note that the first condition implies that the model generally prefers subgoals that are close to the demonstration set — a property that cannot be justified in all cases. For example, the recording of the demonstrations could have simply ended before the expert was able to

reach the goal (Figure 5). Yet, if desired, this proximity property should be more naturally attributed to the subgoal prior model  $p_g(g|\mathbf{s})$ .

Moreover, we observe that the second condition depends primarily on the system dynamics  $T$ , which can be more or less strongly influenced by the actions of the agent, depending on the scenario. In fact, in a pathological example,  $T$  could be even independent of the agent’s decisions, meaning that the agent has no control over its state. An example illustrating this extreme case would be a scenario where the agent gets always driven to the same terminal state, regardless of the executed policy. Although it is somewhat pointless speak of “subgoals” in this context, that terminal state would exhibit a high subgoal likelihood according to the softmax model because the corresponding visitation frequency would be inevitably large. A softened variant of this condition can occur at corner/border states (i.e., states in which the agent experiences fewer degrees of freedom and which are hence more difficult to leave than others) and transition states (i.e., states that must be passed in order to get from certain regions of the space to others), which naturally exhibit an increased visitation frequency due to the characteristics of the environment.

In our example in Figure 6, we can observe the symptoms of both described conditions clearly. In particular, for an upward-directed policy as it is implied by the shown demonstration set, the induced state visitation distribution exhibits increased values at exactly the aforementioned border and corner states (due to the reflections occurring to the agent when hitting the state space boundary) as well as close to the trajectory ending (caused by the proximity condition).

### 3.1.3 THE NORMALIZED LIKELIHOOD MODEL

To address these problems, we modify the likelihood model using a rescaling of the involved Q-values. Let  $Q^\wedge(s|g)$  and  $Q^\vee(s|g)$  denote the maximum and minimum Q-values at state  $s$  for subgoal  $g$ , i.e.,  $Q^\wedge(s|g) := \max_{a \in \mathcal{A}} Q^*(s, a|g)$  and  $Q^\vee(s|g) := \min_{a \in \mathcal{A}} Q^*(s, a|g)$ . We then define the normalized state-action value function  $Q^\bullet : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow [0, 1]$  as

$$Q^\bullet(s, a|g) := \begin{cases} \frac{Q^*(s, a|g) - Q^\vee(s|g)}{Q^\wedge(s|g) - Q^\vee(s|g)} & \text{if } Q^\wedge(s|g) \neq Q^\vee(s|g), \\ \epsilon & \text{otherwise,} \end{cases} \quad (7)$$

where  $\epsilon \in (0, 1]$  is an arbitrary constant that is canceled out in Equation (8). In contrast to the Bellman state-action value function  $Q^*$ , which quantifies the expected return of an action, the normalized function  $Q^\bullet$  assesses the return of that action in relation to the returns of all other actions. This concept is similar to that of the advantage function (Baird, 1993) with the important difference that the values returned by  $Q^\bullet$  are normalized to the range  $[0, 1]$  and thus serve as an indicator for the *relative* quality of actions. Accordingly, the values can be interpreted as *relative advantages* (i.e., relative to the maximum possible advantage among all actions). The normalized subgoal likelihood model is then constructed analogously to the BNIRL likelihood model,

$$\pi^\bullet(a_d | s_d, g\bar{z}_d) \propto \exp \{ \beta Q^\bullet(s_d, a_d | g\bar{z}_d) \}. \quad (8)$$

The key property of this model is that it is invariant to affine transformations of the reward function, as summarized by the following proposition.



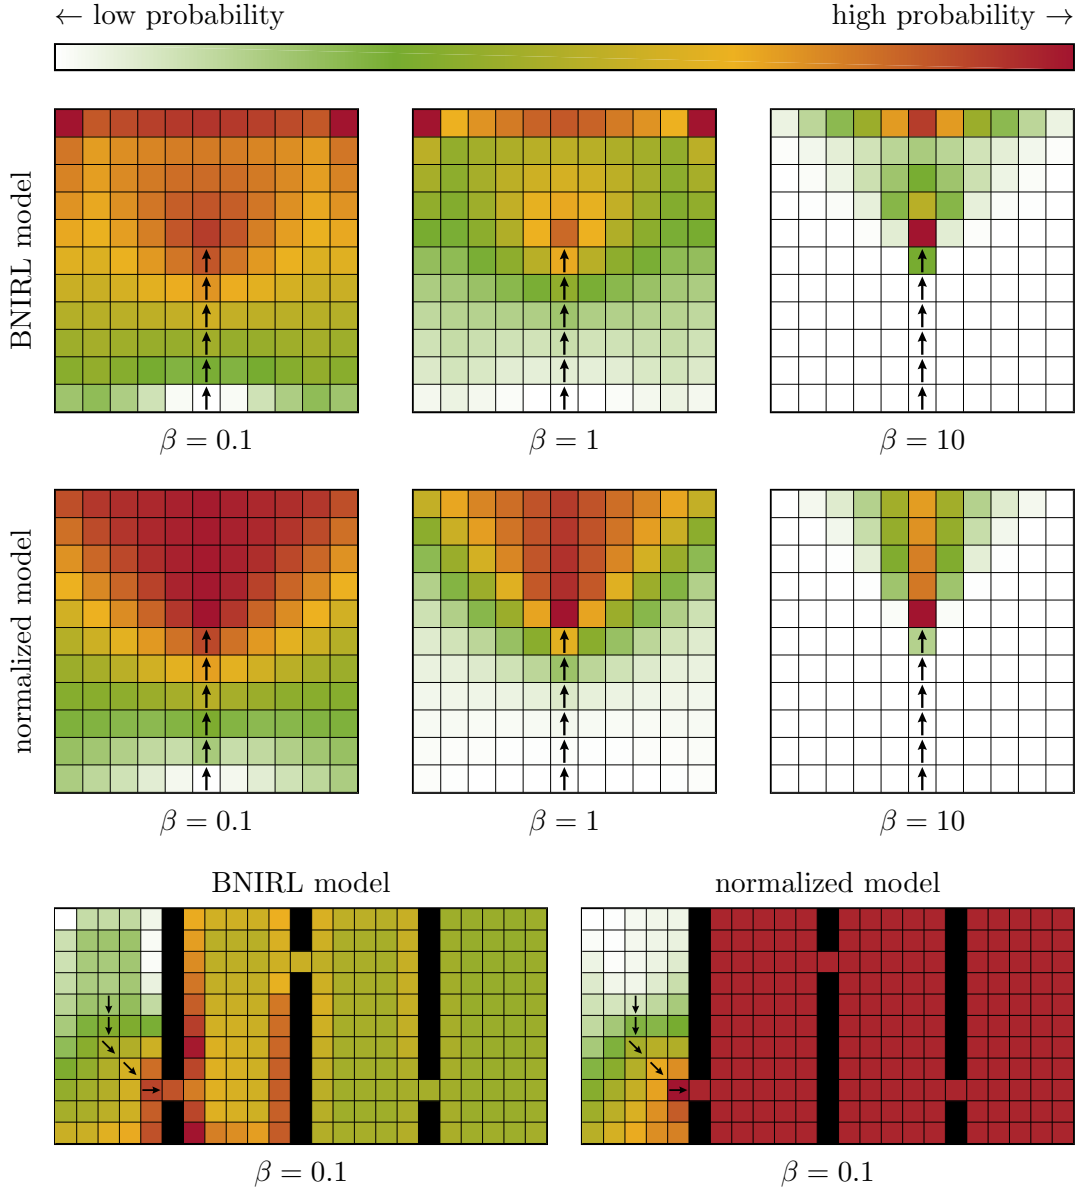


Figure 6: Comparison of the subgoal posterior distributions induced by the original BNIRL likelihood model and by the proposed normalized model, based on the grid world dynamics described in Section 5.1 and a uniform subgoal prior distribution  $p_g$ . The range of the shown color scheme is to be understood per subfigure. Black squares indicate wall states. The BNIRL likelihood model yields unreasonably high subgoal posterior mass at the border states and corners of the state space (due to locally increased state visitation probabilities arising from wall reflections) as well as at trajectory endings (caused by the implicit proximity property of the model) — see Section 3.1.2 for details. Both effects are mitigated by the proposed normalized likelihood model, which describes the action-selection process of the agent using relative advantages of actions instead of absolute returns.

**Proposition 1 (Affine Invariance)** *Consider an MDP with reward function  $R : \mathcal{S} \rightarrow \mathbb{R}$  and let  $Q^*(s, a | R)$  denote the corresponding optimal state-action value function. For the corresponding normalized function  $Q^\bullet$  it holds that  $Q^\bullet(s, a | R) = Q^\bullet(s, a | xR + y) \forall x \in (0, \infty), y \in \mathbb{R}, s \in \mathcal{S}, a \in \mathcal{A}$ . Hence, the subgoal likelihood model in Equation (8) is invariant to affine transformations of  $R$ .*

**Proof** Due to the linear dependence of  $Q^*$  on the reward function  $R$  (Equation 3) it holds that  $Q^*(s, a | xR + y) = xQ^*(s, a | R) + \frac{y}{1-\gamma}$ . Using this relationship in Equation (7), it follows immediately that  $Q^\bullet(s, a | R) = Q^\bullet(s, a | xR + y)$ .  $\blacksquare$

Using the proposed likelihood model offers several advantages. First of all, it enables a more generic choice of the uncertainty coefficient  $\beta$  (Section 3.1.1). This is because the returned  $Q^\bullet$ -values lie in the fixed range  $[0,1]$ , where 0 always indicates the lowest and 1 indicates the highest confidence. For example, setting  $\beta = \log(\beta')$  for some  $\beta' \in (0, \infty)$  always corresponds to the assumption that the expert chooses the optimal action with a probability that is  $\beta'$  times higher than the probability of choosing the least favorable action, irrespective of the underlying system model.

Moreover, as the results in Figure 6 reveal, the induced subgoal posterior distribution is notably closer to our expectation. The reason for this is twofold: first, a likelihood computation based on relative advantages mitigates the influence of the transition dynamics discussed in Section 3.1.2. This is because the described cumulation effect of the state visitation distribution  $\rho^{\pi_g}$  (Equation 6) is present in the returns of all actions and is thus reduced through the proposed normalization. For instance, if the agent in our grid world follows a policy that is all upward directed (as shown in the example), the induced state visitation distribution exhibits increased values at the upper border states of the world, even if we manipulated the *first* action of the agent (as considered in the Bellman Q-function). Accordingly, the original model would indicate an increased subgoal likelihood at those states. The normalized model, by contrast, is less affected as it constructs the likelihood by considering the increased visitation frequencies *relative* to each other.

Second, since the normalization diminishes the effect of the discounting, the subgoal posterior distribution is less concentrated around the trajectory ending and shows significant mass along the extrapolated path of the agent. This property allows us to identify far located states as potential goal locations, which adds more flexibility to the inferred subgoal constellation (compare Figure 4). As an illustrating example, consider the scenario shown in the bottom part of Figure 6. We observe that the normalized model assigns high posterior mass to all states in the right three corridors since any subgoal located in those corridors explains the demonstration set equally well. Here, the difference between the two models is even more pronounced because the transition dynamics have a strong impact on the agent behavior due to the added wall states. For further details, we refer to Section 5.1, where we provide additional insights into the subgoal inference mechanism.

### 3.2 Modeling Time-Invariant Intentions

With our redesigned likelihood model, we now focus on the partitioning structure of the model. Herein, we first consider the case where the intentions of the agent are constant with

respect to time. As explained in Limitation 3, this setting is consistent with the standard MDP formalism in the sense that the optimal policy for the considered task can be described in the form of a state-to-action mapping.

As a first step, to account for this relation, we establish a link between the model partitioning structure and the underlying system state space by replacing the demonstration-based indicators  $\tilde{\mathbf{z}} = \{\tilde{z}_d \in \mathbb{N}\}_{d=1}^D$  with a new set of variables  $\mathbf{z} := \{z_i \in \mathbb{N}\}_{i=1}^{|\mathcal{S}|}$ . Unlike  $\tilde{\mathbf{z}}$ , these new indicators do not operate directly on the data but are instead tied to the elements in  $\mathcal{S}$ . Although they formally represent a new type of variable, we can still imagine that their distribution follows a CRP. This yields an intermediate model of the form

$$p(\mathbf{a}, \mathbf{z}, \mathcal{G} | \mathbf{s}) = p(\mathbf{z}) \prod_{k=1}^{\infty} p_g(g_k | \mathbf{s}) \prod_{d=1}^D \pi^\bullet(a_d | s_d, g_{\mathbf{z}_{s_d}}),$$

whose structure is illustrated in Figure 2b. To see the difference to Equation (4), notice the way the subgoals are indexed in this model.

The intermediate model makes it possible to reason about the policy (or, more suggestively, the underlying state-to-action rule approximated by the expert) at visited parts of the state space. Yet, the model is unable to extrapolate the gathered information to unvisited states, for the reasons explained in Section 2.2. This problem can be solved by replacing the exchangeable prior distribution over subgoal assignments induced by the CRP with a non-exchangeable one, in order to account explicitly for the covariate state information contained in the demonstration set. Based on our insights from Bayesian policy recognition (Šošić et al., 2018b), we use the distance-dependent Chinese restaurant process (ddCRP, Blei and Frazier, 2011) for this purpose, which allows a very intuitive handling of the state context, as explained below. For alternatives, we point to the survey paper by Foti and Williamson (2015).

In contrast to the CRP, which assigns states to partitions, the ddCRP assigns states to other states, based on their pairwise distances. These “to-state” assignments are described by a set of indicators  $\mathbf{c} := \{c_i \in \mathcal{S}\}_{i=1}^{|\mathcal{S}|}$  with prior distribution  $p(\mathbf{c}) = \prod_{i=1}^{|\mathcal{S}|} p(c_i)$ ,

$$p(c_i = j) = \begin{cases} \nu & \text{if } i = j, \\ f(\Delta_{i,j}) & \text{otherwise,} \end{cases} \quad (9)$$

for  $i, j \in \mathcal{S}$ . Herein,  $\nu \in [0, \infty)$  is called the self-link parameter of the process,  $\Delta_{i,j}$  denotes the distance from state  $i$  to state  $j$ , and  $f : [0, \infty) \rightarrow [0, \infty)$  is a monotone decreasing score function. Note that the distances  $\{\Delta_{i,j}\}$  can be obtained via a suitable metric defined on the state space, which may be furthermore used for calibrating the score function  $f$  (see subsequent section). The state partitioning structure itself is then determined by the connected components of the induced ddCRP graph (Figure 7). Our joint distribution, visualized in Figure 2c, thus reads as

$$p(\mathbf{a}, \mathbf{c}, \mathcal{G} | \mathbf{s}) = p(\mathbf{c}) \prod_{k=1}^{\infty} p_g(g_k | \mathbf{s}) \prod_{d=1}^D \pi^\bullet(a_d | s_d, g_{\mathbf{z}(\mathbf{c})|_{s_d}}), \quad (10)$$

where  $\mathbf{z}(\mathbf{c})|_s$  denotes the subgoal label of state  $s$  arising from the considered indicator set  $\mathbf{c}$ . In order to highlight the state dependence of the underlying subgoal mechanism, we refer to this model as ddBNIRL-S.

### 3.2.1 THE CANONICAL STATE METRIC FOR SPATIAL SUBGOAL MODELING

The use of the ddCRP as a prior model for the state partitioning in Equation (10) inevitably requires some notion of distance between any two states of the system, in order to compute the involved function scores  $\{f(\Delta_{i,j})\}$ . When no such distances are provided by the problem setting (see Limitation 2, second point), a suitable (quasi-)metric can be derived from the transition dynamics of the system, which turns out to be the canonical choice for the ddBNIRL-S model. Consider the Markov chain governing the state process  $\{s_{t=n}\}_{n=1}^{\infty}$  of an agent for some specific policy  $\pi$ . For any ordered pair of states  $(i, j)$ , the chain naturally induces a value  $T_{i \rightarrow j}^{\pi}$ , called a *hitting time* (Taylor and Karlin, 1984; Tewari and Bartlett, 2008), which represents the expected number of steps required until the state process, initialized at  $i$ , eventually reaches state  $j$  for the first time,

$$T_{i \rightarrow j}^{\pi} := \mathbb{E}[\min\{n \in \mathbb{N} : s_{t=n} = j\} \mid s_0 = i, \pi].$$

In the context of our subgoal problem, the natural quasi-metric to measure the directed distance between two states  $i$  and  $j$  is thus given by the time it takes to reach the goal state  $j$  from the starting state  $i$  under the corresponding optimal subgoal policy  $\pi_j(s) = \arg \max_{a \in \mathcal{A}} Q^*(s, a \mid j)$ , i.e.,  $\Delta_{i,j} := T_{i \rightarrow j}^{\pi_j}$ . For ddBNIRL-S (as well as for the waypoint method in BNIRL), this choice is particularly appealing since the subgoal policies  $\{\pi_j\}$  are already available within the inference procedure after the state-action values have been computed for the likelihood model (more on this in Section 4.5). The corresponding distances  $\{\Delta_{i,j}\}$  can be obtained efficiently in a single policy evaluation step since  $\Delta_{i,j}$  corresponds to the optimal (negative) expected return at the starting state  $i$  for the special setting where the respective target state  $j$  is made absorbing with zero reward while all other states are assigned a reward of  $-1$ .

### 3.2.2 CHOICE OF THE SCORE FUNCTION

From Equation (9) it is evident that the ddCRP model favors partitioning structures that result from the connection of nearby states. In the context of the subgoal problem, this property translates to the prior assumption that, most likely, each subgoal is approached by the expert from only one specific localized region in the system state space. While this assumption may be reasonable for some tasks, other tasks require that certain target states be approached more than one time, from different regions in the system state space. In such cases, it is beneficial if the model can reuse the same subgoal in various contexts, in order to obtain a more efficient task encoding (Figure 4).

From a mathematical point of view, the prerequisite for learning such encodings is that the score function  $f$  does not shrink to zero at large distance values, so that there remains a non-zero probability of connecting states that are far apart from each other. This can be achieved, for example, by representing  $f$  as a convex combination of a monotone decreasing zero-approaching function  $\bar{f} : [0, \infty) \rightarrow [0, \infty)$  and some constant offset  $\kappa \in (0, 1]$ ,

$$f(\Delta) = (1 - \kappa)\bar{f}(\Delta) + \kappa,$$

where  $\bar{f}$  is chosen, e.g., as a radial basis function (Šošić et al., 2018b). Note that, in order to implement a desired degree of locality in the model, the scale of the decay function  $f$  (or  $\bar{f}$ , respectively) can be further calibrated based on the quantiles of the distribution of the given distances  $\{\Delta_{i,j}\}$ .

### 3.3 Modeling Time-Varying Intentions

For the case of changing expert intentions, we need to keep the flexibility of BNIRL to select a new subgoal at each decision instant, instead of restricting our policy to target a unique subgoal per state (Figure 5). Hence, we retain the basic BNIRL structure in this case and define the subgoal allocation mechanism using a set of data-related indicator variables. However, in contrast to BNIRL, which makes no assumptions about the temporal relationship of the subgoals and thus allows arbitrary changes of the expert’s intentions (Section 2.1), we design our joint distribution in a way that favors smooth action plans in which the expert persistently follows a subgoal over an extended period of time. Again, we can make use of the ddCRP properties to encode the underlying smoothness assumption, but this time using a score function defined on the *temporal distance* between demonstration pairs. For this purpose, we require an additional piece of information, namely the unique timestamp of each demonstration example. Accordingly, we need to assume that our data set is of the form  $\tilde{\mathcal{D}} := \{(s_d, a_d, t_d)\}_{d=1}^D$ , where  $t_d$  denotes the recording time of the  $d$ th demonstration pair  $(s_d, a_d)$ .<sup>2</sup>

The prior distribution over data partitionings can then be written as  $p(\tilde{\mathbf{c}}) = \prod_{d=1}^D p(\tilde{c}_d)$ ,

$$p(\tilde{c}_d = d') \propto \begin{cases} \nu & \text{if } d = d', \\ f(\tilde{\Delta}_{d,d'}) & \text{otherwise,} \end{cases}$$

where the indices  $d, d' \in \{1, \dots, D\}$  range over the size of the demonstration set. Herein,  $\tilde{\Delta}_{d,d'} := |t_d - t_{d'}|$  denotes the temporal distance between the data points  $d$  and  $d'$ . As before, we use the “ $\sim$ ”-notation to distinguish the data-related partitioning variables  $\tilde{\mathbf{c}}$ ,  $\tilde{\mathbf{z}}$  and distances  $\{\tilde{\Delta}_{d,d'}\}$  from their state-space-related counterparts  $\mathbf{c}$ ,  $\mathbf{z}$  and  $\{\Delta_{i,j}\}$  used in ddBNIRL-S. Note, however, that the score function  $f$  is independent of the underlying model type and may be chosen as described in Section 3.2.1, with a scale calibrated to the duration of the demonstrated task.

With that, we obtain our temporal subgoal model as

$$p(\mathbf{a}, \tilde{\mathbf{c}}, \mathcal{G} | \mathbf{s}) = p(\tilde{\mathbf{c}}) \prod_{k=1}^{\infty} p_g(g_k | \mathbf{s}) \prod_{d=1}^D \pi^\bullet(a_d | s_d, g_{\tilde{\mathbf{z}}(\tilde{\mathbf{c}})|_d}), \quad (11)$$

where  $\tilde{\mathbf{z}}(\tilde{\mathbf{c}})|_d$  refers to the subgoal label of the  $d$ th demonstration pair induced by the given assignment  $\tilde{\mathbf{c}}$ . Analogous to our spatial subgoal model, we refer to this model as ddBNIRL-T. The structural differences between all models can be seen from Figure 2.

#### 3.3.1 RELATIONSHIP TO BNIRL

Since the distance-dependent CRP contains the classical CRP as a special case for a specific choice of distance metric and score function (Blei and Frazier, 2011), the ddBNIRL-T model can be considered a strict generalization of the original BNIRL framework (neglecting the likelihood normalization in Section 3.1). In the same way, ddBNIRL-S generalizes

2. Note that the timestamps  $\{t_d\}$  are naturally available if the demonstrations are recorded in trajectory form, where we observe several consecutive state-action pairs. In fact, the temporal information of the data is also required for the waypoint method to work (Limitation 2), even though the authors of BNIRL formally assume to have access to the reduced data set of state-action pairs only.

the intermediate model presented in Section 3.2 (Figure 2). However, although the BNIRL model can be recovered from ddBNIRL, it is important to note that the sampling mechanisms of both frameworks are fundamentally different. Whereas in BNIRL the subgoal assignments are sampled directly, the clustering structure in ddBNIRL is defined *implicitly* via the assignment variables  $\mathbf{c}$  and  $\tilde{\mathbf{c}}$ , respectively. As explained by Blei and Frazier (2011), this has the effect that the Markov chain governing the Gibbs sampler mixes significantly faster because several cluster assignments can be altered in a single step, which effectively realizes a blocked Gibbs sampler (Roberts and Sahu, 1997).

### 3.4 Static versus Dynamic Subgoal Allocation

With the model structures described in Sections 3.2 and 3.3, we have presented two alternative views on the subgoal problem. Naturally, the question arises which of the two approaches is better suited for a particular application scenario. As explained in the previous paragraphs, the main difference between the two models lies in their structure, i.e., in the way subgoals are allocated. While ddBNIRL-S relies on a static assignment mechanism that consistently links the individual states of a system to their corresponding subgoals, ddBNIRL-T allocates its subgoals per demonstration pair. The latter means that different state-action pairs observed at the *same* state can be explained using *different* intentional settings (Figure 5). To answer the above question, we hence need to ask under which conditions an observed decision-making process can be described via a static assignment rule that *uniquely* characterizes each system state, and in which situations we require a more flexible model that allows to take into account additional side information.

From decision-making theory, we know that the optimal solutions for time-invariant MDPs can be formulated as a deterministic *time-invariant* Markov policies (Puterman, 1994), the class of which is fully covered by the static ddBNIRL-S framework.<sup>3</sup> Therefore, assuming that the transition dynamics of our system are constant with respect to time and that the agent acts rationally while having complete knowledge of the environment, there exist only two plausible reasons why we would potentially observe the agent execute a time-variant policy:

- either, the reward model of the agent changes over time,
- or, the observed decision-making process is not Markovian with respect to the assumed state space model (i.e., the agent’s decisions depend on additional context information that is not explicitly captured in our state representation).

Accordingly, if we assume that the Markov property holds (meaning that the chosen state representation is sufficiently rich to capture the decision-making strategy of the agent), the only theoretical justification to prefer a dynamic subgoal model like ddBNIRL-T over a static one such as ddBNIRL-S would be if we assume that the intentions of the agent are truly time-dependent.

Practically speaking, however, there can be several reasons why a given state representation might not fulfill the Markov requirement. One obvious explanation would be that the actual state space of the demonstrator is not perfectly known. This situation occurs, for

---

3. While we omit a rigorous proof here, this can be seen intuitively by noticing that any state-to-action rule that is optimal for a given MDP reward function can be synthesized via ddBNIRL-S by assuming an individual subgoal for each state in the extreme case.

example, if not all state context available to the agent is observable by the modeler. Another potential situation is when the strategy of the agent depends on information that is independent of the system dynamics and hence deliberately excluded from the state variable (i.e., parameters that are unaffected by the actions of the agent, such as the preselection of a specific high-level strategy). A generic framework for such settings is described by Daniel et al. (2016b), where the agent learns multiple sub-policies that are triggered depending on context information that is treated separately from the state.

To an external observer who is unaware of that context information, the resulting policy of the agent would potentially appear time-dependent, in which case the only chance to disentangle the individual sub-policies would be to resort to a dynamic subgoal encoding, such as provided by ddBNIRL-T. However, if the context is known (like the temporal information in Section 3.3 as a particular example), both approaches can be used equivalently and will only differ in the resulting state representation. More specifically, we can either fall back on the static ddBNIRL-S model by augmenting the state variable with the context information accordingly, or we can resort to the dynamic subgoal allocation scheme of ddBNIRL-T, using a distance metric that accounts for the context. Conversely, when considered in a purely time-invariant setting (where the context is described by some other known quantity), ddBNIRL-S and ddBNIRL-T can be regarded as two sides of the same coin, i.e., both can be used to describe the time-invariant policy of an observed demonstrator but they differ in the way the side information is represented.

## 4. Prediction and Inference

Having introduced the ddBNIRL framework, we now explain how it can be used to generalize a given expert behavior. To this end, we first focus on the task of action prediction at a given query state, and then explain in a second step how to extract the necessary information from the demonstration data. Along the way, we also give insights into the implicit intentional model learned through the framework.

**Note:** In order to keep the level of redundancy at a minimum, the following considerations are based on the ddBNIRL-S model. The results for ddBNIRL-T follow straightforwardly; the only change in the equations is the way the subgoals are referenced. To obtain the corresponding expressions, we simply replace the assignment variables  $\mathbf{c}$  with  $\tilde{\mathbf{c}}$  and change the cluster definition in Equation (16) to  $\mathcal{C}_k := \{d \in \{1, \dots, D\} : \tilde{\mathbf{z}}(\tilde{\mathbf{c}})|_d = k\}$ . Accordingly, all occurrences of  $\mathbf{z}(\mathbf{c})|_{s^*}$  change to  $\tilde{\mathbf{z}}(\tilde{\mathbf{c}})|_{d^*}$ ,  $\mathbf{z}(\mathbf{c})|_{s_d}$  becomes  $\tilde{\mathbf{z}}(\tilde{\mathbf{c}})|_d$ , and  $s_d \in \mathcal{C}_k$  is replaced with  $d \in \mathcal{C}_k$ .

### 4.1 Action Prediction

Similar to the work by Abbeel and Ng (2004), we consider the task of predicting an action  $a^* \in \mathcal{A}$  at some query state  $s^* \in \mathcal{S}$  that is optimal with respect to the expert’s *unknown* reward model. However, in contrast to most existing IRL methods, our approach is not based on point estimates of the expert’s reward function but takes into account the entire hypothesis space of reward models. This allows us to obtain the full posterior predictive policy from the expert data. Mathematically, the task is formulated as computing the predictive action distribution  $p(a^* | s^*, \mathcal{D})$ , which captures the full information about the expert

behavior contained in the demonstration set  $\mathcal{D}$ . We start by expanding that distribution with the help of the latent state assignments  $\mathbf{c}$ ,

$$p(a^* | s^*, \mathcal{D}) = \sum_{\mathbf{c} \in \mathcal{S}^{|\mathcal{S}|}} p(a^* | s^*, \mathcal{D}, \mathbf{c}) p(\mathbf{c} | \mathcal{D}).$$

The conditional distribution  $p(a^* | s^*, \mathcal{D}, \mathbf{c})$  can be expressed in terms of the posterior distribution of the subgoal targeted at the query state  $s^*$ ,

$$p(a^* | s^*, \mathcal{D}) = \sum_{\mathbf{c} \in \mathcal{S}^{|\mathcal{S}|}} p(\mathbf{c} | \mathcal{D}) \sum_{i \in \mathcal{S}} p(a^* | s^*, \mathbf{c}, g_{\mathbf{z}(\mathbf{c})|_{s^*}} = i) p(g_{\mathbf{z}(\mathbf{c})|_{s^*}} = i | \mathcal{D}, \mathbf{c}),$$

where we used the fact that the prediction  $a^*$  is conditionally independent of the demonstration set  $\mathcal{D}$  given the state partitioning structure and the corresponding subgoal assigned to  $s^*$  (that is, given  $\mathbf{c}$  and  $g_{\mathbf{z}(\mathbf{c})|_{s^*}}$ ). From the joint distribution in Equation (10), it follows that

$$p(g_k | \mathcal{D}, \mathbf{c}) = \frac{1}{Z_k(\mathcal{D}, \mathbf{c})} p_g(g_k | \mathbf{s}) \prod_{d: \mathbf{z}(\mathbf{c})|_{s_d} = k} \pi(a_d | s_d, g_k), \quad (12)$$

where  $Z_k(\mathcal{D}, \mathbf{c})$  is the corresponding normalizing constant,

$$Z_k(\mathcal{D}, \mathbf{c}) := \sum_{i \in \text{supp}(p_g)} p_g(g_k = i | \mathbf{s}) \prod_{d: \mathbf{z}(\mathbf{c})|_{s_d} = k} \pi(a_d | s_d, g_k = i). \quad (13)$$

Using this relationship, we get

$$\begin{aligned} p(a^* | s^*, \mathcal{D}) &= \sum_{\mathbf{c} \in \mathcal{S}^{|\mathcal{S}|}} \frac{1}{Z_k(\mathcal{D}, \mathbf{c})} p(\mathbf{c} | \mathcal{D}) \sum_{i \in \text{supp}(p_g)} p_g(g_{\mathbf{z}(\mathbf{c})|_{s^*}} = i | \mathbf{s}) \dots \\ &\dots \times \prod_{d: \mathbf{z}(\mathbf{c})|_{s_d} = \mathbf{z}(\mathbf{c})|_{s^*}} \pi(a_d | s_d, g_{\mathbf{z}(\mathbf{c})|_{s^*}} = i) p(a^* | s^*, \mathbf{c}, g_{\mathbf{z}(\mathbf{c})|_{s^*}} = i). \end{aligned}$$

In contrast to the summation over subgoal locations  $i$ , whose computational complexity is determined by the support of the subgoal prior distribution  $p_g$  and which grows at most linearly with the size of  $\mathcal{S}$ , the marginalization with respect to the indicator variables  $\mathbf{c}$  involves the summation of  $|\mathcal{S}|^{|\mathcal{S}|}$  terms and becomes quickly intractable even for small state spaces. Therefore, we approximate this operation via Monte Carlo integration, which yields

$$p(a^* | s^*, \mathcal{D}) \approx \frac{1}{N} \sum_{n=1}^N \sum_{i \in \text{supp}(p_g)} p(g_{\mathbf{z}(\mathbf{c}^{(n)})|_{s^*}} = i | \mathcal{D}, \mathbf{c}^{(n)}) p(a^* | s^*, \mathbf{c}^{(n)}, g_{\mathbf{z}(\mathbf{c}^{(n)})|_{s^*}} = i),$$

where  $\mathbf{c}^{(n)} \sim p(\mathbf{c} | \mathcal{D})$ . The final prediction step can then be performed, for example, via the maximum a posteriori (MAP) policy estimate,

$$\hat{\pi}(s^*) := \arg \max_{a^* \in \mathcal{A}} p(a^* | s^*, \mathcal{D}). \quad (14)$$

The inference task, hence, reduces to the computation of the posterior samples  $\{\mathbf{c}^{(n)}\}$ , which is described in the next section.



## 4.2 Partition Inference

Based on the joint model in Equation (10), we obtain the posterior distribution  $p(\mathbf{c} | \mathcal{D})$  in factorized form as

$$\begin{aligned} p(\mathbf{c} | \mathcal{D}) &= p(\mathbf{c}) \prod_{k=1}^{\infty} \sum_{g_k \in \text{supp}(p_g)} p_g(g_k | \mathbf{s}) \prod_{d=1}^D \pi(a_d | s_d, g_{\mathbf{z}(\mathbf{c})|s_d}) \\ &= p(\mathbf{c}) \prod_{k=1}^{|\mathbf{z}(\mathbf{c})|} \sum_{g_k \in \text{supp}(p_g)} p_g(g_k | \mathbf{s}) \prod_{d: s_d \in \mathcal{C}_k} \pi(a_d | s_d, g_k), \end{aligned} \quad (15)$$

where  $\mathcal{C}_k$  denotes the  $k$ th state cluster induced by the assignment  $\mathbf{c}$ ,

$$\mathcal{C}_k := \{s \in \mathcal{S} : \mathbf{z}(\mathbf{c})|_s = k\}, \quad (16)$$

and  $|\mathbf{z}(\mathbf{c})|$  is the total number of clusters defined by  $\mathbf{c}$ . As explained by Blei and Frazier (2011), the indicator samples  $\{\mathbf{c}^{\{n\}}\}$  can be efficiently generated using a fast-mixing Gibbs chain. Starting from a given ddCRP graph defined by the subset of indicators  $\mathbf{c}_{\setminus i} := \{c_j\} \setminus c_i$ , the insertion of an additional edge  $c_i$  will result in one of three possible outcomes, as illustrated in Figure 7: in the case of adding a self-loop ( $c_i = i$ ), the underlying partitioning structure stays unaffected. Setting  $c_i \neq i$  either leaves the structure unchanged (if the target state is already in the same cluster as state  $i$ ) or creates a new link between two clusters. In the latter case, the involved clusters are merged, which corresponds to a merging of the associated sums in Equation (15). According to these three cases, the conditional distribution for the Gibbs procedure is obtained as

$$p(c_i = j | \mathbf{c}_{\setminus i}, \mathcal{D}) \propto \begin{cases} \nu & \text{if } i = j, \\ f(d_{i,j}) & \text{if no clusters are merged,} \\ f(d_{i,j}) \frac{\mathcal{L}(\mathcal{C}_{z_i} \cup \mathcal{C}_{z_j})}{\mathcal{L}(\mathcal{C}_{z_i}) \cdot \mathcal{L}(\mathcal{C}_{z_j})} & \text{if clusters } \mathcal{C}_{z_i} \text{ and } \mathcal{C}_{z_j} \text{ are merged.} \end{cases} \quad (17)$$

Herein,  $\mathcal{L}(\mathcal{C})$  denotes the marginal action likelihood of all demonstrations accumulated in cluster  $\mathcal{C}$ ,

$$\mathcal{L}(\mathcal{C}) = \sum_{g \in \text{supp}(p_g)} p_g(g | \mathbf{s}) \prod_{d: s_d \in \mathcal{C}} \pi(a_d | s_d, g), \quad (18)$$

which further represents the normalizing constant for the posterior distribution of the cluster subgoal (Equation 13). Accordingly, the fraction in Equation (17) can be interpreted as the likelihood ratio of the partitioning defined by  $\mathbf{c}_{\setminus i}$  and the merged structure after inserting the new edge  $c_i$ .

## 4.3 Subgoal Inference

It is important to note that the inference method described in Sections 4.1 and 4.2 is based on a collapsed sampling scheme where all subgoals of our model are marginalized out. In fact, the ddBNIRL framework differs from BNIRL and other IRL methods in that the reward model of the expert is never made explicit for predicting new actions. Nonetheless, if desired (e.g., for the purpose of analyzing the expert’s intentions), an estimate of the subgoal locations can be obtained in a post-hoc fashion from the subgoal posterior distribution in Equation (12) for any given assignment  $\mathbf{c}$ . Examples are provided in Figure 8.




Figure 7: Insertion of an edge (dashed arrow) to the ddCRP graph. Colors indicate the cluster memberships of the nodes, which are defined implicitly via the connected components of the graph. (a) Adding a self-loop or (b) inserting an edge between two already connected nodes does not alter the clustering structure. (c) Adding an edge between two unconnected components merges the associated clusters.

#### 4.4 Action Inference

As mentioned in Section 2.2, the original BNIRL algorithm requires complete knowledge of the expert’s action record  $\mathbf{a}$ , which limits the range of potential application scenarios. For this reason, we generalize our inference scheme to the case where we have access to state information only, provided in the form of an alternative data set  $\bar{\mathcal{D}} := \{(s_d, \bar{s}_d)\}_{d=1}^D$ , where  $\bar{s}_d$  refers to the state visited by the expert immediately after  $s_d$ . In this setting, inference can be performed by extending the Gibbs procedure with an additional collapsed sampling stage,

$$p(a_d | \mathbf{a}_{\setminus d}, \bar{\mathcal{D}}, \mathbf{c}) \propto T(\bar{s}_d | s_d, a_d) \sum_{i \in \text{supp}(p_g)} p_g(g_{\mathbf{z}(\mathbf{c})|s_d} = i) \prod_{d': \mathbf{z}(\mathbf{c})|s_{d'} = \mathbf{z}(\mathbf{c})|s_d} \pi(a_{d'} | s_{d'}, g_{\mathbf{z}(\mathbf{c})|s_d} = i), \quad (19)$$

which, for a fixed assignment  $\mathbf{c}$ , recovers an estimate of the latent action set  $\mathbf{a}$  from the observed state transitions. Note that knowledge of the transition model  $T$  is required for this step as it provides the necessary link between the expert’s actions and the observed successor states. The same extension is possible for the ddBNIRL-T model, provided that the transition timestamps  $\{t_d\}$  are known (Section 3.3).

#### 4.5 Computational Complexity

As a last point in this section, we would like to discuss the computational complexity of our approach. For this purpose, here a quick reminder on the used notation: we write  $|\mathcal{S}|$  and  $|\mathcal{A}|$  for the cardinalities of the state and action space, respectively, and use the letter  $D$  for the size of the demonstration set. Further, we write  $\mathcal{C}_k$  to refer to the  $k$ th state cluster (ddBNIRL-S) or data cluster (ddBNIRL-T). In the subsequent paragraphs, we additionally use the notation  $N_D(\mathcal{C}_k)$  to access the number of demonstration data points associated with

cluster  $\mathcal{C}_k$ ,  $K$  to indicate the number of clusters in the current iteration,  $N_g := |\text{supp}(p_g)|$  as a shorthand for the size of the support of the subgoal prior distribution, and  $N_c$  for the number of indicator variables, i.e.,  $N_c := |\mathcal{S}|$  for ddBNIRL-S and  $N_c := D$  for ddBNIRL-T.

*Initialization Phase:* Common to all discussed models (including BNIRL) is that they depend on a preceding planning phase, where we compute, potentially in parallel, the state-action value functions (Equation 3) for all  $N_g$  considered subgoals, which allows us to construct the subgoal likelihood model (Equation 2 or 8). The overall computational complexity of this procedure is of order  $\mathcal{O}(N_g \mathbb{C}_{\text{MDP}}(|\mathcal{S}|, |\mathcal{A}|))$ , where  $\mathbb{C}_{\text{MDP}}(x, y)$  denotes the complexity of the used planning routine to (approximately) solve an MDP of size  $x$  with a total number of  $y$  actions. Using a value iteration algorithm, for instance, this can be achieved in  $\mathcal{O}(\mathbb{C}_{\text{MDP}}(|\mathcal{S}|, |\mathcal{A}|)) = \mathcal{O}(|\mathcal{S}|^2 |\mathcal{A}|)$  steps (Littman et al., 1995). If we assume that the expert reaches all subgoals during the demonstration phase (Michini and How, 2012), we can restrict the support of the subgoal prior to the visited states, so that  $N_g$  is upper-bounded by  $\min(|\mathcal{S}|, D)$ . Note that there exist approximation techniques that make the computation tractable in large/continuous state spaces (see discussion in Section 6).

Before we start the sampling procedure, we compute all single-cluster likelihoods  $\{\mathcal{L}(\mathcal{C}_k)\}$  and pairwise likelihoods  $\{\mathcal{L}(\mathcal{C}_k \cup \mathcal{C}_{k'})\}$  according to Equation (18), based on some (random) initial cluster structure. The likelihood computation for the  $k$ th cluster  $\mathcal{C}_k$  involves a product over  $N_D(\mathcal{C}_k)$  data points, which needs to be calculated for each of the  $N_g$  subgoals before taking their weighted average. This step has to be executed (potentially in parallel) for all clusters. However, because each demonstration is associated with exactly one cluster (either directly as in ddBNIRL-T or via the corresponding state variable as in ddBNIRL-S) and hence  $\sum_k N_D(\mathcal{C}_k) = D$ , the total complexity for computing all single-cluster likelihoods is of order  $\mathcal{O}(N_g D)$ , irrespective of the actual cluster structure. A similar line of reasoning applies to the computation of the pairwise likelihoods, yielding the same complexity order. Yet, for the latter we need to consider all possible cluster combinations. Assuming an initial number of  $K$  clusters, there are in total  $K(K - 1)/2$  pairwise likelihoods to be computed. Hence, the overall complexity of the initialization phase can be summarized as  $\mathcal{O}(N_g D K^2)$ .

*Partition Inference:* For the partition inference, the bulk of the computation lies in the repeated construction of the likelihood term in Equation (17), which needs to be updated whenever the cluster structure changes. To analyze the complexity, we consider the sampling step of an individual assignment variable  $c_i$  (or likewise  $\tilde{c}_i$ ). In the worst case, removing the edge that belongs to  $c_i$  from the ddCRP graph divides the associated cluster into two parts (Figure 7), so that two new single-cluster likelihoods need to be computed. With the upper bound  $D$  on the number of data points associated with the cluster before the division, this operation is of worst-case complexity  $\mathcal{O}(N_g D)$  (see initialization phase). Irrespective of whether a division occurs, we then need to compute all pairwise cluster likelihoods with the (new) cluster connected via  $c_i$ . For a total of  $K - 1$  possible choices, this is done in  $\mathcal{O}(N_g D K)$  operations (see initialization phase). After assigning the indicator, we move on to the next variable where the process repeats. If we assume, for simplicity, that the number of clusters stays constant during a full Gibbs cycle, the total complexity of updating all cluster assignments is hence of order  $\mathcal{O}(N_g D K N_c)$ . A (pessimistic) upper bound for the general case can be obtained by assuming that each data point defines its own cluster, in

which case the complexity increases to  $\mathcal{O}(N_g D^2 N_c)$ . Note that, in order to identify the new cluster structure after changing an assignment, we additionally need to track the connected components of the underlying ddCRP graph. As explained by Kapron et al. (2013), this can be done in polylogarithmic worst-case time.

*Action Sampling:* In order to compute the conditional probability distribution of a particular action  $a_d$ , we need to evaluate a product involving all actions that belong to the same cluster as action  $a_d$  (Equation 19). First, we can compute the product over all actions except  $a_d$  itself, where the number of involved terms is again upper-bounded by  $D$ . Appending the term that belongs to  $a_d$  for all possible action choices requires another  $|\mathcal{A}|$  operations. These two steps need to be repeated for all possible subgoals, yielding an upper bound on the complexity of order  $\mathcal{O}(N_g(D + |\mathcal{A}|))$ . For a full Gibbs cycle, which involves sampling all  $D$  action variables, the overall (worst-case) complexity is hence of order  $\mathcal{O}(N_g(D + |\mathcal{A}|)D)$ .

## 5. Experimental Results

In this section, we present experimental results for our framework. The evaluation is separated into four parts:

- (i) a proof of concept and conceptual comparison to BNIRL (Section 5.1),
- (ii) a performance comparison with related algorithms (Section 5.2),
- (iii) a real data experiment conducted on a KUKA robot (Section 5.3) and
- (iv) an active learning task (Section 5.4).

### 5.1 Proof of Concept

To illustrate the conceptual differences to BNIRL and provide additional insights into the latent intentional model learned through our framework, we begin with the motivating data set from Figure 1a, which had been originally presented by Michini and How (2012). The considered system environment, defined by  $|\mathcal{S}| = 20 \times 20 = 400$  grid positions, is again shown in the top left corner of Figure 8. Nine of those positions correspond to inaccessible wall states, marked by the horizontal black bar. At the valid states, the expert can choose from an action set comprising a total of eight actions, each initiating a noisy state transition toward one of the (inter-)cardinal directions. The observed state-action pairs are depicted in the form of arrows, whose colors indicate the MAP partitioning learned through BNIRL. The remaining subfigures show the results of the ddBNIRL framework, which were obtained from a posterior sample returned by the respective algorithm (ddBNIRL-S/T) at a low temperature in a simulated annealing schedule (Kirkpatrick et al., 1983).

Comparing the obtained results, we observe the following main differences to the original approach:

- (i) Unlike BNIRL, the proposed framework allows to choose between a spatial and a temporal encoding of the observed task, providing the possibility to account explicitly for the type of demonstrated behavior (static/dynamic). As explained in Section 3.3.1, the context-unaware (yet in principle dynamic) vanilla BNIRL inference scheme is still included as a special case.

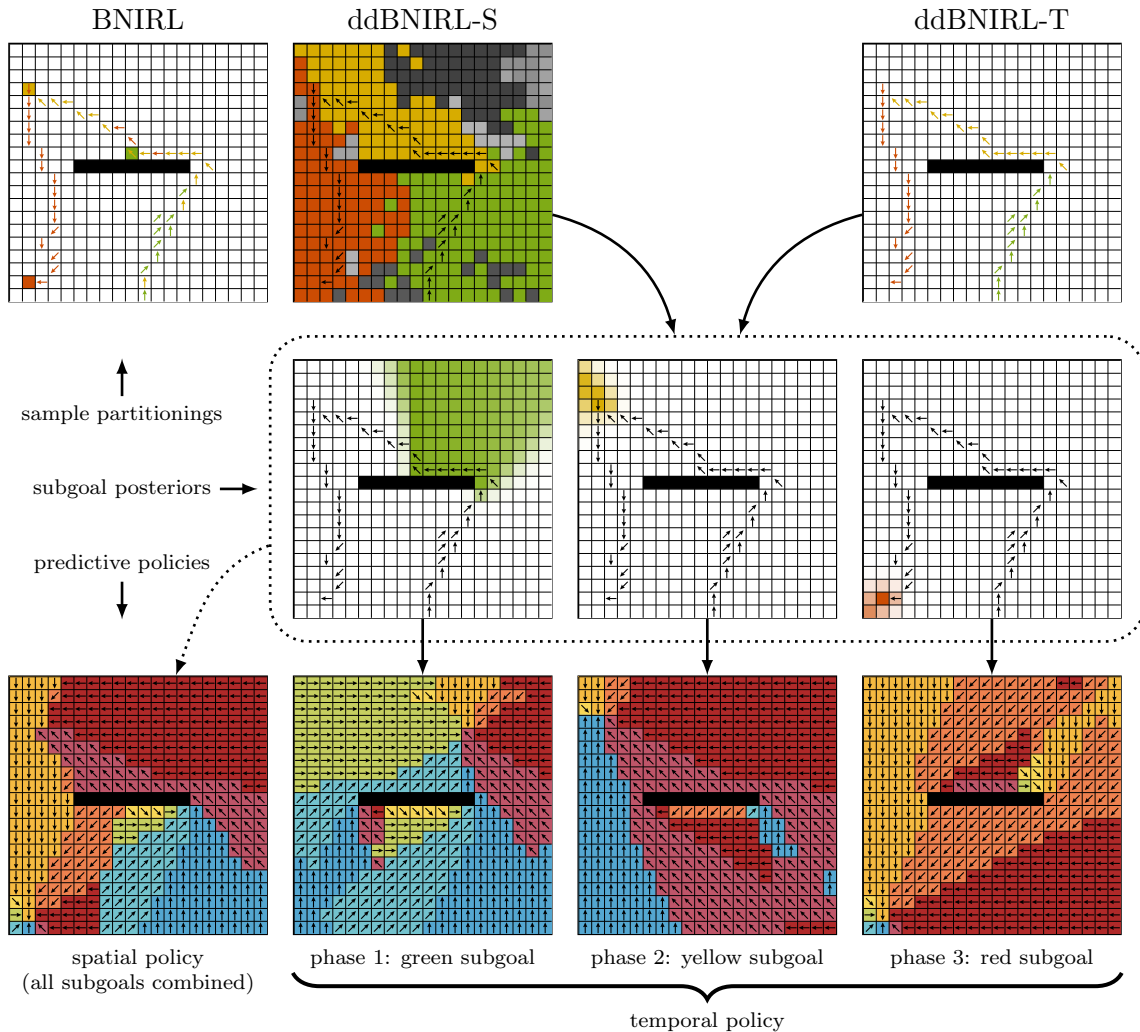


Figure 8: Results on the BNIRL data set (Michini and How, 2012). **Top row:** demonstration data and sample partitionings generated by the different inference algorithms. **Center row:** subgoal posterior distributions associated with the partitions found by ddBNIRL-S and ddBNIRL-T. For a clearer overview, the corresponding BNIRL distributions are omitted (see Figure 6 for a comparison). **Bottom row:** time-invariant ddBNIRL-S policy model synthesized from all three detected subgoals (left) and temporal phases identified by ddBNIRL-T (right). The background colors have no particular meaning and were added only to highlight the structures of the policies. Because of its missing generalization mechanism, BNIRL does not itself provide a reasonable predictive policy model (Limitation 1).

- (ii) Exploiting the spatial/temporal context of the data, the ddbNIRL solution is inherently robust to demonstration noise, giving rise to notably smoother partitioning structures (top row). This effect is particularly pronounced in the case of real data, as we shall see later in Section 5.3.2.
- (iii) For each state partition or trajectory segment, we obtain an implicit representation of the associated subgoal in the form of a posterior distribution, without the need of assigning point estimates (center row). It is striking that the posterior distribution corresponding to the green state partition has a comparably large spread on the upper side of the wall. This can be explained intuitively by the fact that any subgoal located in this high posterior region could have potentially caused the green state sequence, which circumvents the wall from the right. At the same time, the green area of high posterior values exhibits a sharp boundary on the left side since a subgoal located in the upper left region of the state space would have more likely resulted in a trajectory approaching from the left.
- (iv) In contrast to BNIRL, which has no built-in generalization mechanism (Limitation 1), our method returns a predictive policy model comprising the full posterior action information at all states. Note that we only show the resulting MAP policy estimates here (bottom row), computed according to Equation (14). Additional results concerning the posterior uncertainty are provided in Sections 5.3 and 5.4.

The example illustrates how the synthesis of the predictive policy differs between ddbNIRL-S (bottom left) and ddbNIRL-T (bottom row, rightmost three subfigures). While ddbNIRL-T uses a set of (conditionally) independent policy models to describe the different identified behavioral phases, ddbNIRL-S maps the entire subgoal schedule onto a single time-invariant policy representation. Looking closer at the learned models, we recognize that the ddbNIRL-S solution in fact realizes a spatial combination of the three temporal ddbNIRL-T components, where each component is activated in the corresponding cluster region of the state space. This gives us two alternative interpretations of the same behavior.

## 5.2 Random MDP Scenario

Our next experiment is designed to provide insights into the generalization abilities of the framework. For this purpose, we consider a class of randomly generated MDPs similar to the Garnet problems (Bhatnagar et al., 2009). The transition dynamics  $\{T(\cdot | s, a)\}$  are sampled independently from a symmetric Dirichlet distribution with a concentration parameter of 0.01, where we choose  $|\mathcal{S}| = 100$  and  $|\mathcal{A}| = 10$ . For each repetition of the experiment,  $N_R$  states are selected uniformly at random and assigned rewards that are, in turn, sampled uniformly from the interval  $[0, 1]$ . All other states contain zero reward. Next, we compute an optimal deterministic MDP policy  $\pi^*$  with respect to a discount factor of  $\gamma = 0.9$  and generate a number of expert trajectories of length 10. Herein, we let the expert select the optimal action with probability 0.9 and a random, suboptimal action with probability 0.1. The obtained state sequences are passed to the algorithms and we compute the normalized value loss of the reconstructed policies according to

$$L(\pi^*, \hat{\pi}) := \frac{\|\mathbf{V}^* - \mathbf{V}^{\hat{\pi}}\|_2}{\|\mathbf{V}^*\|_2}, \quad (20)$$

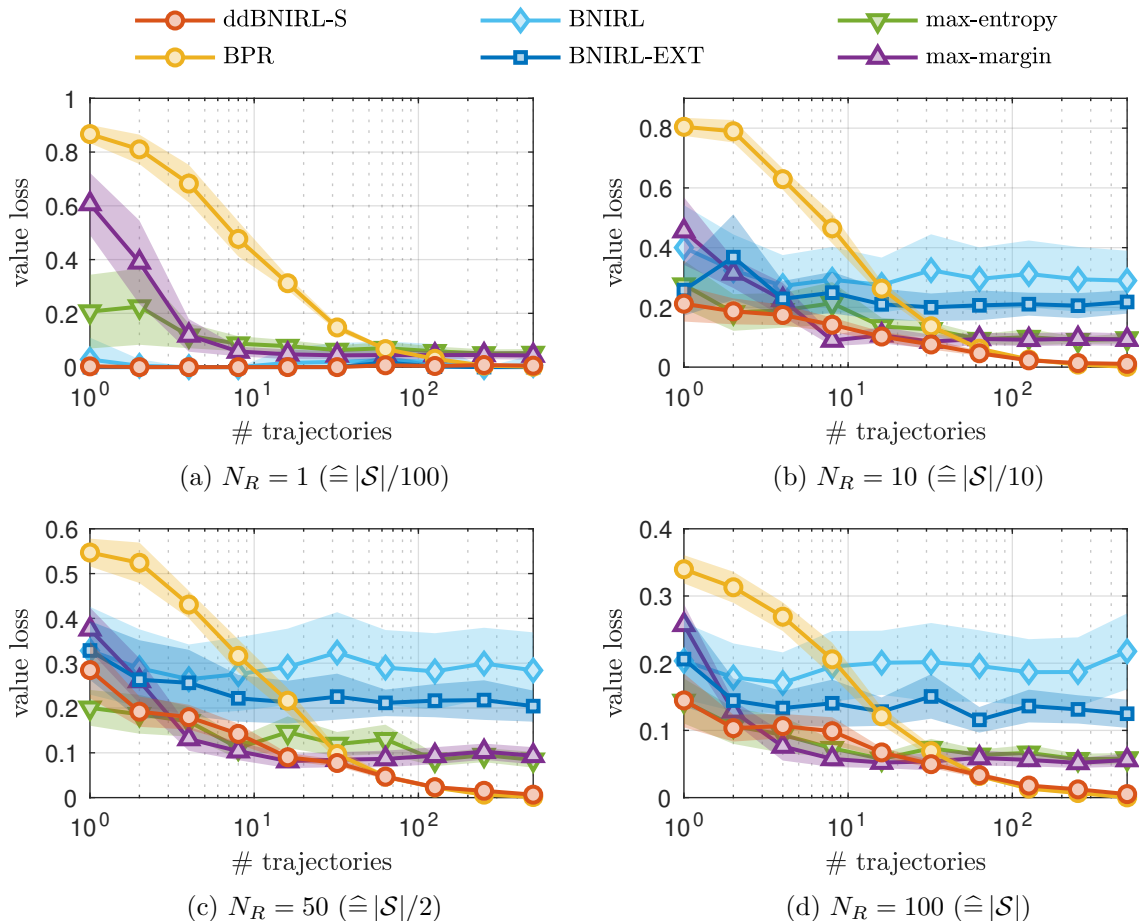


Figure 9: Comparison of all inference methods in the random MDP scenario for different reward densities. Shown are the empirical mean values and standard deviations of the resulting value losses, obtained from 100 Monte Carlo runs. The graphs show a clear difference between BNIRL, BNIRL-EXT and ddBNIRL-S, which illustrates the importance of considering the spatial context for subgoal extraction.

where  $\mathbf{V}^*$  and  $\mathbf{V}^{\hat{\pi}}$  represent, respectively, the vectorized value functions of the optimal policy  $\pi^*$  and the reconstruction  $\hat{\pi}$ .

Since the considered system belongs to the class of time-invariant MDPs, ddBNIRL-S lends itself as the natural choice to model the expert behavior. As baseline methods, we adopt our subintentional Bayesian policy recognition framework (BPR, Šošić et al., 2018b), as well as maximum-margin IRL (Abbeel and Ng, 2004), maximum-entropy IRL (Ziebart et al., 2008), and vanilla BNIRL. Due to the missing generalization abilities of BNIRL (Limitation 1) and because the waypoint method (Section 2.2) does not straightforwardly apply to the considered scenario of multiple unaligned trajectories, we further compare our algorithm to an extension of BNIRL, which we refer to as BNIRL-EXT. Mimicking the ddBNIRL-S principle, the method accounts for the spatial context of the demonstrations

by assigning each state to the BNIRL subgoal that is targeted by the closest (see metric in Section 3.2.1) state-action pair— however, these assignments are made *after* the actual subgoal inference. When compared to ddBNIRL-S, this provides a reference of how much can be gained by considering the spatial relationship of the data *during* the inference. For the experiment, both ddBNIRL-S and BNIRL(-EXT) are augmented with their corresponding action sampling stages (Section 4.4) since the action sequences of the expert are discarded from the data set, in order to enable a fair comparison to the remaining algorithms.

Figure 9 shows the value loss over the size of the demonstration set for different reward settings. For small  $N_R$ , both BNIRL(-EXT) and ddBNIRL-S significantly outperform the reference methods. This is because the sparse reward structure allows for an efficient subgoal-based encoding of the expert behavior, which enables the algorithms to reconstruct the policy even from minimal amounts of demonstration data. However, the BNIRL(-EXT) solutions drastically deteriorate for denser reward structures. In particular, we observe a clear difference in performance between the cases where

- (i) we do not account for the spatial information in the partitioning model (BNIRL),
- (ii) include it in a post-processing step (BNIRL-EXT), and
- (iii) exploit it during the inference itself (ddBNIRL-S),

which demonstrates the importance of processing the context information. Most tellingly, ddBNIRL-S outperforms the baseline methods even in the dense reward regimes, although the subgoal-based encoding loses its efficiency here. In fact, the results reveal that the proposed approach combines the merits of both model types, i.e., the sample efficiency of the intentional models (max-margin/max-entropy) required for small data set sizes, as well as the asymptotic accuracy and fully probabilistic nature of the subintentional Bayesian framework (BPR).<sup>4</sup>

### 5.3 Robot Experiment

In the next experiment, we test the ddBNIRL framework on various real data sets, which we recorded on a KUKA lightweight robotic arm (Figure 10) via kinesthetic teaching. Videos of all demonstrated tasks can be found at <http://www.spg.tu-darmstadt.de/jmlr2018>.

The system has seven degrees of freedom, corresponding to the seven joints of the arm. Each joint is equipped with a torque sensor and an angle encoder, providing recordings of joint angles, velocities and accelerations. For our experiments, we only consider the xy-Cartesian coordinates spanning the transverse plane, which we computed from the raw measurements using a forward kinematic model. The data was recorded at a sampling rate of 50 Hz and further downsampled by a factor of 10, yielding an effective sample rate of 5 Hz, which provided a sufficient temporal resolution for the considered scenario.

The goal of the experiment is to learn a set of high-level intentional models for the recorded behavior types by partitioning the data sets into meaningful parts that can be used to predict the desired motion direction of the expert. For simplicity and to demonstrate the

---

4. The comparably large loss of BPR for small data set sizes can be explained by the fact that the framework is based on a more general policy model in which the expert behavior is assumed to be *inherently* stochastic, in contrast to the here considered setting where stochasticity arises merely a consequence of suboptimal decision-making.






Figure 10: KUKA lightweight robotic arm.

algorithm’s robustness to modeling errors, we adopt the simplistic transition model from Section 5.1 with the same action set containing the eight (inter-)cardinal motion directions. The high measurement accuracy of the end-effector position allows us to extract these high-level actions directly from the raw data, i.e., by selecting the directions with the smallest angular deviations from the ground truth (see example in Figure 11a). The underlying state space is obtained by discretizing the part of the coordinate range that is covered by the measurements into blocks of predefined size (see next sections for details). Apart from this discretization step and the aforementioned data downsampling, no preprocessing is applied.

### 5.3.1 SPATIAL PARTITIONING

First, we consider a case where the expert behavior can be described using a time-invariant policy model, which we aspire to capture via ddBNIRL-S. For our example, we consider the “Cycle” task shown in the video and in Figure 14. The same setting is analyzed using the time-variant ddBNIRL-T model in Section 5.3.2, which allows a direct comparison of the two approaches. The task consists in approaching a number of target positions, indicated by a set of markers (see video), before eventually returning to the initial state. The setting can be regarded as a real-world version of the “Loop” problem described by Michini and How (2012). As explained in their paper, classical IRL algorithms that rely on a global state-based reward model (such as max-margin IRL and max-entropy IRL) completely fail on this problem, due to the periodic nature of the task.

Figure 11a shows the downsampled and discretized data set (black arrows) obtained from four expert trajectories (white lines). For visualization purposes, the discretization block size is chosen as  $2\text{ cm} \times 2\text{ cm}$ , giving rise to a total of  $18 \times 24 = 432$  states. As in the top row of Figure 8 (ddBNIRL-S), the coloring of the background indicates the learned partitioning structure, computed from a low-temperature posterior sample. We observe that the found state clusters clearly reveal the modular structure of the task, providing an intuitive and interpretable explanation of the data. However, although the induced policy

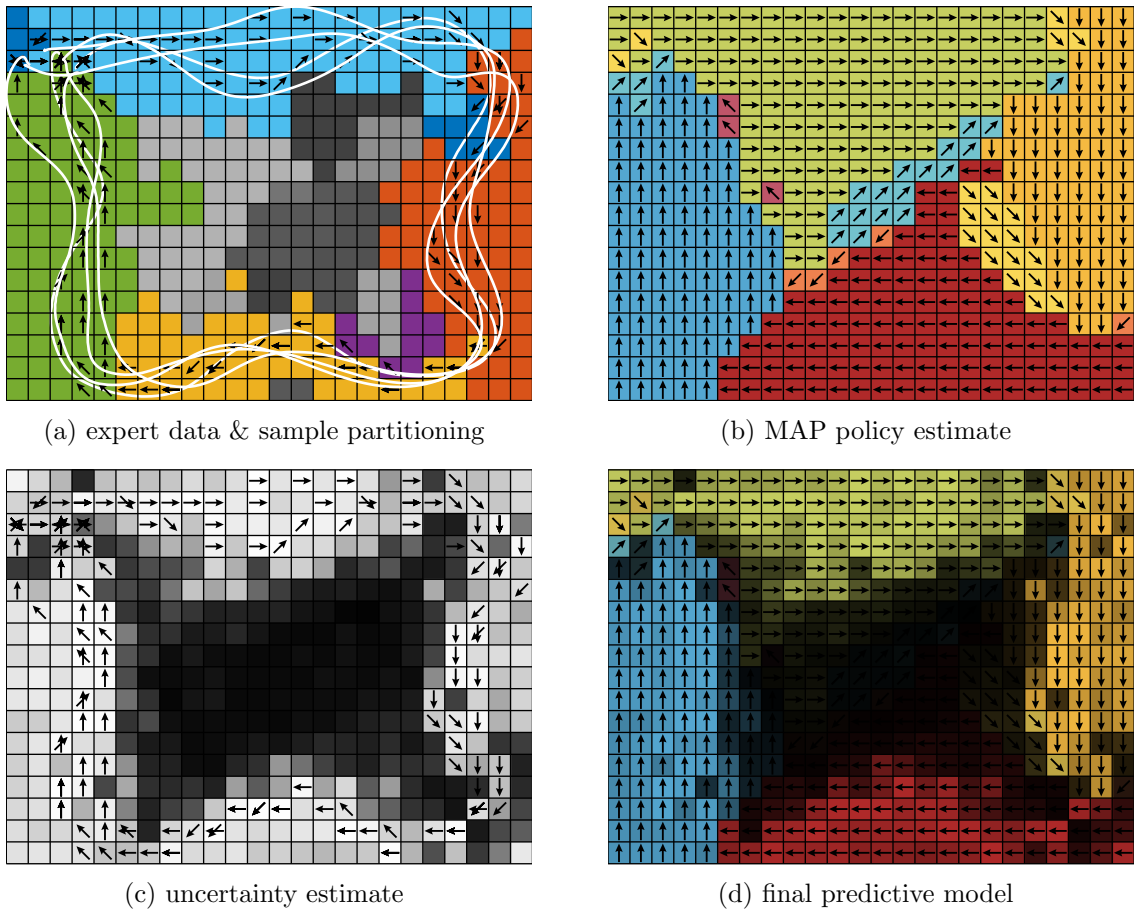


Figure 11: Results of ddBNIRL-S on the “Cycle” task. (a) Raw measurements (white lines) and discretized demonstration data (black arrows). The coloring of the background indicates a partitioning structure obtained from a low-temperature posterior sample. (b) Maximum a posteriori policy estimate. (c) Visualization of the model’s prediction uncertainty at all system states, represented by the entropies of the corresponding posterior predictive action distributions. Dark background indicates high uncertainty. (d) Illustration of the final predictive model, comprising both the action information and the prediction uncertainty.

model (Figure 11b) smoothly captures the cyclic nature of the task, we cannot expect to obtain trustworthy predictions in the center region of the state space, due to the lack of additional demonstration data that would be required to unveil the expert’s true intention in that region. Clearly, a point estimate such as the shown MAP policy cannot reflect this prediction uncertainty since it does not carry any confidence information. Yet, following a Bayesian approach, we can naturally quantify the prediction uncertainty at any query state  $s^*$  based on the shape of the corresponding posterior predictive action distribution  $p(a^* | s^*, \mathcal{D})$ . A straightforward option is, for example, to consider the prediction entropy,

defined as

$$H(s^*) := \sum_{a^* \in \mathcal{A}} p(a^* | s^*, \mathcal{D}) \log p(a^* | s^*, \mathcal{D}).$$

In order to obtain an unbiased approximation of the true *non-tempered* predictive distribution  $p(a^* | s^*, \mathcal{D})$ , we run a second Gibbs chain with unaltered temperature in parallel to the tempered chain. The resulting entropy estimates are summarized in an uncertainty map (Figure 11c), which we overlaid on the original prediction result to produce the final figure shown at the bottom right. Note that the obtained posterior uncertainty information of the model can be further used in an active learning setting, as demonstrated in Section 5.4.

### 5.3.2 TEMPORAL PARTITIONING

Next, we turn our attention to the ddBNIRL-T model, which we test against the vanilla BNIRL approach. For this purpose, we consider the full collection of tasks shown in the supplementary video, which comprises different time-dependent expert behaviors of varying complexity. In order to obtain a quantitative performance measure for our evaluation, we conducted a manual segmentation of all recorded trajectories, thereby creating a set of ground truth subgoal labels for all observed decision times. The result of this segmentation step is depicted in the appendix (Figure 14, center column). Note that the ground truth subgoals are assumed immediately at the ends of the corresponding segments.

The left and right column of Figure 14 show, respectively, the partitioning structures found by BNIRL and ddBNIRL-T, based on a uniform subgoal prior distribution with support at the visited states. The underlying state discretization block size is chosen as 1 cm  $\times$  1 cm, as indicated by the regular grid in the background. A simple visual comparison of the learned structures reveals the clear superiority of ddBNIRL-T over vanilla BNIRL on this problem set.

For our quantitative comparison, we consider the instantaneous subgoal localization errors of the two models over the entire course of a demonstration (Figure 12). Herein, the instantaneous localization error for a given state-action pair is measured in terms of the Euclidean distance between the grid location of the ground truth subgoal associated with the pair and the corresponding subgoal location predicted by the model. Note that the predictions of both models are based on the *entire* trajectory data of an experiment, considering the full posterior information after completing the demonstration. For ddBNIRL-T, which does not directly return a subgoal location estimate but instead provides access to the full subgoal posterior distribution, the error is computed with respect to the MAP subgoal locations  $\{\hat{g}_k\}$ ,

$$\hat{g}_k := \arg \max_{g_k \in \text{supp}(p_g)} p(g_k | \mathcal{D}, \tilde{\mathbf{c}}),$$

using the ddBNIRL-T version of Equation (12)—see note at the beginning of Section 4.

The black dots in the figure indicate the time instants where the ground truth annotations change. At those time instants, we observe significantly increased localization errors for both models, which can be explained by the fact that the ground truth annotation is somewhat subjective around the switching points (see labeling in Figure 14). Also, we notice a comparably high error at the beginning and the end of some trajectories, which stems from the imperfect synchronization between the recording interval and the execution

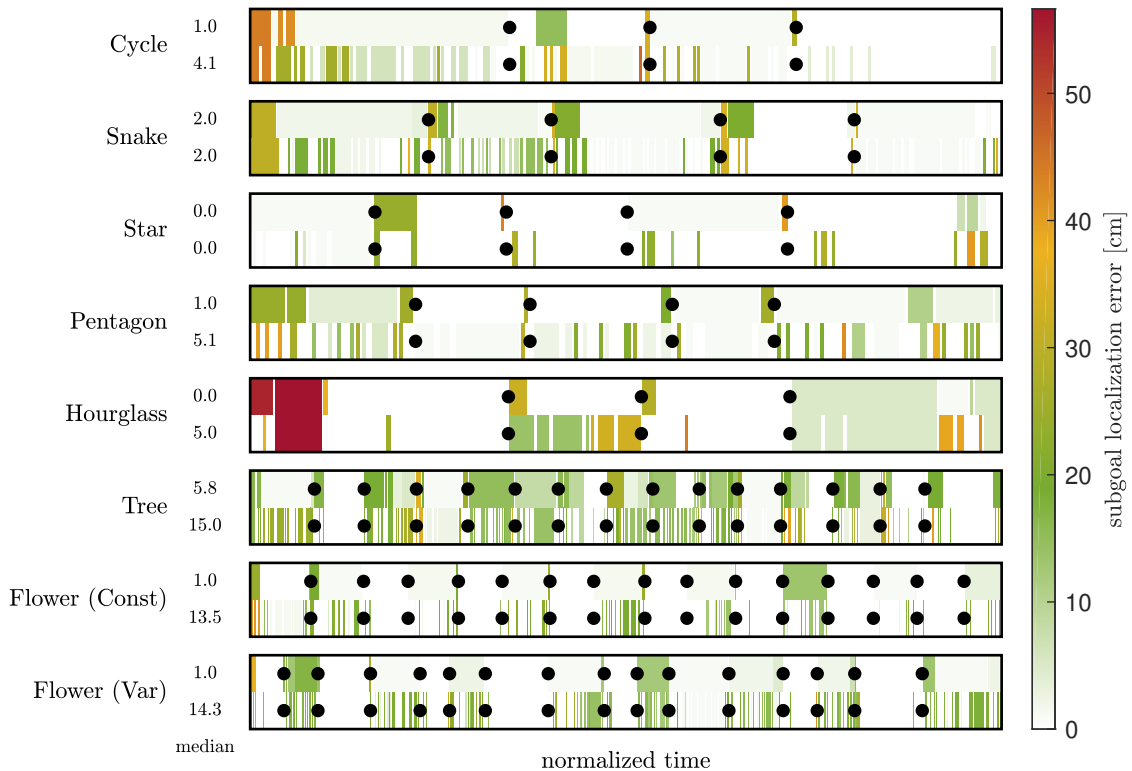


Figure 12: Instantaneous subgoal localization errors of ddBNIRL-T (upper rows) and BNIRL (lower rows) for the eight recorded data sets. The black dots indicate the subgoal switching times in the corresponding ground truth subgoal annotation, depicted in the center column of Figure 14. On average, the localization error of ddBNIRL-T is significantly lower compared to the BNIRL approach, as indicated by the median values shown on the left. For a qualitative comparison of the underlying partitioning structures, see Appendix A.

of the task (recall that we skipped the corresponding data preprocessing step). Hence, to capture the accuracy in a single figure of performance, we consider the median localization error of each time series, as it masks out these outliers and provides a more realistic error quantification than the sample mean. The obtained values are shown next to the error plots in Figure 12, indicating that the ddBNIRL-T localization error is in the range of the discretization interval in most cases. *Compared to BNIRL, the proposed method yields an error reduction of more than 70% on average.*

## 5.4 Active Learning

In Section 5.3, we saw that the posterior predictive action distribution  $p(a^* | s^*, \mathcal{D})$  provides a natural way to quantify the prediction uncertainty of our model at any given query state  $s^*$ . This offers the opportunity to apply the framework in an active learning setting, since the




Figure 13: Comparison between random data acquisition and active learning in the random MDP scenario. Shown are the empirical mean value losses of the obtained policy models over the number of data queries, obtained from 200 Monte Carlo runs.

induced uncertainty map (see example in Figure 11c) indicates in which parts of the state space the trained model can process further instructions from the expert most effectively.

To demonstrate the basic procedure, we reconsider the random MDP problem from Section 5.2 in an active learning context, where we compare different active strategies with the previously used random data acquisition scheme. As an initialization for the learning procedure, we request a single state-action pair  $(s_1, a_1)$  from the demonstrator, which we store in the initial data set  $\mathcal{D}_1 := \{(s_1, a_1)\}$ . Herein, the state  $s_1$  is drawn uniformly at random from  $\mathcal{S}$  and the action  $a_1 \sim \pi_E(a | s_1)$  is generated according to the noisy expert policy  $\pi_E : \mathcal{A} \times \mathcal{S} \rightarrow [0, 1]$  described in Section 5.2. Continuing from this point, each of the considered active learning algorithms requests a series of subsequent demonstrations  $((s_2, a_2), (s_3, a_3), \dots)$ , inducing a sequence of data sets  $(\mathcal{D}_1, \mathcal{D}_2, \mathcal{D}_3, \dots)$ , where the next query state  $s_{d+1}$  is chosen according to the specific data acquisition criterion  $f_{\text{acq}}$  of the algorithm evaluated on the current predictive model,

$$\begin{aligned} \mathcal{D}_{d+1} &= \mathcal{D}_d \cup \{(s_{d+1}, a_{d+1})\} \\ s_{d+1} &= \arg \max_{s^* \in \mathcal{S}} f_{\text{acq}}[p(a^* | s^*, \mathcal{D}_d)] \\ a_{d+1} &\sim \pi_E(a | s_{d+1}). \end{aligned}$$

The purpose of the acquisition criterion is to assess the uncertainty of the model at all possible query states, so that the next demonstration can be requested in the high uncertainty region of the state space (see *uncertainty sampling*, Settles, 2010). For our experiment, we consider the following three common choices,

- highest entropy:  $f_{\text{acq}}(p) := - \sum_{a \in \mathcal{A}} p(a) \log p(a)$ ,
- least confidence:  $f_{\text{acq}}(p) := 1 - \max_{a \in \mathcal{A}} p(a)$ ,
- smallest margin:  $f_{\text{acq}}(p) := p(\hat{a}_2) - p(\hat{a}_1)$ ,

where  $\hat{a}_1$  and  $\hat{a}_2$  denote, respectively, the most likely and second most likely action according to the considered distribution  $p$ , i.e.,  $\hat{a}_1 := \arg \max_{a \in \mathcal{A}} p(a)$  and  $\hat{a}_2 := \arg \max_{a \in \mathcal{A} \setminus \hat{a}_1} p(a)$ . At each iteration, we compute the value losses (Equation 20) of the induced policy models and compare them with the corresponding loss obtained from random data acquisition. The resulting curves are delineated in Figure 13. As expected, the learning speed of the model is significantly improved under all active acquisition schemes, which reduces the number of expert demonstrations required to successfully learn the observed task.

## 6. Conclusion

Building upon the principle of Bayesian nonparametric inverse reinforcement learning, we proposed a new framework for data-efficient IRL that leverages the context information of the demonstration set to learn a predictive model of the expert behavior from small amounts of training data. Central to our framework are two model architectures, one designed for learning spatial subgoal plans, the other to capture time-varying intentions. In contrast to the original BNIRL model, both architectures explicitly consider the covariate information contained in the demonstration set, giving rise to predictive models that are inherently robust to demonstration noise. While the original BNIRL model can be recovered as a special case of our framework, the conducted experiments show a drastic improvement over the vanilla BNIRL approach in terms of the achieved subgoal localization accuracy, which stems from both an improved likelihood model and a context-aware clustering of the data. Most notably, our framework outperforms all tested reference methods in the analyzed benchmark scenarios while it additionally captures the full posterior information about the learned subgoal representation. The resulting prediction uncertainty about the expert behavior, reflected by the posterior predictive action distribution, provides a natural basis to apply our method in an active learning setting where the learning system can request additional demonstration data from the expert.

The current limitation of our approach is that both presented architectures require an MDP model with discrete state and action space. While the subgoal principle carries over straightforwardly to continuous metric spaces, the construction of the likelihood model becomes difficult in these environments as it requires knowledge of the optimal state-action value functions for all potential subgoal locations. However, for BNIRL, there exist several ways to approximate the likelihood in these cases (Michini et al., 2013) and the same concepts apply equally to ddBNIRL. Thus, an interesting future study would be to compare the efficacy of both model types on larger problems involving continuous spaces, where it appears even more natural to follow a distance-based approach.

## Acknowledgments

This project has received funding from the European Union’s Horizon 2020 research and innovation program under grant agreement No #713010 (GOALRobots) and No #640554 (SKILLS4ROBOTS).

Appendix A. Robot experiment

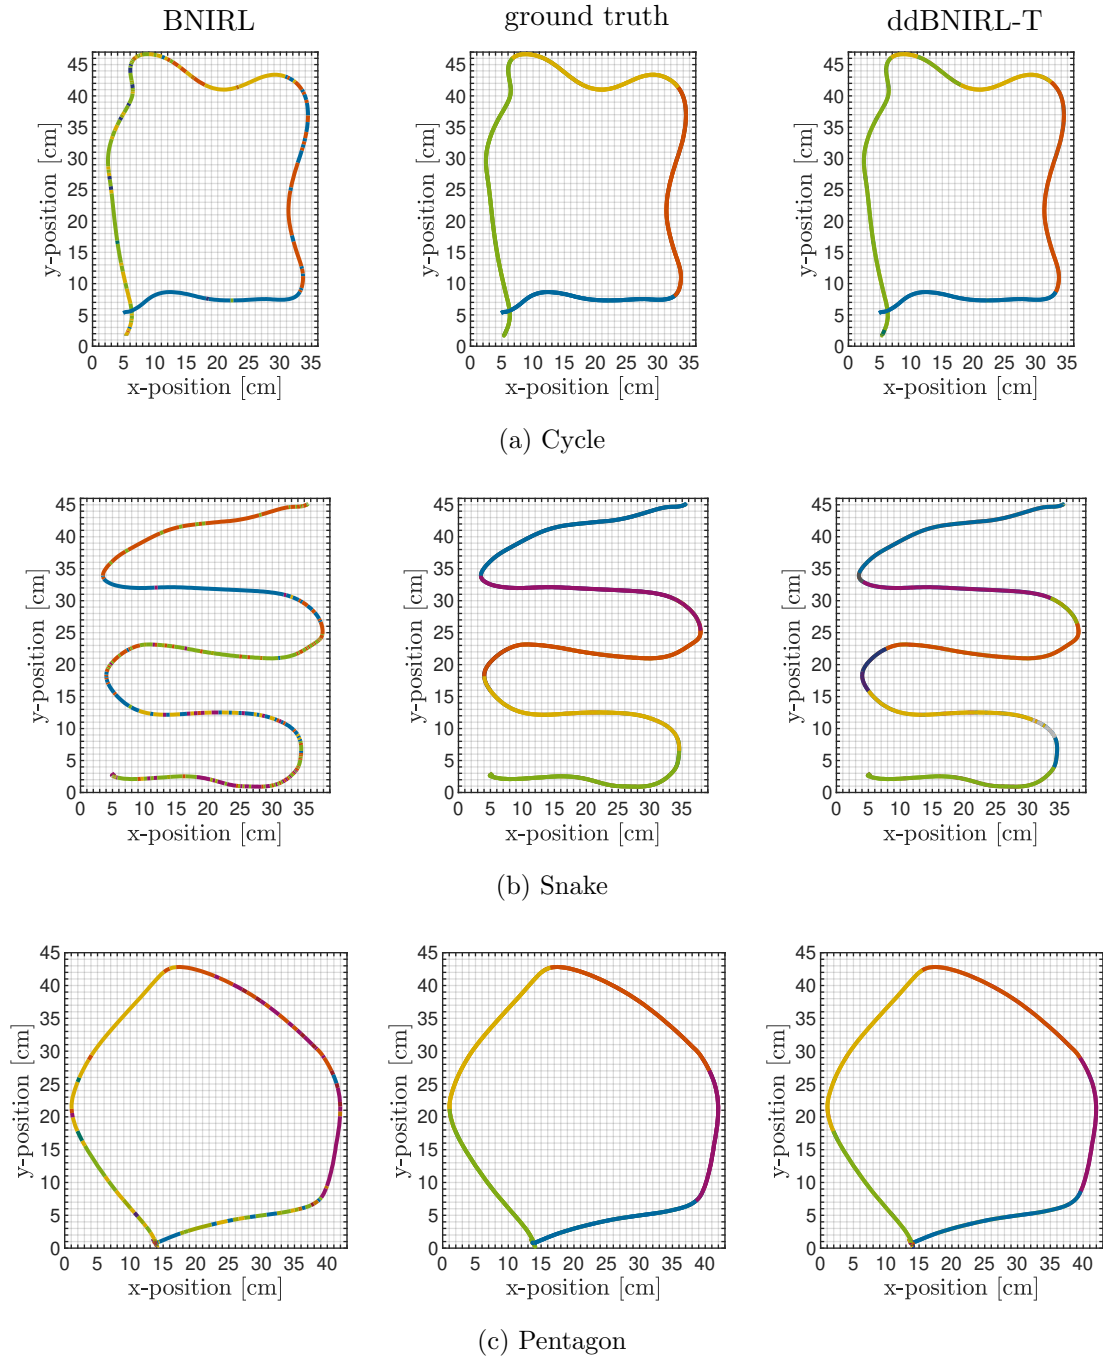


Figure 14: Motion sequences without trajectory crossings, which can be represented using a spatial subgoal pattern.

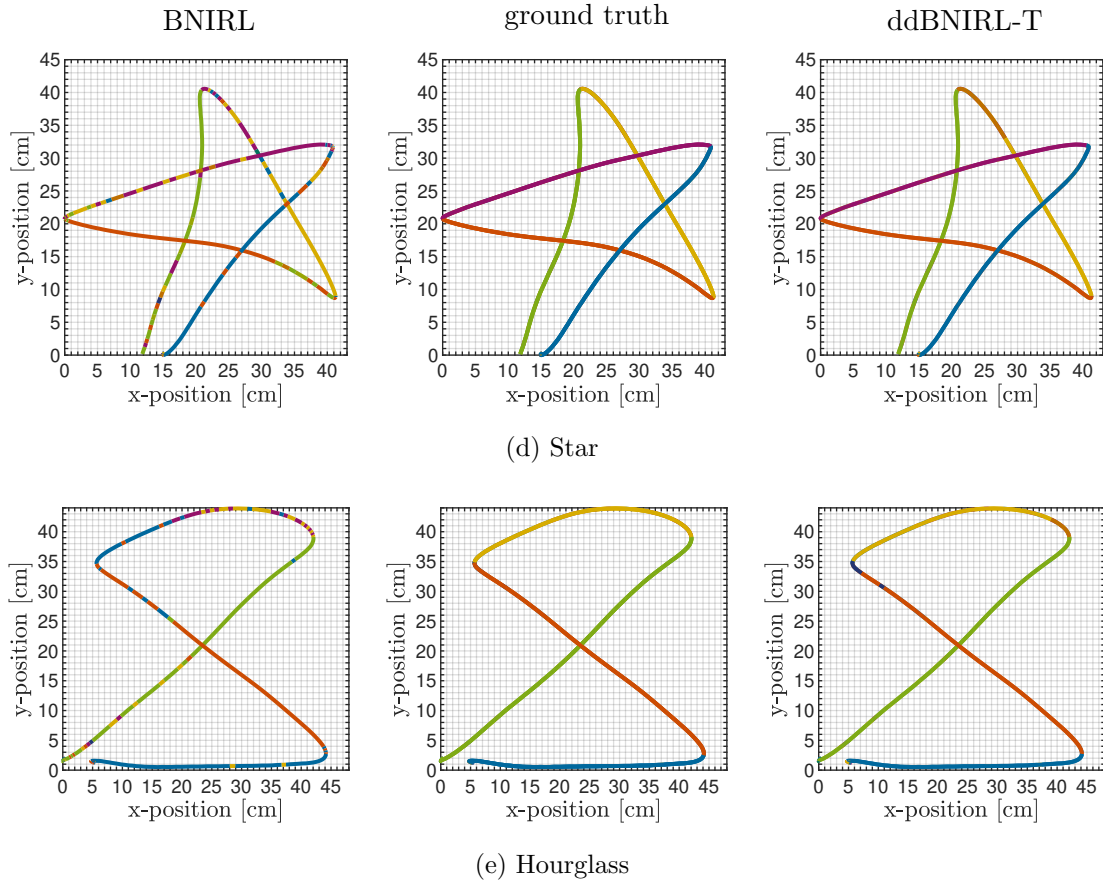


Figure 14 (continued): Motion sequences with few trajectory crossings, requiring a time-varying subgoal representation.



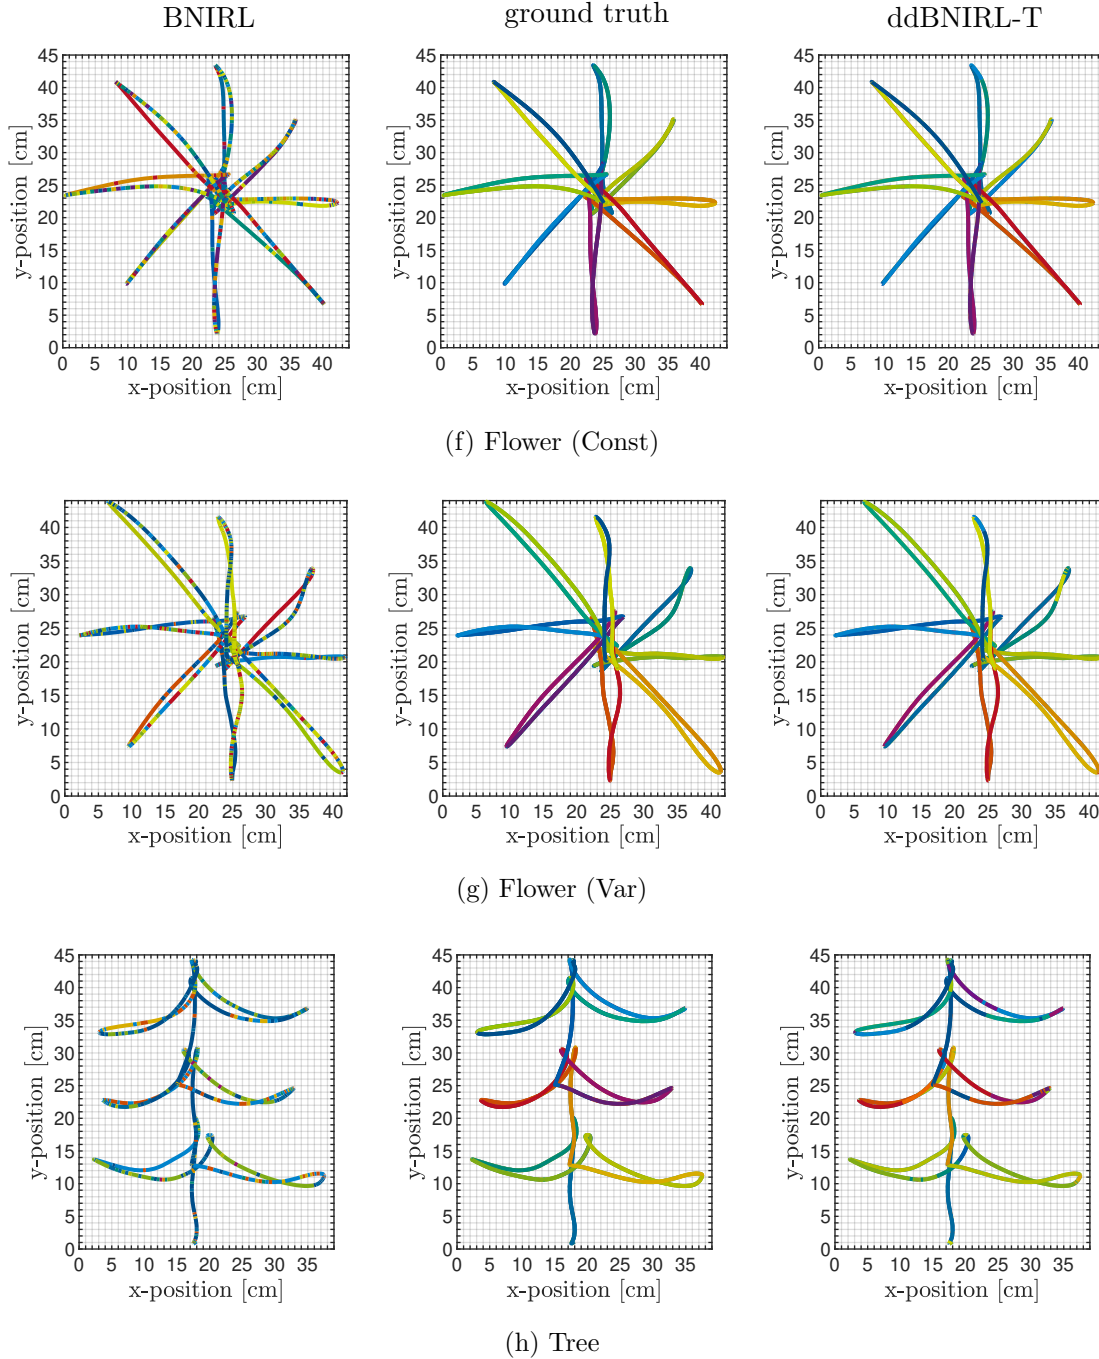


Figure 14 (continued): Long motion sequences comprising a large number of sub-patterns with overlapping parts that can be only separated by considering the temporal context. Flower (Const): all strokes are performed with the same absolute velocity. Flower (Var): the individual strokes are performed with alternating velocity.

## References

- P. Abbeel and A. Y. Ng. Apprenticeship learning via inverse reinforcement learning. In *International Conference on Machine Learning*, page 1, 2004.
- M. Al-Emran. Hierarchical reinforcement learning: a survey. *International Journal of Computing and Digital Systems*, 4(2), 2015.
- S. V. Albrecht and P. Stone. Autonomous agents modelling other agents: a comprehensive survey and open problems. [arXiv:1709.08071 \[cs.AI\]](https://arxiv.org/abs/1709.08071), 2017.
- D. J. Aldous. *Exchangeability and Related Topics*. Springer, 1985.
- B. D. Argall, S. Chernova, M. Veloso, and B. Browning. A survey of robot learning from demonstration. *Robotics and Autonomous Systems*, 57(5):469–483, 2009.
- M. Babeş-Vroman, V. Marivate, K. Subramanian, and M. Littman. Apprenticeship learning about multiple intentions. In *International Conference on Machine Learning*, pages 897–904, 2011.
- L. C. Baird. Advantage updating. Technical report, Wright Lab, 1993.
- S. Bhatnagar, R. Sutton, M. Ghavamzadeh, and M. Lee. Natural actor-critic algorithms. *Automatica*, 45(11), 2009.
- D. M. Blei and P. I. Frazier. Distance dependent Chinese restaurant processes. *Journal of Machine Learning Research*, 12(Nov):2461–2488, 2011.
- D. M. Blei, A. Y. Ng, and M. I. Jordan. Latent Dirichlet allocation. *Journal of Machine Learning Research*, 3(Jan):993–1022, 2003.
- M. M. Botvinick. Hierarchical reinforcement learning and decision making. *Current Opinion in Neurobiology*, 22(6):956–962, 2012.
- S. J. Bradtke and M. O. Duff. Reinforcement learning methods for continuous-time Markov decision problems. In *Advances in Neural Information Processing Systems*, pages 393–400, 1994.
- N. Cesa-Bianchi, C. Gentile, G. Neu, and G. Lugosi. Boltzmann exploration done right. In *Advances in Neural Information Processing Systems*, pages 6275–6284, 2017.
- J. Choi and K.-E. Kim. Nonparametric Bayesian inverse reinforcement learning for multiple reward functions. In *Advances in Neural Information Processing Systems*, pages 305–313, 2012.
- C. Daniel, H. Van Hoof, J. Peters, and G. Neumann. Probabilistic inference for determining options in reinforcement learning. *Machine Learning*, 104(2-3):337–357, 2016a.
- Christian Daniel, Gerhard Neumann, Oliver Kroemer, and Jan Peters. Hierarchical relative entropy policy search. *Journal of Machine Learning Research*, 17(1):3190–3239, 2016b.

- C. Dimitrakakis and C. A. Rothkopf. Bayesian multitask inverse reinforcement learning. In *European Workshop on Reinforcement Learning*, pages 273–284, 2011.
- N. J. Foti and S. A. Williamson. A survey of non-exchangeable priors for Bayesian non-parametric models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(2):359–371, 2015.
- M. Ghavamzadeh, S. Mannor, J. Pineau, and A. Tamar. Bayesian reinforcement learning: a survey. *Foundations and Trends in Machine Learning*, 8(5-6):359–483, 2015.
- B. M. Kapron, V. King, and B. Mountjoy. Dynamic graph connectivity in polylogarithmic worst case time. In *Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1131–1142, 2013.
- S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi. Optimization by simulated annealing. *Science*, 220(4598):671–680, 1983.
- D. Koller and N. Friedman. *Probabilistic Graphical Models: Principles and Techniques*. MIT Press, 2009.
- G. Konidaris, S. Kuindersma, R. Grupen, and A. Barto. Robot learning from demonstration by constructing skill trees. *International Journal of Robotics Research*, 31(3):360–375, 2012.
- S. Krishnan, A. Garg, R. Liaw, L. Miller, F. T. Pokorny, and K. Goldberg. HIRL: hierarchical inverse reinforcement learning for long-horizon tasks with delayed rewards. [arXiv:1604.06508 \[cs.R0\]](https://arxiv.org/abs/1604.06508), 2016.
- S. Levine, Z. Popovic, and V. Koltun. Nonlinear inverse reinforcement learning with Gaussian processes. In *Advances in Neural Information Processing Systems*, pages 19–27, 2011.
- R. Lioutikov, G. Neumann, G. Maeda, and J. Peters. Learning movement primitive libraries through probabilistic segmentation. *International Journal of Robotics Research*, 36(8):879–894, 2017.
- M. L. Littman, T. L. Dean, and L. P. Kaelbling. On the complexity of solving Markov decision problems. In *Conference on Uncertainty in Artificial Intelligence*, pages 394–402, 1995.
- B. Michini and J. P. How. Bayesian nonparametric inverse reinforcement learning. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 148–163, 2012.
- B. Michini, M. Cutler, and J. P. How. Scalable reward learning from demonstration. In *IEEE International Conference on Robotics and Automation*, pages 303–308, 2013.
- B. Michini, T. J. Walsh, A.-A. Agha-Mohammadi, and J. P. How. Bayesian nonparametric reward learning from demonstration. *IEEE Transactions on Robotics*, 31(2):369–386, 2015.

- G. Neu and C. Szepesvári. Apprenticeship learning using inverse reinforcement learning and gradient methods. In *Conference on Uncertainty in Artificial Intelligence*, pages 295–302, 2007.
- A. Y. Ng and S. J. Russell. Algorithms for inverse reinforcement learning. In *International Conference on Machine Learning*, pages 663–670, 2000.
- A. Y. Ng, D. Harada, and S. Russell. Policy invariance under reward transformations: Theory and application to reward shaping. In *International Conference on Machine Learning*, pages 278–287, 1999.
- Q. P. Nguyen, B. K. H. Low, and P. Jaillet. Inverse reinforcement learning with locally consistent reward functions. In *Advances in Neural Information Processing Systems*, pages 1747–1755, 2015.
- S. Niekum, S. Osentoski, G. Konidaris, and A. G. Barto. Learning and generalization of complex tasks from unstructured demonstrations. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 5239–5246, 2012.
- A. Panella and P. Gmytrasiewicz. Interactive POMDPs with finite-state models of other agents. *Autonomous Agents and Multi-Agent Systems*, pages 861–904, 2017.
- M. L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, Inc., 1994.
- D. Ramachandran and E. Amir. Bayesian inverse reinforcement learning. *International Joint Conference on Artificial Intelligence*, pages 2586–2591, 2007.
- P. Ranchod, B. Rosman, and G. Konidaris. Nonparametric Bayesian reward segmentation for skill discovery using inverse reinforcement learning. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 471–477, 2015.
- G. O. Roberts and S. K. Sahu. Updating schemes, correlation structure, blocking and parameterization for the Gibbs sampler. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 59(2):291–317, 1997.
- C. A. Rothkopf and C. Dimitrakakis. Preference elicitation and inverse reinforcement learning. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 34–48, 2011.
- E. Rueckert, G. Neumann, M. Toussaint, and W. Maass. Learned graphical models for probabilistic planning provide a new class of movement primitives. *Frontiers in Computational Neuroscience*, 6:97, 2013.
- S. Schaal, J. Peters, J. Nakanishi, and A. Ijspeert. Learning movement primitives. *Robotics Research*, pages 561–572, 2005.
- B. Settles. Active learning literature survey. Technical report, University of Wisconsin-Madison, 2010.

- Ö Şimşek, A. P. Wolfe, and A. G. Barto. Identifying useful subgoals in reinforcement learning by local graph partitioning. In *International Conference on Machine Learning*, pages 816–823, 2005.
- A. Šošić, A. M. Zoubir, and H. Koepl. Inverse reinforcement learning via nonparametric subgoal modeling. In *AAAI Spring Symposium on Data-Efficient Reinforcement Learning*, 2018a.
- A. Šošić, A. M. Zoubir, and H. Koepl. A Bayesian approach to policy recognition and state representation learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(6):1295–1308, 2018b.
- M. Stolle and D. Precup. Learning options in reinforcement learning. In *International Symposium on Abstraction, Reformulation, and Approximation*, pages 212–223, 2002.
- A. Surana and K. Srivastava. Bayesian nonparametric inverse reinforcement learning for switched markov decision processes. In *IEEE International Conference on Learning and Applications*, pages 47–54, 2014.
- R. S. Sutton and A. G. Barto. *Reinforcement Learning: An Introduction*. MIT Press, 1998.
- R. S. Sutton, D. Precup, and S. Singh. Between MDPs and semi-MDPs: a framework for temporal abstraction in reinforcement learning. *Artificial Intelligence*, 112(1):181–211, 1999.
- M. Tamassia, F. Zambetta, W. Raffe, and X. Li. Learning options for an MDP from demonstrations. In *Australasian Conference on Artificial Life and Computational Intelligence*, pages 226–242, 2015.
- H. M. Taylor and S. Karlin. *An Introduction to Stochastic Modeling*. Academic Press, 1984.
- A. Tewari and P. L. Bartlett. Optimistic linear programming gives logarithmic regret for irreducible MDPs. In *Advances in Neural Information Processing Systems*, pages 1505–1512, 2008.
- J. Yang, Y.-G. Jiang, A. G. Hauptmann, and C.-W. Ngo. Evaluating bag-of-visual-words representations in scene classification. In *International Workshop on Multimedia Information Retrieval*, pages 197–206, 2007.
- S. Zhifei and E. M. Joo. A survey of inverse reinforcement learning techniques. *International Journal of Intelligent Computing and Cybernetics*, 5(3):293–311, 2012.
- B. D. Ziebart, A. L. Maas, J. A. Bagnell, and A. K. Dey. Maximum entropy inverse reinforcement learning. In *AAAI Conference on Artificial Intelligence*, pages 1433–1438, 2008.