

Efficient Learning of High Level Plans from Play

Núria Armengol Urpí¹, Marco Bagatella¹, Otmar Hilliges¹, Georg Martius² and Stelian Coros¹

Abstract—Real-world robotic manipulation tasks remain an elusive challenge, since they involve both fine-grained environment interaction, as well as the ability to plan for long-horizon goals. Although deep reinforcement learning (RL) methods have shown encouraging results when planning end-to-end in high-dimensional environments, they remain fundamentally limited by poor sample efficiency due to inefficient exploration, and by the complexity of credit assignment over long horizons. In this work, we present Efficient Learning of High-Level Plans from Play (ELF-P), a framework for robotic learning that bridges motion planning and deep RL to achieve long-horizon complex manipulation tasks. We leverage task-agnostic play data to learn a discrete behavioral prior over object-centric primitives, modeling their feasibility given the current context. We then design a high-level goal-conditioned policy which (1) uses primitives as building blocks to scaffold complex long-horizon tasks and (2) leverages the behavioral prior to accelerate learning. We demonstrate that ELF-P has significantly better sample efficiency than relevant baselines over multiple realistic manipulation tasks and learns policies that can be easily transferred to physical hardware.

I. INTRODUCTION

One of the collective visions of robotics is a world where robots help humans with daily household chores, such as setting up a table or emptying the dishwasher. An immediate challenge is that, to operate in the physical world, a robot must be able to reason in a mixed decision space. That is, it must combine decisions relating continuous motions with discrete subtasks in order to accomplish complex tasks. For example, the task of emptying the dishwasher requires the robot to first go to the dishwasher, then open it, grab a single dish, open the cupboard and place the dish inside, and repeat until the dishwasher is empty. Failing to perform a single subtask or ordering them incorrectly would ultimately lead to a failure. The ability to reason over long horizons, we argue, is thus a crucial milestone in achieving the above-mentioned vision.

Despite significant improvements in several fields of robotics, ranging from motion planning to robust control [1]–[4], long-horizon problems involving several tasks still presents a considerable challenge. While planning algorithms can efficiently search well-defined state spaces and plan over relatively long horizons, they generally struggle with complex, stochastic dynamics, and high-dimensional systems, in particular featuring narrow passages (i.e. low-measure regions which have to be traversed to reach a

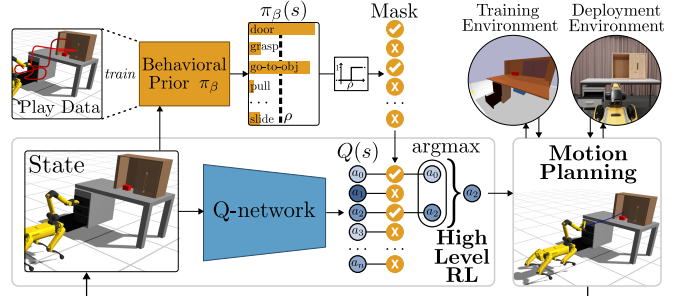


Fig. 1. Overview of ELF-P. A discrete behavioral prior learned on play data is used to discard infeasible actions at every state, enabling a high-level RL agent to search over a smaller space of behaviors. Each action is mapped to a motion primitive, which is executed through classical motion planning methods. After training in simulation, the policies can be easily deployed on real hardware.

goal). Moreover, they often require accurate models of the environment or handcrafted distance measures [5]–[7].

Learning methods have thus emerged as a more scalable approach for handling the high-dimensionality of the real world. While capable of obtaining hard-to-engineer behaviors [8]–[11], learning-based methods are similarly limited to short horizons tasks in general, since longer plans involving multi-level reasoning aggravate the challenges of exploration and temporal credit assignment [12].

In this work, we introduce ELF-P, a novel approach for robotic learning of long-horizon manipulation tasks, bridging motion planning and deep reinforcement learning (RL). In a two-level hierarchy, high-level planning is entrusted to a goal-conditioned policy, trained through deep RL and designed to deal with the high-dimensionality and complexity of real-world dynamics. This policy operates at a coarser time scale and controls the execution of primitives, which lie at the lower level in the hierarchy (see Figure 1). Crucially, we sidestep the instability of jointly learning high and low-level policies and thus rely on a predefined library of object-centric primitives (such as grasping or reaching a target configuration in a collision-free path) obtained through classical motion planning.

While planning over primitives substantially reduces the burden of the higher level, uniform exploration may still be insufficient in long-horizon tasks with sparse rewards, as the space of possible behaviors grows rapidly with the size of the action space and number of steps to the goal [13]. Although humans excel at inferring the set of behaviors enabled by a situation, uninformed RL agents tend to repeatedly attempt futile actions, resulting in poor sample efficiency and potentially unsafe exploration.

We thus additionally propose to prune the high-level action space to a subspace of feasible primitives, inferred dynam-

¹Department of Computer Science, ETH Zurich, Switzerland {nuria.armengolurpi, mbagatella, scoros, otmar.hilliges}@inf.ethz.ch

²Max Planck Institute for Intelligent Systems, Tübingen, Germany georg.martius@tuebingen.mpg.de

ically at each step. In practice, this is modelled through a discrete behavioral prior learned from task-agnostic play data.

While existing works on behavioral priors largely rely on soft integration schemes, we instead propose to discretize the prior into a hard, binary constraint. We then define an algorithm that only explores and exploits over the set of feasible actions, effectively performing Q-learning in a reduced Markov Decision Process (MDP). Theoretically, a near-optimal solution to this MDP can be retrieved more efficiently and under mild assumptions, can be generalized to the original MDP. Empirically, we observe improved performance in a variety of complex manipulation tasks and highlight the ease of transferring the policy to real-world hardware.

Our approach represents a promising step towards learning to solve long-horizon tasks in a hierarchical manner by combining task-agnostic play data, RL and motion planning algorithms. Our contributions can be summarized as follows: (a) we explore the integration of RL with classic motion planning algorithms and show its feasibility in a simulated manipulation environment; (b) we highlight how the exploration problem can be further addressed by introducing a prior model learned from play data; (c) we evaluate our proposed algorithm both theoretically and empirically, while comparing it against relevant baselines; (d) we show how learned policies can be deployed to physical hardware. Videos and additional information can be found at nuria95.github.io/elf-p.

II. RELATED WORK

A. Planning Algorithms

Classical motion planning methods to move a robot in a collision-free path are based on sampling [16]–[18] or constrained optimization [19], [20] and can efficiently find long-horizon paths. However, they do not allow a robot to alter the world. Task-and-motion planning (TAMP) methods are used when planning for a robot that operates in environments containing a large number of objects, taking actions to navigate the world as well as to change the state of the objects [21]–[25]. Despite performing well in long-horizon planning, they are limited by the need of a world model, their task representations, the dimensionality of the search space and the inability to execute high-dimensional, complex tasks robustly.

B. Model-Free Learning

Model-free learning offers a promising alternative when dealing with unknown environments but relies on a reward function defining the task at hand. Additionally, model-free methods are sample inefficient [26] in general and have difficulties reasoning over long horizons. Hierarchical RL (HRL) [27]–[31] methods have been proposed as a scalable solution that directly leverages temporal abstraction but in practice, they struggle with sample complexity and suffer from brittle generalization [32]. Several HRL algorithms rely on goal-conditioned policies [33]–[36] for low-level control

but they tend to make training unstable (see discussion in [12]). To address this, several works have adopted classical motion planning techniques to replace the high-level policy (e.g. tree search methods [5], [6]). In contrast, we sidestep the problem by replacing the low-level policy with a predefined set of task-primitives and train a model-free high-level policy to control their execution. Closely related to our work, [32], [37], [38] use Parameterized Action MDPs [39] in which the agent executes a parameterized primitive at each decision-making step. These methods also rely on a specified library of primitives but, unlike ours, they are parameterized. However, this versatility in primitive instantiation comes at a low efficiency cost since the agent needs to explore large amounts of parameters to solve long horizon tasks. Other methods also use motion planners as a low-level controller and learn an RL policy in a high-level action space [40]–[42].

C. Guiding Exploration in Large Action Spaces

As exploration remains a fundamental challenge for RL agents, previous works have attempted to ease the burden of planning in complex spaces by pruning the actions available at each step. Invalid action masking has been proposed in large strategy games to restrict sampling to a fixed subset of the action space [43]–[45]. However, these methods assume that the set of illegal actions is given a priori. The case in which a random subset of all actions is available at each step was also formally studied by [46]. When prior information is not available, several methods naturally propose to learn which actions are suitable. In the action elimination literature, this is achieved with [47] or without [48] assuming the availability of additional supervision from the environment, although the latter case remains constrained to tabular settings. Another relevant framework is that of affordances [49]–[51], which measure the possibility for single actions to achieve desired future configurations. Affordances can again be learned from an explicit signal from the environment [50] or utilize prior information. Our method is perhaps more closely related to behavioral priors, which can be learned from offline trajectories [52]–[54] or online interaction [55] and can be used to direct exploration [54] or regularize learned policies [55]. Behavioral priors can be modeled as conditional probability distributions over the action space. Our method also learns a state-conditional model of possible actions from data and does not require explicit supervision (i.e. validity labels for the agent’s actions). However, while behavioral prior literature largely focuses on continuous actions and soft constraints (e.g. KL-regularization), our method learns over a discrete set of actions and relies on a hard integration of the prior into the training process.

III. PROBLEM STATEMENT

We model the environment as a goal-conditioned Markov Decision Process represented as an 8-tuple $(\mathcal{S}, \mathcal{G}, \mathcal{A}, P, R, \rho_0, \rho_g, \gamma)$ with possibly continuous states $s \in \mathcal{S}$ and goals $g \in \mathcal{G}$, discrete actions $a \in \mathcal{A}$, transition kernel $P(\cdot|s, a)$, reward function $R(s, g)$, initial state distribution ρ_0 , goal distribution ρ_g and discount factor γ . We focus on

sparse reward signals: $R(\cdot|s, g) = \mathbb{1}_{\{|f(s)-g|_d < \epsilon\}}$, where $|\cdot|_d$ is an arbitrary distance metric, $f: \mathcal{S} \rightarrow \mathcal{G}$ is a projection and ϵ is a threshold. We additionally denote by $\pi: \mathcal{S} \times \mathcal{G} \rightarrow \mathcal{A}$ a goal-conditioned stationary policy.

Crucially, we assume access to a fixed dataset collected through *play*, following an unknown behavior policy π_β . Play data [56] can be inexpensively collected by a human operator controlling a robot to achieve arbitrary environment configurations while interacting with the objects at hand. This is useful to extract *affordances* [49], representing the subset of possible actions which are feasible in a current situation. For example, a drawer can be pushed or pulled, but not moved sideways. In our framework, *play* happens at a high level, more specifically at the level of motion primitives or skills, where a human can choose from a predefined set of primitives, which are generally available (e.g. precoded behaviors for real-word hardware such as for Spot robot [57] or as in [58]). Another solution, which could involve recording low-level action sequences, and segmenting them into meaningful skills as in [59]–[61], is left for future works. The resulting play dataset consists of N high-level state-action pairs $\mathcal{D} = \{(s_i, a_i)\}_{i=1}^N$ where a high level action a_i represents the primitive enacted by the play agent in state s_i .

Following the goal-conditioned framework [62], we intend to find a policy $\pi: \mathcal{S} \times \mathcal{G} \rightarrow \mathcal{A}$ that maximizes the expected discounted return

$$J(\pi) = \mathbb{E}_{g \sim \rho_g, \mu^\pi} \left[\sum_{t=1}^{\infty} \gamma^{t-1} R(s_t, g) \right], \quad (1)$$

under the trajectory distribution $\mu^\pi(\tau|g) = \rho_0(s_0) \prod_{t=0}^{\infty} P(s_{t+1}|s_t, a_t)$ with $a_t = \pi(s_t, g)$.

We remark that through this paper, the notation of atomic actions $a \in \mathcal{A}$ is used to represent non-parametric motion primitives. We effectively abstract away the complexity of low-level control and let the action space \mathcal{A} be composed of a finite set of high-level behaviors which are, in practice, executed through established motion planning methods. As a consequence, a single timestep in the MDP corresponds to the execution of a single primitive. We finally remark that, while the method described in Section IV is designed and deployed in this particular setting, it remains generally applicable (e.g. in MDPs with low-level actions).

IV. EFFICIENT LEARNING OF HIGH LEVEL PLANS FROM PLAY (ELF-P)

We now present our algorithm ELF-P for solving long-horizon tasks with motion primitives and play-guided RL. To address exploration, we first learn a discrete behavioral prior that eliminates infeasible actions from the set of primitives and hence prunes the search space (Section IV-A, Figure 2). Next, in Section IV-C we propose and motivate an integration scheme for the learned prior onto Q-learning. Consequently, our agent can focus on learning Q-values for feasible state-action pairs only, as the prior generally lifts the burden of learning to avoid infeasible actions.

A. Learning a Prior from Play

Let us start by considering the play dataset $\mathcal{D} = \{(s_1, a_1), \dots, (s_N, a_N)\}$ introduced in Section III. While lacking explicit exploitative behaviors, play data inherently favors actions that are *feasible*: while not necessarily desirable for a given goal, those actions are likely to be successfully executed given the current state of the environment. We aim to extract this information by estimating a (goal-independent) behavioral prior $\pi^\beta(\cdot|s)$, modeled as a conditional categorical distributions over primitives, which associates feasible primitives to higher likelihoods.

We parameterize π^β through a neural network with learnable parameters ω , which can be trained via standard mini-batch first order techniques to minimize the negative log-likelihood $\mathcal{L}_{NLL}(\omega)$:

$$\mathcal{L}_{NLL}(\omega) = \mathbb{E}_{B \sim \mathcal{D}} \left[\frac{1}{|B|} \sum_{(s,a) \in B} -\log \pi_\omega^\beta(a|s) \right], \quad (2)$$

where B represents a batch of state-action pairs sampled uniformly from the dataset \mathcal{D} .

B. Selecting Feasible Actions

Given the learned behavioral prior $\pi_\beta(\cdot|s)$, we propose to turn its soft probability distribution into hard, binary constraints on the action space. We thus define a threshold-based selection operator $\alpha: \mathcal{S} \rightarrow \mathcal{P}(\mathcal{A})$:

$$\alpha(s) = \{a \in \mathcal{A} \mid \pi_\beta(a|s) > \rho\}, \quad (3)$$

where $\mathcal{P}(\cdot)$ represents a powerset. Ideally, an action $a \notin \alpha(s)$ would not be chosen by an optimal goal-conditioned policy in state s . We refer to $\alpha(s)$ as the set of *feasible* actions for state s . See Figure 2 for a visualization.

C. Learning in a Reduced MDP

The learned selection operator $\alpha(s)$ enables the definition of an auxiliary MDP \mathcal{M}' , which we refer to as *reduced* MDP. The definition and solution of the reduced MDP lay at the core of our method. We model \mathcal{M}' through a generalized definition of MDPs [63] in which available actions depend on the current state: in the state s , the action space is restricted to a subset $\alpha(s) \subseteq \mathcal{A}$.




Fig. 2. The trained behavioral prior π^β learns to estimate the set of feasible primitives in different environment configurations. (Left) When the agent is at the center of the space, the prior favors primitives that involve reaching elements through a free-collision path (go to *door*, *drawers* or above the *object*), but prevents actions that involve object manipulation. (Right) When being close to the door, the prior learns correct object affordances such as *grasp* or *slide*.

Definition 1 (Reduced MDP): Given an MDP $\mathcal{M} = (\mathcal{S}, \mathcal{G}, \mathcal{A}, P, R, \rho_0, \rho_g, \gamma)$ and a selection operator $\alpha : \mathcal{S} \rightarrow \mathcal{P}(\mathcal{A})$ such that for all $s \in \mathcal{S}$, $\alpha(s) \neq \emptyset$, the reduced MDP \mathcal{M}' is defined as the 9-tuple $(\mathcal{S}, \mathcal{G}, \mathcal{A}, P, R, \rho_0, \rho_g, \alpha, \gamma)$.

where the assumption on $\alpha(s)$ ensures that there exist a feasible action in each state and Q-values can be well-defined. Intuitively, \mathcal{M}' encodes the same environment as \mathcal{M} but restricts the set of action-state pairs.

We note that in PAC RL settings, the analysis of sample complexity [64] (i.e. the number of steps for which a learned policy is not ϵ -optimal with high probability) produces upper bounds that are directly dependent on the number of state-action pairs [65]. In particular, model-free PAC-MDP algorithms can attain a sample complexity that is $\tilde{O}(N)$, where $N \leq |\mathcal{S}||\mathcal{A}|$ is the number of state-action pairs, and $\tilde{O}(\cdot)$ represents $O(\cdot)$ where logarithmic factors are ignored [66]. Learning in \mathcal{M}' instead of \mathcal{M} is thus desirable and could lead to near-linear improvements in sample efficiency as the number of infeasible actions grows. Crucially, under mild assumptions, the optimal policy for \mathcal{M}' can not only be retrieved more efficiently but also attains optimality in the original MDP \mathcal{M} (see Appendix A.5)¹.

We thus propose a practical modified Q-learning iteration on the original MDP \mathcal{M} , which is equivalent to performing Q-learning directly in the reduced MDP \mathcal{M}' . Given a transition (s, a, s', g, r) :

$$Q(s, a, g) \leftarrow (1 - \delta)Q(s, a, g) + \delta(r + \gamma \max_{a' \in \alpha(s)} Q(s', a', g)), \quad (4)$$

where the value of the next-state s' is only computed over feasible actions and δ is the learning rate.

Under common assumptions (i.e. infinite visitation of each state-action pair and well-behaved learning rate in tabular settings [67]), this algorithm converges to $Q_{\mathcal{M}'}^*$, from which we can easily extract $\pi_{\mathcal{M}'}^*(s, g) = \pi_{\mathcal{M}'}^*(s, g) = \arg \max_{a \in \alpha(s)} Q_{\mathcal{M}'}^*(s, a, g)$. For simplicity, we will from now on refer to $Q_{\mathcal{M}'}$ as Q .

In practice, following the goal-conditioned framework [62], we scale this algorithm by parameterizing $Q_{\theta}(s, a, g)$ through a neural network. Inspired by recent success in scaling Q-learning [68] to high-dimensional spaces [69]–[71] while reducing overestimation bias, we learn the parameters θ of the Q-function using Clipped Double Q-learning [71], which minimizes the temporal difference (TD) loss:

$$\mathcal{L}(\theta_j) = \mathbb{E}_{(s, a, s', g, r) \sim \mathcal{B}} [(y_j - Q_{\theta_j}(s_t, a_t, g))^2], \text{ with } \quad (5)$$

$$y_j = r + \gamma \min_{i=1,2} Q_{\theta_i}(s', \arg \max_{a_{t+1} \in \alpha(s_{t+1})} Q_{\theta_j}(s_{t+1}, a_{t+1}, g), g),$$

where $j \in \{1, 2\}$, and θ_j, θ'_j are the parameters for Q and target Q-networks respectively and where (s, a, s', g, r) tuples are sampled uniformly from an experience replay buffer [72] \mathcal{B} exploiting the off-policy nature of Q-learning. We collect experience following an ϵ -greedy exploration mechanism

on the feasible action set $\alpha(s) \subseteq \mathcal{A}$. We summarize our approach in Algorithm 1.

Algorithm 1: ELF-P

input Trained prior π_w^β , randomly initialized Q_θ and Q-target $Q_{\theta'}$ with $\theta = \theta'$, probability threshold ρ , learning rate η , replay buffer $\mathcal{D} = \emptyset$, soft update parameter μ .
for episode = 1, ... N **do**
 Sample $s \sim \rho_0, g \in \rho_g$.
 for step = 1, ... T **do**
 Compute feasible action set $\alpha(s_t)$ in (3), compute $Q_\theta(s_t, a, g)$ for each $a \in \mathcal{A}$.
 With probability ϵ sample $a_t \sim \mathcal{U}\{\alpha(s_t)\}$, else select $a_t = \arg \max_{a \in \alpha(s_{t+1})} Q_\theta(s_t, a, g)$.
 Execute a_t and store transition $(s_t, a_t, r_t, s_{t+1}, g)$ in \mathcal{D} .
 Sample minibatch of transitions $(s_j, a_j, r_j, s_{j+1}, g)$ uniformly from \mathcal{D} .
 Compute TD loss $\mathcal{L}_Q(\theta)$ in (5).
 Gradient step $\theta \leftarrow \theta - \eta \nabla \mathcal{L}_Q(\theta)$.
 Perform soft-update on $\theta' \leftarrow \mu\theta + (1 - \mu)\theta'$.
 end for
end for

V. RESULTS

In this Section, we empirically evaluate the performance of ELF-P against several baselines. We investigate whether ELF-P can retrieve optimal policies more efficiently than existing approaches that leverage prior data, and whether the number of infeasible actions that are attempted throughout the training process is reduced. We evaluate ELF-P on a variety of manipulation tasks in a simulated environment and then deploy the learned policies on real hardware to evaluate the ease of transfer. For further details on the environment, additional experiments and extended results, see Appendix A.1, A.2, A.3 and A.4.

We compare ELF-P with an unmodified DDQN [69], with DDQN with Hindsight Experience Replay (DDQN+HER) [34], with two other state-of-the-art methods that can leverage prior data, namely DDQNfD [73] (with 1-step returns and unprioritized experience replay for a fair comparison),




Fig. 3. An example of a task. Starting from an arbitrary configuration, the agent needs to place the block, which is hidden behind the cabinet door, in the mid drawer, which is blocked by the drawer above. The sequence of 19 actions required to achieve the goal is shown below.

¹See our website: nuria95.github.io/elf-p/ for all Supplementary Materials.




Fig. 4. Success rate and number of infeasible actions attempts for the Medium (left) and Hard (right) task variants. (Results are averaged across 10 independent random seeds, shaded area represents standard deviation).




Fig. 5. Comparison between ELF-P and Soft-ELF-P. ELF-P. Success rate and number of infeasible actions attempts for the Hard task variant. (Play dataset size 10^4). Results are averaged across 10 independent random seeds, shaded area represents standard deviation)

DDQN+Prefill, which initializes its replay buffer with the task-agnostic data \mathcal{D} , and with SPIRL [52], which also uses play data to constrain exploration, but adopts an actor-critic framework with soft prior regularization, which we adapted to operate over discrete action spaces. Extended details are in Appendix A.3.3.

A. Simulation Experiments

We train and evaluate the algorithm on a variety of simulated long-horizon manipulation tasks in which a robot needs to interact with a realistic desk environment featuring several household items. Possible behaviors include sliding a cabinet door, opening several drawers and moving a wooden block (see Figure 3). The tasks are defined in the same environment that play data was obtained in. While play data can be collected by humans, for simplicity we use a scripted policy (details are in Appendix A.3). The goal of each task is to place the block in an arbitrary desired position, which generally also requires manipulating the rest of the items. The episode ends when exceeding a predefined number of steps or when the goal is achieved. The reward function is sparse and is equal to one if and only if the task is completed in time, otherwise a reward of zero is given.

We evaluate on tasks distributions with two levels of difficulty: *Medium* (M) and *Hard* (H). The average number of actions required to solve the (M) tasks for an expert planner is 14, with the longest task requiring 16 steps. For the (H) tasks, the average is 23 steps, with the longest task requiring 29 steps. See Figure 3 for an example.

1) *Performance Analysis*: Figure 4 shows the success rate and the number of infeasible actions attempts averaged over 50 evaluation episodes across the two task distributions. In the (M) tasks ELF-P shows the same sample efficiency as the best baseline, while in the (H) tasks it is significantly better than all competitors. For both (M) and (H) tasks, the number of infeasible actions that are attempted by the agent

is significantly lower than other baselines. We note that the execution of infeasible actions is due to inaccuracies in the trained prior.

While vanilla DDQN cannot master the (M) task and fails to solve the (H) task, Prefill+DDQN and DDQNfD manage to solve the (M) task as efficiently as ELF-P, proving that having access to task-agnostic play data is beneficial for the learning process. However, when the action space grows in size, their performance decreases significantly. We also notice that although managing to reach over 0.5 success rate for both tasks, SPIRL shows a much lower performance than our method and most competitive baselines, hence integrating a soft prior via KL-regularization might not be beneficial in this setting. Finally, we observe that adding HER relabeling helps slightly on the (H) tasks but hurts performance on the (M) tasks. We report that using HER with ELF-P also hurts performance (see Appendix A.2). HER relies on a gradual growth of the frontier of achieved goals, which can be used for relabeling. Since the dynamics in our environment are not smooth (i.e. a single action often leads to large changes in the state), we hypothesize that HER cannot interpolate to unseen goals and hence hurts performance. This behavior was also pointed out by [6].

We note that all methods in this section that can leverage prior data have access to a play dataset with 10^4 datapoints, which simulates an amount of data that could in practice, be collected by a few human operators.

2) Robustness to Play

Dataset Size: We study

the effect of play dataset size on training performance for both ELF-P and DDQN+Prefill (which affects the amount of data used for training the prior for ELF-P and the amount of data used to pre-fill the replay buffer for DDQN+Prefill). In particular, we measure the number of timesteps required to reach a success rate of 0.95 on (H) tasks when using different dataset sizes. We report results in Figure 6. We observe that our method retains most of its performance when the prior is trained on minimal quantities of data whereas DDQN+Prefill struggles in low-data regimes. This shows that our method is more suited for the setting of *play*, and is able to learn efficiently with reasonable amounts




Fig. 6. Effect of play dataset size on sample complexity (DDQN+Prefill fails to solve the task with 1000 datapoints).

of data, i.e., data that could in practice be collected by a few human operators ($\sim 1\text{h}30\text{min}$ of interaction data). When more data is available, we generally find the performance gap between methods to decrease, while ELF-P consistently outperforms baselines in all data regimes. See Appendix A.2.3 for additional results.

3) Sample Complexity:

We further analyze the gains in sample complexity and compare them with the theoretical rates mentioned in Subsection IV-C, which hypothesize a linear dependency on the number of feasible action-state pairs in PAC-RL settings. In particular, we measure the number of timesteps required to reach a success rate of 0.95 on (M) tasks with increasing values for the threshold ρ , thus pruning the state-action space more aggressively. In Figure 7 we plot this quantity against the ratio of infeasible state-action pairs for each value of ρ , computed according to the visitation distribution of a random policy. While our algorithm recovers the performance of DDQN for $\rho = 0$, we indeed observe a linear trend in the decrease of sample complexity as the number of infeasible action-state pairs increases. Further details are in Appendix A.2.2.

4) *Soft Prior Integration*: We implement Soft-ELF-P, an ablation for our method which softens the integration of the prior into learning. This modification could potentially allow Soft-ELF-P to recover from degenerated priors, i.e., when play data has significant distribution mismatch with the target tasks of interest. Soft-ELF-P samples from the prior (instead of sampling from the feasible set α) both during initial exploration phase and while performing ϵ -greedy exploration. Additionally, during exploitation, it multiplies the softmax of Q-values by the prior, thus biasing the greedy action selection towards the prior. Finally, given the soft integration of the prior, it performs Q-learning over the set of all actions, instead of over the reduced set of feasible actions. In Figure 5 we compare the performance of ELF-P with Soft-ELF-P on the (H) task. Results on the (M) task are in the Appendix A.2.5. We observe that while Soft-ELF-P is able to learn both medium and hard tasks, it has lower sample efficiency than ELF-P. The hard prior integration of ELF-P alleviates the Q-network from learning values for all state-action pairs: it can focus on learning Q-values for feasible state-action pairs only and ignore Q-values for unfeasible actions. This shows one of the main contributions of our algorithm. While binarizing the prior results in greater learning efficiency, a soft integration could be useful when dealing with degenerated priors and we reserve further exploration on the topic for future work.

B. Real-world Experiments

In this experiment, we evaluate how ELF-P performs when transferred to physical hardware. Figure 1 (upper right)

depicts our physical robot workspace and further details can be found in the Appendix A.4. We use *Boston Dynamics Spot* robot [57] for all our experiments because of its stability and robustness when following desired end-effector trajectories. The high-level planner, trained in simulation, is used at inference time to predict the required sequence of motion primitives to achieve a desired goal. We then instantiate each of them using established motion planning methods [74] to create optimal control trajectories. For further implementation details, we refer the reader to [75].

We show that ELF-P can be easily transferred to physical hardware while maintaining its goal-reaching capabilities. In Figure 2 we show the agent executing several primitives in order to complete a task. We refer the reader to Appendix A.4 and the video material for extended visualizations.

VI. DISCUSSION AND FUTURE WORK

We present ELF-P, a method that bridges motion planning and deep RL to achieve complex long-horizon manipulation tasks. We show that by integrating a discrete behavioral prior learned from easily collectable play data, we can achieve significant gains in sample efficiency compared to other baselines that leverage prior data. This approach has the added benefit of largely avoiding infeasible actions during training. By planning in a two-level hierarchy, we show how our method allows reasoning over long-horizons in a mixed decision space in an efficient manner. We finally demonstrate that within this framework, ELF-P can be easily transferred to physical hardware without further modifications, showing the potential of combining readily available motion planners with sample efficient RL algorithms.

Despite showing promising results, our method assumes full observability of the state space and perfect execution of the motion primitives. These limitations could be addressed by introducing perception and by querying ELF-P at higher frequencies at inference time to ensure primitive completion. Furthermore, choosing a suitable level of abstractions for skills remains an open question, whose answer could relax the need for providing a predefined set of skills, while still maintaining a low-dimensional parametrization. Future work may also actively learn the behavioral prior instead of leveraging a static play dataset, reaching a compromise between sample complexity and reliance on collected data.

We expect our work to enable future research directions such as a tighter coupling between the training of the high-level planner and the execution of motion primitives. Although introducing motion planning in the training loop is time consuming, we believe that the significant gains in sample efficiency demonstrated in this work can help address this challenge.

ACKNOWLEDGMENT

Núria Armengol and Marco Bagatella are financially supported by the Max Planck ETH Center for Learning Systems. We thank Simon Zimmermann and Oliver Stark for the help throughout the project.




Fig. 7. Effect of ρ on sample complexity. Means across seeds are connected by lines.

REFERENCES

- [1] A. W. Winkler, "Optimization-based motion planning for legged robots," Ph.D. dissertation, ETH Zurich, 2018.
- [2] A. W. Winkler, C. D. Bellicoso, M. Hutter, and J. Buchli, "Gait and trajectory optimization for legged systems through phase-based end-effector parameterization," *IEEE Robotics and Automation Letters*, vol. 3, no. 3, pp. 1560–1567, 2018.
- [3] H. Dai, A. Valenzuela, and R. Tedrake, "Whole-body motion planning with centroidal dynamics and full kinematics," in *2014 IEEE-RAS International Conference on Humanoid Robots*, 2014, pp. 295–302.
- [4] M. Bjelonic, R. Grandia, M. Geilinger, O. Harley, V. S. Medeiros, V. Pajovic, E. Jelavic, S. Coros, and M. Hutter, "Offline motion libraries and online mpc for advanced mobility skills," *The International Journal of Robotics Research*, 2022.
- [5] P. Sermanet, C. Lynch et al., "Broadly-exploring, local-policy trees for long-horizon task planning," in *5th Annual Conference on Robot Learning*, 2021.
- [6] B. Eysenbach, R. R. Salakhutdinov, and S. Levine, "Search on the replay buffer: Bridging planning and reinforcement learning," *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [7] J. Mainprice, N. Ratliff, M. Toussaint, and S. Schaal, "An interior point method solving motion planning problems with narrow passages," in *2020 29th IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*, 2020, pp. 547–552.
- [8] D. Kalashnikov, J. Varley, Y. Chebotar, B. Swanson, R. Jonschkowski, C. Finn, S. Levine, and K. Hausman, "Mt-opt: Continuous multi-task robotic reinforcement learning at scale," *arXiv preprint arXiv:2104.08212*, 2021.
- [9] I. Akkaya, M. Andrychowicz, M. Chociej, M. Litwin, B. McGrew, A. Petron, A. Paino, M. Plappert, G. Powell, R. Ribas et al., "Solving rubik's cube with a robot hand," *arXiv preprint arXiv:1910.07113*, 2019.
- [10] O. M. Andrychowicz, B. Baker, M. Chociej, R. Jozefowicz, B. McGrew, J. Pachocki, A. Petron, M. Plappert, G. Powell, A. Ray et al., "Learning dexterous in-hand manipulation," *The International Journal of Robotics Research*, vol. 39, no. 1, pp. 3–20, 2020.
- [11] D. Kalashnikov, A. Irpan, P. Pastor, J. Ibarz, A. Herzog, E. Jang, D. Quillen, E. Holly, M. Kalakrishnan, V. Vanhoucke et al., "Qt-opt: Scalable deep reinforcement learning for vision-based robotic manipulation (2018)," *arXiv preprint arXiv:1806.10293*, 2018.
- [12] O. Nachum, S. S. Gu, H. Lee, and S. Levine, "Data-efficient hierarchical reinforcement learning," *Advances in neural information processing systems*, vol. 31, 2018.
- [13] A. Levy, R. Platt, and K. Saenko, "Hierarchical reinforcement learning with hindsight," *arXiv preprint arXiv:1805.08180*, 2018.
- [14] E. Coumans and Y. Bai, "Pybullet, a python module for physics simulation for games, robotics and machine learning," <http://pybullet.org>, 2016–2021.
- [15] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [16] L. E. Kavraki, P. Svestka, J.-C. Latombe, and M. H. Overmars, "Probabilistic roadmaps for path planning in high-dimensional configuration spaces," *IEEE transactions on Robotics and Automation*, vol. 12, no. 4, pp. 566–580, 1996.
- [17] S. M. LaValle and J. J. Kuffner Jr, "Randomized kinodynamic planning," *The international journal of robotics research*, vol. 20, no. 5, pp. 378–400, 2001.
- [18] S. Karaman and E. Frazzoli, "Sampling-based algorithms for optimal motion planning," *The international journal of robotics research*, vol. 30, no. 7, pp. 846–894, 2011.
- [19] N. Ratliff, M. Zucker, J. A. Bagnell, and S. Srinivasa, "Chomp: Gradient optimization techniques for efficient motion planning," in *2009 IEEE International Conference on Robotics and Automation*, 2009, pp. 489–494.
- [20] J. Schulman, Y. Duan, J. Ho, A. Lee, I. Awwal, H. Bradlow, J. Pan, S. Patil, K. Goldberg, and P. Abbeel, "Motion planning with sequential convex optimization and convex collision checking," *The International Journal of Robotics Research*, vol. 33, no. 9, pp. 1251–1270, 2014.
- [21] L. P. Kaelbling and T. Lozano-Pérez, "Integrated task and motion planning in belief space," *The International Journal of Robotics Research*, vol. 32, no. 9-10, pp. 1194–1227, 2013.
- [22] M. Toussaint, "Logic-geometric programming: An optimization-based approach to combined task and motion planning," in *Twenty-Fourth International Joint Conference on Artificial Intelligence*, 2015.
- [23] M. Toussaint and M. Lopes, "Multi-bound tree search for logic-geometric programming in cooperative manipulation domains," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*, 2017, pp. 4044–4051.
- [24] C. R. Garrett, R. Chitnis, R. Holladay, B. Kim, T. Silver, L. P. Kaelbling, and T. Lozano-Pérez, "Integrated task and motion planning," *Annual review of control, robotics, and autonomous systems*, vol. 4, pp. 265–293, 2021.
- [25] C. V. Braun, J. Ortiz-Haro, M. Toussaint, and O. S. Oguz, "Rhh-lgp: Receding horizon and heuristics-based logic-geometric programming for task and motion planning," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 13 761–13 768.
- [26] Y. Duan, M. Andrychowicz, B. Stadie, O. Jonathan Ho, J. Schneider, I. Sutskever, P. Abbeel, and W. Zaremba, "One-shot imitation learning," *Advances in neural information processing systems*, vol. 30, 2017.
- [27] A. G. Barto and S. Mahadevan, "Recent advances in hierarchical reinforcement learning," *Discrete event dynamic systems*, vol. 13, no. 1, pp. 41–77, 2003.
- [28] P. Dayan and G. E. Hinton, "Feudal reinforcement learning," *Advances in neural information processing systems*, vol. 5, 1992.
- [29] M. Wiering and J. Schmidhuber, "Hq-learning," *Adaptive Behavior*, vol. 6, no. 2, pp. 219–246, 1997.
- [30] R. S. Sutton, D. Precup, and S. Singh, "Between mdps and semi-mdps: A framework for temporal abstraction in reinforcement learning," *Artificial intelligence*, vol. 112, no. 1-2, pp. 181–211, 1999.
- [31] T. G. Dietterich, "Hierarchical reinforcement learning with the maxq value function decomposition," *Journal of artificial intelligence research*, vol. 13, pp. 227–303, 2000.
- [32] S. Nasiriany, H. Liu, and Y. Zhu, "Augmenting reinforcement learning with behavior primitives for diverse manipulation tasks," in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 7477–7484.
- [33] L. P. Kaelbling, "Learning to achieve goals," in *IN PROC. OF IJCAI-93*, 1993, pp. 1094–1098.
- [34] M. Andrychowicz, F. Wolski, A. Ray, J. Schneider, R. Fong, P. Welinder, B. McGrew, J. Tobin, O. Pieter Abbeel, and W. Zaremba, "Hindsight experience replay," *Advances in neural information processing systems*, vol. 30, 2017.
- [35] A. Gupta, V. Kumar, C. Lynch, S. Levine, and K. Hausman, "Relay policy learning: Solving long horizon tasks via imitation and reinforcement learning," *Conference on Robot Learning (CoRL)*, 2019.
- [36] S. Christen, L. Jendele, E. Aksan, and O. Hilliges, "Learning functionally decomposed hierarchies for continuous control tasks with path planning," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 3623–3630, 2021.
- [37] M. Hausknecht and P. Stone, "Deep reinforcement learning in parameterized action space," in *Proceedings of the International Conference on Learning Representations (ICLR)*, San Juan, Puerto Rico, May 2016. [Online]. Available: <http://www.cs.utexas.edu/users/ai-lab/hausknecht/iclr16>
- [38] M. Dalal, D. Pathak, and R. R. Salakhutdinov, "Accelerating robotic reinforcement learning via parameterized action primitives," *Advances in Neural Information Processing Systems*, vol. 34, pp. 21 847–21 859, 2021.
- [39] W. Masson, P. Ranchod, and G. Konidaris, "Reinforcement learning with parameterized actions," in *Thirtieth AAAI Conference on Artificial Intelligence*, 2016.
- [40] F. Xia, C. Li, R. Martín-Martín, O. Litany, A. Toshev, and S. Savarese, "Relmogen: Integrating motion generation in reinforcement learning for mobile manipulation," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*, 2021, pp. 4583–4590.
- [41] J. Yamada, Y. Lee, G. Salhotra, K. Pertsch, M. Pflueger, G. Sukhatme, J. Lim, and P. Englert, "Motion planner augmented reinforcement learning for robot manipulation in obstructed environments," in *Proceedings of the 2020 Conference on Robot Learning*, ser. Proceedings of Machine Learning Research, J. Kober, F. Ramos, and C. Tomlin, Eds., vol. 155. PMLR, 16–18 Nov 2021, pp. 589–603.
- [42] D. Angelov, Y. Hristov, M. Burke, and S. Ramamoorthy, "Composing diverse policies for temporally extended tasks," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 2658–2665, 2020.
- [43] O. Vinyals, T. Ewalds, S. Bartunov, P. Georgiev, A. Vezhnevets, M. Yeo, A. Makhzani, H. Küttler, J. Agapiou, J. Schrittwieser

- et al., “A new challenge for reinforcement learning,” *arXiv preprint ArXiv:1708.04782*, vol. 5, 2017.
- [44] C. Berner, G. Brockman, B. Chan, V. Cheung, P. Dkebiak, C. Denison, D. Farhi, Q. Fischer, S. Hashme, C. Hesse et al., “Dota 2 with large scale deep reinforcement learning,” *arXiv preprint arXiv:1912.06680*, 2019.
- [45] D. Ye, Z. Liu, M. Sun, B. Shi, P. Zhao, H. Wu, H. Yu, S. Yang, X. Wu, Q. Guo et al., “Mastering complex control in moba games with deep reinforcement learning,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 34, 2020, pp. 6672–6679.
- [46] C. Boutilier, A. Cohen, A. Hassidim, Y. Mansour, O. Meshi, M. Mladenov, and D. Schuurmans, “Planning and learning with stochastic action sets,” in *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI-18*, 2018, pp. 4674–4682.
- [47] T. Zahavy, M. Haroush, N. Merlis, D. J. Mankowitz, and S. Mannor, “Learn what not to learn: Action elimination with deep reinforcement learning,” *Advances in Neural Information Processing Systems*, vol. 31, 2018.
- [48] E. Even-Dar, S. Mannor, and Y. Mansour, “Action elimination and stopping conditions for reinforcement learning,” in *Proceedings of the 20th International Conference on Machine Learning (ICML-03)*, 2003, pp. 162–169.
- [49] J. J. Gibson, “The theory of affordances,” *Hilldale, USA*, vol. 1, no. 2, pp. 67–82, 1977.
- [50] K. Khetarpal, Z. Ahmed, G. Comanici, D. Abel, and D. Precup, “What can i do here? a theory of affordances in reinforcement learning,” in *ICML*, 2020.
- [51] R. Costales, S. Iqbal, and F. Sha, “Possibility before utility: Learning and using hierarchical affordances,” in *International Conference on Learning Representations*, 2022.
- [52] K. Pertsch, Y. Lee, and J. J. Lim, “Accelerating reinforcement learning with learned skill priors,” in *Conference on Robot Learning (CoRL)*, 2020.
- [53] K. Pertsch, Y. Lee, Y. Wu, and J. J. Lim, “Demonstration-guided reinforcement learning with learned skills,” *5th Conference on Robot Learning*, 2021.
- [54] A. Singh, H. Liu, G. Zhou, A. Yu, N. Rhinehart, and S. Levine, “Parrot: Data-driven behavioral priors for reinforcement learning,” in *International Conference on Learning Representations*, 2021.
- [55] D. Tirumala, A. Galashov, H. Noh, L. Hasenclever, R. Pascanu, J. Schwarz, G. Desjardins, W. M. Czarnecki, A. Ahuja, Y. W. Teh et al., “Behavior priors for efficient reinforcement learning,” *Journal of Machine Learning Research*, vol. 23, no. 221, pp. 1–68, 2022.
- [56] C. Lynch, M. Khansari, T. Xiao, V. Kumar, J. Tompson, S. Levine, and P. Sermanet, “Learning latent plans from play,” *Conference on Robot Learning (CoRL)*, 2019.
- [57] “Spot® | Boston Dynamics,” <https://www.bostondynamics.com/products/spot>, 2022.
- [58] M. Ahn, A. Brohan, N. Brown, Y. Chebotar, O. Cortes, B. David, C. Finn, C. Fu, K. Gopalakrishnan, K. Hausman, A. Herzog, D. Ho, J. Hsu, J. Ibarz, B. Ichter, A. Irpan, E. Jang, R. J. Ruano, K. Jeffrey, S. Jesmonth, N. Joshi, R. Julian, D. Kalashnikov, Y. Kuang, K.-H. Lee, S. Levine, Y. Lu, L. Luu, C. Parada, P. Pastor, J. Quiambao, K. Rao, J. Rettinghouse, D. Reyes, P. Sermanet, N. Sievers, C. Tan, A. Toshev, V. Vanhoucke, F. Xia, T. Xiao, P. Xu, S. Xu, M. Yan, and A. Zeng, “Do as i can and not as i say: Grounding language in robotic affordances,” in *arXiv preprint arXiv:2204.01691*, 2022.
- [59] Y. Zhu, P. Stone, and Y. Zhu, “Bottom-up skill discovery from unsegmented demonstrations for long-horizon robot manipulation,” *IEEE Robotics and Automation Letters*, 2022.
- [60] G. Konidaris, S. Kuindersma, R. Grupen, and A. Barto, “Constructing skill trees for reinforcement learning agents from demonstration trajectories,” *Advances in neural information processing systems*, vol. 23, 2010.
- [61] S. Niekum, S. Osentoski, C. G. Atkeson, and A. G. Barto, “Online bayesian changepoint detection for articulated motion models,” in *2015 IEEE International Conference on Robotics and Automation (ICRA)*, 2015, pp. 1468–1475.
- [62] T. Schaul, D. Horgan, K. Gregor, and D. Silver, “Universal value function approximators,” in *International conference on machine learning*, 2015, pp. 1312–1320.
- [63] M. L. Puterman, *Markov Decision Processes*. Wiley, 1994.
- [64] S. M. Kakade, *On the sample complexity of reinforcement learning*. University of London, University College London (United Kingdom), 2003.
- [65] T. Lattimore and M. Hutter, “Pac bounds for discounted mdps,” in *International Conference on Algorithmic Learning Theory*, 2012, pp. 320–334.
- [66] A. L. Strehl, L. Li, E. Wiewiora, J. Langford, and M. L. Littman, “Pac model-free reinforcement learning,” in *Proceedings of the 23rd international conference on Machine learning*, 2006, pp. 881–888.
- [67] D. P. Bertsekas and J. N. Tsitsiklis, *Neuro-dynamic programming*, ser. Optimization and neural computation series. Athena Scientific, 1996, vol. 3.
- [68] C. J. Watkins and P. Dayan, “Q-learning,” *Machine learning*, vol. 8, no. 3, pp. 279–292, 1992.
- [69] H. Van Hasselt, A. Guez, and D. Silver, “Deep reinforcement learning with double q-learning,” in *Proceedings of the AAAI conference on artificial intelligence*, vol. 30, 2016.
- [70] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski et al., “Human-level control through deep reinforcement learning,” *Nature*, vol. 518, no. 7540, pp. 529–533, 2015.
- [71] S. Fujimoto, H. Hoof, and D. Meger, “Addressing function approximation error in actor-critic methods,” in *International conference on machine learning*, 2018, pp. 1587–1596.
- [72] L.-J. Lin, *Reinforcement learning for robots using neural networks*. Carnegie Mellon University, 1992.
- [73] T. Hester, M. Vecerik, O. Pietquin, M. Lanctot, T. Schaul, B. Piot, D. Horgan, J. Quan, A. Sendonaris, I. Osband et al., “Deep q-learning from demonstrations,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 32, 2018.
- [74] S. Zimmermann, M. Busenhardt, S. Huber, R. Poranne, and S. Coros, “Differentiable collision avoidance using collision primitives,” in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 8086–8093.
- [75] S. Zimmermann, R. Poranne, and S. Coros, “Go fetch!-dynamic grasps using boston dynamics spot with external robotic arm,” in *2021 IEEE International Conference on Robotics and Automation (ICRA)*, 2021, pp. 4488–4494.