

Received 16 March 2024, accepted 9 May 2024, date of publication 16 May 2024, date of current version 23 May 2024.

Digital Object Identifier 10.1109/ACCESS.2024.3402089

## RESEARCH ARTICLE

# Bidirectional Progressive Neural Networks With Episodic Return Progress for Emergent Task Sequencing and Robotic Skill Transfer

SUZAN ECE ADA<sup>1</sup>, HANNE SAY<sup>2</sup>, EMRE UGUR<sup>1</sup>, AND ERHAN OZTOP<sup>2,3</sup>, (Member, IEEE)

<sup>1</sup>Department of Computer Engineering, Boğaziçi University, 34342 İstanbul, Turkey

<sup>2</sup>Department of Computer Science, Özyeğin University, 34794 İstanbul, Turkey

<sup>3</sup>SISREC, OTRI, Osaka University, Osaka 565-0871, Japan

Corresponding author: Suzan Ece Ada (ece.ada@bogazici.edu.tr)

This research has been supported by Japan Society for the Promotion of Science (JSPS) KAKENHI Grant Number JP23K24926 and the project JPNP16007 commissioned by the New Energy and Industrial Technology Development Organization (NEDO). In addition, partial support has been given by the INVERSE project (no. 101136067), funded by the European Union.

**ABSTRACT** Human brain and behavior provide a rich venue that can inspire novel control and learning methods for robotics. In an attempt to exemplify such a development by inspiring how humans acquire knowledge and transfer skills among tasks, we introduce a novel multi-task reinforcement learning framework named Episodic Return Progress with Bidirectional Progressive Neural Networks (ERP-BPNN). The proposed ERP-BPNN model 1) learns in a human-like interleaved manner by 2) autonomous task switching based on a novel intrinsic motivation signal and, in contrast to existing methods, 3) allows bidirectional skill transfer among tasks. ERP-BPNN is a general architecture applicable to several multi-task learning settings; in this paper, we present the details of its neural architecture and show its ability to enable effective learning and skill transfer among morphologically different robots in a reaching task. The developed Bidirectional Progressive Neural Network (BPNN) architecture enables bidirectional skill transfer without requiring incremental training and seamlessly integrates with online task arbitration. The task arbitration mechanism developed is based on soft Episodic Return progress (ERP), a novel intrinsic motivation (IM) signal. To evaluate our method, we use quantifiable robotics metrics such as ‘expected distance to goal’ and ‘path straightness’ in addition to the usual reward-based measure of episodic return common in reinforcement learning. With simulation experiments, we show that ERP-BPNN achieves faster cumulative convergence and improves performance in all metrics considered among morphologically different robots compared to the baselines. Overall, our method provides a human-inspired and efficient multi-task reinforcement learning approach with interleaved learning, making it highly suitable for lifelong learning applications.

**INDEX TERMS** Multi-task learning, reinforcement learning, robot learning, intrinsic motivation, knowledge transfer, deep learning, cognitive robotics, transfer learning.

## I. INTRODUCTION

Developing robots capable of autonomous and continual learning effectively requires the exploitation of acquired knowledge without human intervention, which could be described as the main goal of lifelong robot learning [1], [2], [3]. A key feature of human learning is autonomous task

The associate editor coordinating the review of this manuscript and approving it for publication was Alessia Saggese<sup>1</sup>.

switching and interleaved learning, which is not addressed in mainstream machine learning and robotics research [4]. Deep reinforcement learning (RL) is an area of machine learning that leverages deep learning to make decisions by learning from real-time or recorded interactions with the environment. Traditional deep RL methods often do not include mechanisms to exploit the potential benefits of interleaved learning and bidirectional skill transfer based on partial learning. In contrast to prior works, our proposed

model integrates human-like interleaved learning, leverages intrinsic motivation for autonomous task-switching, and includes a novel bidirectional progressive architecture tailored for deep multitask reinforcement learning. In sum, we aim to help fill this gap by developing a multi-task reinforcement learning framework that can sustain and, importantly, benefit from interleaved task learning and allow bidirectional skill transfer among tasks.

In humans, it is shown that interleaved learning yields improved recall of information and better memory retention in the long run rather than blocked learning [5], [6], [7]. Supporting this behavioral data, the human brain is endowed with mechanisms against task interference and forgetting, such as internal rehearsal of experiences and memory consolidation [8], [9]. To enable interleaved learning, the question of when and which task to engage during multi-task learning must be answered. Developmental learning offers some inspiration: during learning, an infant autonomously decides what to do or play without external directions dictating what task (s)he needs to engage in. This behavior of infants is usually associated with the notion of *intrinsic motivation* (IM), which guides behavior through a putative internal reward system [10]. IM can be based on curiosity, novelty, learning progress (LP), or even challenge; as such, it is also used in robotics as LP [11], curiosity [12], and in machine learning as novelty [13], or surprise [14]. In the current study, we adopt an IM approach and propose a novel learning progress (LP) signal for RL tasks that guides task switching. Overall, inspired by the above discussion, we aim to develop a multi-task reinforcement learning (RL) framework with autonomous task switching, which can benefit from interleaved task learning without suffering from task interference. We argue that such a learning system may benefit a wide range of robot learning scenarios ranging from human-like learning for a social robot to skill transfer applications among morphologically different robots.

While significant progress has been made in deep RL for robotics, most approaches [15] have focused on transferring skills between robots with identical action spaces [16], [17] and tasks with pixel-level state inputs [18], [19], [20]. Nevertheless, achieving multi-task reinforcement learning between robots with different physical structures provides insights for human-inspired reinforcement learning and acts as a critical driving force for machine learning research [21]. This is particularly significant due to the different state and action spaces, providing a unique perspective on knowledge generalization. In light of this, we choose learning and transferring skills among morphologically different robots as the target to address with the the developed framework.

One of the key challenges in lifelong learning is catastrophic interference/forgetting, which needs to be considered if a robot is to learn continually. When novel instances to be learned diverge greatly from previously observed ones, new information may overwrite the already acquired knowledge by modifying the representations shared among multiple tasks, leading to catastrophic forgetting. To minimize the

interference while learning a novel task, in the literature, several techniques have been proposed, such as restricting the probable update(s) on the network parameters, dynamic resource allocation, or rehearsing the old task samples while learning the new ones [1]. In this study, similar to the Progressive Neural Networks (PNN) [22], we utilize task-specific resources while learning to prevent possible task interference but also allow bidirectional inter-task connectivity to support positive skill transfer.

In sum, to enable human-like interleaved multi-task learning while avoiding task interference, we develop a multi-task reinforcement learning system composed of (1) a novel architecture called BPNN that improves the PNN [22], [23] and (2) a novel Intrinsic Motivation signal, Episodic Return Progress (ERP), for task-switching. Unlike the PNN architecture, which restricts skill transfer to the forward direction and requires learning previous tasks until convergence before transferring to the next task, our BPNN method enables bidirectional skill transfer during training. This means that skill transfer can happen in midway along multi-task learning among all tasks without the necessity of one task waiting for another to finish. The ERP signal evaluates task progress based on episodic return values, detecting the task that significantly contributes to enhancing overall performance across multiple tasks. The efficacy of the proposed multi-task learning framework is shown by its application to the learning of reaching skill by morphologically different robots, namely two degrees of freedom (2-DoF), 3-DoF, and 4-DoF manipulators (Figure 1). The conducted systemic experiments show that synergistic multi-task learning is possible due to bidirectional inter-task skill transfer provided by the proposed BPNN architecture and the ERP-based task switching.

The rest of the paper is organized as follows. In Section II, we present an overview of the related studies present in the literature. Then, we describe our method in detail in Section III, providing metrics used to evaluate the performance of our proposed method. Section IV details the experiments and presents the results for skill transfer between morphologically different reacher robots. Finally, we discuss the broader impact of our method and future research directions in Section V.




FIGURE 1. (a) 2-DoF, (b) 3-DoF, (c) 4-DoF reacher robot arm environments.

## II. RELATED WORK

### A. HUMAN LEARNING

Numerous examples demonstrate how neuroscience and artificial intelligence have paved the way for each other [24], [25], [26]. In this vein, the ability of humans to acquire multiple skills with ease throughout their lives may guide

machine learning research in multi-task learning and lifelong learning fronts. Human learning, especially during infancy and childhood, is characterized by autonomous engagement in play, i.e., exploration, where no task is learned completely in one sitting. Besides being ecologically unreasonable to focus on learning a single task until mastery, interleaved learning may allow positive skill transfer among tasks from partial learning if adequate mechanisms are engaged. This notion is supported by the contextual-interference (CI) effect studies showing that practicing tasks in an interleaved regime often results in improved learning compared to practicing in a block order [27]. Benefits of CI have been linked to increased brain activity during interleaved practice as opposed to repetitive practice [28], [29]. In addition, the CI effect not only improves information retention but also skill transfer between similar tasks [30].

On the other hand, machine learning settings usually prefer blocked learning to interleaved learning. Multi-task learning settings generally assume that either a random task (or a subset of tasks) is chosen for each training trial or that training proceeds to the next task after one task is mastered with a few exceptions [31], [32]. In the case of continual learning, tasks are learned sequentially as each task arrives [33]. This is in contrast with how humans learn. It is clearly seen that the beneficial impacts of interleaved learning have not been thoroughly examined within the machine learning literature thus far. Investigating these effects has the potential to enhance the alignment of artificial intelligence with the underlying mechanisms of the human brain.

### B. INTRINSIC MOTIVATION

Intrinsic motivation (IM) is a significant topic in infant cognitive development and learning, which refers to motivation originating from innate satisfaction instead of the extrinsic reward gained from the environment [34], [35]. IM has been adopted for enabling open-ended robot learning [36], and improving self-supervised exploration [37], [38]. Intrinsic rewards can arise from exploring novel states, satisfying an ingrained curiosity, or the rate of acquiring skills and knowledge in an environment. Inclusion of intrinsic rewards to the extrinsic rewards from the environment is one way of solving challenging exploration problems [39] and discovering a diverse set of skills [40] in deep reinforcement learning. In our approach, unlike the existing uses of IM in reinforcement learning, we propose to utilize episodic return progress (ERP) as a higher-level IM signal to dynamically switch tasks for learning in an online manner instead of modulating the reward signal guiding RL. Consequently, the dynamic task switching carried out by the task selection mechanism leads to an emergent interleaved multi-task learning regime.

### C. CURRICULUM REINFORCEMENT LEARNING

Curriculum learning methods focus on discovering a goal or task sequencing procedure that can lead to a faster convergence during training or improved performance

compared to random sequencing [41]. Most curriculum learning techniques require a priori domain knowledge of tasks to distinguish task levels to train from easier tasks to more difficult tasks [42], [43], [44]. For example, for predicting the output of a short Python code with Long Short Term Memory (LSTM) networks, the task levels can be identified by the number of nestings and the number of digits in the integers [45]. After manually identifying these task difficulty measures, it can be demonstrated that while training, a combination of a random curriculum and a naive curriculum where tasks are selected in ascending order of difficulty performs better than using only a random or a naive curriculum [45]. However, the difficulty of each task might not be readily available a priori in the robotics domain, and thus an automatic or emergent curriculum formation can be desirable. Deep Q-Networks (DQN) with prioritized experience replay [46] assigns importance to transitions based on their associated temporal difference error, thereby selecting data for more frequent replay in single-task reinforcement learning. On the other hand, our method prioritizes which task network is allowed to learn in an online manner. Hence, it operates at a higher level than DQN with prioritized experience replay's prioritization scheme and creates emergent interleaved task-switching patterns on the fly. As such, the learning scheduling obtained is quite different from what can be obtained through a transition-level curriculum or a usual task curriculum where each task is learned to completion.

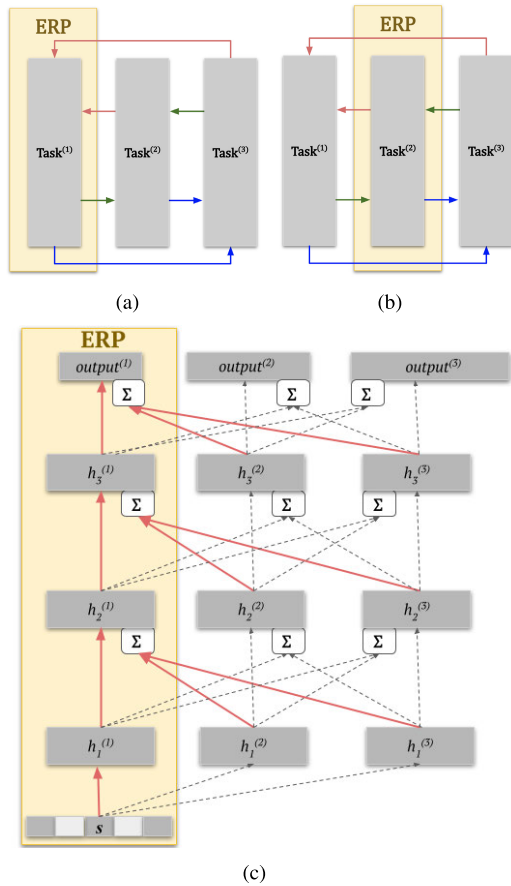
### D. MULTI-TASK REINFORCEMENT LEARNING

Multi-task learning involves sharing skills and knowledge between multiple tasks, where each task is identified as either a source or a target task during training. Typically, tasks share a part of the neural network model, and the model integrates a task conditioning parameter that defines the task during training. PathNet is a technique that uses a tournament selection genetic algorithm to evolve pathways of a neural network for multi-task, lifelong, and forward transfer learning [47]. However, Pathnet is trained consecutively for reinforcement learning tasks, meaning that the source task needs to be trained until convergence before moving on to the target task. This procedure does not allow backward transfer of tasks. Similarly, in PNN [22], training the source task until convergence is required to transfer skills to other tasks connected to the trained task.

Current state-of-the-art methods do not consider human-inspired interleaved learning as a viable approach for multi-task learning, yet there are potential benefits of adopting a human-like learning strategy. In this vein, ERP-BPNN supports bidirectional transfer and does not require convergence in one task to use previously learned representations in other tasks. Hence, ERP-BPNN fits better into the multi-task learning framework where all tasks are source and target tasks throughout training. This is essential because the BPNN architecture allows for the integration of learning progress and bidirectional transfer.

### III. METHOD

We propose a novel multi-task reinforcement learning framework that integrates a bidirectional progressive neural network (BPNN) with a unique architecture and soft task-switching mechanism crafted for RL, inspired by our previous work [4]. The BPNN architecture consists of bidirectional lateral connections among the hidden layers of each fully connected task network to allow skill transfer. By allocating a separate network module for each task, the model avoids negative task transfer at the core level but allows potential positive transfer to take place due to the bidirectional lateral connections among networks.



**FIGURE 2.** A high-level (a,b) ERP-BPNN architecture demonstrates ERP selecting Task 1 (a) then Task 2 (b), showcasing bidirectional flow for skill transfer in a many-to-many fashion among three tasks. Graphical representation of ERP-BPNN framework with ERP task switching, wherein (c) Task 1 is selected to learn. Weight updates are denoted by red (Task 1) arrows. Dashed gray arrows indicate no gradient flow during the learning update. In the current report, Task 1, 2, and 3 refer to RL tasks for 2-Dof, 3-Dof, and 4-Dof Reacher Robot arms as illustrated in Fig. 1 (a), (b), (c) respectively.

#### A. BIDIRECTIONAL PROGRESSIVE NEURAL NETWORKS

We initialize a fully connected neural network module  $m \in M$  with task parameters  $\theta_m$  for each task  $\mathcal{T}_i$  with index  $i$ . As each task corresponds to a module, we will use “network module” and “task” terms interchangeably. Since each robot controls different numbers of joints, each network module has

task-specific output layers with a different number of neurons  $n_{l=L}^m$  where  $l \in \{1..L\}$  refers to the layer index.

During training, each network module receives the input of the task selected for training. In this way, the lateral activations can receive representations of other tasks for skill transfer. Subsequently, we compute the hidden activations  $h_l^{(m)} \in \mathbb{R}^{n_l^m}$

$$h_l^{(m)} = f\left(W_l^{(m)} h_{l-1}^{(m)} + b_l^{(m)} + \sum_{t \neq m} U_l^{(m:t)} h_{l-1}^{(t)} + b_l^{(m:t)}\right) \quad (1)$$

where,  $U_l^{(m:t)}$ ,  $b_l^{(m:t)}$  denote the weight matrix and the bias vector corresponding to the lateral connections of the ordered pair of modules  $(m, t)$  from the previous layer  $h_{l-1}^{(t)}$  of the module  $t$  to the  $l^{\text{th}}$  layer of module  $m$ .  $W_l^{(m)} \in \mathbb{R}^{n_l^m \times n_{l-1}^m}$  and  $b_l^{(m)}$  are the weight matrix and bias vector, respectively, for layer  $l$  of module  $m$ , and  $f$  is the element-wise activation function. In the BPNN architecture, linear layers are utilized to potentially transmit information from other tasks, represented by  $\Sigma$  in Fig. 2(c), functioning as a module for summation. The linear transformations applied to the previous layers of other tasks allow the tuning of the incoming lateral signals. In particular, a negative transfer can be suppressed, and a positive transfer can be enhanced by the tuning of lateral weights through gradient descent-based learning. In this work, we use the activation function  $\tanh(x) = \frac{e^x - e^{-x}}{e^x + e^{-x}}$  following the suggested hyperparameters for Proximal Policy Optimization (PPO) [48] for continuous control. In the initial phase, the lateral connections of each module are frozen, and all tasks are individually trained to jumpstart task-specific learning for a predetermined number of  $K_{init}$  training iterations. In the experiments reported in Section IV, we set  $K_{init} = 20$ . In the subsequent learning steps, the parameters of the task modules that are not selected for training are frozen to avoid negative transfer between tasks. This prevents the gradient flow into the networks associated with the remaining tasks. On the other hand, the parameters of and the lateral connections incoming to the task module selected for training are unfrozen to facilitate task-specific learning as well as the skill transfer from other tasks. At a high level, this training mechanism and architecture enables bidirectional information transfer by allowing gradient flow through a subset of neural network parameters associated with the task and integrating information from other tasks. How the dynamic task selection takes place is described next.

#### B. TASK SWITCHING BY A NOVEL INTRINSIC MOTIVATION SIGNAL: AVERAGE SOFT EPISODIC RETURN PROGRESS

We propose a task-switching mechanism for multi-task reinforcement learning based on a novel Intrinsic Motivation (IM) signal, namely Average Soft Episodic Return Progress (ERP), that captures the learning progress of an agent in the RL context. At each optimization iteration  $k$ , for each task network module  $m$ , we record the expected discounted cumulative reward or average episodic return denoted by  $R_m(k)$ . To compute  $R_m(k)$ , we first normalize the immediate rewards to ensure the exponential moving average of the



rewards has a constant variance and clip them between  $(-10,10)$  to stabilize training [49]. Then we compute the mean of the returns over  $P * \kappa$  trajectories where  $P = 8$  is the number of parallel RL environments, and  $\kappa = 2$  is the number of episodes completed in each RL environment. Then, we take ERP for task  $m$  at step  $k$ ,  $ERP_m(k)$ , as the slope of the line that is fitted to  $R_m(k-w+1), R_m(k-w+2), \dots, R_m(k)$  using least squares, where  $w$  is a predetermined window size ( $w = 5$  in the experiments reported in this paper). The least squares solution is given by

$$ERP_m(k) = \frac{w \left( \sum_{i=0}^{w-1} X_k[i] Y_m(k)[i] \right) - \sum_{i=0}^{w-1} X_k[i] \sum_{i=0}^{w-1} Y_m(k)[i]}{w \left( \sum_{i=0}^{w-1} (X_k[i])^2 \right) - \left( \sum_{i=0}^{w-1} X_k[i] \right)^2} \quad (2)$$

where  $X_k = [k-w+1, k-w+2, \dots, k]$  and  $Y_m(k) = [R_m(k-w+1), R_m(k-w+2), \dots, R_m(k)]$ .

For bootstrapping ERP computation for each task, at the beginning of a multi-task learning session, each task is given an initial ERP bootstrapping run of  $K_{init} > w$  iterations. Then, we compute the ERP for each module individually for the initial run and subsequently at each iteration to select the task that has made the highest recent progress. The objective of the ERP procedure is to select tasks dynamically by identifying the task with the highest recent progress so that the selected task can continue its rapid progress. Selecting the most efficient task for training is particularly important as it can facilitate skill transfer to the other tasks, utilizing the bidirectional lateral connections of BPNN. For instance, the tasks that have made less progress can benefit from the steeper improvement of the selected task. Since ERP is computed for all tasks at each iteration, other tasks can benefit from the top-performing task and increase their chances of getting selected. The window size should be chosen to balance the reduction of noise and the tracking of recent updates. A large window size,  $w$ , can lead to selecting the initially top-performing task for an extended number of iterations. Correspondingly, a small window size can introduce noise during training. Thus, to monitor the recent changes in the progress of all tasks, we empirically selected the window size as five based on a grid search. Consecutive selection of the same task will eventually reach a plateau during training either due to the fact that the task is learned or a local minimum is encountered, where learning in other tasks may help the plateaued task to jumpstart as the reduction in progress of the current task allows other tasks with more progress to be selected. This dynamic task selection regime can be considered analogous to the flow state theory [50], which strives to maintain a balance between challenging and effortless tasks.

### C. MULTI-TASK REINFORCEMENT LEARNING

We initialize two separate BPNN architectures for critic and policy networks to integrate our method into the actor-critic reinforcement learning framework. In this manner, the policy network receives the state for the corresponding task as input and outputs the mean of a diagonal multivariate Gaussian distribution with a learned log standard deviation parameter

independent of the state. Correspondingly, the critic network receives the state and learns the value function. Hence, only the parameters of the training task's actor-critic modules and their corresponding lateral connections are updated using Adam optimizer [51] during task learning. We use the tuned hyperparameters in [52] for the reacher task as there are multiple extensions of PPO [53]. This extended PPO loss  $L_{PPO}^{\mathcal{T}_i}$  with a clipped value function loss and a value function entropy bonus can be formulated [48] as

$$L_{PPO}^{\mathcal{T}_i} = \mathbb{E}_t [L_t^{CLIP\mathcal{T}_i}(\theta) - c_1 L_t^{VF\mathcal{T}_i}(\theta) + c_2 S^{\mathcal{T}_i}(\theta(\pi))(s_t)] \quad (3)$$

where  $L_t^{CLIP\mathcal{T}_i}(\theta)$  is the clipped PPO surrogate objective,  $c_1 L_t^{VF\mathcal{T}_i}(\theta)$  is the clipped value loss with coefficient  $c_1$ ,  $c_2 S^{\mathcal{T}_i}(\theta(\pi))(s_t)$  is the entropy bonus for the actor network with coefficient  $c_2$ . Although the actor and critic neural networks are separate, we can denote them collectively as  $\theta$  as in previous works [48] for brevity. For instance,  $\theta(\pi)$  and  $\theta(\phi)$  refer to the actor-network and critic-networks parameters, respectively.

### D. EVALUATION METRICS

Designing a reward function is crucial and challenging in deep reinforcement learning, primarily because the reward function may not fully encapsulate all attributes expected from a learning agent. One example is in the Reacher environment, where the need often arises to tune the reward function coefficients or introduce additional parameters. However, this requires significant time and resources to adhere to the predetermined metrics and careful tuning of the parameters. In this sense, evaluating how a straightforward, uninformed reward function performs in terms of metrics meaningful for the task domain at hand is important. Therefore, in this section, in addition to the classical RL metric of episodic return, we also present task domain metrics used to evaluate the collective performance of morphologically different agents. To account for early stopping introduced in [54], we evaluate all tasks after each iteration, save the best policies obtained up to that iteration, and use them in our evaluations. Notably, saving the best policy based on the cumulative performance of all tasks for all methods ensures an accurate comparison of the proposed method against baselines and enhances performance across all methods after training is completed.

#### 1) EPISODIC RETURN

Maximum episodic return tracks the best expected discounted cumulative reward achieved across all tasks in an iteration. We report the best expected discounted cumulative reward after each training iteration to ensure a fair comparison with baselines. Crucially, the ERP task-switching procedure provides an inherent early-stopping procedure, provided that we allow a fixed number of cumulative training iterations under resource constraints. Since the task is to reach a given point in space with minimum effort, we follow the reward

**Algorithm 1** ERP-BPNN

**Require:**  $\mathcal{T}_i$ : Task with index  $i$ ,  $\mathcal{T}$ : set of tasks,  $m_{\mathcal{T}_i}$ : network module of  $\mathcal{T}_i$

```

1: Constants:  $\kappa = 2$  (#episodes),  $P = 8$  (#parallel tasks),
    $K_{init} = 20$  (#jumpstart training iterations)
2: for all  $\mathcal{T}_i$  do
3:   Unfreeze module  $m_{\mathcal{T}_i}$ 
4:   for  $k \in \{1, \dots, K_{init}\}$  do
5:     for  $p \in \{1, \dots, P\}$  do
6:       Sample  $\kappa$  episodes by running policy  $\pi_{\theta_{\mathcal{T}_i}^{old}}$ 
7:        $\forall t$ , Compute Advantage estimates  $\hat{A}_t$ 
8:     end for
9:     Record average episodic return  $R_{m_{\mathcal{T}_i}}(k)$ 
10:    Optimize  $L_{PPO}^{\mathcal{T}_i}$  w.r.t.  $\theta_{\mathcal{T}_i}$  by Equation 3
11:     $\theta_{\mathcal{T}_i}^{old} \leftarrow \theta_{\mathcal{T}_i}$ 
12:  end for
13:  Freeze module  $m_{\mathcal{T}_i}$ 
14: end for
15: while not done do
16:   Calculate progress for all tasks using ERP
17:   Choose task  $\mathcal{T}_i$  with maximum ERP
18:   Unfreeze module  $m_{\mathcal{T}_i}$ 
19:   for  $p \in \{1, \dots, P\}$  do
20:     Sample  $\kappa$  episodes by running policy  $\pi_{\theta_{\mathcal{T}_i}^{old}}$ 
21:      $\forall t$ , Compute Advantage estimates  $\hat{A}_t$ 
22:   end for
23:   Record average episodic return  $R_{m_{\mathcal{T}_i}}(k)$ 
24:   Optimize  $L_{PPO}^{\mathcal{T}_i}$  w.r.t.  $\theta_{\mathcal{T}_i}$  by Equation 3
25:    $\theta_{\mathcal{T}_i}^{old} \leftarrow \theta_{\mathcal{T}_i}$ 
26:   Freeze module  $m_{\mathcal{T}_i}$ 
27:   for all  $\mathcal{T}_j \in T - \{\mathcal{T}_i\}$  do
28:     for  $p \in \{1, \dots, P\}$  do
29:       Sample  $\kappa$  episodes by running policy  $\pi_{\theta_{\mathcal{T}_j}^{old}}$ 
30:     end for
31:     Record average episodic return  $R_{m_{\mathcal{T}_j}}(k)$ 
32:   end for
33: end while

```

function definition in [55]

$$\mathbf{r} = -\|p_{\text{fingertip}} - p_{\text{target}}\|_2 - \sum a^2 \quad (4)$$

where  $a$  is the action,  $-\sum a^2$  is the control cost,  $p_{\text{fingertip}}$ , and  $p_{\text{target}}$  are the positions of the fingertip and the target, respectively.

## 2) DISTANCE TO THE GOAL

Distance to the goal is the minimum expected L2-norm distance between the final end-effector position and the goal position obtained over all tasks. The Reacher environment only terminates after 50 timesteps, which is equal to the episode length. Hence, we expect the end-effector of the reacher to stay in the immediate vicinity of the goal until episode termination.

## 3) DEVIATION FROM SHORTEST PATH TO THE GOAL

At a high level, we expect an efficient learning agent to follow the shortest feasible path to the goal, which is also desirable for many robotic applications. Thus, we introduce a straightness metric that measures the minimum expected deviation from the shortest path to the goal taken by the manipulator's end-effector over all tasks. We define the deviation metric,  $\mathcal{D}$  across tasks as

$$\mathcal{D} = E_{\mathcal{T}} \left[ \left( \sum_{t=1} L_2(\mathbf{x}_{\mathcal{T}}(t), \mathbf{x}_{\mathcal{T}}(t-1)) \right) - L_2(\mathbf{g}, \mathbf{x}_{\mathcal{T}}(0)) \right]$$

where  $\mathcal{T}$  is a task,  $\mathbf{x}_{\mathcal{T}}(t)$  is the position of the end-effector for task  $\mathcal{T}$  at timestep  $t$ ,  $\mathbf{g}$  is the goal and  $\mathbf{x}_{\mathcal{T}}(0)$  is the initial end-effector position. Here, we first compute the length of the spatial trajectory by summing the  $L_2$ -norm of the distance between the current and the previous end-effector position. Subsequently, we subtract the  $L_2$ -norm of the distance between the goal position and the initial end-effector position to obtain the deviation from the shortest path to the goal for task  $\mathcal{T}$ . After each iteration, we calculate the path length traveled by all agents for every task and record the smallest expected deviation from the optimal path to the goal as obtained by the learning algorithm.

## E. IMPLEMENTATION DETAILS

Given the morphological diversity among robots results in different state space dimensions, and our BPNN architecture has a static input layer; we have standardized the state input dimensions for robots with 2-DoF and 3-DoF to match the maximum state space dimension, which of the 4-DoF robot, across all tasks. Hence, we apply zero-padding to the dimensions corresponding to the angular velocity and the angle of each missing link. To encourage skill transfer among networks, we limit the available computational resources to the task network modules by setting the number of hidden layers as three and the hidden layer size as two.

For the Reacher tasks, we use PPO [48] as the RL algorithm and adopt the hyperparameter set from [52] for the remaining hyperparameters. After each training iteration, we sample trajectories from each task, excluding the task trained most recently, using the latest BPNN actor and critic parameters. Then, we compute the ERP, maximum episodic return, minimum expected distance to the goal, and minimum expected deviation from the shortest path to the goal for each task. Additionally, we run eight parallel environments per task to reduce simulation time. Each environment runs two episodes, with each episode lasting 50 steps, thereby accumulating a total of 800 timesteps per task.

## IV. EXPERIMENTS

In this section, we first present the environment used for benchmarking multi-task learning for morphologically different robots illustrated in Fig. 1. Then, we elaborate on ERP-based task switching and present results obtained with the evaluation metrics introduced in Section III-D.

## A. SINGLE GOAL MULTI-TASK LEARNING BETWEEN MORPHOLOGICALLY DIFFERENT ROBOTS

To demonstrate skill transfer between morphologically different robots, we modified the Reacher-v2 environment, simulated in MuJoCo [56] using Gymnasium framework [55]. The learning experiments have been conducted on a modern computer equipped with an NVIDIA 3090 GPU and 10-core Intel i9 CPU at 3.7GHz. The experiment setup consists of three different reacher environments, each having distinct action spaces, specifically featuring 2-DoF, 3-DoF, and 4-DoF robot arms. Although each environment has unique environmental dynamics due to its distinct morphologies, it shares the same reward functions. The reward function defined in Eq. 4 comprises a control cost, the negative vector norm of the end-effector of the reacher, and a predetermined goal. We consider the baselines RANDOM-BPNN and random Multi-Layer Perceptron (RANDOM-MLP) to evaluate the performance of our ERP-BPNN approach across various metrics. RANDOM-BPNN abides by random task selection while featuring the same underlying BPNN architecture as ERP-BPNN. Similarly, RANDOM-MLP follows the same random task selection procedure as RANDOM-BPNN but utilizes a separate actor-critic network pair for each task with no lateral connections among different tasks. Accordingly, RANDOM-MLP can be regarded as training a PPO algorithm for each task separately.

## B. RESULTS

In this section, we report the results of our method ERP-BPNN compared to the baselines of RANDOM-BPNN and RANDOM-MLP, based on the evaluation metrics presented in Section III-D. Table 1 shows each method's mean and standard deviation results obtained during 100K episodes across five random seeds for each metric. The results indicate that ERP-BPNN can obtain superior performance compared to the baselines across all metrics with the lowest standard deviation. More importantly, as seen from the Episodic Return plot in Fig. 3 (a), the proposed model, ERP-BPNN, achieves faster convergence than the baselines.


Since the combination of BPNN architecture with ERP-based task selection (ERP-BPNN) yields the best results

**TABLE 1.** Episodic return, distance to goal, and deviation from shortest path to goal results across 5 random seeds using ERP-BPNN, RANDOM-BPNN, and RANDOM-MLP.

Episodic Return	ERP-BPNN	RANDOM-BPNN	RANDOM-MLP
60,000 episodes	$-7.95 \pm 0.48$	$-9.55 \pm 0.94$	$-9.12 \pm 0.82$
75,000 episodes	$-6.26 \pm 0.20$	$-7.32 \pm 0.45$	$-7.24 \pm 1.15$
80,000 episodes	$-6.09 \pm 0.13$	$-6.72 \pm 0.34$	$-6.69 \pm 0.91$
100,000 episodes	$-5.58 \pm 0.08$	$-5.86 \pm 0.17$	$-5.95 \pm 0.69$
Distance to Goal	ERP-BPNN	RANDOM-BPNN	RANDOM-MLP
60,000 episodes	$5.3 \pm 1.79$	$9.13 \pm 2.86$	$7.92 \pm 1.82$
75,000 episodes	$2.25 \pm 0.33$	$4.85 \pm 1.58$	$4.82 \pm 2.89$
80,000 episodes	$2.03 \pm 0.34$	$3.42 \pm 0.97$	$3.57 \pm 2.39$
100,000 episodes	$1.29 \pm 0.18$	$1.55 \pm 0.4$	$2.25 \pm 2.1$
Deviation from Shortest Path	ERP-BPNN	RANDOM-BPNN	RANDOM-MLP
60,000 episodes	$21.25 \pm 2.33$	$24.39 \pm 4.15$	$22.44 \pm 3.01$
75,000 episodes	$16.71 \pm 0.93$	$21.29 \pm 2.32$	$19.3 \pm 1.48$
80,000 episodes	$15.85 \pm 0.73$	$18.03 \pm 1.23$	$17.88 \pm 1.97$
100,000 episodes	$13.94 \pm 0.27$	$14.97 \pm 0.6$	$14.79 \pm 0.8$

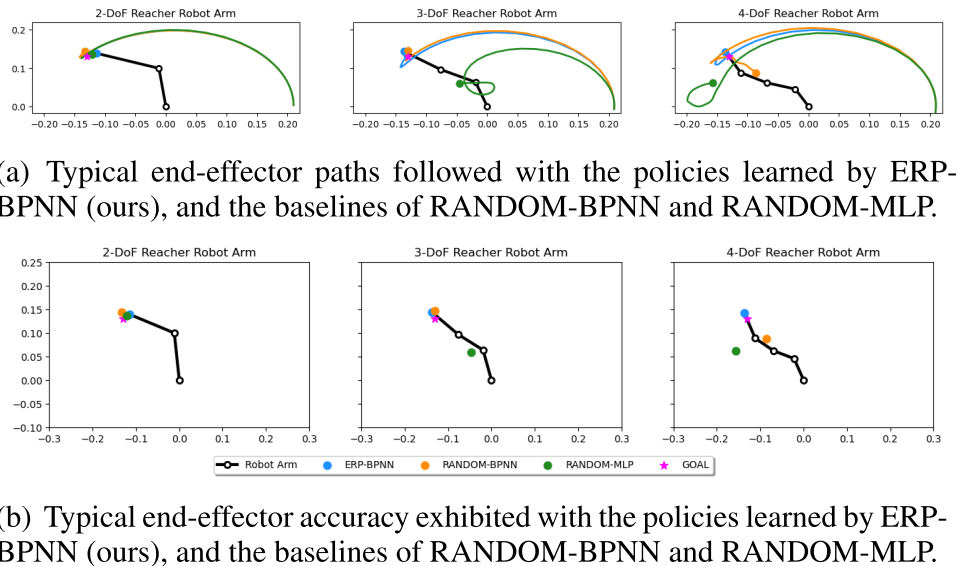
surpassing the random task selection strategy (RANDOM-BPNN), it can be argued that the ERP task selection procedure is essential for successful multi-task learning with positive transfer among tasks. Note that interestingly, the random task selection strategy (RANDOM-BPNN), although inferior to ERP-BPNN, performs better than the RANDOM-MLP baseline in terms of Episodic Return and yields better endpoint accuracy measured as Distance to Goal at the end of the training (Table 1). This indicates the lateral connections among task networks may facilitate limited positive skill transfer even with random task selection. However, this picture is not reflected in the Deviation from the Shortest Path to the Goal measure as random task selection leads to negative transfer for the Deviation from the Shortest Path to the Goal measure (see Fig 3(c), episodes > 55K). In the early stages of learning, the performance of ERP-BPNN and RANDOM-BPNN are better than RANDOM-MLP (episodes < 55K); however, while the proposed ERP-BPNN enjoys positive skill transfer and thus continues to improve its progress, RANDOM-BPNN suffers from negative interference degrading to a performance worse than RANDOM-MLP in this metric (Fig 3(c), episodes > 55K).

In order to get an intuitive understanding of the learned policy performances, we illustrate typical trajectories followed by the end-effector of the Reacher robot in Fig. 4(a) when controlled by policies acquired by the 1750<sup>th</sup> learning iteration, corresponding to approximately 80K episodes, i.e., when approximately 80% of the training is completed. Observe that at this training stage, the proposed ERP-BPNN has already acquired the skill to reach the target accurately through a less curved path compared to the baselines. In particular, the RANDOM-BPNN diverges from the shortest path to the goal for the 4-DoF robot arm, whereas the RANDOM-MLP diverges from the shortest path to the goal for both 3-DoF and 4-DoF Reacher Robot arms.




**FIGURE 3.** Performances of the proposed model, ERP-BPNN, and the two baselines of RANDOM-BPNN and RANDOM-MLP across five random seeds are shown in terms of (a) maximum episodic return, (b) minimum expected final end-effector distance to goal, and (c) minimum expected deviation from the shortest path to the goal.

To have an intuitive grasp of the reaching ability obtained on morphologically different robots by the proposed model and the baselines, a typical goal position and the final end-effector positions reached for each robot are shown in Fig. 4(b). Final end-effector positions reached using each method indicate that the agent trained using ERP-BPNN is consistently closest to the goal across all environments



**FIGURE 4.** The policies obtained by our model and the baselines during multi-task learning in terms of straightness (a) and endpoint accuracy (b) are demonstrated (at 1750<sup>th</sup> policy update).



**FIGURE 5.** Selection frequency plot of task switching by average episodic return progress with an iteration window size of  $v = 35$ .

compared to the baselines. Likewise, the results for Deviation from Shortest Path to Goal in Fig. 4(a) show that the agent trained with ERP-BPNN can reach the goal by taking a shorter path than the baselines in the 4-DoF reacher robot arm environment. Importantly, accurate reaching using a straighter path can be learned significantly faster by our proposed method ERP-BPNN compared to the baselines as evidenced by the distance plots given in Fig. 3(b).

### C. ERP-BASED DYNAMIC TASK SELECTION

To examine the task engagement patterns that emerge with the ERP-based task selection mechanism, we record the selection frequency of each task over a moving learning window, i.e., a fixed number of learning iterations. By using this scheme, the selection frequencies of the tasks are plotted against training iteration counts in Fig. 5. As can be seen, although ERP primarily selects the more straightforward 2-DoF Reacher Robot Arm as the one to learn initially (Fig. 5, time range: 0-10), it gradually starts selecting more complex tasks with 3-DoF and 4-DoF Reacher Robots, respectively (Fig. 5 time range: 10-150). Then, we observe that ERP continues selecting the 4-DoF Reacher Robot for skill transfer more frequently (Fig. 5, time range: 250-500). This task selection behavior suggests that the 4-DoF Reacher Robot has benefited from the skill transfer

of the other tasks and started to learn more rapidly later in the training. Upon further elaboration of the episodic return plots for all tasks in Fig. 3 (a), and standard deviations across iterations in Table 1, ERP-BPNN exhibits a consistent overall improvement across training iterations, achieving a faster convergence with lower standard deviation compared to the baselines. In line with this, faster convergence of ERP-BPNN indicates that consecutive selections of the 4-DoF Reacher task have collectively increased the maximum episodic return.

### V. CONCLUSION

This paper introduces Episodic Return Progress with Bidirectional Progressive Neural Networks (ERP-BPNN), a novel multi-task reinforcement learning approach incorporating a human-like interleaved learning mechanism, and shows its application to skill transfer across morphologically different robots. ERP-BPNN comprises a Bidirectional Progressive Neural Network (BPNN) that enables efficient, bidirectional skill transfer and a soft Episodic Return Progress (ERP) mechanism for dynamic task selection. The BPNN architecture is specifically designed to enable skill transfer through bidirectional lateral connections, drawing inspiration from the human brain's adeptness at retaining existing knowledge while acquiring new skills. Episodic return progress-based (ERP) task selection complements BPNN by autonomously selecting the tasks to engage in learning during training, eliminating the need for prior domain knowledge to evaluate difficulty levels of tasks. We demonstrate that soft ERP-based task selection, with BPNN, achieves higher performance and faster convergence than the baselines across all the tested metrics of Episodic Return, Distance To the Goal, and Deviation from Shortest Path to Goal.

The insights presented in this work are closely related to the field of human brain-inspired reinforcement learning (RL) research. By implementing an Intrinsic Motivation (IM)



signal and leveraging curiosity-driven exploration within the context of multi-task RL, ERP-BPNN framework closely aligns with biological processes observed in the human brain. Specifically, the behavior of dopaminergic neurons, known for their critical role in reward-based learning, mirrors the temporal difference (TD) prediction errors utilized in RL algorithms [24], [57], [58]. BPNN provides a mechanism for a shared multi-task RL architecture that prevents task interference during skill transfer by adjusting knowledge transfer from other tasks during training. In addition, the ERP-guided freezing of lateral connections among task modules prevents interference in the previously acquired knowledge. Furthermore, the bidirectional lateral connections augment noise in the task-specific computational flow akin to noisy computations in the brain [26], alleviating catastrophic forgetting [59] while amplifying plasticity [60]. Analogous to the modular control architecture in cerebellum [61], as corroborated by Functional Magnetic Resonance Imaging studies [62], [63], BPNN adopts a modular architecture. This architecture choice is instrumental in decreasing catastrophic forgetting via instantiating a module for each task.

Overall, we have designed a cognitive architecture applicable to a wide range of lifelong learning scenarios. Yet, an important future work that remains is to extend our architecture to support heterogeneous learning, where each task may require a different type of learning mechanism. For instance, one module may support supervised learning while others may handle reinforcement and unsupervised learning tasks. Other future work involves enhancing the framework to accommodate more complex lifelong RL scenarios, including real-world applications. This can be achieved by increasing the number of tasks and improving the bidirectional skill transfer mechanisms to enhance learning capacity and generalization ability.

## REFERENCES

- [1] E. Oztup and E. Ugur, "Lifelong robot learning," in *Encyclopedia of Robotics*, M. H. Ang, O. Khatib, and B. Siciliano, Eds. Berlin, Germany: Springer, 2021, pp. 1–12.
- [2] G. I. Parisi, R. Kemker, J. L. Part, C. Kanan, and S. Wermter, "Continual lifelong learning with neural networks: A review," *Neural Netw.*, vol. 113, pp. 54–71, May 2019.
- [3] S. Thrun and T. M. Mitchell, "Lifelong robot learning," *Robot. Auton. Syst.*, vol. 15, nos. 1–2, pp. 25–46, Jul. 1995.
- [4] H. Say and E. Oztup, "A model for cognitively valid lifelong learning," in *Proc. IEEE Int. Conf. Robot. Biomimetics (ROBIO)*, Dec. 2023, pp. 1–7.
- [5] C. Lin, M. Chiang, B. J. Knowlton, M. Iacoboni, P. Udompholkul, and A. D. Wu, "Interleaved practice enhances skill learning and the functional connectivity of fronto-parietal networks," *Human Brain Mapping*, vol. 34, no. 7, pp. 1542–1558, Jul. 2013.
- [6] J. Samani and S. C. Pan, "Interleaved practice enhances memory and problem-solving ability in undergraduate physics," *npj Sci. Learn.*, vol. 6, no. 1, p. 32, Nov. 2021.
- [7] J. Park, K. Varma, and S. Varma, "The role of executive function abilities in interleaved vs. blocked learning of science concepts," *Frontiers Psychol.*, vol. 14, Nov. 2023, Art. no. 1199682.
- [8] L. Nadel, A. Hupbach, R. Gomez, and K. Newman-Smith, "Memory formation, consolidation and transformation," *Neurosci. Biobehavioral Rev.*, vol. 36, no. 7, pp. 1640–1645, Aug. 2012.
- [9] N. J. Cepeda, E. Vul, D. Rohrer, J. T. Wixted, and H. Pashler, "Spacing effects in learning: A temporal ridge of optimal retention," *Psychol. Sci.*, vol. 19, no. 11, pp. 1095–1102, Nov. 2008.
- [10] P.-Y. Oudeyer, F. Kaplan, and V. V. Hafner, "Intrinsic motivation systems for autonomous mental development," *IEEE Trans. Evol. Comput.*, vol. 11, no. 2, pp. 265–286, Apr. 2007.
- [11] C. Colas, P. Fournier, M. Chetouani, O. Sigaud, and P.-Y. Oudeyer, "CURIOUS: Intrinsically motivated modular multi-goal reinforcement learning," in *Proc. 36th Int. Conf. Mach. Learn.*, in Proceedings of Machine Learning Research, vol. 97, K. Chaudhuri and R. Salakhutdinov, Eds., Jun. 2019, pp. 1331–1340.
- [12] S. Forestier and P.-Y. Oudeyer, "Curiosity-driven development of tool use precursors: A computational model," in *Proc. 38th Annu. Conf. Cognit. Sci. Soc. (COGSCI)*, 2016, pp. 1859–1864.
- [13] Y. Gatsoulis and T. M. McGinnity, "Intrinsically motivated learning systems based on biologically-inspired novelty detection," *Robot. Auto. Syst.*, vol. 68, pp. 12–20, Jun. 2015.
- [14] J. Achiam and S. Sastry, "Surprise-based intrinsic motivation for deep reinforcement learning," 2017, *arXiv:1703.01732*.
- [15] N. V. Varghese and Q. H. Mahmoud, "A survey of multi-task deep reinforcement learning," *Electronics*, vol. 9, no. 9, p. 1363, Aug. 2020.
- [16] Z. Xu, K. Wu, Z. Che, J. Tang, and J. Ye, "Knowledge transfer in multi-task deep reinforcement learning for continuous control," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 33, H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, Eds. Red Hook, NY, USA: Curran Associates, 2020, pp. 15146–15155.
- [17] Y. Teh, V. Bapst, W. M. Czarnecki, J. Quan, J. Kirkpatrick, R. Hadsell, N. Heess, and R. Pascanu, "Distal: Robust multitask reinforcement learning," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 30, 2017. [Online]. Available: [https://papers.nips.cc/paper\\_files/paper/2017/hash/0abdc563a06105aee3c6136871c9f4d1-Abstract.html](https://papers.nips.cc/paper_files/paper/2017/hash/0abdc563a06105aee3c6136871c9f4d1-Abstract.html)
- [18] J. Zhang, J. T. Springenberg, J. Boedecker, and W. Burgard, "Deep reinforcement learning with successor features for navigation across similar environments," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, Sep. 2017, pp. 2371–2378.
- [19] H. Yin and S. J. Pan, "Knowledge transfer for deep reinforcement learning with hierarchical experience replay," in *Proc. AAAI Conf. Artif. Intell.* Washington, DC, USA: The Association for the Advancement of Artificial Intelligence (AAAI), 2017, pp. 1640–1646.
- [20] L. T. Liu, U. Dogan, and K. Hofmann, "Decoding multitask DQN in the world of minecraft," in *Proc. 13th Eur. Workshop Reinforcement Learn. (EWRL), 11th Women Mach. Learn. Workshop Deep Reinforcement Learn. Workshop*, Dec. 2016.
- [21] N. Roy, I. Posner, T. Barfoot, P. Beaudoin, Y. Bengio, J. Bohg, O. Brock, I. DePatie, D. Fox, D. Koditschek, T. Lozano-Perez, V. Mansinghka, C. Pal, B. Richards, D. Sadigh, S. Schaal, G. Sukhatme, D. Therien, M. Toussaint, and M. Van de Panne, "From machine learning to robotics: Challenges and opportunities for embodied intelligence," 2021, *arXiv:2110.15245*.
- [22] A. A. Rusu, N. C. Rabinowitz, G. Desjardins, H. Soyer, J. Kirkpatrick, K. Kavukcuoglu, R. Pascanu, and R. Hadsell, "Progressive neural networks," 2016, *arXiv:1606.04671*.
- [23] A. A. Rusu, M. T. Rothörl, N. Heess, R. Pascanu, and R. Hadsell, "Sim-to-real robot learning from pixels with progressive nets," in *Proc. 1st Annu. Conf. Robot Learn.*, in Proceedings of Machine Learning Research, vol. 78, Mountain View, CA, USA, 2017, pp. 262–270.
- [24] D. Hassabis, D. Kumaran, C. Summerfield, and M. Botvinick, "Neuroscience-inspired artificial intelligence," *Neuron*, vol. 95, no. 2, pp. 245–258, Jul. 2017.
- [25] Y. Niv, "Reinforcement learning in the brain," *J. Math. Psychol.*, vol. 53, no. 3, pp. 139–154, 2009.
- [26] K. Khetarpal, M. Riemer, I. Rish, and D. Precup, "Towards continual reinforcement learning: A review and perspectives," *J. Artif. Intell. Res.*, vol. 75, pp. 1401–1476, Dec. 2022.
- [27] F. Brady, "The contextual interference effect and sport skills," *Perceptual Motor Skills*, vol. 106, no. 2, pp. 461–472, Apr. 2008.
- [28] E. S. Cross, P. J. Schmitt, and S. T. Grafton, "Neural substrates of contextual interference during motor learning support a model of active preparation," *J. Cogn. Neurosci.*, vol. 19, no. 11, pp. 1854–1871, Nov. 2007.
- [29] C.-H. Lin, B. J. Knowlton, M.-C. Chiang, M. Iacoboni, P. Udompholkul, and A. D. Wu, "Brain-behavior correlates of optimizing learning through interleaved practice," *NeuroImage*, vol. 56, no. 3, pp. 1758–1772, Jun. 2011.
- [30] J. M. Schorn and B. J. Knowlton, "Interleaved practice benefits implicit sequence learning and transfer," *Memory Cognition*, vol. 49, no. 7, pp. 1436–1452, Oct. 2021.
- [31] S. Sharma, A. Jha, P. Hegde, and B. Ravindran, "Learning to multi-task by active sampling," 2017, *arXiv:1702.06053*.
- [32] S. Jean, O. Firat, and M. Johnson, "Adaptive scheduling for multi-task learning," 2019, *arXiv:1909.06434*.

- [33] T. Lesort, V. Lomonaco, A. Stoian, D. Maltoni, D. Filliat, and N. Díaz-Rodríguez, “Continual learning for robotics: Definition, framework, learning strategies, opportunities and challenges,” *Inf. Fusion*, vol. 58, pp. 52–68, Jun. 2020.
- [34] P.-Y. Oudeyer, “What is intrinsic motivation? A typology of computational approaches,” *Frontiers Neurobotics*, vol. 1, p. 108, Nov. 2007.
- [35] K. Kasmarik, G. Baldassarre, V. G. Santucci, and J. Triesch, “Guest editorial special issue on intrinsically motivated open-ended learning (IMOL),” *IEEE Trans. Cogn. Develop. Syst.*, vol. 15, no. 2, pp. 321–324, Jun. 2023.
- [36] V. G. Santucci, P.-Y. Oudeyer, A. Barto, and G. Baldassarre, “Intrinsically motivated open-ended learning in autonomous robots,” *Frontiers Neurobot.*, vol. 13, p. 115, Jan. 2020.
- [37] M. I. Sener, Y. Nagai, E. Oztop, and E. Ugur, “Exploration with intrinsic motivation using object-action-outcome latent space,” *IEEE Trans. Cognit. Develop. Syst.*, vol. 15, no. 2, pp. 325–336, Jun. 2023.
- [38] S. Bugur, E. Oztop, Y. Nagai, and E. Ugur, “Effect regulated projection of robot’s action space for production and prediction of manipulation primitives through learning progress and predictability-based exploration,” *IEEE Trans. Cogn. Develop. Syst.*, vol. 13, no. 2, pp. 286–297, Jun. 2021.
- [39] Y. Burda, H. Edwards, A. Storkey, and O. Klimov, “Exploration by random network distillation,” 2018, *arXiv:1810.12894*.
- [40] B. Eysenbach, A. Gupta, J. Ibarz, and S. Levine, “Diversity is all you need: Learning skills without a reward function,” 2018, *arXiv:1802.06070*.
- [41] Y. Bengio, J. Louradour, and R. Collobert, “Curriculum learning,” in *Proc. 26th Annu. Int. Conf. Mach. Learn.* Montreal, QC, Canada: ACM, Jun. 2009, pp. 41–48.
- [42] S. Narvekar, B. Peng, M. Leonetti, J. Sinapov, M. E. Taylor, and P. Stone, “Curriculum learning for reinforcement learning domains: A framework and survey,” 2020, *arXiv:2003.04960*.
- [43] D. Weinshall, G. Cohen, and D. Amir, “Curriculum learning by transfer learning: Theory and experiments with deep networks,” in *Proc. Int. Conf. Mach. Learn.*, Jul. 2018, pp. 5238–5246.
- [44] E. Ugur and J. Piater, “Emergent structuring of interdependent affordance learning tasks,” in *Proc. 4th Int. Conf. Develop. Learn. Epigenetic Robot.*, Genoa, Italy, Oct. 2014, pp. 489–494.
- [45] W. Zaremba and I. Sutskever, “Learning to execute,” 2014, *arXiv:1410.4615*.
- [46] T. Schaul, J. Quan, I. Antonoglou, and D. Silver, “Prioritized experience replay,” 2015, *arXiv:1511.05952*.
- [47] C. Fernando, D. Banarse, C. Blundell, Y. Zwols, D. Ha, A. A. Rusu, A. Pritzel, and D. Wierstra, “PathNet: Evolution channels gradient descent in super neural networks,” 2017, *arXiv:1701.08734*.
- [48] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, “Proximal policy optimization algorithms,” 2017, *arXiv:1707.06347*.
- [49] S. Huang, R. F. J. Dossa, C. Ye, J. Braga, D. Chakraborty, K. Mehta, and J. G. Araújo, “CleanRL: High-quality single-file implementations of deep reinforcement learning algorithms,” *J. Mach. Learn. Res.*, vol. 23, no. 274, pp. 1–18, 2022. [Online]. Available: <http://jmlr.org/papers/v23/21-1342.html>
- [50] M. Csikszentmihalyi, *Flow: The Psychology of Optimal Experience*. New York, NY, USA: Harper & Row, 1990.
- [51] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” in *Proc. ICLR*, Y. Bengio and Y. LeCun, Eds., 2015.
- [52] A. Raffin. (2020). *RL Baselines3 Zoo*. [Online]. Available: <https://github.com/DLR-RM/rl-baselines3-zoo>
- [53] L. Engstrom, A. Ilyas, S. Santurkar, D. Tsipras, F. Janoos, L. Rudolph, and A. Madry, “Implementation matters in deep RL: A case study on PPO and TRPO,” in *Proc. Int. Conf. Learn. Represent.*, 2020.
- [54] S. E. Ada, E. Ugur, and H. L. Akin, “Generalization in transfer learning: Robust control of robot locomotion,” *Robotica*, vol. 40, no. 11, pp. 3811–3836, Nov. 2022.
- [55] M. Towers, J. K. Terry, A. Kwiatkowski, J. U. Balis, G. de Cola, T. Deleu, M. Goulão, A. Kallinteris, K. G. Arjun, M. Krimmel, R. Perez-Vicente, A. Pierré, S. Schulhoff, J. J. Tai, A. T. J. Shen, and O. G. Younis, Mar. 2023, “Gymnasium,” *Zenodo*, doi: [10.5281/zenodo.8127025](https://doi.org/10.5281/zenodo.8127025).
- [56] E. Todorov, T. Erez, and Y. Tassa, “MuJoCo: A physics engine for model-based control,” in *Proc. IEEE/RSS Int. Conf. Intell. Robots Syst.*, Oct. 2012, pp. 5026–5033.
- [57] J. P. O’Doherty, P. Dayan, K. Friston, H. Critchley, and R. J. Dolan, “Temporal difference models and reward-related learning in the human brain,” *Neuron*, vol. 38, no. 2, pp. 329–337, Apr. 2003.
- [58] W. Schultz, P. Dayan, and P. R. Montague, “A neural substrate of prediction and reward,” *Science*, vol. 275, no. 5306, pp. 1593–1599, Mar. 1997.
- [59] R. Ajemian, A. D’Ausilio, H. Moorman, and E. Bizzi, “A theory for how sensorimotor skills are learned and retained in noisy and nonstationary neural circuits,” *Proc. Nat. Acad. Sci. USA*, vol. 110, no. 52, pp. E5078–E5087, Dec. 2013.
- [60] S. Dohare, R. S. Sutton, and A. R. Mahmood, “Continual backprop: Stochastic gradient descent with persistent randomness,” 2021, *arXiv:2108.06325*.
- [61] D. M. Wolpert, R. C. Miall, and M. Kawato, “Internal models in the cerebellum,” *Trends Cogn. Sci.*, vol. 2, no. 9, pp. 338–347, Sep. 1998.
- [62] H. Imamizu, “Separated modules for visuomotor control and learning in the cerebellum: A functional MRI study,” in *Proc. 3rd Int. Conf. Funct. Mapping Human Brain*, vol. 5. New York, NY, USA: Academic, 1997, p. S598.
- [63] H. Imamizu, S. Miyauchi, T. Tamada, Y. Sasaki, R. Takino, B. Pütz, T. Yoshioka, and M. Kawato, “Human cerebellar activity reflecting an acquired internal model of a new tool,” *Nature*, vol. 403, no. 6766, pp. 192–195, Jan. 2000.



**SUZAN ECE ADA** received the B.S. degree in industrial engineering from Koç University, İstanbul, Turkey, and the M.S. degree from Boğaziçi University, İstanbul, where she is currently pursuing the Ph.D. degree in computer engineering. She was a Researcher with the COLORS Lab, Boğaziçi University, and has taken on the role of a Research and Teaching Assistant with the COLORS Lab, Boğaziçi University. From 2022 to 2023, she was a Research Fellow with the Wearable Flexible Sensor Supported Lower Body Exoskeleton System Project, under the supervision of Assoc. Prof. Emre Ugur. In 2023, she was a Research Fellow with the Symbiotic Intelligent Systems Research Center (SISREC), Osaka University, Japan, under the supervision of Prof. Erhan Oztop. Her research interests include deep reinforcement learning, robotics, lifelong learning, generative models, and meta-learning.



**HANNE SAY** received the B.S. degree in computer science from the TOBB University of Economics and Technology, Ankara, Turkey, in 2022. She is currently pursuing the M.S. degree in artificial intelligence with Özyeğin University, İstanbul, Turkey. She was a Research Fellow with the Symbiotic Intelligent Systems Research Center (SISREC), Osaka University, Osaka, Japan, in 2023, under the supervision of Prof. Erhan Oztop. Her research interests include interleaved learning of multi-tasking using intrinsic motivations, lifelong learning, brain energy consumption, and cognitive science.



**EMRE UGUR** received the Ph.D. degree in computer engineering from Middle East Technical University, Turkey. He was a Research Scientist with ATR, Japan; visited Osaka University, specially appointed as an Assistant and Associate Professor; and a Senior Researcher with the University of Innsbruck. He is currently an Associate Professor with the Department of Computer Engineering, Boğaziçi University, and the Chair of the Cognitive Science MA Program. He is interested in robot learning.



**ERHAN OZTOP** (Member, IEEE) received the Ph.D. degree from the University of Southern California, in 2002. In the same year, he joined Computational Neuroscience Laboratories, Advanced Telecommunications Research Institute International (ATR), Japan. There, he was a Researcher and later a Senior Researcher and the Group Leader, in addition to serving as the Vice Department Head for two research groups and holding a visiting associate professor position with Osaka University, from 2002 to 2011. Currently, he is a Professor with Özyeğin University, Turkey, and a Specially Appointed Professor with Osaka University, Japan.