

01 Jan 2023

Cooperative Deep Q^2 -Learning Framework for Environments Providing Image Feedback

Krishnan Raghavan

Vignesh Narayanan

Sarangapani Jagannathan

Missouri University of Science and Technology, sarangap@mst.edu

Follow this and additional works at: https://scholarsmine.mst.edu/ele_comeng_facwork

 Part of the [Electrical and Computer Engineering Commons](#)

Recommended Citation

K. Raghavan et al., "Cooperative Deep Q^2 -Learning Framework for Environments Providing Image Feedback," *IEEE Transactions on Neural Networks and Learning Systems*, Institute of Electrical and Electronics Engineers, Jan 2023.

The definitive version is available at <https://doi.org/10.1109/TNNLS.2022.3232069>

This Article - Journal is brought to you for free and open access by Scholars' Mine. It has been accepted for inclusion in Electrical and Computer Engineering Faculty Research & Creative Works by an authorized administrator of Scholars' Mine. This work is protected by U. S. Copyright Law. Unauthorized use including reproduction for redistribution requires the permission of the copyright holder. For more information, please contact scholarsmine@mst.edu.

Cooperative Deep Q -Learning Framework for Environments Providing Image Feedback

Krishnan Raghavan¹, *Member, IEEE*, Vignesh Narayanan, *Member, IEEE*,
and Sarangapani Jagannathan², *Fellow, IEEE*

Abstract—In this article, we address two key challenges in deep reinforcement learning (DRL) setting, sample inefficiency and slow learning, with a dual-neural network (NN)-driven learning approach. In the proposed approach, we use two deep NNs with independent initialization to robustly approximate the action-value function in the presence of image inputs. In particular, we develop a temporal difference (TD) error-driven learning (EDL) approach, where we introduce a set of linear transformations of the TD error to directly update the parameters of each layer in the deep NN. We demonstrate theoretically that the cost minimized by the EDL regime is an approximation of the empirical cost, and the approximation error reduces as learning progresses, irrespective of the size of the network. Using simulation analysis, we show that the proposed methods enable faster learning and convergence and require reduced buffer size (thereby increasing the sample efficiency).

Index Terms—Deep Q -learning (DQN), games, images.

I. INTRODUCTION

CONTROL of complex systems is an important research thrust in the domain of science and engineering [1], [2], [3]. In a typical application, sensors measure physical quantities such as position, velocity, pressure, and temperature. These measurements are then utilized to design controllers that solve control/decision-making tasks. However, in many applications, such as the self-driving autonomous cars or control tasks performed by humans, sensory inputs are often encoded as images and/or audio signals. More recently, deep learning has enabled NNs to decode complex abstractions from measurement data. For instance, in the context of reinforcement learning (RL), the deep RL (DRL) facilitates complex

decision-making problems involving high-dimensional state, such as images, audio inputs, and action spaces [4]. Despite decades of research, learning controllers for complex physical systems cannot be designed with performance guarantees when the sensors generate a sequence of images or audio signals [5] as feedback data for learning. One of the key reasons is that deep neural networks (DNNs) with many hidden layers are required to extract complex abstractions from images or audio signals [6]. However, the complex compositional structure of a deep neural network (NN) does not allow for traditional Lyapunov theory-based analysis to provide any suitable conclusion about the performance of the DNN-based controller without limiting assumptions.

A typical RL task is composed of three steps: 1) an agent selects an action and utilizes it to actuate or effect change in the environment; 2) the environment, as a consequence of the action, transitions from its current state to a new state; and 3) the environment then provides a scalar reinforcement/reward signal as feedback. By performing these three steps repeatedly, the agent explores the environment and accumulates experiences or knowledge regarding the environment. Based on these experiences collected over time, the agent learns the best sequence of actions from any situation/state [7] to achieve a desired objective/goal.

The desired objective is typically described as a function of the temporal difference (TD) error, and this objective is estimated through an RL agent (Q -network parameterized by NN) and its clone (a target) [4]. A typical DRL learning scheme such as the one introduced in [4] aims to successively nudge the target network toward the optimal Q -function, while the Q -network is trained to be close to the target. The key distinction between the target in a DRL task with that of supervised learning task is that the targets in a DRL problem are nonstationary [4]. Consequentially, efficient learning requires accurate approximation of the target, that is, the optimal Q -function and efficient learning are pivotal on two key conditions: 1) efficient use of experiences and 2) sufficient exploration of the parameter or weight space.

To enhance the process of collecting and efficiently utilizing experiences within high-dimensional state (such as images and audio inputs) and action spaces [8], [9], [10], [11], an experience replay strategy [12], [13], [14], [15] is usually employed in RL schemes. However, the memory space to store the experiences may be limited or expensive in many applications, and therefore, a mechanism to omit irrelevant

Manuscript received 19 January 2021; revised 8 September 2021 and 16 June 2022; accepted 14 December 2022. This work was supported in part by the Department of the Navy, Office of Naval Research, under Grant N00014-21-1-2232; in part by the Department of Army Cooperative Research under Agreement W911NF-21-2-0260 and Agreement W911NF-22-2-0185; and in part by the Intelligent Systems Center at the Missouri University of Science and Technology, Rolla. (Corresponding author: Krishnan Raghavan.)

Krishnan Raghavan is with the Mathematics and Computer Science Division, Argonne National Laboratory, Lemont, IL 60439 USA (e-mail: krm9c@outlook.com).

Vignesh Narayanan is with the AI Institute and the Department of Computer Science and Engineering, University of South Carolina, Columbia, SC 63130 USA.

Sarangapani Jagannathan is with the Department of Electrical and Computer Engineering, Missouri University of Science and Technology, Rolla, MO 65401 USA.

This article has supplementary material provided by the authors and color versions of one or more figures available at <https://doi.org/10.1109/TNNLS.2022.3232069>.

Digital Object Identifier 10.1109/TNNLS.2022.3232069

2162-237X © 2023 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission.
See <https://www.ieee.org/publications/rights/index.html> for more information.

experiences is desirable and have been considered [15], [16], [17]. Despite appropriate sample selection strategies, useful experiences may be discarded. For instance, experiences collected initially, which are near the initial states, may be useful, but when some experiences are omitted based on time or reward, these experiences may be discarded. Even when these initial experiences are not discarded, they have the tendency to bias the learning procedure [4], [18] as the agent may spend more time in exploring states dictated by the initial experiences but these experiences do not contribute positively to the overall objective of the agent. Despite several successful applications of DRL, these strategies of discarding or utilizing experiences, in general, are typically sample-inefficient [19], [20].

On the other hand, a tractable and efficient training strategy to use these experiences and update the network parameters is essential [21], [22]. The most common strategy for training DNNs is the back-propagation/stochastic gradient descent (SGD) algorithm [4]. Despite promising results, SGD suffers from issues, such as vanishing gradient problem and slow learning [23], [24], [25]. Many alternative approaches have been introduced in the literature to target these issues [24], [26], [27], [28], [29]. One prominent method is a direct error-driven learning (EDL) approach that was proposed to train DNNs [24] in the context of supervised learning. However, the EDL approach was not evaluated in DRL applications with image pixels as the NN inputs, and a theoretical analysis of the EDL algorithm is a missing component in [24].

To address these issues, we introduce a dual-NN-driven exploratory learning approach, wherein we improve sample efficiency using two NNs that are updated alternatively using the EDL update rule. We refer to this dual-NN-driven alternative update mechanism as the *cooperative (coop) update strategy*. In particular, the two DNNs estimate and minimize the desired objective (a function of TD errors). Each DNN alternatively behaves as the target and the Q -network; therefore, each NN is allowed to exploit, while behaving as target, and explore, while behaving as Q -network.

In this setting, we adapt and expand the direct EDL rule proposed in [24] to the case of DRL. In the EDL update rule, we transform the TD error onto each layer in the deep NN through a nonsingular matrix (linear transformation). Since the choice of the transformation matrix is critical, we construct this matrix for the purpose of increased exploration. We show that the solutions obtained with our scheme are as good as the typical SGD. We show that our coop approach: 1) improves sample efficiency; 2) enables better exploration; 3) results in faster convergence; and 4) provides better performance in the context of online games using DRL. The contributions of this article include: 1) alternative optimization-driven online learning approach for DRL involving high-dimensional state and action spaces to address sample efficiency; 2) extension of direct error-driven to RL with the cooperative learning strategy in order to mitigate vanishing gradient problem; and 3) convergence analysis.

II. BACKGROUND

In this section, we begin with a brief background on the RL problem and then introduce the notations used in the article along with DNNs. Finally, we formally introduce the problem

considered in this work. Additional background details are also provided in the Appendix in the Supplementary Material.

A. RL Preliminaries

The RL problems can be formulated as Markov decision processes (MDPs) composed of a state-space \mathcal{S} , an action space \mathcal{A} , a state-transition function $T : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{S}$, and a reward function R [30]. During the learning process, an RL agent perceives the environmental information, described by the state $x(k) \in \mathcal{S}$, and takes an action $u(k) \in \mathcal{A}$ so that the environment transits from the current state to the next state $x(k+1)$ based on the transition function T . The action $u(k)$ is, in general, based on a policy $\pi : \mathcal{S} \rightarrow \mathcal{A}$, and in this process, the agent receives an external scalar reward/reinforcement signal $r(k+1) \in \mathcal{R}$ from the environment. The set composed of $\{x(k), u(k), x(k+1), \text{and } r(k+1)\}$ is collected as experience by the RL agent to learn the value function or the Q -function. In this work, we utilize a DNN-based architecture to learn and approximate Q -functions in control/decision making problems concerned with systems/environments using data (from the environment) in the form of images.

B. DNNs and Notations

Consider a feed-forward NN with d layers denoted using a parametric map $y(\mathbf{z}; \boldsymbol{\theta})$, where \mathbf{z} is the input and $\boldsymbol{\theta}$ denotes the weight parameters. In particular, let $\boldsymbol{\theta} = [\mathbf{W}^{(1)}, \dots, \mathbf{W}^{(d)}]$ denotes the ideal weights of the DNN, such that the parametric map $y(\mathbf{z}; \boldsymbol{\theta})$ is an approximation to a smooth nonlinear function $\mathcal{F}(\mathbf{z})$ with compact support and an approximation error ε [31] such that $\mathcal{F}(\mathbf{z}) = y(\mathbf{z}, \boldsymbol{\theta}) + \varepsilon$. The term $f^{(i)}$, for $i = 1, \dots, d$ denotes the layerwise activation functions (applied componentwise to the vector input) corresponding to the d layers. In our proposed framework, we will employ such NNs to approximate the unknown function $\mathcal{F}(\mathbf{z})$. Let the estimated weights be denoted by $\hat{\boldsymbol{\theta}} = [\hat{\mathbf{W}}^{(1)}, \dots, \hat{\mathbf{W}}^{(d)}]$, such that the estimated net is denoted by $\hat{y}(\mathbf{z}, \hat{\boldsymbol{\theta}})$, where the symbol $\hat{(\cdot)}$ is used to denote the estimated quantities. Additionally, throughout the article, we use standard mathematical notations. Specifically, we use \mathbb{N} and \mathbb{R} to denote the set of natural numbers and real numbers, respectively. In the analysis presented in this article, we use $\|\cdot\|$ to denote the Euclidean norm for vectors and Frobenius norm for matrices. In the following, we formally introduce the problem considered in this article.

C. Problem Description

We consider an RL problem, where the dynamics or the transition function, T , of the environment is given as follows:

$$\mathbf{x}(k+1) = T(\mathbf{x}(k), \mathbf{u}(k), \boldsymbol{\omega}(k)) \quad (1)$$

where $\mathbf{x}(k) \sim p(\mathcal{S})$ is the state of the system at the k th sampling instant, $\mathbf{u}(k)$ (actions) is the input to the system, and $\boldsymbol{\omega}(k)$ is some internal fluctuation, e.g., noise. We will define a compact set \mathcal{S} to describe all the states that can be observed by the agent, and $p(\mathcal{S})$ describes a distribution on the states. Typically, the exact dynamics of the system are unknown at any sampling instant k . We consider applications where the system output (i.e., the feedback data) is a sequence of images

or snapshots. The input to the (learning) model are raw images that are preprocessed. We follow the preprocessing procedure detailed in [4], and the input to our network is four images of size 84×84 . One key difference between standard learning schemes and DRL is that the images are not the states of the system. These images encode the states, and it is the job of the NN to extract the states. Hence, the exact states are unknown to the model and must be extracted from the images. For instance, consider the game *breakout*, where the state of the game is represented by the position of the bar at the bottom of the screen. Since the exact position of the bar is unknown, the NN must extract it from the images of the screen. With this feedback information, we seek to train an RL agent to learn a sequence of actions that yields maximum cumulative reward. We denote the reward functions by $r(\mathbf{x}(k), \mathbf{u}(k))$, which is a function of both the state and the action or control inputs such that for a given (state, action pair) $(\mathbf{x}(k), \mathbf{u}(k))$ and an action policy π , the Q -function is given as follows:

$$Q^\pi(\mathbf{x}(k), \mathbf{u}(k)) = \sum_{k=0}^{\infty} \gamma^k r(\mathbf{x}(k), \mathbf{u}(k)) \quad (2)$$

where $\gamma \in (0, 1)$ is the discount factor. In a typical DRL environment, the reward and images are obtained from an emulator. Our goal is to utilize the state and reward information from the environment to develop a learning framework to iteratively learn the optimal sequence of actions for a given task. Since the optimality for a sequence of actions is determined by the Q -function, we seek to learn the optimal Q -function using NN; and since the feedback is in the form of images, we incorporate DNNs in our architecture.

III. PROPOSED METHOD

In this section, we first introduce the optimization problem and the overall learning architecture. Later, we introduce the learning scheme to train the DNNs.

A. RL Approximation of Q -Values

The RL agent receives feedback information from the environment in the form of images, the RL agent must learn a map between the inputs and the Q -values corresponding to a fixed policy (a greedy policy is considered in this article). To update the Q -values, we use the Bellman's principle of optimality [30] such that the optimal Q -function satisfies

$$Q^*(\mathbf{x}(k), \mathbf{u}(k)) = \max_{\mathbf{u}(k)} (r(\mathbf{x}(k), \mathbf{u}(k)) + \gamma^k (Q^*(\mathbf{x}(k+1), \mathbf{u}(k+1)))) \quad (3)$$

where Q^* is the optimal Q -function. Let $\rho(\mathcal{A})$ describes a distribution over \mathcal{A} providing

$$\mathbf{u}^*(\mathbf{x}(k)) = \operatorname{argmax}_{\mathbf{u}(k) \sim \rho} (Q^\rho(\mathbf{x}(k), \mathbf{u}(k))) \quad (4)$$

where we use the notation ρ to denote $\rho(\mathcal{A})$. When the dynamics are known as in a linear system, the evaluation of the optimal input to the system is well-known [32]. However, in most general cases, the system/environment dynamics are often unknown. Therefore, the optimal Q -function and, in turn, the control policy have to be approximated. In this

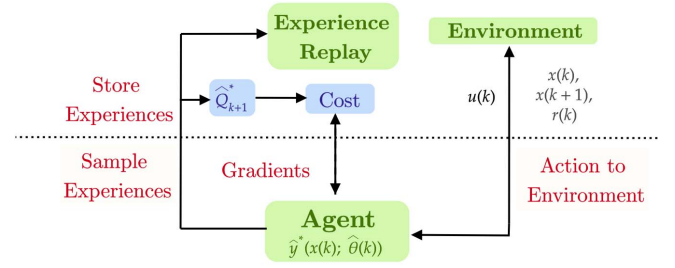


Fig. 1. Deep Q network (DQN) Algorithm.

work, we will approximate the Q -function value for each action in the set \mathcal{A} by defining a generic parametric map, $y(\mathbf{x}(k), \mathbf{u}(k); \theta)$ with an ideal set of parameters θ such that

$$\mathbf{u}^*(\mathbf{x}(k)) = \operatorname{argmax}_{\mathbf{u}(k) \in \mathcal{A}} (y(\mathbf{x}(k), \mathbf{u}(k); \theta)). \quad (5)$$

The learning problem is to find an approximation \hat{y} with parameters $\hat{\theta}(k)$ for the function y . We employ DNNs as function approximators to learn \hat{y} . To train the weights of the DNN, the cost is defined as the expected value over all the states and all the actions such that $\mathbb{E}_{\mathbf{x}(k) \sim p(\mathcal{S}), \mathbf{u}(k) \sim p(\mathcal{A})} [J(\hat{\theta}(k))]$, where

$$J(\hat{\theta}(k)) = \frac{1}{2} (y(\mathbf{x}(k), \mathbf{u}(k); \theta(k)) - \hat{y}(\mathbf{x}(k), \mathbf{u}(k); \hat{\theta}(k)))^2 \quad (6)$$

with target $y(\mathbf{x}(k), \mathbf{u}(k); \theta(k)) = r(\mathbf{x}(k), \mathbf{u}(k)) + \gamma^k \hat{Q}^*(\mathbf{x}(k+1), \mathbf{u}(k+1))$, coming from the Bellman equation with \hat{Q}^* denoting the approximation of the optimal Q -value obtained by extrapolating the current Q -function estimate. Therefore, we may write

$$\hat{\theta}(k) = \arg \min_{\hat{\theta}(k)} \mathbb{E}_{\mathbf{x}(k) \in \mathcal{S}, \mathbf{u}(k) \sim p(\mathcal{A})} [J_E(\mathbf{x}(k), \mathbf{u}(k), \hat{y}(\mathbf{x}(k), \mathbf{u}(k); \hat{\theta}(k)))] \quad (7)$$

with

$$\begin{aligned} J_E(\mathbf{x}(k), \mathbf{u}(k), \hat{y}(\mathbf{x}(k), \mathbf{u}(k); \hat{\theta}(k))) &= \frac{1}{2} [r_k + \gamma^k \hat{Q}_{k+1}^* - \hat{y}_k]^2 \\ &= \frac{1}{2} [\epsilon_k]^2. \end{aligned}$$

To simplify notations, we have used r_k to denote $r(\mathbf{x}(k), \mathbf{u}(k))$ and denote $\hat{Q}^*(\mathbf{x}(k+1), \mathbf{u}(k+1))$ by \hat{Q}_{k+1}^* and $\hat{y}(k; \hat{\theta}(k))$ by \hat{y}_k . Note that $r_k + \gamma^k \hat{Q}_{k+1}^* - \hat{y}_k$ is the TD error denoted as ϵ_k from here on. The empirical cost which we seek to minimize is the squared TD error.

B. Deep Q -Learning

A learning algorithm to solve the optimization problem in (7) was provided in [4]. In [4], the total learning process was split into episodes. For each episode, there are several plays, each play is comprised of five basic steps (refer Fig. 1): 1) current state is inferred from images provided by the environment; 2) this state information is utilized by the Q -network to generate actions; 3) the actions are then presented to the environment; 4) the environment evolves and provides the rewards and the next state via images; 5) the tuple of current state, next state, action, and reward is stored into an experience replay array in batches; and 6) finally, batches are sampled from the experience replay array to update the network parameters using the gradient of $J_E(\mathbf{x}(k), \mathbf{u}(k), \hat{y}(\mathbf{x}(k); \hat{\theta}(k)))$.

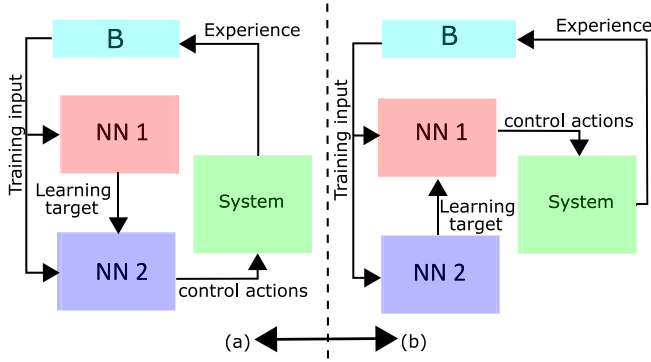


Fig. 2. Schematic illustrating the coop dual DQN architecture. During learning, the NNs (NN 1 and NN 2) switch roles as shown in (a) and (b). (a) NN 2 is trained using the experience stored in the memory buffer with targets given by NN 1 and (b) NNs swap their roles.

In a typical DRL algorithm [4], the TD error ($\epsilon_k = r_k + \gamma^k \hat{Q}_{k+1}^* - \hat{y}_k$) is computed by extrapolating the current Q -network to “guess” the Q -values of future states (the target). This is typically achieved by defining two networks (Q -network and target) and periodically copying the weights of the current Q -network to the target. Then, the current state is used as input to the Q -network to generate \hat{y}_k , while the future state is used as an input to the target to generate \hat{Q}_{k+1}^* . Using the ϵ_k , we update the weights of $\hat{y}(x(k); \hat{\theta}(k))$ during the next C plays.

C. Cooperative Dual Network Architecture (Coop)

Our coop strategy utilizes the algorithm from [4] (also described in Section III-B) but incorporates two DNNs initialized at random with different hyperparameters (weights and biases) instead of the periodically cloning the Q -network.

Our coop strategy is comprised of two steps. First, we designate one of the NNs, say NN 1, to generate the targets (future state is given as input to NN 1) and NN 2 for generating action (current state is given as input), where we update the weights of NN 2. Subsequently, in the following iteration, NN 2 generates the target and NN 1 generates the action. We always update the network that generates the action and with alternative updates, the two NNs play the role of learning the Q -function and the target, respectively (see Fig. 2).

The rationale behind the proposed learning scheme, where the NNs have cooperative functionality, is twofold. We expect that the NNs encode the training data from the samples much more efficiently while reducing the memory footprint and necessity of raw data stored in the memory buffer, thereby improving sampling efficiency. Second, since each NN obtains target from the other NN during training, these NNs are expected to converge asymptotically while improving the robustness of the learning. In the numerical experiments presented in Section IV, we show that equipped with an exploratory EDL rule, the proposed learning architecture is robust, improves the sample utilization, performs better with lesser buffer memory, and converges faster when compared with the existing results [4].

D. Error Driven TD Learning

To minimize the empirical cost, gradient-based update rules are commonly employed [33], [34]. These update rules are

typically written as follows:

$$\hat{\mathbf{W}}^{(i)}(k+1) = \hat{\mathbf{W}}^{(i)}(k) + \alpha \Delta^{(i)}(k) \quad (8)$$

where $\alpha > 0$ is the learning rate, i is the layer number, k is the sampling instant, and $\Delta^{(i)}(k)$ is the parameter adjustment made along a descent direction by using the gradient of the error propagated backward from the output layer of the network [25], [33]. To simplify notations, we switch to a subscript notation instead of expressing the sampling instance in parenthesis, so that (8) is rewritten as $\hat{\mathbf{W}}_{k+1}^{(i)} = \hat{\mathbf{W}}_k^{(i)} + \alpha \Delta_k^{(i)}$. From here on, for brevity, the notation of the expected value operator is suppressed, and we refer to the expected value of the cost as just the cost value. Furthermore, we add a regularization term to the cost function such that the revised cost is written as follows:

$$H(\hat{\theta}_k) = \left[J_E(\hat{\theta}_k) + \sum_{i=1}^d \lambda^{(i)} R_k(\hat{\mathbf{W}}_k^{(i)}) \right] \quad (9)$$

where $\lambda^{(i)} > 0$ is the decay coefficient, $R_k(\hat{\mathbf{W}}_k^{(i)})$ denotes the function of regularization applied on the weights $\hat{\mathbf{W}}_k^{(i)}$, and the optimal cost $H^*(\hat{\theta}_k)$ satisfies

$$H^*(\hat{\theta}_k) = \min_{\hat{\theta} \in \Omega} \left[J_E(\hat{\theta}_k) + \sum_{i=1}^d \lambda^{(i)} R_k(\hat{\mathbf{W}}_k^{(i)}) \right] \quad (10)$$

where $J_E(\hat{\theta}_k)$ is used in place of $J_E(x(k), u(k); \hat{\theta}(k))$. The weight adjustment at each iteration k is given as follows:

$$\Delta_k^{(i)} = -\left[\nabla_{\hat{\mathbf{W}}_k^{(i)}} H(\hat{\theta}_k) \right] = -\left[\delta_k^{(i)} + \lambda^{(i)} \nabla_{\hat{\mathbf{W}}_k^{(i)}} R(\hat{\mathbf{W}}_k^{(i)}) \right] \quad (11)$$

where the term $\nabla_{\hat{\mathbf{W}}_k^{(i)}}(\cdot)$ denotes the gradient of (\cdot) with respect to the NN weight $\hat{\mathbf{W}}_k^{(i)}$, the second term in this update rule depends on the choice of $R(\hat{\mathbf{W}}_k^{(i)})$, and $\delta_k^{(i)}$ is obtained by applying the chain rule to compute the gradient of the cost with respect to weight. Specifically, applying the chain rule to compute layerwise gradient through error backpropagation [23], [24], a generalized expression for $\delta_k^{(i)}$ can be derived as follows:

$$\delta_k^{(i)} = f^{(i-1)}(x) \epsilon_k \left[\prod_{j=d}^{i+1} \text{diag}(\nabla f^{(j)}(x)) \hat{\mathbf{W}}_k^{(j)} \right] \times \text{diag}(\nabla f^{(i)}(x)) \mathbf{I}^{(i)}. \quad (12)$$

We now denote $\prod_{j=d}^{i+1} (\text{diag}(\nabla f^{(j)}(x)) \hat{\mathbf{W}}_k^{(j)}) \text{diag}(\nabla f^{(i)}(x))$ as $\mathcal{T}^{(i)}(x)$ and simplify (12) to get

$$\delta_k^{(i)} = f^{(i-1)}(x) \epsilon_k \mathcal{T}^{(i)} \mathbf{I}^{(i)} \quad (13)$$

where $\epsilon_k = y_k - \hat{y}_k = r_k + \gamma^k \hat{Q}_{k+1}^* - \hat{y}_k$ is the TD error. Observe that the TD error ϵ_k propagates through transformation $\mathcal{T}^{(i)}$ to influence learning. Moreover, for a fixed x , $\mathcal{T}^{(i)}(x)$ is a linear approximation of the gradients in the neighborhood of the weights. Since this approximation of the gradient is a matrix, the singular vectors of $\mathcal{T}^{(i)}$ dictate the directions of learning. Their magnitude is described by the singular values of $\mathcal{T}^{(i)}$. Each diagonal element in $\mathcal{T}^{(i)}$ is the product of the derivative of the layerwise activation functions. In other words, the singular values of $\mathcal{T}^{(i)}$ would also tend

toward zero as the number of layers in the DNN increase, otherwise known as the vanishing gradients issue (see [35] for more details).

Therefore, to address this issue, we introduce a user-defined feedback matrix $\mathbf{B}_k^{(i)}$ such that $\mathbf{B}_k^{(i)}(\mathbf{B}_k^{(i)})^T$ is positive definite, to take the role of $\mathcal{T}_k^{(i)}$ in (13). The new feedback, denoted as $\sigma_k^{(i)}$, is then given as follows:

$$\sigma_k^{(i)} = [f^{(i-1)}(\mathbf{x})] \epsilon_k \mathbf{B}_k^{(i)}. \quad (14)$$

With this definition of feedback, the new layerwise cost function denoted with $\mathcal{H}^{(i)}(\mathbf{x}(k), \mathbf{u}(k); \hat{\mathbf{W}}_k^{(i)})$ is defined to be

$$\mathcal{H}^{(i)}(\mathbf{x}(k), \mathbf{u}(k); \hat{\mathbf{W}}_k^{(i)}) = \frac{1}{2} \left[\text{tr} \left(\left(\sigma_k^{(i)} \right)^T \mathbf{P} \hat{\mathbf{W}}_k^{(i)} \right) + \lambda^{(i)} R(\hat{\mathbf{W}}_k^{(i)}) \right] \quad (15)$$

where $\text{tr}(\cdot)$ is the trace operator and \mathbf{P} is a positive definite symmetric matrix of choice which replaces the identity matrix in the update for the gradient descent rule. Note that for each layer i , $\sigma_k^{(i)}$ can be understood as the feedback provided by the overall cost $J_E(\hat{\theta})$ toward controlling the layer i of the DNN. Therefore, the layerwise cost can be interpreted as the minimization of the correlation between $\sigma_k^{(i)}$ and $\hat{\mathbf{W}}_k^{(i)}$ under the constraint that $\|\hat{\mathbf{W}}_k^{(i)}\|$ is bounded (due to the regularization). Finally, the overall cost denoted as $\mathcal{H}(\mathbf{x}(k), \mathbf{u}(k); \hat{\theta}(k))$ may be written as the sum of layerwise costs as follows:

$$\mathcal{H}(\mathbf{x}(k), \mathbf{u}(k); \hat{\theta}(k)) = \sum_{i=1}^d \mathcal{H}^{(i)}(\mathbf{x}(k), \mathbf{u}(k); \hat{\mathbf{W}}_k^{(i)}). \quad (16)$$

With the cost $\mathcal{H}(\hat{\theta}_k)$ defined in (16), the optimization problem for the direct EDL can be rewritten as follows:

$$\hat{\theta}(k) = \arg \min_{\hat{\theta}(k)} \mathbb{E}_{\mathbf{x}(k) \in \mathcal{S}, \mathbf{u}(k) \sim \rho(\mathcal{A})} [\mathcal{H}(\mathbf{x}(k), \mathbf{u}(k); \hat{\theta}(k))]. \quad (17)$$

The weight updates for each layer are defined as $\Delta_k^{(i)} = -\nabla_{\hat{\mathbf{W}}_k^{(i)}} \mathcal{H}(\mathbf{x}(k), \mathbf{u}(k); \hat{\theta}(k))$, which results in the update rule as defined in (11) with $\delta_k^{(i)}$ replaced by $\sigma_k^{(i)}$ and defined as in (14). For simplicity, from here on, we will denote $\mathcal{H}(\mathbf{x}(k), \mathbf{u}(k); \hat{\theta}(k))$ as $\mathcal{H}(\hat{\theta}(k))$ and $\mathcal{H}^{(i)}(\mathbf{x}(k), \mathbf{u}(k); \hat{\mathbf{W}}_k^{(i)})$ as $\mathcal{H}^{(i)}(\hat{\mathbf{W}}_k^{(i)})$.

Remark 1: The direct EDL-based weight updates proposed in [24] can be interpreted as an exploratory update rule, as the descent direction in the weight updates is assigned through $\mathbf{B}_k^{(i)}$ matrix. To define the matrix $\mathbf{B}_k^{(i)}$, we decompose it as $\mathbf{B}_k^{(i)} = \mathbf{U}_k^{(i)}(\Sigma_k^{(i)} + s \times \mathbf{I}^{(i)})\mathbf{V}_k^{(i)}$, where $\mathbf{U}_k^{(i)}\Sigma_k^{(i)}\mathbf{V}_k^{(i)}$ is the SVD of $\mathcal{T}_k^{(i)}$ for a fixed \mathbf{x} with $\mathbf{I}^{(i)}$ being an identity matrix of appropriate dimensions, and s is a chosen perturbation.

In the following, we analyze the EDL scheme and extract some insights on the role of the matrix $\mathbf{B}_k^{(i)}$ introduced in the update rule to train NNs of d layers.

E. Analysis of EDL for Training DNN

In this article, we approximate the optimization problem associated with (10) using (17) and introduce the idea of exploration in the learning approach through the choice of perturbations in the design matrix $\mathbf{B}^{(i)}$. In this section, we seek

to analyze different components of this learning methodology. In order to proceed further, we make the following assumptions.

Assumption 1: The objective function is a smooth function in its domain, and in particular, $J_E \in C^1$. In addition, for all $\mathbf{x} \in \mathcal{S}$, $\mathbf{u} \in \mathcal{A}$, and $\hat{\theta} \in \Omega$, the empirical cost and its gradient are bounded, i.e., $J_E(\cdot, \cdot, \hat{\theta}) \leq L$ and $(\partial J_E / \partial \hat{\theta}) \leq M$. There exists a positive constant θ_B such that for any $\hat{\theta}, \theta \in \Omega$, we have $\max(\|\hat{\theta}_k\|, \|\theta\|) \leq \theta_B$. The activation function f of the NNs is chosen such that for any $x \in \mathbb{R}$, $\|f(x)\| \leq 1$. Finally, $s \sim \mathcal{N}(0, 1)$, where \mathcal{N} denotes the normal distribution.

Lemma 1: Let Assumption 1 be true. Consider the empirical cost with regularization given as

$$H(\hat{\theta}_k) = \left[J_E(\hat{\theta}_k) + \lambda^{(i)} \sum_{i=1}^d R_k(\hat{\mathbf{W}}_k^{(i)}) \right].$$

Then, the value of $H(\hat{\theta}_k)$ can be rewritten as

$$H(\hat{\theta}_k) = \sum_{i=1}^d \frac{1}{2} \left[\text{tr} \left(\left(\delta_k^{(i)} \right)^T \hat{\mathbf{W}}_k^{(i)} \right) + \lambda^{(i)} R(\hat{\mathbf{W}}_k^{(i)}) \right] + \zeta$$

where $\zeta = J_E(\theta) - ((\partial J_E(\hat{\theta}_k) / \partial \hat{\theta}_k))^T \theta - \text{H.O.T}$, where H.O.T represents the higher order terms of $(\hat{\theta}_k - \theta)$.

The preceding lemma provides the link between the regular DRL learning cost and the error driven cost.

Theorem 1: Let Assumption 1 and Lemma 1 hold true. Consider the original optimization problem defined in (9) with the cost denoted as $H(\hat{\theta}_k)$ and consider the modified optimization problem defined in (16) with the cost denoted as $\mathcal{H}(\hat{\theta}_k)$. Let the feedback matrix be defined as $\mathbf{B}_k^{(i)} = \mathbf{U}_k^{(i)}(\Sigma_k^{(i)} + s \times \mathbf{I}^{(i)})\mathbf{V}_k^{(i)}$, where $\mathbf{U}_k^{(i)}\Sigma_k^{(i)}\mathbf{V}_k^{(i)}$ is the SVD of $\mathcal{T}_k^{(i)}$ with $\mathbf{I}^{(i)}$ being an identity matrix of appropriate dimensions. Let $s \sim \mathcal{N}(0, 1)$ be a chosen perturbation, where \mathcal{N} describes the normal distribution. Consider the conditions for $i = 1, \dots, d$ that $\|\hat{\mathbf{W}}_k^{(i)}\| \leq W_B$ and $\|f^{(i-1)}(\mathbf{x})\| < (\eta)^{1/2}$, where η is the total number of neurons at each layer. Choose \mathbf{P} such that $\|\mathbf{P}\| \leq 1$. The square difference between original cost [see (9)] and the new cost [see (16)] is then given as

$$\|H(\hat{\theta}_k) - \mathcal{H}(\hat{\theta}_k)\| \leq d \frac{|s|}{2} \|\epsilon_k\| W_B \sqrt{\eta} + \|\zeta\| \quad (18)$$

and $E\|H(\hat{\theta}_k) - \mathcal{H}(\hat{\theta}_k)\| \leq \|\zeta\|$.

Proof: See Appendix in the Supplementary Material. \square

In the above theorem, we have shown that the direct EDL problem is an approximation to the original optimization problem with residuals controlled by the choice of the perturbations. In other words, the residual will be zero when the choice of perturbations is uniformly zero.

Theorem 2: Let Assumption 1 be true. Let the gradient of the empirical cost be given as $(\partial J_E(\hat{\theta}_k) / \partial \hat{\theta}_k) = [f^{(1)}(\mathbf{x})\epsilon_k \mathcal{T}^{(1)}, \dots, f^{(d)}(\mathbf{x})\epsilon_k \mathcal{T}^{(d)}]$. Consider the weight update rule as $\hat{\theta}_{k+1} = \hat{\theta}_k + \Delta_k$, $\Delta_k = -\alpha(k) \times (\partial \mathcal{H}_k(\hat{\theta}_k) / \partial \hat{\theta}_k)$, with $\alpha(k) > 0$ and $(\partial \mathcal{H}_k(\hat{\theta}_k) / \partial \hat{\theta}_k) = [f^{(1)}(\mathbf{x})\epsilon_k \mathbf{B}^{(1)} + (\partial \lambda^{(i)} R(\hat{\mathbf{W}}^{(1)}) / \partial \hat{\mathbf{W}}^{(1)}), \dots, f^{(d)}(\mathbf{x})\epsilon_k \mathbf{B}^{(d)} + (\partial \lambda^{(i)} R(\hat{\mathbf{W}}^{(d)}) / \partial \hat{\mathbf{W}}^{(d)})]$ with the feedback matrix defined as $\mathbf{B}_k^{(i)} = \mathbf{U}_k^{(i)}(\Sigma_k^{(i)} + s \times \mathbf{I}^{(i)})\mathbf{V}_k^{(i)}$, where $\mathbf{U}_k^{(i)}\Sigma_k^{(i)}\mathbf{V}_k^{(i)}$ is the SVD of $\mathcal{T}_k^{(i)}(\mathbf{x})$ for a given batch of data \mathbf{x} with $\mathbf{I}^{(i)}$ being

an identity matrix of appropriate dimensions and s is a known perturbation. Let $B^{(i)}, i = 1, \dots, d$ be chosen with nonzero singular values then J_E converges asymptotically in the mean.

Proof: See Appendix in the Supplementary Material. \square

Remark 2: The EDL learning rule was introduced in [24], and it was shown to perform very well in learning applications. In this article, we attempt to uncover the reasons for the performance of the EDL. In this context, the analysis presented in this section reveals the following. The results of Lemma 1 and Theorem 1 reveal that the cost function which is minimized by the EDL is a first-order approximation of the empirical cost, and that the difference between these costs will reduce as the weights of the NN converge. Second, from Theorem 1, we show that the parameters of the EDL update rule can be chosen such the original cost function asymptotically converges, so that the training error (TD error) goes to zero asymptotically.

Remark 3: In the cooperative learning strategy, two NNs are utilized to estimate the objective (J_E) which is a function of TD error. Therefore, the asymptotic convergence of J_E (Theorem 2) demonstrates the convergence of the TD error. The convergence of the TD error also implies the two NNs converge to each other. In this article, we claim that this convergence provides a Q -function estimate that is sample efficient. While a theoretical proof is not considered in this article, we show the intuition with an example in Section IV and using a simulation example.

F. Algorithm

Our algorithm is summarized in Algorithm 1. The training process is performed for a total of M episodes. At the start of episode 1, we initialize the two NNs, that is, NN_1 and NN_2 . In each episode, there are K plays. In each play, we provide an action to the environment and we receive a tuple comprised of reward, state, and next state. This tuple and the action are stored in the experience replay buffer (denoted as \mathcal{B} with size N).

In this training strategy, each episode comprises repetition of two phases. *Phase 1: we first choose NN_2 to be target and NN_1 to be the actor.* For each play, the actor network is used to gather actions that are provided to the emulator. At the end of each play, the reward and the states are obtained from the emulator, and NN_1 is updated using EDL update rule. Once, the first phase is completed, we initiate phase 2. *Phase 2: we designate NN_2 to be the actor and NN_1 to be the target.* We alternatively switch between phase 1 and phase 2 for every C plays and update the two networks. For each update, we sample a batch of data from the experience replay and use it to evaluate the error and update the weights.

In the cooperative learning strategy, the two NNs are utilized to estimate the objective which is a function of TD error. Therefore, the asymptotic convergence of J_E (Theorem 2) demonstrates the convergence of the TD error. To gain more insight into the convergence of the coop strategy, consider two cars (Cars 1 and 2) starting at two different points in a city. These cars do not know their destination but are able to communicate with each other. Furthermore, both cars are provided a reward value which describes the goodness of their

Algorithm 1 Coop-RL

```

Initialize  $\hat{\theta}_1 (Q_1)$ ,  $\hat{\theta}_2 (Q_2)$ , and buffer  $\mathcal{B}$  of capacity  $N$ 
for episodes = 1, 2, 3, ..., do
  Set flag = True
  for  $k = 1, 2, \dots$ , do
    Observe image  $x_k$ , and preprocess it
    if flag == True then
      With probability  $\epsilon$  select a random action  $a_k$ ;
      otherwise  $a_k = \text{argmax}(Q_1(x_k))$ ;
      Execute  $a_k$  in emulator and observe reward  $r_k$ 
      and image  $x_{k+1}$ .
      Store the tuple  $(x_k, a_k, x_{k+1}, r_k)$  into  $\mathcal{B}$  after
      preprocessing.
      Sample batch of tuples from  $\mathcal{B}$ .
      Calculate the error  $r + \gamma Q_2 - Q_1$  and update  $\hat{\theta}_1$ .
    else
      With probability  $\epsilon$  select a random action  $a_k$ 
      otherwise  $a_k = \text{argmax}(Q_2(x_k))$ .
      Execute  $a_k$  in emulator and observe reward  $r_k$ 
      and image  $x_{k+1}$ .
      Store the tuple  $(x_k, a_k, x_{k+1}, r_k)$  into  $\mathcal{B}$  after
      preprocessing.
      Sample batch of tuples from  $\mathcal{B}$ .
      Calculate the error  $r + \gamma Q_1 - Q_2$  and update  $\hat{\theta}_2$ .
    for every  $C$  steps toggle flag.

```

location [(0) no need to move, (1) good direction, and (−1) bad direction]. The goal of these cars is to minimize the TD error, which is defined as a function of target, actor, and the reward value. Each car can either be the target or an actor. When a car is target (say Car 1), it broadcasts its location but does not move. On the other hand, the actor car (say Car 2) exploits the targets' location and the reward value (reinforcement for the actors' current location) to determine its movement. Car 2 seeks to move toward Car 1 (guided by location) while progressing toward the destination (guided by the reward). After Car 2 has moved many times, it finds a new location that is closer (depending on the initialization) to the destination than the older location. The new and better location is used to guide Car 1, and therefore, Car 2 becomes the target and Car 1 becomes the actor. Using the reward and location, Car 1 can now move many times. Next, Car 1 becomes the target and Car 2 becomes the actor and the process continues. This process of learning reaches its culmination when the TD error is small, that is, Car 1 and Car 2 have both reached a common location where the cost is zero. At this point, there is no incentive for moving (cost 0) and the cooperative strategy has converged.

IV. IMPLEMENTATION AND SIMULATION RESULTS

To substantiate the efficacy of our approach, we consider the application of playing games as a way to verify our theoretical claims. We choose a total of four games, Cartpole, Breakout, Pong, and Space Invaders. The game environments (emulators) for our analysis are provided by OpenAI and their details can be availed from the OpenAI website [36]. For all these games, we consider two NNs with equal hyperparameters.




Fig. 3. RL problem: input-output flow diagram.

Our network architecture is depicted in Fig. 3, where we use two convolutional layers and two feed-forward layers with *relu* activation function. The input of the network are images from the emulator, whereas the output of the network are the Q -values corresponding to each action. For additional details, please refer to [4].

For comparisons, we consider three realizations. First, we consider the standard DQN with one NN (target is a copy of the action network) and gradient driven updates, denoted as DQN. Second, we consider error-DQL (EDQL) which is an NN (target is a copy of the action network) with error driven updates [24]. The training strategy for these approaches is identical to the one proposed in [4]. Third, we consider the two network setups with standard gradient driven updates (G-coop). For G-coop, we follow Algorithm 1; however, we use the Gradient-driven Adam optimizer instead of the error driven update rule. Finally, we use the two network setups with error-driven update and the coop strategy. As the two NNs are used to approximate function, from here on, we will refer to output of NN_1 as Q_1 and NN_2 as Q_2 .

To record performance, we track the progression of the average reward (instantaneous reward averaged over 100 episodes) and the cumulative rewards (the average output of Q_1 and Q_2) with respect to the episodes. To record these values, we do the following at the end of each episode: we measure the performance of the network by recording the reward obtained by the network at the end of that episode. We also record the Q -values for each action at the end of the episode. We report the mean and standard deviation of these quantities from the last 100 episodes.

In these results on games presented below, the model is trained for a total of 100 000 episodes, and each episode is composed of 1000 plays where each play provides a total of four images. Therefore, the NN has to process a total of 400 million frames. To provide insights into the inner workings of our approach, we choose the simpler, cartpole example as it is memory efficient. We start by discussing performance on the cart pole example, and the NNs variants are implemented in Python 3.6 using Pytorch libraries. All simulations are performed using NVIDIA Tesla V100 SXM2 with 32 GB HBM2 and NVIDIA Tesla K80 with dual GPUs provided by the Argonne Leadership Computing Facility.

A. Example 1—Cartpole

One of most common examples in RL is that of cartpole [37]. The problem is to prevent a vertical pole from falling by moving the cart left or right. Additional details can be found in the supplementary files. In our experiments, we execute

our networks for a total of 10 000 episodes. We record the Q -function values and the cumulative reward in each episode.

We start by demonstrating that for our *coop* approach, the average reward over the last 100 episodes is 200 (which is the best reward for the cartpole problem) with a standard deviation of 0.2123, refer Table I. For this experiment, we choose the hyperparameters $C = 30$, $\text{len}(\mathcal{B}) = 5000$ with an exploration rate of $s = 0.1$ that decays exponentially. In Fig. 4(a), we plot the progression of cumulative reward achieved by the *coop* method with the exploration rate of 0.1, as a function of the total number of episodes. Cumulative rewards [refer to *coop* (0.1)] steadily increase and reaches its peak value near 200 around 2000 episodes. There are three crucial components that affect the performance of our methodology. First, we use two NNs Q_1 and Q_2 . Second, the length of the experience replay buffer, and third, the exploration rate. We start by analyzing the behavior of the two NNs, Q_1 and Q_2 . Next, we will analyze, how these components affect the performance of our methodology. We start by analyzing the convergence of the two NNs.

1) *Convergence of the Two Networks*: In our approach, we independently initialize the two networks and update them alternatively. Provided, an infinite number of iterations are provided, and the Q -function values predicted by the two networks are expected to converge to the optimal Q -function. To observe this behavior, we plot $\|Q_1 - Q_2\|$ as a function of the episodes in Fig. 4. We point out that the normed difference between Q_1 and Q_2 values come to a value around 0.005 around episode 2000. We note that 2000 episodes are where our approach reached the best cumulative reward values.

2) *Different Buffer Sizes*: Next, we analyze we compare the two NN approach (*coop*) with one NN approach (EDQL), both for various buffer sizes. In Fig. 4(c)–(f), we plot the trend of rewards with respect to different episodes for buffer sizes, 1000, 2000, 3000, and 4000. In Fig. 4(c)–(f), the markers indicate the mean of the cumulative reward over last 100 episodes. Note that the best performance is observed for both EDQL, and *coop* is obtained when the buffer size is 45 000.

However, we point that when the buffer size is reduced, the drop in performance for EDQL is much larger than *coop*. To observe this, note the trends from Fig. 5(c)–(f). We see a steady deterioration in performance. The plot shows that the use of the second NN allows to learn a representation for the data and thus provides convergence benefits. To further substantiate this, we observe Fig. 5(b). We note from Fig. 5(b) that the trend for *coop* (observe that this curve is plotted in Fig. 5(a) as well) is consistently above the curve for the EDQL. The size of the buffer is a huge memory load in practical applications (both for read/write and storage). The use of a second NN allows us to reduce this load by providing good performance for smaller buffer sizes.

3) *Effect of Exploration*: Next, we analyze the effect of different exploration rates for the *coop* methodology and compare it to that of G-coop. First, we note that *coop* with exploration rate of 0 is equivalent to G-coop. The comparison is provided in Fig. 4(b) where even a small exploration of 0.1 improves the performance of G-coop drastically. However, for a large exploration rate (1), *coop* converges quickly to




Fig. 4. (a) Mean rewards for different exploration rates (exploration rate = 0) are the standard SGD, referred as G – coop. We use ADAM optimizer for this implementation. (b) Difference between the Q -values generated by the two networks as a function of episodes.

200 but oscillates significantly more in comparison to G -coop (no exploration). On the other hand, coop (0.01) is more stable compared to coop (0.5). These results concur with the common idea that larger exploration provides for a fast but unstable learning process. On the other hand, reasonable exploration can strike the right balance between exploration and exploitation.

B. RL Games

For this part of analysis, we choose a total of three games: breakout, space invaders, and pong, details can be found in [38]. For comparisons, we choose DQN, EDQL, G -coop, and coop. The hyperparameters for this analysis are taken from [4] and kept consistent across different games. We record the mean and standard deviation values of the average reward in Table I. Note that for all these games, the network architecture is identical to that proposed in DQN paper [4]. Therefore, the input to the network is images. At the onset, we point that DQN is the baseline and provides reasonable cumulative rewards on all the games. The cumulative rewards obtained for DQN in our experiments are similar to the performance reported in [4]. Note from the table that we observe mean cumulative rewards with standard deviation 287 (78) for breakout, 1139 (158) for space invaders, 193.6 (0.71) for cartpole, and 18.9 (2.12) for pong. A slight improvement is observed with EDQL everywhere except space invaders where we observe a 14% drop in performance. The improvement is expected as EDQL introduces an inherent exploration strategy which provides convergence benefits. However, one has to note that, even with an improved exploration strategy, the best performance often depends on appropriate hyperparameter tuning which has not been done in this article.

With the introduction of the second NN for learning the Q -function, we observe a consistent improvement in performance indicated by better cumulative rewards in both G -coop and coop. For coop, we observe an improvement of 9% for breakout, 16% for space invaders, 3.3% for cartpole, and 0.2% for pong. Similarly, we observe an improvement of 8% for breakout, 16% for space invaders, 0.2% for cartpole, and 0.2% for pong.

C. Discussion

A common approach in the literature is the use of back-propagation/SGD [4]. Typically, in such an approach, labels have to be defined. However, in RL-based design, the only feedback from the environment is the reward signals. To work with a definition of error, in most DRL schemes including ours, synthetic labels in the form of an optimal Q -function are defined as a function of the reward signal. Since the optimal Q -function is unknown and must be approximated using samples from the history. The error signal is a function of an approximated optimal Q -function, which provides imprecise information early in the learning phase. If this imprecise and small error is utilized as is done with the SGD method, the vanishing gradient problem will stagnate the learning. However, our approach guarantees nonvanishing leaning signals. Furthermore, we introduce an exploratory strategy resulting in improved performance.

However, the improvement is contingent upon the idea that Q^* is well approximated. The canonical way of obtaining this approximation is as follows: we start from the current value of Q (that is, the output of the RL agent). Since the optimal Q -value can only be greater than the current value in a maximization problem, we run multiple updates starting from the current value. These updates are essentially equivalent to running a Markov chain starting from the current value to obtain a new estimate of the Q^* value. This approximated Q^* can be utilized as the target. Once new data are observed, one can repeat the procedure to get another approximation. To describe the error signal, we repeatedly replace the Q^* with this running approximation of Q^* through Q which is essentially the output of the NN. However, the main drawback of these canonical approaches is that, the process of approximating the optimal function is imperfect. To introduce enough exploration in the network for efficient exploration, one needs to maintain a large experience replay buffer.

Although, to stabilize this approximation, approaches were introduced in [39] and [40], they rely on the learning process of a single network that can still be biased by its own initialization. These constraints can be removed with the use of a second NN. If we run the same approximation procedure with two independent initialization, the solution is more robust as the optimal value is obtained only when the two networks converge to each other which is in contrast with [39] and [40]. This type of approach allows for the steady improvement in cumulative rewards we observed from Table I results. Another crucial component of our methodology is the use of buffer and the use of the second network paired with the exploration strategy thus reducing the need for a large replay buffer which we demonstrated in Fig. 4.




Fig. 5. (a) and (b) Mean rewards (with error bars) depicting variance across different buffer sizes. Buffer sizes considered for this plot are 500, 1000, 1500, 2000, 3000, 3500, 4000, and 5000. As expected, the average rewards across different buffer sizes are better for the double NN-driven method instead of a single-NN-driven method. (c)–(f) Trend of rewards with respect to different buffer sizes.

TABLE I
MEAN SCORE (STANDARD DEVIATION) OVER 100 EPISODES

| Method | Breakout | Space Invaders | Cartpole | Pong |
|--------|----------|----------------|-------------|-------------|
| DQN | 287 (78) | 1139 (158) | 193.6(0.71) | 18.9(2.12) |
| EDQL | 290 (55) | 980 (163) | 199.6(0.82) | 19.11(1.88) |
| G-coop | 314 (15) | 1174 (163) | 200(0.82) | 18.91(1.32) |
| coop | 313 (25) | 1144 (121) | 200(0.2123) | 19.15(2.1) |

An analysis of computational overhead, relative to the standard DQN approach, can be performed as follows: let $o[g]$ be the complexity of updating a NN once using gradients. Since each NN is alternatively updated, at any point of time, only one network is being updated at a given time. Let the worst case complexity of evaluating the SVD of the gradient matrix be denoted by $o[S]$. Therefore, for every batch of data, the worst case complexity of our approach is $o[g+S]$. In comparison, the complexity of a standard DQN [4] is $o[g]$.

V. CONCLUSION

In this article, we present the two NN-driven exploratory EDL approaches for addressing DRL tasks. This approach presented the following advantages.

- 1) We reduced the need for a large buffer sizes with the use of two NNs.
- 2) We demonstrated that the error driven strategy improves convergence by disentangling the learning across different layers in a deep NN. We provided both theoretical and experimental evidence to support this claim.
- 3) Our approach allows for methodical analysis of the learning problem and is extendable to diverse applications where a notion of error can be defined for an RL problem involving a dynamical process.

Although significant improvements have been shown in this article, there are major challenges that are still present. For instance, in this article, we consider a typical reinforcement episodic learning setup where the RL agent attempts to learn from rewards due to both successes and failures. In real-time control applications such as controlling autonomous cars, failures are not acceptable as each failure can lead to enormous loss of revenue and worse and the loss of human life.

Despite the reduction in requirement of buffer size, the size of experience replay buffer is still fixed at one million. Furthermore, the larger the size of the experience array, the better

the performance of the RL agent. There is a significant need of resources to achieve the required precision in learning optimal policies for a given task. Therefore, prior to deployment in real life control applications, the DRL methods must improve to the extent that learning is sufficiently precise with guaranteed convergence given limited time and resources.

ACKNOWLEDGMENT

Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the Office of Naval Research and Department of Army.

REFERENCES

- [1] M. Mahmud, M. S. Kaiser, A. Hussain, and S. Vassanelli, "Applications of deep learning and reinforcement learning to biological data," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 29, no. 6, pp. 2063–2079, Jun. 2018.
- [2] N. C. Luong et al., "Applications of deep reinforcement learning in communications and networking: A survey," *IEEE Commun. Surveys Tuts.*, vol. 21, no. 4, pp. 3133–3174, 4th Quart., 2019.
- [3] T. T. Nguyen, N. D. Nguyen, and S. Nahavandi, "Deep reinforcement learning for multiagent systems: A review of challenges, solutions, and applications," *IEEE Trans. Cybern.*, vol. 50, no. 9, pp. 3826–3839, Sep. 2020.
- [4] V. Mnih et al., "Human-level control through deep reinforcement learning," *Nature*, vol. 518, no. 7540, pp. 529–533, 2015.
- [5] K. Muhammad, A. Ullah, J. Lloret, J. D. Ser, and V. H. C. De Albuquerque, "Deep learning for safe autonomous driving: Current challenges and future directions," *IEEE Trans. Intell. Transp. Syst.*, vol. 22, no. 7, pp. 1–21, Jul. 2020.
- [6] H. Mhaskar, Q. Liao, and T. Poggio, "Learning functions: When is deep better than shallow," 2016, *arXiv:1603.00988*.
- [7] K. Iwata, "Extending the peak bandwidth of parameters for softmax selection in reinforcement learning," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 28, no. 8, pp. 1865–1877, Aug. 2017.
- [8] M. Leonetti, L. Iocchi, and P. Stone, "A synthesis of automated planning and reinforcement learning for efficient, robust decision-making," *Artif. Intell.*, vol. 241, pp. 103–130, Dec. 2016.
- [9] T. Hester and P. Stone, "Generalized model learning for reinforcement learning in factored domains," in *Proc. The 8th Int. Conf. Auto. Agents Multiagent Syst.*, vol. 2, 2009, pp. 717–724.

- [10] T. Hester, M. Quinlan, and P. Stone, "Generalized model learning for reinforcement learning on a humanoid robot," in *Proc. IEEE Int. Conf. Robot. Autom.*, May 2010, pp. 2369–2374.
- [11] A. M. Farahmand, A. Shademan, M. Jagersand, and C. Szepesvari, "Model-based and model-free reinforcement learning for visual servoing," in *Proc. IEEE Int. Conf. Robot. Autom.*, May 2009, pp. 2917–2924.
- [12] L.-J. Lin, "Self-improving reactive agents based on reinforcement learning, planning and teaching," *Mach. Learn.*, vol. 8, nos. 3–4, pp. 293–321, May 1992.
- [13] S. Kalyanakrishnan and P. Stone, "Batch reinforcement learning in a complex domain," in *Proc. 6th Int. Joint Conf. Auto. Agents Multiagent Syst. (AAMAS)*, 2007, p. 94.
- [14] F. Ruelens, B. J. Claessens, S. Vandaal, B. De Schutter, R. Babuska, and R. Belmans, "Residential demand response of thermostatically controlled loads using batch reinforcement learning," *IEEE Trans. Smart Grid*, vol. 8, no. 5, pp. 2149–2159, Sep. 2017.
- [15] M. Pieters and M. A. Wiering, "Q-learning with experience replay in a dynamic environment," in *Proc. IEEE Symp. Comput. Intell. (SSCI)*, Dec. 2016, pp. 1–8.
- [16] Y.-H. Wang, T.-H.-S. Li, and C.-J. Lin, "Backward Q-learning: The combination of Sarsa algorithm and Q-learning," *Eng. Appl. Artif. Intell.*, vol. 26, no. 9, pp. 2184–2193, Oct. 2013.
- [17] D. Horgan et al., "Distributed prioritized experience replay," 2018, *arXiv:1803.00933*.
- [18] S. Thrun and A. Schwartz, "Issues in using function approximation for reinforcement learning," in *Proc. 4th Connectionist Models Summer School*. Hillsdale, NJ, USA, 1993, pp. 255–263.
- [19] T. Schaul, J. Quan, I. Antonoglou, and D. Silver, "Prioritized experience replay," 2015, *arXiv:1511.05952*.
- [20] W.-C. Jiang, V. Narayanan, and J.-S. Li, "Model learning and knowledge sharing for cooperative multiagent systems in stochastic environment," *IEEE Trans. Cybern.*, vol. 51, no. 12, pp. 5717–5727, Dec. 2021.
- [21] Z. Zhang, D. Zhao, J. Gao, D. Wang, and Y. Dai, "FMRQ—A multiagent reinforcement learning algorithm for fully cooperative tasks," *IEEE Trans. Cybern.*, vol. 47, no. 6, pp. 1367–1379, Jun. 2017.
- [22] F. L. Da Silva, R. Glatt, and A. H. R. Costa, "MOO-MDP: An object-oriented representation for cooperative multiagent reinforcement learning," *IEEE Trans. Cybern.*, vol. 49, no. 2, pp. 567–579, Feb. 2019.
- [23] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, p. 436, Feb. 2015.
- [24] R. Krishnan, S. Jagannathan, and V. A. Samaranayake, "Direct error driven learning for deep neural networks with applications to bigdata," *Proc. Comput. Sci.*, vol. 144, pp. 89–95, Jan. 2018.
- [25] M. Hardt, B. Recht, and Y. Singer, "Train faster, generalize better: Stability of stochastic gradient descent," in *Proc. Int. Conf. Mach. Learn.*, 2016, pp. 1225–1234.
- [26] D.-H. Lee, S. Zhang, A. Fischer, and Y. Bengio, "Difference target propagation," in *Proc. Joint Eur. Conf. Mach. Learn. Knowl. Discovery Databases*. Cham, Switzerland: Springer, 2015, pp. 498–515.
- [27] A. Nøkland, "Direct feedback alignment provides learning in deep neural networks," in *Proc. Adv. Neural Inf. Process. Syst.*, 2016, pp. 1037–1045.
- [28] D. L. Donoho and C. Grimes, "Hessian eigenmaps: Locally linear embedding techniques for high-dimensional data," *Proc. Nat. Acad. Sci. USA*, vol. 100, no. 10, pp. 5591–5596, 2003.
- [29] S. Reed, H. Lee, D. Anguelov, C. Szegedy, D. Erhan, and A. Rabinovich, "Training deep neural networks on noisy labels with bootstrapping," 2014, *arXiv:1412.6596*.
- [30] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*. Cambridge, MA, USA: MIT Press, 2018.
- [31] S. Sonoda and N. Murata, "Neural network with unbounded activation functions is universal approximator," *Appl. Comput. Harmon. Anal.*, vol. 43, no. 2, pp. 233–268, Sep. 2017.
- [32] F. L. Lewis, D. Vrabie, and V. L. Syrmos, *Optimal Control*. Hoboken, NJ, USA: Wiley, 2012.
- [33] F. Lewis, S. Jagannathan, and A. Yesildirak, *Neural Network Control of Robot Manipulators and Non-Linear Systems*. Boca Raton, FL, USA: CRC Press, 1998.
- [34] D. E. Rumelhart, G. E. Hinton, and R. J. Williams, "Learning representations by back-propagating errors," *Nature*, vol. 323, no. 6088, pp. 533–536, Oct. 1986.
- [35] R. Pascanu, T. Mikolov, and Y. Bengio, "On the difficulty of training recurrent neural networks," in *Proc. Int. Conf. Mach. Learn.*, 2013, pp. 1310–1318.
- [36] G. Brockman et al., "OpenAI gym," 2016, *arXiv:1606.01540*.
- [37] S. Kumar, "Balancing a CartPole system with reinforcement learning—A tutorial," 2020, *arXiv:2006.04938*.
- [38] R. Mall, R. Langone, and J. A. K. Suykens, "Self-tuned kernel spectral clustering for large scale networks," in *Proc. IEEE Int. Conf. Big Data*, Oct. 2013, pp. 385–393.
- [39] Z. Wang, T. Schaul, M. Hessel, H. van Hasselt, M. Lanctot, and N. De Freitas, "Dueling network architectures for deep reinforcement learning," 2015, *arXiv:1511.06581*.
- [40] H. van Hasselt, A. Guez, and D. Silver, "Deep reinforcement learning with double Q-learning," 2015, *arXiv:1509.06461*.



Krishnan Raghavan (Member, IEEE) received the B.E. degree in instrumentation and control engineering from the University of Mumbai, Mumbai, India, and the master's and Ph.D. degrees from the Missouri University of Science and Technology, Rolla, MO, USA, in 2014 and 2019, respectively.

He conducted a post-doctoral research at the Mathematics and Computer Science (MCS) Division, Argonne National Laboratory (ANL), Lemont, IL, USA, from 2019 to 2022. He is currently an Assistant Computational Mathematician at the MCS Department, ANL, and has been primarily focused on machine learning and artificial intelligence. He is also affiliated with the laboratory for applied mathematics and numerical software at MCS Division. His research interests include dynamical systems, data science, machine learning, and statistics.



Vignesh Narayanan (Member, IEEE) received the B.Tech. degree in electrical and electronics engineering from SASTRA University, Thanjavur, India, in 2012, the M.Tech. degree in electrical engineering from the National Institute of Technology, Kurukshetra, India, in 2014, and the Ph.D. degree from the Missouri University of Science and Technology, Rolla, MO, USA, in 2017.

He was a Post-Doctoral Research Associate at Washington University in St. Louis, St. Louis, MO, USA. He is an Assistant Professor at the AI Institute and the Department of Computer Science and Engineering, University of South Carolina (USC), Columbia, SC, USA. He is also with the Carolina Autism and Neurodevelopment Center, USC. His research interests include dynamical systems and networks, data science, learning, and computational neuroscience.



Sarangapani Jagannathan (Fellow, IEEE) is a Professor and the Rutledge-Emerson Chair of electrical and computer engineering at the Missouri University of Science and Technology (former University of Missouri-Rolla), Rolla, MO, USA. He is also a Professor (courtesy) at the Department of Computer Science. He served as the Director for the NSF Industry/University Cooperative Research Center on Intelligent Maintenance Systems, Rolla, for 13 years. He has coauthored 189 peer-reviewed journal articles, 292 refereed IEEE conference papers, written several book chapters, authored/coauthored six books, and received 21 U.S. patents and one patent defense publication. He graduated 32 Ph.D. and 31 M.S. thesis students. His research interests include learning, adaptation and control, secure human-cyber-physical systems, prognostics, and autonomous systems/robotics.

Dr. Jagannathan is a fellow of the National Academy of Inventors, the Institute of Measurement and Control, U.K., the Institution of Engineering and Technology (IET), U.K., and the Asia-Pacific Artificial Intelligence Association. He received many awards, including the 2020 Best Associate Editor Award, the 2018 IEEE CSS Transition to Practice Award, the 2007 Boeing Pride Achievement Award, the 2001 Caterpillar Research Excellence Award, and the 2000 NSF Career Award. He was a coeditor for the IET book series on control from 2010 to 2013 and served on many editorial boards, including IEEE TRANSACTIONS ON SYSTEMS, MAN, AND CYBERNETICS, and has been on organizing committees of several IEEE conferences. He is currently an Associate Editor of IEEE TRANSACTIONS ON NEURAL NETWORKS AND LEARNING SYSTEMS and others.