Catastrophic Interference in Reinforcement Learning: A Solution Based on Context Division and Knowledge Distillation

This paper was downloaded from TechRxiv (<https://www.techrxiv.org>).

LICENSE

CC BY 4.0

SUBMISSION DATE / POSTED DATE

04-08-2021 / 09-08-2021

CITATION

Zhang, Tiantian; Wang, Xueqian; Liang, Bin; Yuan, Bo (2021): Catastrophic Interference in Reinforcement Learning: A Solution Based on Context Division and Knowledge Distillation. TechRxiv. Preprint.
<https://doi.org/10.36227/techrxiv.15105492.v1>

DOI

[10.36227/techrxiv.15105492.v1](https://doi.org/10.36227/techrxiv.15105492.v1)

Catastrophic Interference in Reinforcement Learning: A Solution Based on Context Division and Knowledge Distillation

Tiantian Zhang, Xueqian Wang, *Member, IEEE*,
Bin Liang, *Senior Member, IEEE*, and Bo Yuan*, *Senior Member, IEEE*

Abstract—The powerful learning ability of deep neural networks enables *reinforcement learning* (RL) agents to learn competent control policies directly from high-dimensional and continuous environments. In theory, to achieve stable performance, neural networks assume *i.i.d.* inputs, which unfortunately does not hold in the general RL paradigm where the training data is temporally correlated and non-stationary. This issue may lead to the phenomenon of “catastrophic interference” (a.k.a. “catastrophic forgetting”) and the collapse in performance as later training is likely to overwrite and interfere with previously learned good policies. In this paper, we introduce the concept of “context” into the single-task RL and develop a novel scheme, termed as *Context Division and Knowledge Distillation* (CDaKD) driven RL, to divide all states experienced during training into a series of contexts. Its motivation is to mitigate the challenge of aforementioned catastrophic interference in deep RL, thereby improving the stability and plasticity of RL models. At the heart of CDaKD is a value function, parameterized by a neural network feature extractor shared across all contexts, and a set of output heads, each specializing on an individual context. In CDaKD, we exploit online clustering to achieve context division, and interference is further alleviated by a knowledge distillation regularization term on the output layers for learned contexts. In addition, to effectively obtain the context division in high-dimensional state spaces (e.g., image inputs), we perform clustering in the lower-dimensional representation space of a randomly initialized convolutional encoder, which is fixed throughout training. Our results show that, with various replay memory capacities, CDaKD can consistently improve the performance of existing RL algorithms on classic OpenAI Gym tasks and the more complex high-dimensional Atari tasks, incurring only moderate computational overhead.

Index Terms—Reinforcement Learning, Catastrophic Interference, Context Division, Knowledge Distillation.

I. INTRODUCTION

IN recent years, the successful application of *deep neural networks* (DNNs) in *reinforcement learning* (RL) [1] has provided a new perspective to boost its performance on high-dimensional continuous problems. With the powerful function approximation and representation learning capabilities of

Tiantian Zhang, and Bo Yuan are with the Intelligent Computing Lab, Division of Informatics, Shenzhen International Graduate School, Tsinghua University, 518055 Shenzhen, P.R. China (e-mail: ztt19@mails.tsinghua.edu.cn; boyuan@ieee.org).

Xueqian Wang is with the Center for Artificial Intelligence and Robotics, Shenzhen International Graduate School, Tsinghua University, 518055 Shenzhen, P.R. China (e-mail: wang.xq@sz.tsinghua.edu.cn).

Bin Liang is with the Research Center for Navigation and Control, Department of Automation, Tsinghua University, 100084 Beijing, P.R. China (e-mail: bliang@tsinghua.edu.cn).

*Corresponding Author.

Fig. 1. An illustration of the catastrophic interference in the single-task RL. (a) The drift of data distributions during learning, where $P1$ – $P3$ are different data distributions and ①–④ represent distribution transitions. The agent experiences the following data distribution transitions during learning: $P1 \xrightarrow{①} P2 \xrightarrow{②} P3 \xrightarrow{③} P1 \xrightarrow{④} P3$. (b) The stability-plasticity trade-off in DNNs considers plasticity *w.r.t.* the current learning and how it degrades the effect of the old learning. (c) The learning curves where the solid line corresponds to the data distribution transitions in (a) and the dashed line shows the training performance. Before $T3$, the data distribution is gradually drifted from $P1$ to $P2$ and to $P3$. When the model fits to $P3$, the learned policies on $P1$ and $P2$ are interfered, resulting in catastrophically degraded performance when the agent encounters states from $P1$ again. Therefore, the model needs to be retrained on $P1$ (in the time period $T3 \rightarrow T4$). The same problem occurs in the time period $T4 \rightarrow T5$.

DNNs, deep RL is regarded as a milestone towards constructing autonomous systems with a higher level of understanding of the physical world. Currently, deep RL has demonstrated great potential on complex tasks, from learning to play video games directly from pixels [2], [3] to making immediate decisions on robot behavior from camera inputs [4]–[6]. However, these successes are limited and prone to catastrophic interference due to the inherent issue of DNNs in face of the non-stationary data distributions, and they rely heavily on a combination of various subtle strategies, such as experience

replay [2] and fixed target networks [3], or distributed training architecture [7]–[9].

Catastrophic interference is the primary challenge for many neural-network-based machine learning systems when learning over a non-stationary stream of data [10]. It is normally investigated in multi-task *continual learning* (CL), mainly including *supervised continual learning* (SCL) for classification tasks [11]–[16] and *continual reinforcement learning* (CRL) [11], [12], [17]–[19] for decision tasks. In the multi-task CL, the agent continually faces new tasks and the neural network may quickly fit to the data distribution of the current task, while potentially overwriting the information related to learned tasks, leading to catastrophic forgetting of the solutions of old tasks. The underlying reason behind this phenomenon is the global generalization and overlapping representation of neural networks [20], [21]. Neural networks training normally assumes that the inputs are identically and independently distributed (*i.i.d.*) from a fixed data distribution and the output targets are sampled from a fixed conditional distribution. Only when this assumption is satisfied, can positive generalization be ensured among different batches of stochastic gradient descent. However, when the data distribution is drifted during training, the information learned from old tasks may be negatively interfered or even overwritten by the newly updated weights, resulting in catastrophic interference.

Deep RL is essentially a CL problem due to its learning mode of exploring while learning [19], and it is particularly vulnerable to catastrophic interference, even within many single-task settings (such as Atari 2600 games, or even simpler classic OpenAI Gym environments) [22]–[25]. The non-stationarity of data distributions in the single-task RL is mainly attributed to the following properties of RL. Firstly, the inputs of RL are sequential observations received from the environment, which are temporally correlated. Secondly, in the progress of learning, the agent’s decision making policy changes gradually, which makes the observations non-stationary. Thirdly, RL methods rely heavily on bootstrapping, where the RL agent uses its own estimated value function as the target, making the target outputs also non-stationary. In addition, as noted in [23], replay buffers with prioritized experience replay [26] that preferentially sample experiences with higher *temporal-difference* (TD) errors will also exasperate the non-stationarity of training data. Once the distribution of training data encounters notable drift, catastrophic interference and a chain reaction are likely to occur, resulting in a sudden deterioration of the training performance, as shown in Fig. 1.

Currently, there are two major strategies for dealing with catastrophic interference in the single-task RL training: experience replay [2], [3] and local optimization [20], [24], [25]. The former usually exhibits extreme sensitivity to key parameters (*e.g.*, replay buffer size) and often requires maintaining a large experience storage memory. Furthermore, when faced with imbalanced state data, interference may still occur even if the memory is sufficiently large. The latter advocates local network updating for the data with a specific distribution instead of global generalization to reduce the representation overlap among different data distributions. The major issues are that some methods are limited in the capability of model transfer

among differently distributed data [24], [25], or can not scale to high-dimensional complex tasks [20], or require pretraining and may not be suitable for the online and incremental setting [25].

In this paper, we focus on the catastrophic interference problem caused by state distribution drift in the single-task RL. We propose a novel scheme with low buffer-size sensitivity called *Context Division and Knowledge Distillation* (CDaKD) that estimates the value function online and incrementally for each state distribution by minimizing the weighted sum of the original loss function of RL algorithms and the regularization term regarding the interference among different groups of states. The schematic architecture is shown in Fig. 3.

In order to mitigate the interference among different state distributions during model training, we introduce the concept of “*context*” into the single-task RL, and propose a novel context division strategy based on online clustering. We show that it is essential to decouple the correlations among different state distributions with this strategy to divide the state space into a series of independent contexts (each context is a set of states distributed close to each other, conceptually similar to “*task*” in the multi-task CRL). To achieve efficient and adaptive partition, we employ *Sequential K-Means Clustering* [27] to process the states encountered during training in real time. Then, we parameterize the value function by a neural network with multiple output heads commonly used in multi-task learning [18], [28], [29] in which each output head specializes on a specific context, and the feature extractor is shared across all contexts. In addition, we apply knowledge distillation as a regularization term in the objective function for value function estimation, which can preserve the learned policies while the RL agent is trained on states guided by the current policy, to further avoid the interference caused by the shared low-level representation. Furthermore, to ease the curse of dimensionality in high-dimensional state spaces, we employ a random encoder as its low-dimensional representation space can effectively capture the information about the similarity among states without any representation learning [30]. Clustering is then performed in the low-dimensional representation space of the randomly initialized convolutional encoder.

Finally, to validate the efficacy of CDaKD, we conduct extensive experiments on several OpenAI Gym standard benchmark environments containing 4 basic MDP control environments [31] and 6 high-dimensional complex Arcade Learning Environments (ALE) [32]. Compared to existing experience-replay-based and local-optimization-based methods, CDaKD features state-of-the-art performance on classic control tasks (*CartPole-v0*, *Pendulum-v0*, *CartPole-v1*, *Acrobot-v1*) and Atari tasks (*Pong*, *Breakout*, *Carnival*, *Free-way*, *Tennis*, *FishingDerby*). The main contributions of this paper are summarized as follows:

- 1) A novel context division strategy is proposed for the single-task RL. It is essential as the widely studied multi-task CRL methods cannot be used directly to reduce interference in the single-task RL due to the lack of pre-defined task boundaries. This strategy can detect contexts adaptively online, so that each context can be regarded as a task in multi-task settings. In this way, we bridge the

gap between the multi-task CRL and the single-task RL in terms of the catastrophic interference problem.

- 2) A novel RL training scheme called CDaKD based on multi-head neural networks is proposed following the context division strategy. By incorporating the knowledge distillation loss into the objective function, our method can better alleviate the interference suffered in the single-task RL than existing methods in an online and incremental manner.
- 3) A high-dimensional state representation method based on the random convolutional encoder is introduced, which further boosts the performance of CDaKD on high-dimensional complex RL tasks.
- 4) The CDaKD framework is highly flexible and can be incorporated into various RL models. Experiments on classic control tasks and high-dimensional complex tasks demonstrate the overall superiority of our method over baselines in terms of plasticity and stability, confirming the effectiveness of the context division strategy.

The rest of this paper is organized as follows. Section II reviews the relevant strategies for alleviating catastrophic interference as well as context detection and identification. Section III introduces the nature of RL in terms of continuous learning and gives a definition of catastrophic interference in the single-task RL. The details of CDaKD are shown in Section IV, and experimental results and analyses are presented in Section V. Finally, this paper is concluded in Section VI with some discussions and directions for future work.

II. RELATED WORK

Catastrophic interference within the single-task RL is a special case of CRL, which involves not only the strategies to mitigate interference but also the context detection and identification techniques.

A. Multi-task Continual Reinforcement Learning

Multi-task CRL has been an active research area with the development of a variety of RL architectures [33]. In summary, existing methods mainly consist of three categories: experience replay methods, regularization-based methods, and parameter isolation methods.

The core idea of experience replay is to store samples of previous tasks in raw format (*e.g.*, Selective Experience Replay (SER) [34], Meta Experience Replay (MER) [17], Continual Learning with Experience And Replay (CLEAR) [35]) or generate pseudo-samples from a generative model (*e.g.*, Reinforcement Pseudo Rehearsal (RePR) [36]) to maintain the knowledge about the past in the model. These previous task samples are replayed while learning a new task to alleviate interference, in the form of either being reused as model inputs for rehearsal [34], [36] or constraining the optimization of the new task loss [17], [35]. As a result, experience replay has become a very successful approach to tackling interference in CRL. However, experience replay in raw format may result in significant storage requirements for more complex CRL settings. Although the generative model can be exempted

from a replay buffer, it is still difficult to capture the overall distribution of previous tasks.

Regularization-based methods avoid storing raw inputs and thus alleviate the memory requirements, by introducing an extra regularization term into the loss function to consolidate previous knowledge while learning on new tasks. The regularization term includes penalty computing and knowledge distillation. The former focuses on reducing the chance of weights being modified. For example, Elastic Weight Consolidation (EWC) [11] and UNCertainty guided Continual LEARNING (UNCLEAR) [18] use Fisher matrix to measure the importance of weights and protect important weights on new tasks. The latter is a form of knowledge transfer [37], which expects that the model trained on a new task can still perform well on the old ones. It is often used for policy transfer from one model to another (*e.g.*, Policy Distillation [38], Genetic Policy Optimization (GPO) [39], Distillation for Continual Reinforcement learning (DisCoRL) [40]). This family of solutions is easy to implement and tends to perform well on a small number of tasks, but still faces challenges as the number of tasks increases.

Parameter isolation methods dedicate different model parameters to each task, to prevent any possible interference among tasks. Without the constraints on the size of neural networks, one can grow new branches for new tasks, while freezing previous task parameters (*e.g.*, Progressive Natural Networks (PNN) [41]). Alternatively, the architecture remains static, with fixed parts being allocated to each task. For instance, PathNet [12] uses a genetic algorithm to find a path from input to output for each task in the neural network and isolates the used network parts in parameter level from the new task training. These methods typically require networks with enormous capacity, especially when the number of tasks is large, and there is often unnecessary redundancy in the network structure, bringing a great challenge to model storage and efficiency.

B. Single-task Reinforcement Learning

Compared with the multi-task CRL, catastrophic interference in the single-task RL remains an emerging research area, which has been relatively under-explored. There are two primary aspects of previous studies: one is finding supporting evidence to confirm that catastrophic interference is indeed prevalent within a specific RL task, and the other is proposing effective strategies for dealing with it.

Researchers in DeepMind studied the learning dynamics of the single-task RL and developed a hypothesis that the characteristic coupling between learning and data generation is the main cause of interference and performance plateaus in deep RL systems [22]. Recent studies further confirmed the above hypothesis and its universality in the single-task RL through large-scale empirical studies (called *Memento experiments*) in Atari 2600 games [23]. However, none of these studies has suggested any practical solution for tackling the interference.

In order to mitigate interference, many deep RL algorithms such as DQN [3] and its variants (*e.g.*, Double DQN [42],

Rainbow [43]) employ experience replay and fixed target networks to produce approximately *i.i.d.* training data, which may quickly become intractable in terms of memory requirement as task complexity increases. Furthermore, even with sufficient memory, it is still possible to suffer from catastrophic interference due to the imbalanced distribution of experiences.

In recent studies [20], [24], [25], researchers proposed some methods based on local representation and optimization of neural networks, which showed that interference can be reduced by promoting the local updating of weights while avoiding global generalization. Sparse Representation Neural Network (SRNN) [25] induces sparse representations in neural networks by introducing a distributional regularizer, which requires a large batch of data generated by a fixed policy that covers the space for pretraining. Dynamic Self-Organizing Map (DSOM) [24] with neural networks introduces a DSOM module to induce such locality updates. These methods can reduce interference to some extent, but they may inevitably suffer from the lack of positive transfer in the representation layer, which is not desirable in complex tasks. Recently, discretizing (D-NN) and tile coding (TC-NN) were used to remap the input observations to a high-dimensional space to sparsify input features, reducing the activation overlap [20]. However, tile coding increases the dimension of inputs to a neural network, which can lead to scalability issues for spaces with high dimensionality.

C. Context Detection and Identification

Context detection and identification is a fundamental step for learning task relatedness in CL. Most multi-task CL methods mentioned above rely on well-defined task boundaries, and are usually trained on a sequence of tasks with known labels or boundaries. Existing context detection approaches commonly leverage statistics or Bayesian inference to detect task boundaries.

On the one hand, some methods tend to be reactive to a changing distribution by finding change points in the pattern of state-reward tuples (e.g., Context QL [44]), or tracking the difference between the short-term and long-term moving average rewards (e.g., CRL-Unsup [45]), or splitting a game into contexts using the undiscounted accumulated game score as a task contextualization [46]. These methods can be agile in responding to scenarios with abrupt changes among contexts or tasks, but are insensitive to smooth transitions from one context to another.

On the other hand, some more ambitious approaches try to learn a belief of the unobserved context state directly from the history of environment interactions, such as Forget-me-not Process (FMN) [47] for piecewise-repeating data generating sources, and Continual Unsupervised Representation Learning (CURL) [48] for task inference without any knowledge about task identity. However, they both need to be pretrained with the complete data when applied to CL problems, and CURL itself also requires additional techniques to deal with catastrophic interference.

Furthermore, Ghosh *et al.* [49] proposed to partition the initial state space into a finite set of contexts by performing

a K-Means clustering procedure, which can decompose more complex tasks, but cannot completely decouple the correlations among different state distributions from the perspective of interference prevention.

III. PRELIMINARIES AND PROBLEM STATEMENT

To better characterize the problem studied in this paper, some key definitions and glossaries of CRL problems are introduced in this section.

A. Definitions and Glossaries

Some important definitions of RL relevant to this paper are presented as follows.

Definition 1 (RL Paradigm [1]). *A RL problem is regarded as a Markov Decision Process (MDP), which is defined as a tuple $M = \langle \mathcal{S}, \mathcal{A}, \mathcal{P}, \mathcal{R}, \gamma \rangle$, where $\mathcal{S}$ is the set of states; $\mathcal{A}$ is the set of actions; $\mathcal{P} : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow [0, 1]$ is the environment transition probability function; $\mathcal{R} : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow \mathbb{R}$ is the reward function, and $\gamma \in [0, 1]$ is the discount factor.*

According to Definition 1, at each time step $t \in \mathbb{N}$, the agent moves from S_t to S_{t+1} with probability $P(S_{t+1}|S_t, A_t)$ after taking action A_t , and receives reward $R(S_t, A_t)$. Based on this definition, the optimization objective of value-based RL models is defined as follow:

Definition 2 (RL Optimization Objective [19]). *The optimization objective of the value-based RL is to learn a policy $\pi(a|s)$ with internal parameter $\theta \in \Theta$ that maximizes the expected long-term discounted returns for each (s, a) in time, also known as the value function:*

$$J(\pi) = Q_\pi(s, a) = \mathbb{E}_{\mathcal{P}, \pi} \left[\sum_{k=0}^{\infty} \gamma^k R(S_{t+k}, A_{t+k}) \middle| S_t = s, A_t = a \right]. \quad (1)$$

Here, the expectation is over the process that generates a history using $\mathcal{P}$ and decides actions from π until the end of the agent's lifetime.

The optimization objective in Definition 2 does not just concern itself with the current state, but also the full expected future distribution of states. As such, it is possible to overcome the catastrophic interference for RL over non-stationary data distributions. However, much of the recent work in RL has been in the so called *episodic environments*, which optimizes the episodic RL objective:

Definition 3 (Episodic RL Optimization Objective [19]). *Given some future horizon H , find a policy $\pi(a|s)$, optimizing the expected discounted returns:*

$$J_{\text{episodic}}(\pi) = Q_\pi(s, a) = \mathbb{E}_{\mathcal{P}, \pi} \left[\sum_{k=0}^{H-1} \gamma^k R(S_{t+k}, A_{t+k}) \middle| S_t = s, A_t = a \right]. \quad (2)$$

Here, to ensure the feasibility and ease of implementation of optimization, the objective is only optimized over a future horizon H until the current episode terminates.

It is clear that the episodic objective in Definition 3 is biased towards the current episode distribution while ignoring the possibly far more important future episode distributions over the agent's lifetime. Plugging in such an objective directly into

the non-stationary RL settings leads to biased optimization, which is likely to cause catastrophic interference effects.

For large scale domains, the value function is often approximated with a member of the parametric function class, such as a neural network with parameter $\theta \in \Theta$, expressed as $Q(s, a; \theta)$, which is fit online using experience samples of the form (s, a, r, s') . This experience is typically collected into a buffer $\mathcal{B}$ from which batches are later drawn at random to form a stochastic estimate of the loss:

$$\mathcal{L}(\theta) = \mathbb{E}_{\mu} \left[L \left(r + \gamma \max_{a' \in \mathcal{A}} Q(s', a'; \theta^-) - Q(s, a; \theta) \right) \right], \quad (3)$$

where $L: \mathbb{R} \rightarrow \mathbb{R}$ is the agent's loss function, and $\mu \in \mathcal{P}(\mathcal{B})$ is the distribution that defines its sampling strategy. In general, the parameter θ^- used to compute the target $Q(s', a'; \theta^-)$ is a prior copy of that used for action selection (as the settings of DQN [3]).

In addition, it is necessary to clarify some important glossaries in relation to CL.

1) *Non-stationary* [50]: a process whose state or probability distribution changes with time.

2) *Interference* [21]: a type of *influence* between two gradient-based processes with objectives J_1, J_2 , sharing parameter θ . Interference is often characterized in the first order by the inner product of their gradients:

$$\rho_{1,2} = \nabla_{\theta} J_1^T \nabla_{\theta} J_2, \quad (4)$$

and can be seen as being constructive ($\rho > 0$, *transfer*) or destructive ($\rho < 0$, *interference*), when applying a gradient update using $\nabla_{\theta} J_1$ on the value of J_2 .

3) *Catastrophic Interference* [50]: a phenomenon observed in neural networks where learning a new task significantly degrades the performance on previous tasks.

B. Problem Statement

The interference within the single-task RL can be approximately measured by the difference in TD errors before and after model update under the current policy, referred to as *Approximate Expected Interference* (AEI) [51]:

$$AEI = \mathbb{E}_{\hat{d}} \left[\delta(s, a, r, s'; \theta_t)^2 - \delta(s, a, r, s'; \theta_{t-1})^2 \right], \quad (5)$$

where $\hat{d}$ is the distribution of (s, a, r, s') under the current policy and $\delta(s, a, r, s'; \theta) = r + \gamma \max_{a' \in \mathcal{A}} Q(s', a'; \theta) - Q(s, a; \theta)$ is the TD error.

To illustrate the interaction between interference and the agent's performance during the single-task RL training, we run an experiment on *CartPole* using the DQN implemented in OpenAI Baselines¹, and set the replay buffer size N to 100, a small capacity to trigger interference to highlight its effect. We trained the agent for 300K environment steps and approximated $\hat{d}$ with a buffer containing recent transitions of capacity 10K to evaluate the AEI value according to Eq. (5) after each update. Fig. 2 shows two segments of the interference and performance curves during training from

Fig. 2. The interference (blue) and training performance (yellow) curve segments of a DQN agent on *CartPole* ($N = 100$). The interference is measured as the expectation in Eq. (5) and the performance is evaluated by the sum of discounted reward per episode. (a) Phase I with $t = 100K \sim 200K$; (b) Phase II with $t = 200K \sim 300K$.

which we can see that the performance started to oscillate when AEI started to increase (e.g., $t \approx 118K$, $t \approx 143K$, and $t \approx 175K$ in Fig. 2(a), and $t \approx 230K$ in Fig. 2(b)). In general, the performance of the agent tends to drop significantly in the presence of increasing interference. This result provides direct evidence that interference is correlated closely with the stability and plasticity of the single-task RL model.

From the analysis above, we state the problem investigated in this paper as: proposing a novel and effective training scheme for the single-task RL, to alleviate catastrophic interference and reduce performance oscillation during training, improving stability and plasticity simultaneously.

IV. THE PROPOSED METHOD

In this section, we give a detailed description of our CDaKD scheme whose architecture is shown in Fig. 3. CDaKD consists of three main components, which are jointly optimized to mitigate catastrophic interference in the single-task RL: context division, knowledge distillation, and the collaborative training of the multi-headed neural network. On the basis of CDaKD, we further propose CDaKD-RE with a random encoder for the efficient contextualization of high-dimensional state spaces.

As mentioned before, catastrophic interference is an undesirable byproduct of global updates to the neural network weights on data whose distribution changes over time. A rational solution to this issue is to estimate an individual value function for each distribution, instead of using a single value function for all distributions. When an agent updates its value estimation of a state, the update should only affect the states within the same distribution. With this intuition in mind, we

¹OpenAI Baselines is a set of high-quality implementations of RL algorithms implemented by OpenAI: <https://github.com/openai/baselines>.

s : current state $\omega(s)$: the context s belongs to $\mathcal{C} = \{c_1, \dots, c_k\}$: the context centroids set $\Omega = \{\omega_1, \dots, \omega_k\}$: the contexts set
 θ_s : feature extractor parameters θ_T : output layer parameters for context that fits current s θ_F : output layer parameters for other contexts
 Φ_R : shared feature extractor Φ_A : action scorer $\mathcal{L}_{ori}$: original loss $\mathcal{L}_D$: distillation loss $\mathcal{L}$: joint optimization loss $\oplus$: concatenation

Fig. 3. An overview of the CDaKD scheme. This framework consists of three components: 1) Context division, including state assignment and centroid update. Adaptive context division is achieved using *Sequential K-Means Clustering* online; 2) Knowledge distillation. The knowledge distillation loss ($\mathcal{L}_D$) is incorporated into the objective function ($\mathcal{L}_{ori}$) to avoid interference among contexts due to the shared feature extractor; 3) Joint optimization with a multi-head neural network, which aims to estimate the value for each $(s, a, \omega(s))$ with the joint optimization loss ($\mathcal{L} = \mathcal{L}_{ori} + \lambda\mathcal{L}_D$). In summary, our method can improve the performance by decoupling the correlations among differently distributed states and intentionally preserving the learned policies.

adopt a multi-head neural network with shared representation layers to parameterize the distribution specific value functions.

The CDaKD scheme proposed in this paper can be incorporated into any existing value-based RL methods to train a piecewise Q -function for the single-task RL. The neural network is parameterized by a shared feature extractor and a set of linear output heads, corresponding to each context. As shown in Fig. 3, the set of weights of the Q -function is denoted by $\theta = \{\theta_s, \theta_T, \theta_F\}$, where θ_s is a set of shared parameters while θ_T and θ_F are both context specific parameters: θ_T is for the context that corresponds to the current input state s , and θ_F is for others. In this section, we take the combination of CDaKD and DQN as an illustrative example.

A. Context Division

In MDPs, states (or “observations”) represent the most comprehensive information regarding the environment. To better understand the states of different distributions, we define a variable ω for a set of states that are close to each other in the state space, referred to as “context”. Formally,

$$\begin{aligned} \Omega &= (\omega_i)_{i=1}^k, \\ \mathcal{S} &= \bigcup_{i=1}^k \mathcal{S}_i, \end{aligned} \quad (6)$$

where Ω is a finite set of contexts and k is the number of contexts. For an arbitrary MDP, we partition its state space into k contexts, and all states within each context follow approximately the same distribution, to decouple the correlations among states against distribution drift. More precisely, for a partition of $\mathcal{S}$ in Eq. (6), we associate a context ω_i with each

set $\mathcal{S}_i$, so that for $s \in \mathcal{S}_i$, $\omega(s) = \omega_i$, where $\omega(s)$ can be thought of as a function of state s .

The inherent learning-while-exploring feature of RL agents leads to the fact that the agent generally does not experience all possible states of the environment while searching for the optimal policy. Thus, it is unnecessary to process the entire state space. Based on this fact, in CDaKD, we only perform context division on states experienced during training. In this paper, we employ *Sequential K-Means Clustering* [27] to achieve context detection adaptively (See Appendix A for more details).

In Fig. 3, k centroids $\mathcal{C} = \{c_1, c_2, \dots, c_k\}$ are initialized at random in the entire state space. In each subsequent time step t , we execute *State Assignment* and *Centroid Update* steps for each incoming state received from the environment, and store its corresponding transition $\{s_t, \omega(s_t), a_t, r_t, s_{t+1}, \omega(s_{t+1})\}$ into the replay buffer $\mathcal{B}$. Accordingly, in the training phase, we randomly sample a batch of transitions from $\mathcal{B}$ and train the shared feature extractor Φ_R and the specific output head Φ_A corresponding to the input state simultaneously, while conducting fine-tuning of other output heads to avoid interference on learned policies. Since we store the context label of each state in the replay buffer, there are no additional state assignments required at every update step².

Note that it is also possible to conduct context division based on the initial state distribution [49]. By contrast, we show that the partition of all states experienced during training can produce more accurate and effective context division

²In CDaKD, we only need to perform state assignment once for each state.

Fig. 4. Measuring the interference among contexts by clustering all experienced states when the agent is trained on *CartPole-v0* and *Pendulum-v0* for 400K environment steps ($k = 3$), respectively. We record the relative changes in Huber loss for all contexts when the agent is trained on a particular context. It is clear that training on a particular context generally reduces the loss on itself and increases the losses on all other contexts.

results, as the trajectories starting from the initial states within different contexts have a high likelihood of overlapping in subsequent time steps (See Appendix B for more detailed experiments and analysis).

Interference Among Contexts: We investigate the interference among contexts obtained by our context division method in details. Specifically, we measure the Huber loss of TD errors in different contexts of the game as the agent learns in other contexts, and then record the relative changes in loss before and after the agent’s learning, as shown in Fig. 4. The results show that, long-term training on any context may lead to negative generalization on all other contexts, even in such simple RL tasks (*i.e.*, *CartPole-v0* and *Pendulum-v0*).

Computational Complexity: Assuming a d -dimensional environment of k contexts, the time and space complexities of our proposed context division module to process T environment steps are $\mathcal{O}(Tkd)$ and $\mathcal{O}(kd)$, respectively.

B. Knowledge Distillation

The shared low-level representation can cause the learning in new contexts to interfere with previous learning results, leading to catastrophic interference. A relevant technique to address this issue is knowledge distillation [37], which works well for encouraging the outputs of one network to approximate the outputs of another. The concept of distillation was originally used to transfer knowledge from a complex ensemble of networks to a relatively simpler network to reduce model complexity and facilitate deployment. In CDaKD, we use it as a regularization term in value function estimation to preserve the previously learned information.

When training the model on a specific context, we need to consider two aspects of the loss function: the general loss of the current training context (denoted by $\mathcal{L}_{ori}$), and the distillation loss of other contexts (denoted by $\mathcal{L}_{\mathcal{D}}$). The former encourages the model to adapt to the current context to ensure plasticity, while the latter encourages the model to keep the memory of other contexts, preventing interference.

To incorporate CDaKD into the DQN framework, we rewrite the original loss function of DQN with the context variable ω as:

$$\mathcal{L}_{ori}(\theta_S, \theta_T) = \mathbb{E}_{\mu} \left[L \left(Q_{\tau} - Q(s, a, \omega(s); \theta_S, \theta_T) \right) \right], \quad (7)$$

where

$$Q_{\tau} = r + \gamma \max_{a' \in \mathcal{A}} Q(s', a', \omega(s'); \theta_S^-, \theta_T^-)$$

is the estimated target value of $Q(s, a, \omega(s); \theta_S, \theta_T)$ and μ is the distribution of samples, *i.e.*, $\{s, \omega(s), a, r, s', \omega(s')\} \sim \mu$, and L is the agent’s loss function.

For each of the other contexts that the environment contains, we expect the output value for each pair of (s, a) to be close to the recorded output from the original network. In knowledge distillation, we regard the learned Q -function before the current update step as the *teacher network*, expressed as $Q_t = Q(s, a, \omega_i; \theta_S^-, \theta_T^-)$, and the current network to be trained as the *student network*, expressed as $Q_s = Q(s, a, \omega_i; \theta_S, \theta_T)$, where $\omega_i \in \Omega$ except the current context $\omega(s)$. Thus, the distillation loss is defined as:

$$\mathcal{L}_{\mathcal{D}}(\theta_S, \theta_T) = \mathbb{E}_{\mu} \sum_{\omega_i \neq \omega(s), \omega_i \in \Omega} L_{\omega_i}(\theta_S, \theta_T), \quad (8)$$

Fig. 5. Illustration of CDaKD-RE. The context division is performed in the low-dimensional representation space of a random encoder. A separate RL encoder is used to work with the MLP layers to estimate the value function.

where

$$L_{\omega_i}(\theta_S, \theta_F) = L(Q_t - Q_s)$$

is the distillation loss function of the output head corresponding to context ω_i .

C. Joint Optimization Procedure

To optimize a Q -function that can guide the agent to make proper decisions on each context without being adversely affected by catastrophic interference, we combine Eqs. (7) and (8) to form a joint optimization framework. Namely, we solve the catastrophic interference problem by the following optimization objective:

$$\min_{\theta_S, \theta_T, \theta_F} \mathcal{L}_{ori}(\theta_S, \theta_T) + \lambda \mathcal{L}_D(\theta_S, \theta_F), \quad (9)$$

where $\lambda \in [0, 1]$ is a coefficient to control the trade-off between the plasticity and stability of the neural network.

The complete procedure is described in Algorithm 1. The proposed method learns the value function and context division policy in parallel. For network training, to reduce the correlations with the target and ensure the stability of model training, the target network parameter θ^- is only updated by the Q -network parameter θ every C steps and is held fixed between individual updates, as in DQN [3]. Similarly, we also adopt fixed target context centroids ($\hat{\mathcal{C}}$) to avoid the instability of RL training caused by constantly updated context centroids ($\mathcal{C}$). To simplify the model implementation, we set the updating frequency of the target context centroids to be consistent with the target network.

D. Random Encoders for High-dimensional State Space

For high-dimensional state spaces, we propose to use random encoders for efficient context division, which can map high-dimensional inputs into low-dimensional representation spaces, easing the curse of dimensionality. Although the original RL model already contains an encoder module, it is constantly updated and directly performing clustering in its representation space may introduce extra instability into context division. Therefore, on the basis of CDaKD, we exploit a dedicated random encoder module for dimension reduction. Fig. 5 gives an illustration of this updated framework called CDaKD-RE in which the structure of the random encoder $f_{\theta_{re}}$ is consistent with the underlying RL encoder, but its parameter

Algorithm 1 CDaKD: Context Division and Knowledge Distillation-Driven Reinforcement Learning

Input: Initial replay buffer $\mathcal{B}$ with capacity N ;
 Initial Q -function f_θ with random weights θ ;
 Initial target $\hat{Q}$ -function f_{θ^-} with weights $\theta^- = \theta$;
 Initial context centroids $\mathcal{C} = \{c_1, c_2, \dots, c_k\}$;
 Initial target context centroids $\hat{\mathcal{C}} = \mathcal{C}$.

Parameter: Total training steps T , the number of contexts k ,
 target update period C , learning rate α .

Output: Updated $\mathcal{C}$ and f_θ .

```

1: Initial state  $s$ ;
2: for  $t = 1, T$  do
3:   Interact with environment to obtain  $\{s_t, a_t, r_t, s_{t+1}\}$ .
4:   States assignment:  $\omega(s_t) \xleftarrow{\hat{\mathcal{C}}} s_t, \omega(s_{t+1}) \xleftarrow{\hat{\mathcal{C}}} s_{t+1}$ .
5:   Store transition  $\{s_t, \omega(s_t), a_t, r_t, s_{t+1}, \omega(s_{t+1})\}$  in  $\mathcal{B}$ .
6:   Context centroids update:  $\mathcal{C} \leftarrow SKM(s_t, \mathcal{C})$ .
7:   Joint optimization:
      Sample mini-batch  $\{s_i, \omega(s_i), a_i, r_i, s'_i, \omega(s'_i)\}_{i=1}^m$ ;
      Calculate  $\mathcal{L}_{ori}, \mathcal{L}_D$  according to Eqs. (7) and (8);
      Update parameter:  $\theta \leftarrow \theta - \alpha \nabla_\theta (\mathcal{L}_{ori} + \lambda \mathcal{L}_D)$ .
8:   if  $t \bmod C == 0$  then
9:      $\theta^- = \theta, \hat{\mathcal{C}} = \mathcal{C}$ .
10:  end if
11: end for
```

θ_{re} is randomly initialized and fixed throughout training. We provide the full procedure of CDaKD-RE in Appendix C.

The main motivation of using random encoders arises from the observation that distances in the representation space of random encoder are adequate for finding similar states without any representation learning [30]. That is, the representation space of a random encoder can effectively capture the information about the similarity among states without any representation learning (See Appendix C).

V. EXPERIMENTS AND EVALUATIONS

In this section, we conduct comprehensive experiments on several standard benchmarks from OpenAI Gym³ containing 4 classic control tasks and 6 high-dimensional complex Atari games to demonstrate the effectiveness of our method⁴. Several state-of-the-art methods such as experience replay based methods (e.g., DQN [3], Rainbow [43]), local optimization based methods (e.g., SRNN [25]), and other context division techniques (e.g., Game Score (GS) [46], Initial States (IS) [49]) are employed as our baselines.

A. Datasets

Classic Control [31] contains 4 classic control tasks: *CartPole-v0*, *Pendulum-v0*, *CartPole-v1*, *Acrobot-v1*, where the dimensions of state spaces are in the range of 3 to 6. The maximum time steps are 200 for *CartPole-v0* and *Pendulum-v0*, and 500 for *CartPole-v1* and *Acrobot-v1*. Meanwhile, the

³OpenAI Gym is a publicly available released implementation repository of RL environments: <https://github.com/openai/gym>.

⁴Experiment code: please refer to supplementary material.

reward thresholds used to determine the success of tasks are 195.0, 475.0 and -100.0 for *CartPole-v0*, *CartPole-v1* and *Acrobot-v1*, respectively, while the threshold for *Pendulum-v0* is not available. We choose these domains as they are well-understood and relatively simple, suitable for highlighting the mechanism and verifying the effectiveness of our method in a straightforward manner.

Atari Games [32] contain 6 image-level complex tasks: *Pong*, *Breakout*, *Carnival*, *Freeway*, *Tennis*, *FishingDerby*, where the observation is the screenshot represented by an RGB image of size $210 \times 160 \times 3$. We choose these domains to further demonstrate the scalability of our method on high-dimensional complex tasks.

B. Implementation

Network Structure. For the 4 classic control tasks, we employ a fully-connected layer as the feature extractor and a fully-connected layer as the multi-head action scorer, following the network configuration for this type of tasks in OpenAI Baseline. For the 6 Atari games, we employ the similar convolution neural network as [43], [52] for feature extracting (*i.e.*, RL Encoder in Fig. 5) and two fully-connected layers as the multi-head action scorer. In addition, for random encoders, we use convolutional neural networks with the same structure as the underlying RL methods, but we do not update its randomly initialized parameters during the training process [30]. Details of the networks can be found in Appendix D.

Parameter Setting. In CDaKD, there are two key parameters: λ and k . To simplify parameter setting, we set λ in accordance with the exploration proportion ϵ of the agent in all experiments: $\lambda = 1 - \epsilon$, due to the inverse relationship between them in training. During the early training, ϵ is close to 1, and the model is normally inaccurate with little interference, and a small λ (close to 0) can promote plasticity construction of the neural network. Then, ϵ gradually approaches 0 during the subsequent training, and the model has accumulated more and more useful information, while interference is also likely to occur. Consequently, smoothly increasing λ is needed to ensure stability and avoid catastrophic interference. Meanwhile, we set k to 3 for all classic control tasks, and 4 for all Atari games. It is suggested that larger k values should be used for complex environments to achieve reasonable state decoupling effect. In CDaKD-RE, there is an extra parameter: the output dimension d of the random encoder. We set d to 50 as in [30], which has been shown to be both efficient and effective. More details of the parameter setting can be found in Appendix D. For classic control tasks, we evaluate the training performance using the average episode returns every 10K time steps for *CartPole-v0*, *Pendulum-v0*, and 20K time steps for *CartPole-v1*, *Acrobot-v1*. For Atari games, the time step range for performance evaluation is 200K. All experiment results reported are the average episode returns over 5 independent runs.

C. Baselines

We compare our method with five state-of-the-art methods, including DQN [3] and Rainbow [43] based on experience replay, SRNN [25] based on local optimization for alleviating

catastrophic interference, and two context division techniques (*i.e.*, Game Score (GS) [46], Initial States (IS) [49]) for efficient contextualization. We briefly introduce these baselines in the following.

- **DQN** [3] is a representative algorithm of Deep RL, which reduces catastrophic interference using experience replay and fixed target networks. We use the DQN agent implemented in OpenAI Baselines.
- **Rainbow** [43] is the upgraded version of DQN containing six extensions, including a prioritized replay buffer [26], n-step returns [1] and distributional learning [8] for stable RL training. The Rainbow agent is implemented in Google’s Dopamine framework⁵.
- **SRNN** [25] employs a distributional regularizer to induce sparse representations in neural networks to avert catastrophic interference in the single-task RL.
- **CDaKD-GS** is our extension to DQN using the CDaKD scheme where the context division is based on the undiscounted accumulated game score in [46] instead of all experienced states.
- **CDaKD-IS** is another extension to DQN using the CDaKD scheme where the context division is based on the initial states distribution as in [49] instead of all experienced states.

D. Evaluation Metrics

Following the convention in previous studies [7]–[9], [23], [43], we employ the average training episode returns $\mathcal{R}_T$ to evaluate our method during training:

$$\mathcal{R}_T = \frac{1}{M} \sum_{i=1}^M \sum_{j=0}^{J_i} R_{ij} \quad (10)$$

where M is the number of episodes experienced within each evaluation period; J_i is the total time steps in episode i ; R_{ij} is the reward received at time step j in episode i .

E. Results

The results on 4 control tasks are presented in Fig. 6, showing the learning curves during training for each task with three levels of replay buffer capacity. Note that CDaKD-ES is our proposed scheme where the context division is based on all experienced states⁶. In general, CDaKD-ES is clearly superior to all baselines in terms of plasticity and stability, especially when the replay buffer capacity is small (*e.g.*, $N = 100$) or even without experience replay (*i.e.*, $N = 1$). In most tasks, CDaKD-ES achieves near optimal performance as well as good stability even without any experience replay. For *Pendulum-v0* and *Acrobot-v1*, a large replay buffer (*e.g.*, $N = 50,000$) can help DQN and SRNN escape from catastrophic interference. However, this is not the case for two *CartPole* tasks where the agents exhibit fast initial learning but then encounter collapse in performance.

⁵Dopamine is a research framework developed by Google for fast prototyping of reinforcement learning algorithms: <https://github.com/google/dopamine>.

⁶In this paper, all appearances of CDaKD refer to CDaKD-ES unless otherwise stated.

Fig. 6. Learning curves on classic control tasks with different replay buffer capacities N . The solid lines and shaded regions represent the mean and standard deviation, respectively, across five runs. In summary, CDaKD-ES significantly outperforms baselines in terms of plasticity and stability.

Fig. 7. Learning curves on Atari games with different replay buffer capacities N . The solid lines and shaded regions represent the mean and standard deviation, respectively, across five runs. It is worth noting that the green line and the red line of Freeway in (a) and the blue line and orange line of Freeway in (b) are overlapping, respectively.

TABLE I
STATISTICS OF THE HIGHEST CUMULATIVE SCORES ACHIEVED DURING TRAINING ON ATARI GAMES
(BASED ON THE AVERAGE PERFORMANCE OF FIVE RUNS IN FIG. 7).

Method	DQN		Rainbow		DQN+CDaKD-RE		Rainbow+CDaKD-RE	
N	10,000	1,000,000	10,000	1,000,000	10,000	1,000,000	10,000	1,000,000
<i>Pong</i>	16.2	15.0	17.1	19.1	16.7	16.3	17.7	19.0
<i>Breakout</i>	71.6	80.9	59.3	62.1	90.5	97.4	67.6	66.7
<i>Carnival</i>	4327.0	4589.0	4529.7	4788.0	5073.1	4610.8	4928.2	5289.0
<i>Freeway</i>	0.5	19.5	25.7	33.2	0.5	25.8	32.3	33.3
<i>Tennis</i>	-13.2	-1.3	0.0	0.0	-10.7	-0.4	0.0	0.0
<i>FishingDerby</i>	-23.8	9.7	-2.7	29.4	-23.5	7.5	-2.4	30.8

TABLE II
STATISTICS OF THE MAXIMUM DETERIORATION RATIOS SUFFERED DURING TRAINING ON ATARI GAMES
(BASED ON THE AVERAGE PERFORMANCE OF FIVE RUNS IN FIG. 7).

Method	DQN		Rainbow		DQN+CDaKD-RE		Rainbow+CDaKD-RE	
N	10,000	1,000,000	10,000	1,000,000	10,000	1,000,000	10,000	1,000,000
<i>Pong</i>	0.041	0.017	0.026	0.005	0.016	0.016	0.013	0.005
<i>Breakout</i>	0.009	0.084	0.071	0.256	0.035	0.094	0.071	0.250
<i>Carnival</i>	0.012	0.007	0.041	0.080	0.009	0.018	0.045	0.076
<i>Freeway</i>	1.000	0.176	0.100	0.002	1.000	0.046	0.067	0.003
<i>Tennis</i>	0.586	0.118	0.493	0.176	0.370	0.181	0.325	0.188
<i>FishingDerby</i>	0.096	0.107	0.042	0.019	0.113	0.082	0.030	0.014
Average	0.291	0.085	0.129	0.090	0.257	0.073	0.092	0.089

The learning curves on 6 Atari games are shown in Fig. 7. Moreover, the highest cumulative returns (indicator of plasticity) achieved during the training of each task are summarized in Table I, and the maximum deterioration ratios (indicator of stability) compared to the previous maximal episode returns are given in Table II. Overall, for high-dimensional image inputs, the training performance of the original RL algorithms can be noticeably improved with our CDaKD-RE scheme. In Fig. 7, our method significantly outperforms DQN on 7 out of 12 tasks, being comparable with DQN on the rest 5 tasks. Similarly, our method outperforms Rainbow on 8 out of 12 tasks, being comparable with Rainbow on the rest 4 tasks. Furthermore, as shown in Table I and Table II, CDaKD-RE achieves higher maximum cumulative scores and less performance degradation in most tasks compared to its counterparts. Among the 24 training settings, only 2 maximum cumulative scores achieved by CDaKD-RE are slightly lower than those of baselines. In terms of the average performance degradation ratio, DQN and Rainbow incorporated with CDaKD-RE surpass the original RL methods by 3.4% and 3.7% ($N = 10,000$), respectively. Note that, even with a large memory ($N = 1,000,000$), CDaKD-RE still shows certain advantages over the baselines.

Moreover, we observe that: 1) DQN and SRNN agents exhibit high sensitivity to the replay buffer capacity. They generally perform well with a large buffer (except on *CartPole-v1*), but their performance deteriorates significantly when the buffer capacity is reduced. DQN performs worst among all baselines when there is no experience replay, as both DQN and SRNN may face severe data drift when the buffer is small, and approximately *i.i.d.* training data are required to avoid possible interference in training. 2) CDaKD-GS and CDaKD-IS leverage the cumulative game score and initial state distribution to partition contexts, respectively. However, neither the game score nor the initial state distribution is a

perfect determining factor for context boundaries, making it difficult to achieve the necessary decoupling of differently distributed states. Furthermore, they require prior knowledge about the scores for each level of the game or the initial state distribution of the environment.

In summary, the proposed techniques containing context division based on the clustering of all experienced states, and knowledge distillation in multi-head neural networks can effectively eliminate catastrophic interference caused by data drift in the single-task RL. In addition, our method leverages a fixed randomly initialized encoder to characterize the similarity among states in the low-dimensional representation space, which can be used to partition contexts effectively for high-dimensional environments.

F. Analysis

1) *Ablation Study*: Since our method can be regarded as an extension to existing RL methods (e.g., DQN [3]), with three novel components (i.e., adaptive context division by online clustering, knowledge distillation, and the multi-head neural network), the ablation experiments are designed as follows:

- **No clustering** means using a random partition of the raw state space instead of adaptive context division by online clustering.
- **No distillation** means removing the distillation loss function $\mathcal{L}_{\mathcal{D}}(\theta_S, \theta_F)$ from the joint objective function in Eq. (9) (i.e., $\lambda = 0$).
- **No multi-head** means removing the context division module and optimize the neural network with a single-head output (i.e., $k = 1$). Here, the distillation term is represented as the distillation of the network before each update of the output head.

The results of ablation experiments are shown in Fig. 8, using classic control tasks for the convenience of validation. From Fig. 8, the following facts can be observed: 1) Across

Fig. 8. Comparisons of CDaKD (red) with DQN (green) and three different ablations (other colors), on each individual task. The solid lines and shaded regions represent the mean and standard deviation, respectively, across five runs. In summary, the removal of either online clustering or multi-head outputs reduces performance across most tasks. Furthermore, removing the distillation regularizer results in performance degradation on all tasks.

all settings, the overall performance of DQN is the worst, showing the effectiveness of the three components introduced for coping with catastrophic interference in the single-task RL, although the contribution of each component varies substantially per task; 2) Removing online clustering from the context division module is likely to damage the performance in most cases; 3) Without the multi-head component, our model is equivalent to a DQN with an extra distillation loss, which can result in better performance in general than DQN; 4) Removing knowledge distillation makes the performance deteriorate on almost all tasks, indicating that knowledge distillation is a key element in our method.

2) *Parameter Analysis*: There are two critical parameters in CDaKD: λ and k . By its nature, λ is related to the training process. Since we need to preserve the learned good policies during training, it is intuitive to gradually increase λ until its value reaches 1. The reason is that, in early-stage training, the model has not learned any sufficiently useful information, so the distillation constraint can be ignored. With the progress of training, the model starts to acquire more and more valuable information and needs to pay serious attention to interference to protect the learned good policies while learning further. In our experiments, we recommend to set λ to be inversely proportional to the exploration proportion ϵ , and the results in Figs. 6 and 7 have demonstrated the simplicity and effectiveness of this setting.

To investigate the effect of k , which is related to the complexity of learning tasks, we conduct experiments with different k values ($k \in \{1, 2, 3, 4, 5\}$) and the results are shown in Fig. 9. In our experiments, $k = 3$ is a reasonably good choice for *CartPole-v0* and *CartPole-v1*, while $k = 1$ is good for CDaKD on *Acrobot-v1*. It is worth noting that, on *Pendulum-v0*, our method achieves similar performance with k set to 2, 3, 4, and 5, respectively, but without any satisfactory performance. A possible explanation is that the agents failed

Fig. 9. Parameter sensitivity analysis w.r.t the number of contexts k . Experiments are conducted on 4 classic control tasks with different k values ($N = 1$). The solid lines and shaded regions represent the mean and standard deviation, respectively, across five runs.

to learn any useful information due to the limited state space explored in the early training, leading to the failure of further learning.

In summary, we can make the following statements: 1) The performance of DQN combined with CDaKD is significantly better than the original DQN regardless of the specific k value, confirming the effectiveness of our CDaKD scheme; 2) For $k > 1$, better performance of CDaKD can be expected. However, large k values are not always desirable as larger k values will result in more fine-grained context divisions and more complex neural networks with a large amount of output heads, making the model unlikely to converge satisfactorily within a limited number of training steps. Thus, we recommend to set the value of k by taking into consideration the state-space structure of specific tasks.

Fig. 10. Training curves tracking the agent’s average loss and average predicted action-value for 400K environment steps in *Pendulum-v0* ($N = 100$ and $k = 3$, see Fig. 6(b) for corresponding curves). (a) Each point is the average loss achieved per training iteration; (b) Average maximum predicted action-value of agents on a held-out set of states; (c) Average maximum predicted action-value of each output head in CDaKD.

(a) Training time

(b) Computational complexity

Fig. 11. Comparison of computational efficiency: (a) Training time of each agent to achieve its performance for 400K environment steps in *Pendulum-v0* ($N = 100$, see Fig. 6(b) for corresponding learning curves); (b) Number of FLOPs used by each agent at 10M environment steps in *Breakout*. Here, we only take into account forward and backward passes through neural network layers (See Fig. 7 for corresponding learning curves).

3) *Convergence Analysis*: To analyze convergence, we track the agent’s average loss and average predicted action-value during training progress. According to Fig. 10, we can conclude that: 1) Our method has better convergence and stability in face of interference compared with original RL algorithms (See Fig. 10(a) and Fig. 10(b)); 2) For a held-out set of states, the average maximum predicted action-value of each output head reflects the difference as expected, and the final output of the CDaKD agent is synthesised based on all output heads.

4) *Computational Efficiency*: Our methods are computationally efficient in that: 1) In each time step, the extra context division module only needs to compute the distances between the current state and k context centroids, which is negligible compared with the computational cost of RL itself; 2) Only $k - 1$ extra output heads are added to the neural network, in which the increased computation is acceptable compared with the representation layer; 3) There are no gradient updates through the random encoder; 4) There is no unnecessary distance computation for finding the corresponding context at every update step as the context label for each state is stored in the replay buffer. Fig. 11 shows the training time of each agent on *Pendulum* and the floating point operations (FLOPs) executed by agents on *Breakout*, respectively.

VI. CONCLUSION AND FUTURE WORK

In this paper, we propose a competent scheme CDaKD to tackle the inherent challenge of catastrophic interference in the single-task RL. The core idea is to partition all states experienced during training into a set of contexts using online clustering techniques and simultaneously estimate the context-specific value function with a multi-head neural network as well as a knowledge distillation loss to mitigate the interference across contexts. Furthermore, we introduced a random convolutional encoder to enhance the context division for high-dimensional complex tasks. Our method can effectively decouple the correlations among differently distributed states and can be easily incorporated into various value-based RL models. Experiments on several benchmark tasks show that our method can significantly outperform state-of-the-art RL methods and dramatically reduce the memory requirement of existing RL methods.

In the future, we aim to incorporate our method into policy-based RL models to reduce the interference during training by applying weight or functional regularization on policies. Furthermore, we will investigate a more challenging setting called continual RL in non-stationary environments [45]. This setting is a more realistic representation of the real-world scenarios and includes abrupt changes or smooth transitions on dynamics, or even the dynamics itself is shuffled.

ACKNOWLEDGMENT

The work presented in this paper was supported by the National Natural Science Foundation of China (U1713214).

REFERENCES

- [1] R. S. Sutton and A. G. Barto, *Reinforcement learning: An introduction*. MIT press, 2018.
- [2] V. Mnih, K. Kavukcuoglu, D. Silver, A. Graves, I. Antonoglou, D. Wierstra, and M. Riedmiller, “Playing atari with deep reinforcement learning,” in *Proceedings of the 27th Conference on Neural Information Processing Systems*, 2013.
- [3] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski et al., “Human-level control through deep reinforcement learning,” *Nature*, vol. 518, no. 7540, pp. 529–533, 2015.
- [4] A. Faust, K. Oslund, O. Ramirez, A. Francis, L. Tapia, M. Fiser, and J. Davidson, “Prm-rl: Long-range robotic navigation tasks by combining reinforcement learning and sampling-based planning,” in *Proceedings of the International Conference on Robotics and Automation*, 2018, pp. 5113–5120.
- [5] H.-T. L. Chiang, A. Faust, M. Fiser, and A. Francis, “Learning navigation behaviors end-to-end with autorl,” *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 2007–2014, 2019.
- [6] A. Francis, A. Faust, H.-T. L. Chiang, J. Hsu, J. C. Kew, M. Fiser, and T.-W. E. Lee, “Long-range indoor navigation with prm-rl,” *IEEE Transactions on Robotics*, vol. 36, no. 4, pp. 1115–1134, 2020.
- [7] V. Mnih, A. P. Badia, M. Mirza, A. Graves, T. Lillicrap, T. Harley, D. Silver, and K. Kavukcuoglu, “Asynchronous methods for deep reinforcement learning,” in *Proceedings of the 33rd International Conference on Machine Learning*, vol. 48, 2016, pp. 1928–1937.
- [8] M. G. Bellemare, W. Dabney, and R. Munos, “A distributional perspective on reinforcement learning,” in *Proceedings of the 32nd International Conference on Machine Learning*, 2017, pp. 449–458.
- [9] L. Espeholt, H. Soyer, R. Munos, K. Simonyan, V. Mnih, T. Ward, Y. Doron, V. Firoiu, T. Harley, I. Dunning et al., “Impala: Scalable distributed deep-rl with importance weighted actor-learner architectures,” in *Proceedings of the 35th International Conference on Machine Learning*. PMLR, 2018, pp. 1407–1416.
- [10] M. McCloskey and N. J. Cohen, “Catastrophic interference in connectionist networks: The sequential learning problem,” in *Psychology of learning and motivation*. Elsevier, 1989, vol. 24, pp. 109–165.
- [11] J. Kirkpatrick, R. Pascanu, N. Rabinowitz, J. Veness, G. Desjardins, A. A. Rusu et al., “Overcoming catastrophic forgetting in neural networks,” in *Proceedings of the National Academy of Sciences*, vol. 114, no. 13, 2017, pp. 3521–3526.
- [12] 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,” *arXiv preprint arXiv:1701.08734*, 2017.
- [13] D. Lopez-Paz and M. Ranzato, “Gradient episodic memory for continual learning,” in *Proceedings of the 31st Conference on Neural Information Processing Systems*, 2017, pp. 6467–6476.
- [14] S.-A. Rebuffi, A. Kolesnikov, G. Sperl, and C. H. Lampert, “icarl: Incremental classifier and representation learning,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 2001–2010.
- [15] A. Mallya and S. Lazebnik, “Packnet: Adding multiple tasks to a single network by iterative pruning,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 7765–7773.
- [16] M. Delange, R. Aljundi, M. Masana, S. Parisot, X. Jia, A. Leonardis, G. Slabaugh, and T. Tuytelaars, “A continual learning survey: Defying forgetting in classification tasks,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021.
- [17] M. Riemer, I. Cases, R. Ajemian, M. Liu, I. Rish et al., “Learning to learn without forgetting by maximizing transfer and minimizing interference,” in *Proceedings of the 7th International Conference on Learning Representations*, 2019.
- [18] S. Kessler, J. Parker-Holder, P. Ball, S. Zohren, and S. J. Roberts, “Unclear: A straightforward method for continual reinforcement learning,” in *Proceedings of the 37th International Conference on Machine Learning*, 2020.
- [19] K. Khetarpal, M. Riemer, I. Rish, and D. Precup, “Towards continual reinforcement learning: A review and perspectives,” *arXiv preprint arXiv:2012.13490*, 2020.
- [20] S. Ghiassian, B. Rafiee, Y. L. Lo, and A. White, “Improving performance in reinforcement learning by breaking generalization in neural networks,” in *Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems*, 2020.
- [21] E. Bengio, J. Pineau, and D. Precup, “Interference and generalization in temporal difference learning,” in *Proceedings of the 37th International Conference on Machine Learning*, 2020, pp. 767–777.
- [22] T. Schaul, D. Borsa, J. Modayil, and R. Pascanu, “Ray interference: a source of plateaus in deep reinforcement learning,” *arXiv preprint arXiv:1904.11455*, 2019.
- [23] W. Fedus, D. Ghosh, J. D. Martin, M. G. Bellemare, Y. Bengio, and H. Larochelle, “On catastrophic interference in atari 2600 games,” *arXiv preprint arXiv:2002.12499*, 2020.
- [24] Y. L. Lo and S. Ghiassian, “Overcoming catastrophic interference in online reinforcement learning with dynamic self-organizing maps,” in *Proceedings of the 33rd Conference on Neural Information Processing Systems*, 2019.
- [25] V. Liu, R. Kumaraswamy, L. Le, and M. White, “The utility of sparse representations for control in reinforcement learning,” in *Proceedings of the 33rd AAAI Conference on Artificial Intelligence*, vol. 33, no. 01, 2019, pp. 4384–4391.
- [26] T. Schaul, J. Quan, I. Antonoglou, and D. Silver, “Prioritized experience replay,” in *Proceedings of the 4th International Conference on Learning Representations*, 2016.
- [27] J. G. Dias and M. J. Cortinhal, “The skm algorithm: A k-means algorithm for clustering sequential data,” in *Proceedings of the Ibero-American Conference on Artificial Intelligence*, 2008, pp. 173–182.
- [28] F. Zenke, B. Poole, and S. Ganguli, “Continual learning through synaptic intelligence,” in *Proceedings of the 34th International Conference on Machine Learning*, 2017, pp. 3987–3995.
- [29] S. Golkar, M. Kagan, and K. Cho, “Continual learning via neural pruning,” *arXiv preprint arXiv:1903.04476*, 2019.
- [30] Y. Seo, L. Chen, J. Shin, H. Lee, P. Abbeel, and K. Lee, “State entropy maximization with random encoders for efficient exploration,” *arXiv preprint arXiv:2102.09430*, 2021.
- [31] G. Brockman, V. Cheung, L. Pettersson, J. Schneider, J. Schulman, J. Tang, and W. Zaremba, “Openai gym,” *arXiv preprint arXiv:1606.01540*, 2016.
- [32] M. G. Bellemare, Y. Naddaf, J. Veness, and M. Bowling, “The arcade learning environment: An evaluation platform for general agents,” *Journal of Artificial Intelligence Research*, vol. 47, pp. 253–279, 2013.
- [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,” *Information Fusion*, vol. 58, pp. 52–68, 2020.
- [34] D. Isele and A. Cosgun, “Selective experience replay for lifelong learning,” in *Proceedings of the 32nd AAAI Conference on Artificial Intelligence*, vol. 32, no. 1, 2018.
- [35] D. Rolnick, A. Ahuja, J. Schwarz, T. P. Lillicrap, and G. Wayne, “Experience replay for continual learning,” in *Proceedings of the 33th Conference on Neural Information Processing Systems*, 2019.
- [36] C. Atkinson, B. McCane, L. Szymanski, and A. Robins, “Pseudo-rehearsal: Achieving deep reinforcement learning without catastrophic forgetting,” *Neurocomputing*, vol. 428, pp. 291–307, 2021.
- [37] G. Hinton, O. Vinyals, and J. Dean, “Distilling the knowledge in a neural network,” in *Workshop of the Conference on Neural Information Processing Systems*, 2015.
- [38] A. A. Rusu, S. G. Colmenarejo, C. Gulcehre, G. Desjardins, J. Kirkpatrick, R. Pascanu et al., “Policy distillation,” in *Proceedings of the 5th International Conference on Learning Representations*, 2016.
- [39] T. Gangwani and J. Peng, “Policy optimization by genetic distillation,” in *Proceedings of the 7th International Conference on Learning Representations*, 2018.
- [40] R. Traoré, H. Caselles-Dupré, T. Lesort, T. Sun, G. Cai et al., “Discorl: Continual reinforcement learning via policy distillation,” in *Workshop of the Conference on Neural Information Processing Systems*, 2019.
- [41] A. A. Rusu, N. C. Rabinowitz, G. Desjardins, H. Soyer, J. Kirkpatrick, K. Kavukcuoglu, R. Pascanu, and R. Hadsell, “Progressive neural networks,” *arXiv preprint arXiv:1606.04671*, 2016.
- [42] H. Van Hasselt, A. Guez, and D. Silver, “Deep reinforcement learning with double q-learning,” in *Proceedings of the 30th AAAI Conference on Artificial Intelligence*, vol. 30, no. 1, 2016.
- [43] M. Hessel, J. Modayil, H. Van Hasselt, T. Schaul, G. Ostrovski, W. Dabney, D. Horgan, B. Piot, M. Azar, and D. Silver, “Rainbow: Combining improvements in deep reinforcement learning,” in *Proceedings of the 32nd AAAI Conference on Artificial Intelligence*, vol. 32, no. 1, 2018.

- [44] S. Padakandla, K. Prabuchandran, and S. Bhatnagar, "Reinforcement learning algorithm for non-stationary environments," *Applied Intelligence*, vol. 50, no. 11, pp. 3590–3606, 2020.
- [45] V. Lomonaco, K. Desai, E. Culurciello, and D. Maltoni, "Continual reinforcement learning in 3d non-stationary environments," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 2020, pp. 248–249.
- [46] V. Jain, W. Fedus, H. Larochelle, D. Precup, and M. G. Bellemare, "Algorithmic improvements for deep reinforcement learning applied to interactive fiction," in *Proceedings of the 34th AAAI Conference on Artificial Intelligence*, vol. 34, no. 04, 2020, pp. 4328–4336.
- [47] K. Milan, J. Veness, J. Kirkpatrick, M. Bowling, A. Koop, and D. Hassabis, "The forget-me-not process," in *Proceedings of the 30th Conference on Neural Information Processing Systems*, vol. 29, 2016, pp. 3702–3710.
- [48] D. Rao, F. Visin, A. A. Rusu, Y. W. Teh, R. Pascanu, and R. Hadsell, "Continual unsupervised representation learning," in *Proceedings of the 33th Conference on Neural Information Processing Systems*, 2019.
- [49] D. Ghosh, A. Singh, A. Rajeswaran, V. Kumar, and S. Levine, "Divide-and-conquer reinforcement learning," in *Proceedings of the 7th International Conference on Learning Representations*, 2018.
- [50] R. Hadsell, D. Rao, A. A. Rusu, and R. Pascanu, "Embracing change: Continual learning in deep neural networks," *Trends in Cognitive Sciences*, 2020.
- [51] V. Liu, A. White, H. Yao, and M. White, "Towards a practical measure of interference for reinforcement learning," *arXiv preprint arXiv:2007.03807*, 2020.
- [52] P. S. Castro, S. Moitra, C. Gelada, S. Kumar, and M. G. Bellemare, "Dopamine: A Research Framework for Deep Reinforcement Learning," *arXiv preprint arXiv:1812.06110*, 2018.

Tiantian Zhang is currently working toward her Ph.D degree in control science and engineering with the Department of Automation, Tsinghua University, Beijing, China. She received her B.Sc degree in automation from the Department of Information Science and Technology, Central South University, Chasha, China, in 2015, and the M.Sc. degree in control engineering with the Department of Automation, Tsinghua University, Beijing, China, in 2018. Her research interests include data science and reinforcement learning.

Xueqian Wang (Member, IEEE) received his M.Sc and Ph.D. degrees in control science and engineering from Harbin Institute of Technology (HIT), Harbin, China, in 2005 and 2010, respectively. From June 2010 to February 2014, he was a Postdoctoral Researcher at HIT. From March 2014 to November 2019, he was an Associate Professor with the Division of Informatics, Shenzhen International Graduate School, Tsinghua University, Shenzhen, China. He is currently a Professor and the leader of the Center for Artificial Intelligence and Robotics, Shenzhen International Graduate School, Tsinghua University. His Research interests include Robot dynamics and control, teleoperation, intelligent decision-making and game-playing, and fault diagnosis.

Bin Liang (Senior Member, IEEE) is currently a Professor with the Research Center for Navigation and Control, Department of Automation, Tsinghua University, Beijing, China. He received his B.Sc. and M.Sc. degrees in control engineering from the Honors College of Northwestern Polytechnical University, Xi'an, China, in 1989 and 1991, respectively, and Ph.D. degree in control engineering from the Department of Precision Instrument, Tsinghua University, Beijing, China, in 1994. From 1994 to 2003, he held his positions as a Postdoctoral Researcher, an Associate Researcher, and a Researcher with China Academy of Space Technology (CAST), Beijing, China. From 2003 to 2007, he held his positions as a researcher and an Assistant Chief Engineer with the China Aerospace Science and Technology Corporation, Beijing. His research interests include modeling and control of intelligent robotic systems, teleoperation, intelligent sensing technology.

Bo Yuan (Senior Member, IEEE) received the B.E. degree in computer science from the Nanjing University of Science and Technology, Nanjing, China, in 1998, and the M.Sc. and Ph.D. degrees in computer science from The University of Queensland (UQ), St Lucia, QLD, Australia, in 2002 and 2006, respectively. From 2006 to 2007, he was a Research Officer on a project funded by the Australian Research Council, UQ. He is currently an Associate Professor with the Division of Informatics, Shenzhen International Graduate School, Tsinghua University, Shenzhen, China. He has authored or coauthored more than 110 articles in refereed international conferences and journals. His research interests include data science, evolutionary computation, and reinforcement learning.

APPENDIX A SEQUENTIAL K-MEANS CLUSTERING

The process of *Sequential K-Means Clustering* for the current state s is shown in Algorithm 2. Each context centroid c_i is the average of all of the states s closest to c_i . In order to get a better initialization of $\mathcal{C}$, we can perform offline *K-Means* clustering on all states experienced before training starts and set the results of centroids as the initial $\mathcal{C}$. Then, *Sequential K-Means Clustering* is performed in subsequent time steps.

Algorithm 2 SKM: Sequential K-Means Clustering

Input: Current state s ;
Initial context centroids $\mathcal{C} = \{c_1, c_2, \dots, c_k\}$.
Output: Updated centroids $\mathcal{C}$.

- 1: Count the number of samples in each contexts:
 $\mathcal{N} = \{n_1, n_2, \dots, n_k\}$
- 2: **if** s is closest to centroid c_i **then**
- 3: Increment n_i : $n_i = n_i + 1$;
- 4: Update c_i : $c_i = c_i + (1/n_i) * (s - c_i)$;
- 5: **end if**
- 6: **return** $\mathcal{C}$.

Fig. 12. Two-dimensional t-SNE results of context division according to different kinds of clustering objects on *CartPole-v0*: (a) ISC vs (b) ESC. Here, the colors represent different contexts of which the points represent the states within the corresponding context.

APPENDIX B INITIAL STATES VS ALL EXPERIENCED STATES

We compare the effects of two kinds of context division techniques:

- **ISC** (Initial States Clustering) means performing context division using *K-Means* on the samples sampled from the initial states distribution, referring to [49];
- **ESC** (Experienced States Clustering) is our proposed performing context division using *Sequential K-Means Clustering* on all states experienced during training process.

We run the DQN agent incorporated with CDaKD with the above two clustering techniques on *CartPole-v0*, respectively, and visualize the two-dimensional t-SNE results of context division in Fig. 12. It can be clearly observed that the contexts divided by ESC are relatively independent and there is almost no overlapping area among contexts, achieving effective decoupling among states with different distributions.

Algorithm 3 CDaKD-RE: Context Division with Random Encoder and Knowledge Distillation-Driven Reinforcement Learning

Input: Initial replay buffer $\mathcal{B}$ with capacity N ;
Initial Q -function f_θ with random weights θ ;
Initial target $\hat{Q}$ -function f_{θ^-} with weights $\theta^- = \theta$;
Initial random encoder $f_{\theta_{re}}$ with weights θ_{re} ;
Initial context centroids $\mathcal{C} = \{c_1, c_2, \dots, c_k\}$;
Initial target context centroids $\hat{\mathcal{C}} = \mathcal{C}$.

Parameter: Total training steps T , the number of contexts k , the output dimension of random encoder d , target update period C , learning rate α .

Output: Updated $\mathcal{C}$ and f_θ .

- 1: Initial state s ;
- 2: **for** $t = 1, T$ **do**
- 3: Interact with environment to obtain $\{s_t, a_t, r_t, s_{t+1}\}$.
- 4: State encoding: get the fixed representation for s_t and s_{t+1} , $y_t = f_{\theta_{re}}(s_t)$, $y_{t+1} = f_{\theta_{re}}(s_{t+1})$.
- 5: States assignment: $\omega(s_t) \xleftarrow{\hat{\mathcal{C}}} y_t$, $\omega(s_{t+1}) \xleftarrow{\hat{\mathcal{C}}} y_{t+1}$.
- 6: Store transition $\{s_t, \omega(s_t), a_t, r_t, s_{t+1}, \omega(s_{t+1})\}$ in $\mathcal{B}$.
- 7: Context centroids update: $\mathcal{C} \leftarrow SKM(y_t, \mathcal{C})$.
- 8: Joint optimization:
 Sample mini-batch $\{s_i, \omega(s_i), a_i, r_i, s'_i, \omega(s'_i)\}_{i=1}^m$;
 Calculate $\mathcal{L}_{ori}$, $\mathcal{L}_{\mathcal{D}}$ according to Eqs. (7) and (8);
 Update parameter: $\theta \leftarrow \theta - \alpha \nabla_\theta (\mathcal{L}_{ori} + \lambda \mathcal{L}_{\mathcal{D}})$.
- 9: **if** $t \bmod C == 0$ **then**
- 10: $\theta^- = \theta$, $\hat{\mathcal{C}} = \mathcal{C}$
- 11: **end if**
- 12: **end for**

By contrast, there are obvious overlapping areas among the contexts divided by ISC, since the trajectories starting from the initial states within different contexts have a high likelihood of overlapping in subsequent time steps, which is not desirable for reducing interference on neural network training among differently distributed states.

APPENDIX C RANDOM ENCODER

We find the *K-nearest neighbors* of some specific states by measuring the distances in the low-dimensional representation space produced by a randomly initialized encoder (Random Encoder) on *Breakout*. The results are shown in Fig. 13 from which we can observe that the raw images corresponding to the *K-nearest neighbors* of the source image in the representation space of the random encoder demonstrate significant similarities. We provide the full procedure of CDaKD with the random encoder in Algorithm 3.

APPENDIX D IMPLEMENTATION DETAILS

To ensure the fairness of comparison, our results compare the agents based on the underlying RL model with the same hyperparameters and neural network architecture. We provide a full list of neural networks architecture of the underlying RL models in Table III and summarize our choices for common key hyperparameters in Table IV.

Fig. 13. Two-dimensional t-SNE visualization of K -nearest neighbors of states found by measuring distances in the representation space of a Random Encoder on *Breakout*. We observe that the representation space of a randomly initialized encoder effectively captures information about similarity between states.

TABLE III
THE NEURAL NETWORKS ARCHITECTURE OF UNDERLYING RL MODELS USED IN EXPERIMENTS.

Tasks	Layer	Input	Filter size	Stride	Num filters	Activation	Output
Classic Control Tasks	FC1	Dimension of state space	-	-	64	Tanh	64
	FC2	64	-	-	Number of actions	Linear	Number of actions
	Conv1	$84 \times 84 \times 4$	8×8	4	32	ReLU	$20 \times 20 \times 32$
Atari Games	Conv2	$20 \times 20 \times 32$	4×4	2	64	ReLU	$9 \times 9 \times 64$
	Conv3	$9 \times 9 \times 64$	3×3	1	64	ReLU	$7 \times 7 \times 64$
	FC4	$7 \times 7 \times 64$	-	-	512	ReLU	512
	FC5	512	-	-	18	Linear	18

TABLE IV
THE COMMON HYPERPARAMETERS OF EACH METHOD USED IN EXPERIMENTS.

Hyperparameter	Classic Control Tasks	Atari Games
Training time step	400,000 steps for <i>CartPole-v0</i> and <i>Pendulum-v0</i> 1,000,000 steps for <i>CartPole-v1</i> and <i>Acrobot-v1</i>	10,000,000 steps
Training ϵ	0.02	0.01
ϵ decay schedule	40,000 steps for <i>CartPole-v0</i> and <i>Pendulum-v0</i> 100,000 steps for <i>CartPole-v1</i> and <i>Acrobot-v1</i>	250,000 steps
Min. history to start learning	1,000 steps	20,000 steps
Target network update frequency	1,000 steps	8,000 steps
Batch size	32	32
Learning rate α	0.0005	0.00025 for DQN and DQN+CDaKD-RE 0.0000625 for Rainbow and Rainbow+CDaKD-RE

APPENDIX E CALCULATION OF MODEL COMPLEXITY

Computational Complexity of Context Division. At each time step, *SKM* only needs to calculate the distances between the current state and k context centroids. Given a d -dimensional state space and T -step environment interactions, the time complexity of context division is $\mathcal{O}(Tkd)$. At the same time, since only k additional context centroids need to be stored for clustering, the space complexity is $\mathcal{O}(kd)$.

Calculation of Floating Point Operations. We obtain the number of operations per forward pass for all layers in the encoder (denoted by E) and the number of operations per forward pass for all MLP layers in each output head (denoted by

M), as in <https://openai.com/blog/ai-and-compute/>. Therefore, the number of FLOPs of CDaKD-RE is:

$$bI(E + kM) + 2bI(E + kM) + T(E + kM) + TE$$

where b is the batch size; T is the number of environment steps; I is the number of training updates. The first two terms are for the forward and backward passes required in training updates, respectively. The latter two terms are for the forward passes required to compute the policy action and obtain the low-dimensional representation from the random encoder, respectively. In our experiments: $b = 32$, $T = 1e7$, $I = 0.25e7$, $k = 4$, $E = 28.582$ MFLOPs, $M = 3.420$ MFLOPs for Rainbow and 3.215 MFLOPs for DQN.