

IMPLICIT UNDER-PARAMETERIZATION INHIBITS DATA-EFFICIENT DEEP REINFORCEMENT LEARNING

Aviral Kumar^{*1,2}, Rishabh Agarwal^{*2,3}, Dibya Ghosh¹, Sergey Levine^{1,2}

¹UC Berkeley, ²Google Research, ³MILA, Université de Montréal

ABSTRACT

We identify an implicit under-parameterization phenomenon in value-based deep RL methods that use bootstrapping: when value functions, approximated using deep neural networks, are trained with gradient descent using iterated regression onto target values generated by previous instances of the value network, more gradient updates decrease the expressivity of the current value network. We characterize this loss of expressivity via a drop in the rank of the learned value network features, and show that this typically corresponds to a performance drop. We demonstrate this phenomenon on Atari and Gym benchmarks, in both offline and online RL settings. We formally analyze this phenomenon and show that it results from a pathological interaction between bootstrapping and gradient-based optimization. We further show that mitigating implicit under-parameterization by controlling rank collapse can improve performance.

1 INTRODUCTION

Many pervasive deep reinforcement learning (RL) algorithms estimate value functions using bootstrapping, that is, by sequentially fitting value functions to target value estimates generated from the value function learned in the previous iteration. Despite high-profile achievements (Silver et al., 2017), these algorithms are highly unreliable due to poorly understood optimization issues. Although a number of hypotheses have been proposed to explain these issues (Achiam et al., 2019; Bengio et al., 2020; Fu et al., 2019; Igl et al., 2020; Liu et al., 2018; Kumar et al., 2020a), a complete understanding remains elusive.

We identify an “implicit under-parameterization” phenomenon that emerges when value networks are trained using gradient descent combined with bootstrapping. This phenomenon manifests as an excessive aliasing of features learned by the value network across states, which is exacerbated with more gradient updates. While the supervised deep learning literature suggests that some feature aliasing is desirable for generalization (e.g., Gunasekar et al., 2017; Arora et al., 2019), implicit under-parameterization exhibits more pronounced aliasing than in supervised learning. This over-aliasing causes an otherwise expressive value network to *implicitly* behave as an *under-parameterized* network, often resulting in poor performance.

Implicit under-parameterization becomes aggravated when the rate of data re-use is increased, restricting the sample efficiency of deep RL methods. In online RL, increasing the number of gradient steps in between data collection steps for *data-efficient* RL (Fu et al., 2019; Fedus et al., 2020b) causes the problem to emerge more frequently. In the extreme case when no additional data is collected, referred to as *offline* RL (Lange et al., 2012; Agarwal et al., 2020; Levine et al., 2020), implicit under-parameterization manifests consistently, limiting the viability of offline methods.

We demonstrate the existence of implicit under-parameterization in common value-based deep RL methods, including Q-learning (Mnih et al., 2015; Hessel et al., 2018) and actor-critic (Haarnoja et al., 2018), as well as neural fitted-Q iteration (Riedmiller, 2005; Ernst et al., 2005). To isolate the issue, we study the effective rank of the features in the penultimate layer of the value network (Section 3). We observe that after an initial learning period, the rank of the learned features drops steeply. As the rank decreases, the ability of the features to fit subsequent target values and the optimal value function generally deteriorates and results in a sharp decrease in performance (Section 3.1).

^{*}Equal Contribution. Correspondence to Aviral Kumar < aviralk@berkeley.edu > and Rishabh Agarwal < rishabhagarwal@google.com >.




Figure 1: **Implicit under-parameterization.** Schematic diagram depicting the emergence of an *effective rank* collapse in deep Q-learning. Minimizing TD errors using gradient descent with deep neural network Q-function leads to a collapse in the effective rank of the learned features Φ , which is exacerbated with further training.

To better understand the emergence of implicit under-parameterization, we formally study the dynamics of Q-learning under two distinct models of neural net behavior (Section 4): kernel regression (Jacot et al., 2018; Mobahi et al., 2020) and deep linear networks (Arora et al., 2018). We corroborate the existence of this phenomenon in both models, and show that implicit under-parameterization stems from a pathological interaction between bootstrapping and the implicit regularization of gradient descent. Since value networks are trained to regress towards targets generated by a previous version of the same model, this leads to a sequence of value networks of potentially decreasing expressivity, which can result in degenerate behavior and a drop in performance.

The main contribution of this work is the identification of implicit under-parameterization in deep RL methods that use bootstrapping. Empirically, we demonstrate a collapse in the rank of the learned features during training, and show it typically corresponds to a drop in performance in the Atari (Bellemare et al., 2013) and continuous control Gym (Brockman et al., 2016) benchmarks in both the offline and data-efficient online RL settings. We verify the emergence of this phenomenon theoretically and characterize settings where implicit under-parameterization can emerge. We then show that mitigating this phenomenon via a simple penalty on the singular values of the learned features improves performance of value-based RL methods in the offline setting on Atari.

2 PRELIMINARIES

The goal in RL is to maximize long-term discounted reward in a Markov decision process (MDP), defined as a tuple $(\mathcal{S}, \mathcal{A}, R, P, \gamma)$ (Puterman, 1994), with state space \mathcal{S} , action space \mathcal{A} , a reward function $R(s, a)$, transition dynamics $P(s'|s, a)$ and a discount factor $\gamma \in [0, 1)$. The Q-function $Q^\pi(s, a)$ for a policy $\pi(a|s)$, is the expected long-term discounted reward obtained by executing action a at state s and following $\pi(a|s)$ thereafter, $Q^\pi(s, a) := \mathbb{E}[\sum_{t=0}^{\infty} \gamma^t R(s_t, a_t)]$. $Q^\pi(s, a)$ is the fixed point of the Bellman operator \mathcal{T}^π , $\forall s, a: \mathcal{T}^\pi Q(s, a) := R(s, a) + \gamma \mathbb{E}_{s' \sim P(\cdot|s, a), a' \sim \pi(\cdot|s')} [Q(s', a')]$, which can be written in vector form as: $\mathbf{Q}^\pi = \mathbf{R} + \gamma P^\pi \mathbf{Q}^\pi$. The optimal Q-function, $Q^*(s, a)$, is the fixed point of the Bellman optimality operator \mathcal{T} : $\mathcal{T}Q(s, a) := R(s, a) + \gamma \mathbb{E}_{s' \sim P(\cdot|s, a)} [\max_{a'} Q(s', a')]$.

Practical Q-learning methods (e.g., Mnih et al., 2015; Hessel et al., 2018; Haarnoja et al., 2018) convert the Bellman equation into an bootstrapping-based objective for training a Q-network, Q_θ , via gradient descent. This objective, known as mean-squared temporal difference (TD) error, is given by: $L(\theta) = \sum_{s, a} (R(s, a) + \gamma \bar{Q}_\theta(s', a') - Q_\theta(s, a))^2$, where \bar{Q}_θ is a delayed copy of the Q-function, typically referred to as the *target network*. These methods train Q-networks via gradient descent and slowly update the target network via Polyak averaging on its parameters. We refer the output of the penultimate layer of the deep Q-network as the learned *feature matrix* Φ , such that $Q(s, a) = \mathbf{w}^T \Phi(s, a)$, where $\mathbf{w} \in \mathbb{R}^d$ and $\Phi \in \mathbb{R}^{|\mathcal{S}||\mathcal{A}| \times d}$.

For simplicity of analysis, we abstract deep Q-learning methods into a generic fitted Q-iteration (FQI) framework (Ernst et al., 2005). We refer to FQI with neural nets as neural FQI (Riedmiller, 2005). In the k -th fitting iteration, FQI trains the Q-function, Q_k , to match the target values, $\mathbf{y}_k = \mathbf{R} + \gamma P^\pi \mathbf{Q}_{k-1}$ generated using previous Q-function, Q_{k-1} (Algorithm 1). Practical methods can be instantiated as variants of FQI, with different target update styles, different optimizers, etc.

Algorithm 1 Fitted Q-Iteration (FQI)

- 1: Initialize Q-network Q_θ , buffer μ .
- 2: **for** fitting iteration k in $\{1, \dots, N\}$ **do**
- 3: Compute $Q_\theta(s, a)$ and target values $y_k(s, a) = r + \gamma \max_{a'} Q_{k-1}(s', a')$ on $\{(s, a)\} \sim \mu$ for training
- 4: Minimize TD error for Q_θ via $t = 1, \dots, T$ gradient descent updates, $\min_{\theta} (Q_\theta(s, a) - y_k)^2$
- 5: **end for**

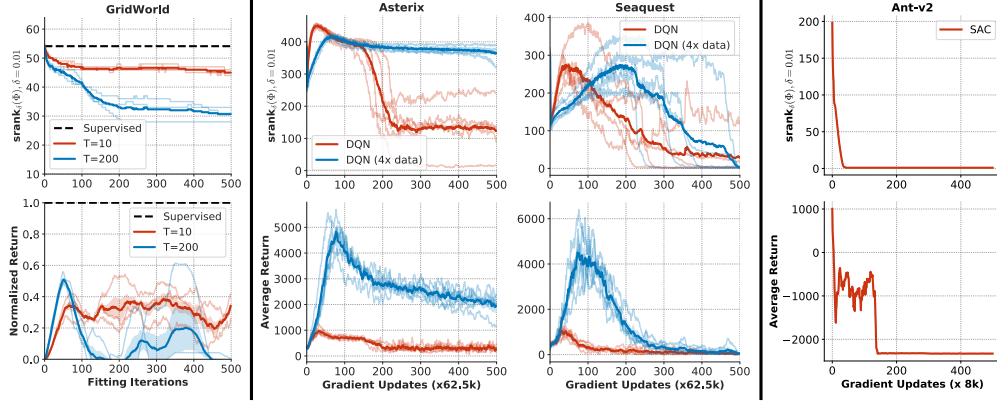


Figure 2: **Offline RL.** $\text{srnk}_\delta(\Phi)$ and performance of neural FQI on gridworld, DQN on Atari and SAC on Gym environments in the offline RL setting. Note that low rank (top row) generally corresponds to worse policy performance (bottom row). Rank collapse is worse with more gradient steps per fitting iteration ($T=10$ vs. 200 on gridworld). Even when a larger, high coverage dataset is used, marked as DQN (4x data), rank collapse occurs (for Asterix also see Figure A.2 for a complete figure with a larger number of gradient updates).

3 IMPLICIT UNDER-PARAMETERIZATION IN DEEP Q-LEARNING

In this section, we empirically demonstrate the existence of implicit under-parameterization in deep RL methods that use bootstrapping. We characterize implicit under-parameterization in terms of the *effective rank* (Yang et al., 2019) of the features learned by a Q-network. The effective rank of the feature matrix Φ , for a threshold δ (we choose $\delta = 0.01$), denoted as $\text{srnk}_\delta(\Phi)$, is given by $\text{srnk}_\delta(\Phi) = \min \left\{ k : \frac{\sum_{i=1}^k \sigma_i(\Phi)}{\sum_{i=1}^d \sigma_i(\Phi)} \geq 1 - \delta \right\}$, where $\{\sigma_i(\Phi)\}$ are the singular values of Φ in decreasing order, i.e., $\sigma_1 \geq \dots \geq \sigma_d \geq 0$. Intuitively, $\text{srnk}_\delta(\Phi)$ represents the number of “effective” unique components of the feature matrix Φ that form the basis for linearly approximating the Q-values. When the network maps different states to orthogonal feature vectors, then $\text{srnk}_\delta(\Phi)$ has high values close to d . When the network “aliases” state-action pairs by mapping them to a smaller subspace, Φ has only a few active singular directions, and $\text{srnk}_\delta(\Phi)$ takes on a small value.

Definition 1. Implicit under-parameterization refers to a reduction in the effective rank of the features, $\text{srnk}_\delta(\Phi)$, that occurs implicitly as a by-product of learning deep neural network Q-functions.

While rank decrease also occurs in supervised learning, it is usually beneficial for obtaining generalizable solutions (Gunasekar et al., 2017; Arora et al., 2019). However, we will show that in deep Q-learning, an interaction between bootstrapping and gradient descent can lead to more aggressive rank reduction (or rank collapse), which can hurt performance.

Experimental setup. To study implicit under-parameterization empirically, we compute $\text{srnk}_\delta(\Phi)$ on a minibatch of state-action pairs sampled i.i.d. from the training data (i.e., the dataset in the offline setting, and the replay buffer in the online setting). We investigate offline and online RL settings on benchmarks including Atari games (Bellemare et al., 2013) and Gym environments (Brockman et al., 2016). We also utilize gridworlds described by Fu et al. (2019) to compare the learned Q-function against the oracle solution computed using tabular value iteration. We evaluate DQN (Mnih et al., 2015) on gridworld and Atari and SAC (Haarnoja et al., 2018) on Gym domains.

Data-efficient offline RL. In offline RL, our goal is to learn effective policies by performing Q-learning on a fixed dataset of transitions. We investigate the presence of rank collapse when deep Q-learning is used with broad state coverage offline datasets from Agarwal et al. (2020). In the top row of Figure 2, we show that after an initial learning period, $\text{srnk}_\delta(\Phi)$ decreases in all domains (Atari, Gym and the gridworld). The final value of $\text{srnk}_\delta(\Phi)$ is often quite small – e.g., in Atari, only 20-100 singular components are active for 512-dimensional features, implying significant underutilization of network capacity. Since under-parameterization is *implicitly* induced by the learning process, even high-capacity value networks behave as low-capacity networks as more training is performed with a bootstrapped objective (e.g., mean squared TD error).

On the gridworld environment, regressing to Q^* using supervised regression results in a much higher $\text{srnk}_\delta(\Phi)$ (black dashed line in Figure 2(left)) than when using neural FQI. On Atari, even when a 4x larger offline dataset with much broader coverage is used (blue line in Figure 2), rank collapse still persists, indicating that implicit under-parameterization is not due to limited offline dataset size.

Figure 2 (2nd row) illustrates that policy performance generally deteriorates as $\text{srnk}(\Phi)$ drops, and eventually collapses simultaneously with the rank collapse. While we do not claim that implicit under-parameterization is the only issue in deep Q-learning, the results in Figure 2 show that the emergence of this under-parameterization is *strongly* associated with poor performance.

To prevent confounding from the distribution mismatch between the learned policy and the offline dataset, which often affects the performance of Q-learning methods, we also study CQL (Kumar et al., 2020b), an offline RL algorithm designed to handle distribution mismatch. We find a similar degradation in effective rank and performance for CQL (Figure A.3), implying that under-parameterization does not stem from distribution mismatch and arises even when the resulting policy is within the behavior distribution (though the policy may not be exactly pick actions observed in the dataset). We provide more evidence in Atari and Gym domains in Appendix A.1.

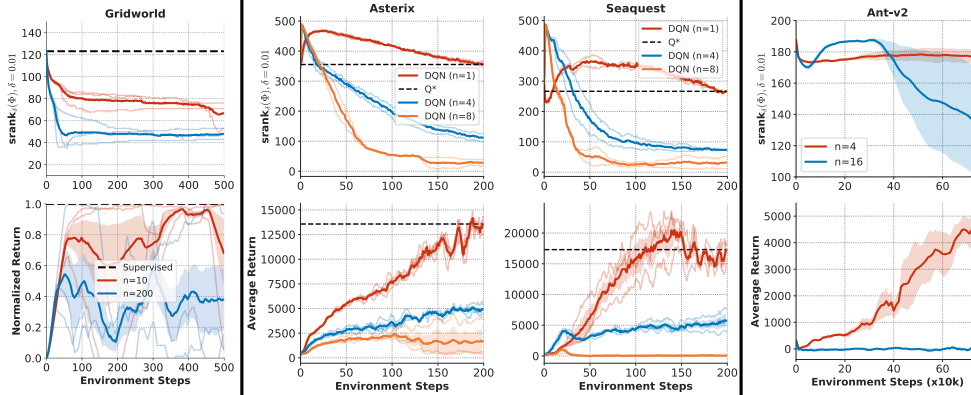


Figure 3: **Data Efficient Online RL.** $\text{srnk}_\delta(\Phi)$ and performance of neural FQI on gridworld, DQN on Atari and SAC on Gym domains in the online RL setting, with varying numbers of gradient steps per environment step (n). Rank collapse happens earlier with more gradient steps, and the corresponding performance is poor.

Data-efficient online RL. Deep Q-learning methods typically use very few gradient updates (n) per environment step (e.g., DQN takes 1 update every 4 steps on Atari, $n = 0.25$). Improving the sample efficiency of these methods requires increasing n to utilize the replay data more effectively. However, we find that using larger values of n results in higher levels of rank collapse as well as performance degradation. In the top row of Figure 3, we show that larger values of n lead to a more aggressive drop in $\text{srnk}_\delta(\Phi)$ (red vs. blue/orange lines), and that rank continues to decrease with more training. Furthermore, the bottom row illustrates that larger values of n result in worse performance, corroborating Fu et al. (2019); Fedus et al. (2020b). We find similar results with the Rainbow algorithm (Hessel et al., 2018) (Appendix A.2). As in the offline setting, directly regressing to Q^* via supervised learning does not cause rank collapse (black line in Figure 3).

3.1 UNDERSTANDING IMPLICIT UNDER-PARAMETERIZATION AND ITS IMPLICATIONS

How does implicit under-parameterization degrade performance? Having established the presence of rank collapse in data-efficient RL, we now discuss how it can adversely affect performance. As the effective rank of the network features Φ decreases, so does the network’s ability to fit the subsequent target values, and eventually results in inability to fit Q^* . In the gridworld domain, we measure this loss of expressivity by measuring the error in fitting oracle-computed Q^* values via a linear transformation of Φ . When rank collapse occurs, the error in fitting Q^* steadily increases during training, and the consequent network is not able to predict Q^* at all by the end of training (Figure 4a) – this entails a drop in performance. In Atari domains, we do not have access to Q^* , and so we instead measure TD error, that is, the error in fitting the target value estimates, $\mathbf{R} + \gamma P^\pi \mathbf{Q}_k$. In SEAQUEST, as rank decreases, the TD error increases (Figure 4b) and the value function is unable to fit the target values, culminating in a performance plateau (Figure 3). This observation is consistent across other environments; we present further supporting evidence in Appendix A.4.

Does bootstrapping cause implicit under-parameterization? We perform a number of controlled experiments in the gridworld and Atari environments to isolate the connection between rank collapse and bootstrapping. We first remove confounding issues of poor network initialization (Fedus et al., 2020a) and non-stationarity (Igl et al., 2020) by showing that rank collapse occurs even when the Q-network is re-initialized from scratch at the start of each fitting iteration (Figure 4c). To show that

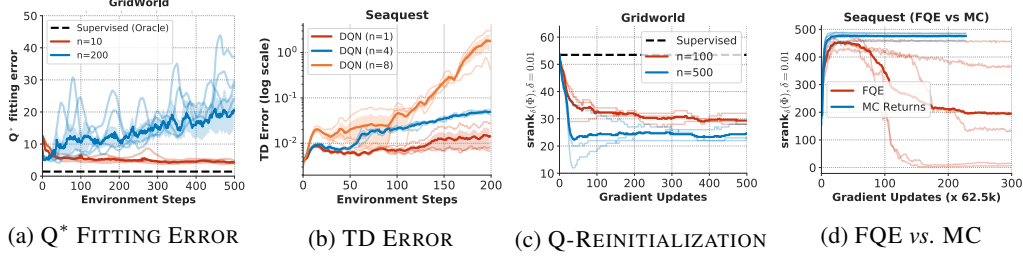


Figure 4: **(a)** Fitting error for Q^* prediction for $n = 10$ vs $n = 200$ steps in Figure 3 (left). Observe that rank collapse inhibits fitting Q^* as the fitting error rises over training while rank collapses. **(b)** TD error for varying values of n for SEAQUEST in Figure 3 (middle). TD error increases with rank degradation. **(c)** Q-network re-initialization in each fitting iteration on gridworld. **(d)** Trend of $\text{srnk}_\delta(\Phi)$ for policy evaluation based on bootstrapped updates (FQE) vs Monte-Carlo returns (no bootstrapping). Note that rank-collapse still persists with reinitialization and FQE, but goes away in the absence of bootstrapping.

the problem is not isolated to the control setting, we show evidence of rank collapse in the policy evaluation setting as well. We trained a value network using fitted Q-evaluation for a fixed policy π (i.e., using the Bellman operator \mathcal{T}^π instead of \mathcal{T}), and found that rank drop still occurs (FQE in Figure 4d). Finally, we show that by removing bootstrapped updates and instead regressing directly to Monte-Carlo (MC) estimates of the value, the effective rank *does not* collapse (MC Returns in Figure 4d). These results, along with similar findings on other Atari environments (Appendix A.3), our analysis indicates that bootstrapping is at the core of implicit under-parameterization.

4 THEORETICAL ANALYSIS OF IMPLICIT UNDER-PARAMETERIZATION

In this section, we formally analyze implicit under-parameterization and prove that training neural networks with bootstrapping reduces the effective rank of the Q-network, corroborating the empirical observations in the previous section. We focus on policy evaluation (Figure 4d and Figure A.9), where we aim to learn a Q-function that satisfies $\mathbf{Q} = \mathbf{R} + \gamma P^\pi \mathbf{Q}$ for a fixed π , for ease of analysis. We also presume a fixed dataset of transitions, \mathcal{D} , to learn the Q-function.

4.1 ANALYSIS VIA KERNEL REGRESSION

We first study bootstrapping with neural networks through a mathematical abstraction that treats the Q-network as a kernel machine, following the neural tangent kernel (NTK) formalism (Jacot et al., 2018). Building on prior analysis of self-distillation (Mobahi et al., 2020), we assume that each iteration of bootstrapping, the Q-function optimizes the squared TD error to target labels y_k with a kernel regularizer. This regularizer captures the inductive bias from gradient-based optimization of TD error and resembles the regularization imposed by gradient descent under NTK (Mobahi et al., 2020). The error is computed on $(s_i, a_i) \in \mathcal{D}$ whereas the regularization imposed by a universal kernel u with a coefficient of $c \geq 0$ is applied to the Q-values at *all* state-action pairs as shown in Equation 1. We consider a setting $c > 0$ for all rounds of bootstrapping, which corresponds to the solution obtained by performing gradient descent on TD error for a small number of iterations with early stopping in each round (Suggala et al., 2018) and thus, resembles how the updates in Algorithm 1 are typically implemented in practice.

$$\mathbf{Q}_{k+1} \leftarrow \arg \min_{\mathbf{Q} \in \mathcal{Q}} \sum_{s_i, a_i \in \mathcal{D}} (Q(s_i, a_i) - y_k(s_i, a_i))^2 + c \sum_{(s, a)} \sum_{(s', a')} u((s, a), (s', a')) Q(s, a) Q(s', a'). \quad (1)$$

The solution to Equation 1 can be expressed as $Q_{k+1}(s, a) = \mathbf{g}_{(s, a)}^T (c\mathbf{I} + \mathbf{G})^{-1} \mathbf{y}_k$, where \mathbf{G} is the Gram matrix for a special positive-definite kernel (Duffy, 2015) and $\mathbf{g}_{(s, a)}$ denotes the row of \mathbf{G} corresponding to the input (s, a) (Mobahi et al., 2020, Proposition 1). A detailed proof is in Appendix C. When combined with the fitted Q-iteration recursion, setting labels $y_k = \mathbf{R} + \gamma P^\pi \mathbf{Q}_{k-1}$, we recover a recurrence that relates subsequent value function iterates

$$\mathbf{Q}_{k+1} = \mathbf{G}(c\mathbf{I} + \mathbf{G})^{-1} \mathbf{y}_k = \underbrace{\mathbf{G}(c\mathbf{I} + \mathbf{G})^{-1}}_{\mathbf{A}} [\mathbf{R} + \gamma P^\pi \mathbf{Q}_k] = \mathbf{A} \left(\sum_{i=1}^k \gamma^{k-i} (P^\pi \mathbf{A})^{k-i} \right) \mathbf{R} := \mathbf{A} \mathbf{M}_k \mathbf{R}. \quad (2)$$

Equation 2 follows from unrolling the recurrence and setting the algorithm-agnostic initial Q-value, \mathbf{Q}_0 , to be $\mathbf{0}$. We now show that the sparsity of singular values of the matrix \mathbf{M}_k generally increases over fitting iterations, implying that the effective rank of \mathbf{M}_k diminishes with more iterations. For this result, we assume that the matrix \mathbf{S} is normal, i.e., the norm of the (complex) eigenvalues of \mathbf{S} is equal to its singular values. We will discuss how this assumption can be relaxed in Appendix A.7.

Theorem 4.1. Let \mathbf{S} be a shorthand for $\mathbf{S} = \gamma P^\pi \mathbf{A}$ and assume \mathbf{S} is a normal matrix. Then there exists an infinite, strictly increasing sequence of fitting iterations, $(k_l)_{l=1}^\infty$ starting from $k_1 = 0$, such that, for any two singular-values $\sigma_i(\mathbf{S})$ and $\sigma_j(\mathbf{S})$ of \mathbf{S} with $\sigma_i(\mathbf{S}) < \sigma_j(\mathbf{S})$,

$$\forall l \in \mathbb{N} \text{ and } l' \geq l, \quad \frac{\sigma_i(\mathbf{M}_{k_{l'}})}{\sigma_j(\mathbf{M}_{k_{l'}})} < \frac{\sigma_i(\mathbf{M}_{k_l})}{\sigma_j(\mathbf{M}_{k_l})} \leq \frac{\sigma_i(\mathbf{S})}{\sigma_j(\mathbf{S})}. \quad (3)$$

Hence, $\text{srnk}_\delta(\mathbf{M}_{k_{l'}}) \leq \text{srnk}_\delta(\mathbf{M}_{k_l})$. Moreover, if \mathbf{S} is positive semi-definite, then $(k_l)_{l=1}^\infty = \mathbb{N}$, i.e., srnk continuously decreases in each fitting iteration.

We provide a proof of the theorem above as well as present a stronger variant that shows a gradual decrease in the effective rank for fitting iterations outside this infinite sequence in Appendix C. As k increases along the sequence of iterations given by $k = (k_l)_{l=1}^\infty$, the effective rank of the matrix \mathbf{M}_k drops, leading to low expressivity of this matrix. Since \mathbf{M}_k linearly maps rewards to the Q-function (Equation 2), drop in expressivity results of \mathbf{M}_k in the inability to model the actual \mathbf{Q}^π .

Summary of our analysis. Our analysis of bootstrapping and gradient descent from the view of regularized kernel regression suggests that rank drop happens with more training (i.e., with more rounds of bootstrapping). In contrast to self-distillation (Mobahi et al., 2020), rank drop may not happen in every iteration (and rank may increase between two consecutive iterations occasionally), but srnk_δ exhibits a generally decreasing trend.

4.2 ANALYSIS WITH DEEP LINEAR NETWORKS UNDER GRADIENT DESCENT

While Section 4.1 demonstrates rank collapse will occur in a kernel-regression model of Q-learning, it does not illustrate *when* the rank collapse occurs. To better specify a point in training when rank collapse emerges, we present a complementary derivation for the case when the Q-function is represented as a deep linear neural network (Arora et al., 2019), which is a widely-studied setting for analyzing implicit regularization of gradient descent in supervised learning (Gunasekar et al., 2017; 2018; Arora et al., 2018; 2019). Our analysis will show that rank collapse can emerge as the generated target values begin to approach the previous value estimate, in particular, when in the vicinity of the optimal Q-function.

Proof strategy. Our proof consists of two steps: (1) We show that the effective rank of the feature matrix decreases within one fitting iteration (for a given target value) due to the low-rank affinity, (2) We show that this effective rank drop is “compounded” as we train using a bootstrapped objective. Proposition 4.1 explains (1) and Proposition 4.2, Theorem 4.2 and Appendix D.2 discuss (2).

Additional notation and assumptions. We represent the Q-function as a deep linear network with at ≥ 3 layers, such that $Q(\mathbf{s}, \mathbf{a}) = \mathbf{W}_N \mathbf{W}_\phi[\mathbf{s}; \mathbf{a}]$, where $N \geq 3$, $\mathbf{W}_N \in \mathbb{R}^{1 \times d_{N-1}}$ and $\mathbf{W}_\phi = \mathbf{W}_{N-1} \mathbf{W}_{N-2} \cdots \mathbf{W}_1$ with $\mathbf{W}_i \in \mathbb{R}^{d_i \times d_{i-1}}$ for $i = 1, \dots, N-1$. \mathbf{W}_ϕ maps an input $[\mathbf{s}; \mathbf{a}]$ to corresponding penultimate layer’s features $\Phi(\mathbf{s}, \mathbf{a})$. Let $\mathbf{W}_j(k, t)$ denotes the weight matrix \mathbf{W}_j at the t -th step of gradient descent during the k -th fitting iteration (Algorithm 1). We define $\mathbf{W}_{k,t} = \mathbf{W}_N(k, t) \mathbf{W}_\phi(k, t)$ and $L_{N,k+1}(\mathbf{W}_{k,t})$ as the TD error objective in the k -th fitting iteration. We study $\text{srnk}_\delta(\mathbf{W}_\phi(k, t))$ since the rank of features $\Phi = \mathbf{W}_\phi(k, t)[\mathcal{S}, \mathcal{A}]$ is equal to rank of $\mathbf{W}_\phi(k, t)$ provided the state-action inputs have high rank.


We assume that the evolution of the weights is governed by a continuous-time differential equation (Arora et al., 2018) *within* each fitting iteration k . To simplify analysis, we also assume that all *except the last-layer* weights follow a “balancedness” property (Equation D.4), which suggests that the weight matrices in the consecutive layers in the deep linear network share the same singular values (but with different permutations). However, note that we do not assume balancedness for the last layer which trivially leads to rank-1 features, making our analysis strictly more general than conventionally studied deep linear networks. In this model, we can characterize the evolution of the singular values of the feature matrix $\mathbf{W}_\phi(k, t)$, using techniques analogous to Arora et al. (2019):

Proposition 4.1. The singular values of the feature matrix $\mathbf{W}_\phi(k, t)$ evolve according to:

$$\dot{\sigma}_r(k, t) = -N \cdot (\sigma_r^2(k, t))^{1 - \frac{1}{N-1}} \cdot \left\langle \mathbf{W}_N(k, t)^T \frac{dL_{N,k+1}(\mathbf{W}_{k,t})}{d\mathbf{W}}, \mathbf{u}_r(k, t) \mathbf{v}_r(k, t)^T \right\rangle, \quad (4)$$

for $r = 1, \dots, \min_{i=1}^{N-1} d_i$, where $\mathbf{u}_r(k, t)$ and $\mathbf{v}_r(k, t)$ denote the left and right singular vectors of the feature matrix, $\mathbf{W}_\phi(k, t)$, respectively.

Solving the differential equation (4) indicates that larger singular values will evolve at an exponentially faster rate than smaller singular values (as we also formally show in Appendix D.1) and the difference in their magnitudes disproportionately increase with increasing t . This behavior also occurs empirically, illustrated in the figure on the right (also see Figure D.1), where larger singular values are orders of magnitude larger than smaller singular values. Hence, the effective rank, $\text{srank}_\delta(\mathbf{W}_\phi(k, t))$, will decrease with more gradient steps *within* a fitting iteration k .



Evolution of singular values of \mathbf{W}_ϕ on SEQUEST

Abstract optimization problem for the low-rank solution. Building on Proposition 4.1, we note that the final solution obtained in a bootstrapping round (*i.e.*, fitting iteration) can be equivalently expressed as the solution that minimizes a weighted sum of the TD error and a data-dependent implicit regularizer $h_{\mathcal{D}}(\mathbf{W}_\phi, \mathbf{W}_N)$ that encourages disproportionate singular values of \mathbf{W}_ϕ , and hence, a low effective rank of \mathbf{W}_ϕ . While the actual form for h is unknown, to facilitate our analysis of bootstrapping, we make a simplification and express this solution as the minimum of Equation 5.

$$\min_{\mathbf{W}_\phi, \mathbf{W}_N \in \mathcal{M}} \|\mathbf{W}_N \mathbf{W}_\phi [\mathbf{s}; \mathbf{a}] - \mathbf{y}_k(\mathbf{s}, \mathbf{a})\|^2 + \lambda_k \text{srank}_\delta(\mathbf{W}_\phi). \quad (5)$$

Note that the entire optimization path may not correspond to the objective in Equation 5, but the Equation 5 represents the final solution of a given fitting iteration. \mathcal{M} denotes the set of constraints that \mathbf{W}_N obtained via gradient optimization of TD error must satisfy, however we do not need to explicitly quantify \mathcal{M} in our analysis. λ_k is a constant that denotes the strength of rank regularization. Since srank_δ is always regularized, our analysis assumes that $\lambda_k > 0$ (see Appendix D.1).

Rank drop within a fitting iteration “compounds” due to bootstrapping. In the RL setting, the target values are given by $y_k(\mathbf{s}, \mathbf{a}) = r(\mathbf{s}, \mathbf{a}) + \gamma P^\pi Q_{k-1}(\mathbf{s}, \mathbf{a})$. First note that when $r(\mathbf{s}, \mathbf{a}) = 0$ and $P^\pi = \mathbf{I}$, *i.e.*, when the bootstrapping update resembles self-regression, we first note that just “copying over weights” from iteration $k-1$ to iteration k is a feasible point for solving Equation 5, which attains *zero* TD error with no increase in srank_δ . A better solution to Equation 5 can thus be obtained by incurring non-zero TD error at the benefit of a decreased srank , indicating that in this setting, $\text{srank}_\delta(\mathbf{W}_\phi)$ drops in each fitting iteration, leading to a compounding rank drop effect.

We next extend this analysis to the full bootstrapping setting. Unlike the self-training setting, $y_k(\mathbf{s}, \mathbf{a})$ is not directly expressible as a function of the previous $\mathbf{W}_\phi(k, T)$ due to additional reward and dynamics transformations. Assuming closure of the function class (Assumption D.1) under the Bellman update (Munos & Szepesvári, 2008; Chen & Jiang, 2019), we reason about the compounding effect of rank drop across iterations in Proposition 4.2 (proof in Appendix D.2). Specifically, srank_δ can increase in each fitting iteration due to \mathbf{R} and P^π transformations, but will decrease due to low rank preference of gradient descent. This change in rank then compounds as shown below.

Proposition 4.2. Assume that the Q -function is initialized to $\mathbf{W}_\phi(0)$ and $\mathbf{W}_N(0)$. Let the Q -function class be closed under the backup, *i.e.*, $\exists \mathbf{W}_N^P, \mathbf{W}_\phi^P$, s.t. $(\mathbf{R} + \gamma P^\pi \mathbf{Q}_{k-1})^T = \mathbf{W}_N^P(k) \mathbf{W}_\phi^P(k) [\mathcal{S}; \mathcal{A}]^T$, and assume that the change in srank due to dynamics and reward transformations is bounded: $\text{srank}_\delta(\mathbf{W}_\phi^P(k)) \leq \text{srank}_\delta(\mathbf{W}_\phi(k-1)) + c_k$. Then,

$$\text{srank}_\delta(\mathbf{W}_\phi(k)) \leq \text{srank}_\delta(\mathbf{W}_\phi(0)) + \sum_{j=1}^k c_j - \sum_{j=1}^k \frac{\|\mathbf{Q}_j - \mathbf{y}_j\|}{\lambda_j}.$$

Proposition 4.2 provides a bound on the value of srank after k rounds of bootstrapping. srank decreases in each iteration due to non-zero TD errors, but potentially increases due to reward and bootstrapping transformations. To instantiate a concrete case where rank clearly collapses, we investigate c_k as the value function gets closer to the Bellman fixed point, which is a favourable initialization for the Q -function in Theorem 4.2. In this case, the learning dynamics begins to resemble the self-training regime, as the target values approach the previous value iterate $\mathbf{y}_k \approx \mathbf{Q}_{k-1}$, and thus, as we show next, the potential increase in srank (c_k in Proposition 4.2) converges to 0.

Theorem 4.2. Suppose target values $\mathbf{y}_k = \mathbf{R} + \gamma P^\pi \mathbf{Q}_{k-1}$ are close to the previous value estimate \mathbf{Q}_{k-1} , *i.e.* $\forall \mathbf{s}, \mathbf{a}, y_k(\mathbf{s}, \mathbf{a}) = Q_{k-1}(\mathbf{s}, \mathbf{a}) + \varepsilon(\mathbf{s}, \mathbf{a})$, with $|\varepsilon(\mathbf{s}, \mathbf{a})| \ll |Q_{k-1}(\mathbf{s}, \mathbf{a})|$. Then, there is a constant ϵ_0 depending upon \mathbf{W}_N and \mathbf{W}_ϕ , such that for all $\|\varepsilon\| < \epsilon_0$, $c_k = 0$. Thus, srank decreases in iteration k : $\text{srank}_\delta(\mathbf{W}_\phi(k)) \leq \text{srank}_\delta(\mathbf{W}_\phi(k-1)) - \|\mathbf{Q}_k - \mathbf{y}_k\|/\lambda_k$.

We provide a complete form, including the expression for ϵ_0 and a proof in Appendix D.3. To empirically show the **consequence of Theorem 4.2** that a decrease in $\text{srank}_\delta(\mathbf{W}_\phi)$ values can lead to an increase in the distance to the fixed point in a neighborhood around the fixed point, we performed a controlled experiment on a deep linear net shown in Figure 5 that measures the relationship between of $\text{srank}_\delta(\Phi)$ and the error to the projected TD fixed point $|\mathbf{Q} - \mathbf{Q}^*|$. Note that a drop in $\text{srank}_\delta(\Phi)$ corresponds to a increased value of $|\mathbf{Q} - \mathbf{Q}^*|$ indicating that rank drop when \mathbf{Q} get close to a fixed point can affect convergence to it.




Figure 5: Trend of $\text{srank}_\delta(\Phi)$ v.s. error on log scale to the projected TD fixed point. A drop in $\text{srank}_\delta(\Phi)$ (shown as blue and yellow circles) corresponds to a corresponding increase in distance to the fixed point.

5 MITIGATING UNDER-PARAMETRIZATION IMPROVES DEEP Q-LEARNING

We now show that mitigating implicit under-parameterization by preventing rank collapse can improve performance. We place special emphasis on the offline RL setting in this section, since it is particularly vulnerable to the adverse effects of rank collapse.

We devise a penalty (or a regularizer) $\mathcal{L}_p(\Phi)$ that encourages higher effective rank of the learned features, $\text{srank}_\delta(\Phi)$, to prevent rank collapse. The effective rank function $\text{srank}_\delta(\Phi)$ is non-differentiable, so we choose a simple surrogate that can be optimized over deep networks. Since effective rank is maximized when the magnitude of the singular values is roughly balanced, one way to increase effective rank is to minimize the largest singular value of Φ , $\sigma_{\max}(\Phi)$, while simultaneously maximizing the smallest singular value, $\sigma_{\min}(\Phi)$. We construct a simple penalty $\mathcal{L}_p(\Phi)$ derived from this intuition, given by:

$$\mathcal{L}_p(\Phi) = \sigma_{\max}^2(\Phi) - \sigma_{\min}^2(\Phi). \quad (6)$$

$\mathcal{L}_p(\Phi)$ can be computed by invoking the singular value decomposition subroutines in standard automatic differentiation frameworks (Abadi et al., 2016; Paszke et al., 2019).

We estimate the singular values over the feature matrix computed over a minibatch, and add the resulting value of \mathcal{L}_p as a penalty to the TD error objective, with a tradeoff factor $\alpha = 0.001$.

Does $\mathcal{L}_p(\Phi)$ address rank collapse? We first verify whether controlling the minimum and maximum singular values using $\mathcal{L}_p(\Phi)$ actually prevents rank collapse. When using this penalty on the gridworld problem (Figure 6a), the effective rank does not collapse, instead gradually decreasing at the onset and then plateauing, akin to the evolution of effective rank in supervised learning. In Figure 6b, we plot the evolution of effective rank on two Atari games in the offline setting (all games in Appendix A.5), and observe that using \mathcal{L}_p also generally leads to increasing rank values.

Does mitigating rank collapse improve performance? We now evaluate the performance of the penalty using DQN (Mnih et al., 2015) and CQL (Kumar et al., 2020b) on Atari dataset from Agarwal et al. (2020) (5% replay data), used in Section 3. Figure 7 summarizes the relative improvement from using the penalty for 16 Atari games. Adding the penalty to DQN improves performance on all 16/16 games with a median improvement of **74.5%**; adding it to CQL, a state-of-the-art offline algorithm, improves performance on 11/16 games with median improvement of **14.1%**. Prior work has discussed that standard Q-learning methods designed for the online setting, such as DQN, are generally ineffective with small offline datasets (Kumar et al., 2020b; Agarwal et al., 2020). Our results show that mitigating rank collapse makes even such simple methods substantially more effective in this setting, suggesting that rank collapse and the resulting implicit under-parameterization may be an crucial piece of the puzzle in explaining the challenges of offline RL.

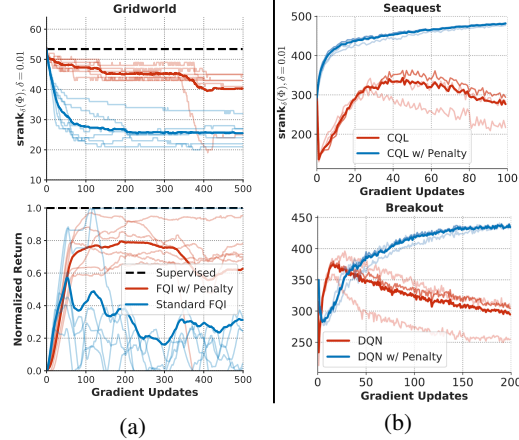


Figure 6: **(a)**: $\text{srank}_\delta(\Phi)$ (top) and performance (bottom) of FQI on gridworld in the offline setting with 200 gradient updates per fitting iteration. Note reduced rank collapse and higher performance with the regularizer $\mathcal{L}_p(\Phi)$. **(b)**: $\mathcal{L}_p(\Phi)$ mitigates the rank collapse in DQN and CQL in the offline RL setting on Atari.

We also evaluated the regularizer $\mathcal{L}_p(\Phi)$ in the data-efficient online RL setting, with results in Appendix A.6. This variant achieved median improvement of 20.6% performance with Rainbow (Hessel et al., 2018), however performed poorly with DQN, where it reduced median performance by 11.5%. Thus, while our proposed penalty is effective in many cases in offline and online settings, it does not solve the problem fully, i.e., it does not address the root cause of implicit under-parameterization and only addresses a symptom, and a more sophisticated solution may better prevent the issues with implicit under-parameterization. Nevertheless, our results suggest that mitigation of implicit under-parameterization can improve performance of data-efficient RL.




Figure 7: DQN and CQL with $\mathcal{L}_p(\Phi)$ penalty vs. their standard counterparts in the 5% offline setting on Atari from Section 3. \mathcal{L}_p improves DQN on 16/16 and CQL on 11/16 games.

6 RELATED WORK

Prior work has extensively studied the learning dynamics of Q-learning with tabular and linear function approximation, to study error propagation (Munos, 2003; Farahmand et al., 2010) and to prevent divergence (De Farias, 2002; Maei et al., 2009; Sutton et al., 2009; Dai et al., 2018), as opposed to deep Q-learning analyzed in this work. Q-learning has been shown to have favorable optimization properties with certain classes of features (Ghosh & Bellemare, 2020), but our work shows that the features learned by a neural net when minimizing TD error do not enjoy such guarantees, and instead suffer from rank collapse. Recent theoretical analyses of deep Q-learning have shown convergence under restrictive assumptions (Yang et al., 2020; Cai et al., 2019; Zhang et al., 2020; Xu & Gu, 2019), but Theorem 4.2 shows that implicit under-parameterization appears when the estimates of the value function approach the optimum, potentially preventing convergence. Xu et al. (2005; 2007) present variants of LSTD (Boyan, 1999), which model the Q-function as a kernel-machine but do not take into account the regularization from gradient descent, as done in Equation 1, which is essential for implicit under-parameterization. Igl et al. (2020); Fedus et al. (2020a) argue that non-stationarity arising from distribution shift hinders generalization and recommend periodic network re-initialization. Under-parameterization is not caused by this distribution shift, and we find that network re-initialization does little to prevent rank collapse (Figure 4c). Luo et al. (2020) proposes a regularization similar to ours, but in a different setting, finding that more expressive features increases performance of on-policy RL methods. Finally, Yang et al. (2019) study the effective rank of the Q^* -values when expressed as a $|\mathcal{S}| \times |\mathcal{A}|$ matrix in online RL and find that low ranks for this Q^* -matrix are preferable. We analyze a fundamentally different object: the learned features (and illustrate that a rank-collapse of features can hurt), not the Q^* -matrix, whose rank is upper-bounded by the number of actions (e.g., 24 for Atari).

7 DISCUSSION

We identified an implicit under-parameterization phenomenon in deep RL algorithms that use bootstrapping, where gradient-based optimization of a bootstrapped objective can lead to a reduction in the expressive power of the value network. This effect manifests as a collapse of the rank of the features learned by the value network, causing aliasing across states and often leading to poor performance. Our analysis reveals that this phenomenon is caused by the implicit regularization due to gradient descent on bootstrapped objectives. We observed that mitigating this problem by means of a simple regularization scheme improves performance of deep Q-learning methods.

While our proposed regularization provides some improvement, devising better mitigation strategies for implicit under-parameterization remains an exciting direction for future work. Our method explicitly attempts to prevent rank collapse, but relies on the emergence of useful features solely through the bootstrapped signal. An alternative path may be to develop new auxiliary losses (e.g., Jaderberg et al., 2016) that learn useful features while passively preventing underparameterization. More broadly, understanding the effects of neural nets and associated factors such as initialization, choice of optimizer, etc. on the learning dynamics of deep RL algorithms, using tools from deep learning theory, is likely to be key towards developing robust and data-efficient deep RL algorithms.

ACKNOWLEDGEMENTS

We thank Lihong Li, Aaron Courville, Aurick Zhou, Abhishek Gupta, George Tucker, Ofir Nachum, Wesley Chung, Emmanuel Bengio, Zafarali Ahmed, and Jacob Buckman for feedback on an earlier version of this paper. We thank Hossein Mobahi for insightful discussions about self-distillation and Hanie Sedghi for insightful discussions about implicit regularization and generalization in deep networks. We additionally thank Michael Janner, Aaron Courville, Dale Schuurmans and Marc Bellemare for helpful discussions. AK was partly funded by the DARPA Assured Autonomy program, and DG was supported by a NSF graduate fellowship and compute support from Amazon.

REFERENCES

- Martin Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, Manjunath Kudlur, Josh Levenberg, Rajat Monga, Sherry Moore, Derek G. Murray, Benoit Steiner, Paul Tucker, Vijay Vasudevan, Pete Warden, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. Tensorflow: A system for large-scale machine learning. In *12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16)*, pp. 265–283, 2016. URL <https://www.usenix.org/system/files/conference/osdi16/osdi16-abadi.pdf>.
- Joshua Achiam, Ethan Knight, and Pieter Abbeel. Towards characterizing divergence in deep q-learning. *ArXiv*, abs/1903.08894, 2019.
- Rishabh Agarwal, Dale Schuurmans, and Mohammad Norouzi. An optimistic perspective on offline reinforcement learning. In *International Conference on Machine Learning (ICML)*, 2020.
- Sanjeev Arora, Nadav Cohen, and Elad Hazan. On the optimization of deep networks: Implicit acceleration by overparameterization. *arXiv preprint arXiv:1802.06509*, 2018.
- Sanjeev Arora, Nadav Cohen, Wei Hu, and Yuping Luo. Implicit regularization in deep matrix factorization. In *Advances in Neural Information Processing Systems*, pp. 7413–7424, 2019.
- Marc G. Bellemare, Yavar Naddaf, Joel Veness, and Michael Bowling. The arcade learning environment: An evaluation platform for general agents. *J. Artif. Int. Res.*, 47(1):253–279, May 2013. ISSN 1076-9757.
- Marc G Bellemare, Will Dabney, and Rémi Munos. A distributional perspective on reinforcement learning. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pp. 449–458. JMLR. org, 2017.
- Emmanuel Bengio, Joelle Pineau, and Doina Precup. Interference and generalization in temporal difference learning. *arXiv preprint arXiv:2003.06350*, 2020.
- Justin A Boyan. Least-squares temporal difference learning. In *ICML*, pp. 49–56. Citeseer, 1999.
- Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. Openai gym, 2016.
- Qi Cai, Zhuoran Yang, Jason D Lee, and Zhaoran Wang. Neural temporal-difference and q-learning provably converge to global optima. *arXiv preprint arXiv:1905.10027*, 2019.
- Jinglin Chen and Nan Jiang. Information-theoretic considerations in batch reinforcement learning. *ICML*, 2019.
- Bo Dai, Albert Shaw, Lihong Li, Lin Xiao, Niao He, Zhen Liu, Jianshu Chen, and Le Song. Sbeed: Convergent reinforcement learning with nonlinear function approximation. In *International Conference on Machine Learning*, pp. 1133–1142, 2018.
- Daniela Pucci De Fariás. *The linear programming approach to approximate dynamic programming: Theory and application*. PhD thesis, 2002.
- Dean G Duffy. *Green’s functions with applications*. CRC Press, 2015.

- Damien Ernst, Pierre Geurts, and Louis Wehenkel. Tree-based batch mode reinforcement learning. *Journal of Machine Learning Research*, 6(Apr):503–556, 2005.
- Amir-massoud Farahmand, Csaba Szepesvári, and Rémi Munos. Error propagation for approximate policy and value iteration. In *Advances in Neural Information Processing Systems (NIPS)*, 2010.
- William Fedus, Dibya Ghosh, John D Martin, Marc G Bellemare, Yoshua Bengio, and Hugo Larochelle. On catastrophic interference in atari 2600 games. *arXiv preprint arXiv:2002.12499*, 2020a.
- William Fedus, Prajit Ramachandran, Rishabh Agarwal, Yoshua Bengio, Hugo Larochelle, Mark Rowland, and Will Dabney. Revisiting fundamentals of experience replay. *arXiv preprint arXiv:2007.06700*, 2020b.
- Justin Fu, Aviral Kumar, Matthew Soh, and Sergey Levine. Diagnosing bottlenecks in deep q-learning algorithms. In *Proceedings of the 36th International Conference on Machine Learning*. PMLR, 2019.
- Dibya Ghosh and Marc G Bellemare. Representations for stable off-policy reinforcement learning. *arXiv preprint arXiv:2007.05520*, 2020.
- Caglar Gulcehre, Ziyu Wang, Alexander Novikov, Tom Le Paine, Sergio Gómez Colmenarejo, Konrad Zolna, Rishabh Agarwal, Josh Merel, Daniel Mankowitz, Cosmin Paduraru, et al. R1 unplugged: Benchmarks for offline reinforcement learning. 2020.
- Suriya Gunasekar, Blake E Woodworth, Srinadh Bhojanapalli, Behnam Neyshabur, and Nati Srebro. Implicit regularization in matrix factorization. In *Advances in Neural Information Processing Systems*, pp. 6151–6159, 2017.
- Suriya Gunasekar, Jason D Lee, Daniel Soudry, and Nati Srebro. Implicit bias of gradient descent on linear convolutional networks. In *Advances in Neural Information Processing Systems*, pp. 9461–9471, 2018.
- Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. *CoRR*, abs/1801.01290, 2018. URL <http://arxiv.org/abs/1801.01290>.
- Matteo Hessel, Joseph Modayil, Hado Van Hasselt, Tom Schaul, Georg Ostrovski, Will Dabney, Dan Horgan, Bilal Piot, Mohammad Azar, and David Silver. Rainbow: Combining improvements in deep reinforcement learning. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- Edmund Hlawka, Rudolf Taschner, and Johannes Schoißengeier. *The Dirichlet Approximation Theorem*, pp. 1–18. Springer Berlin Heidelberg, Berlin, Heidelberg, 1991. ISBN 978-3-642-75306-0. doi: 10.1007/978-3-642-75306-0_1. URL https://doi.org/10.1007/978-3-642-75306-0_1.
- Maximilian Igl, Gregory Farquhar, Jelena Luketina, Wendelin Boehmer, and Shimon Whiteson. The impact of non-stationarity on generalisation in deep reinforcement learning. *arXiv preprint arXiv:2006.05826*, 2020.
- Arthur Jacot, Franck Gabriel, and Clement Hongler. Neural tangent kernel: Convergence and generalization in neural networks. In *Advances in Neural Information Processing Systems 31*. 2018.
- Max Jaderberg, Volodymyr Mnih, Wojciech Marian Czarnecki, Tom Schaul, Joel Z Leibo, David Silver, and Koray Kavukcuoglu. Reinforcement learning with unsupervised auxiliary tasks. *arXiv preprint arXiv:1611.05397*, 2016.
- Robert Johnson. Approximate irrational numbers by rational numbers, 2016. URL <https://math.stackexchange.com/questions/1829743/>.
- Aviral Kumar, Abhishek Gupta, and Sergey Levine. Discor: Corrective feedback in reinforcement learning via distribution correction. *arXiv preprint arXiv:2003.07305*, 2020a.

- Aviral Kumar, Aurick Zhou, George Tucker, and Sergey Levine. Conservative q-learning for offline reinforcement learning. *arXiv preprint arXiv:2006.04779*, 2020b.
- Sascha Lange, Thomas Gabel, and Martin Riedmiller. Batch reinforcement learning. In *Reinforcement learning*, pp. 45–73. Springer, 2012.
- Sergey Levine, Aviral Kumar, George Tucker, and Justin Fu. Offline reinforcement learning: Tutorial, review, and perspectives on open problems. *arXiv preprint arXiv:2005.01643*, 2020.
- Vincent Liu, Raksha Kumaraswamy, Lei Le, and Martha White. The utility of sparse representations for control in reinforcement learning. *CoRR*, abs/1811.06626, 2018. URL <http://arxiv.org/abs/1811.06626>.
- Xufang Luo, Qi Meng, Di He, Wei Chen, and Yunhong Wang. I4r: Promoting deep reinforcement learning by the indicator for expressive representations. In Christian Bessiere (ed.), *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI-20*, pp. 2669–2675. International Joint Conferences on Artificial Intelligence Organization, 7 2020. doi: 10.24963/ijcai.2020/370. URL <https://doi.org/10.24963/ijcai.2020/370>. Main track.
- Hamid R. Maei, Csaba Szepesvári, Shalabh Bhatnagar, Doina Precup, David Silver, and Richard S. Sutton. Convergent temporal-difference learning with arbitrary smooth function approximation. In *Proceedings of the 22nd International Conference on Neural Information Processing Systems*, 2009.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, Stig Petersen, Charles Beattie, Amir Sadik, Ioannis Antonoglou, Helen King, Dharmashan Kumaran, Daan Wierstra, Shane Legg, and Demis Hassabis. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, feb 2015. ISSN 0028-0836.
- Hossein Mobahi, Mehrdad Farajtabar, and Peter L Bartlett. Self-distillation amplifies regularization in hilbert space. *arXiv preprint arXiv:2002.05715*, 2020.
- Rémi Munos. Error bounds for approximate policy iteration. In *Proceedings of the Twentieth International Conference on International Conference on Machine Learning, ICML’03*, pp. 560–567. AAAI Press, 2003. ISBN 1577351894.
- Rémi Munos and Csaba Szepesvári. Finite-time bounds for fitted value iteration. *Journal of Machine Learning Research*, 9(May):815–857, 2008.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. dAlché-Buc, E. Fox, and R. Garnett (eds.), *Advances in Neural Information Processing Systems 32*, pp. 8026–8037. Curran Associates, Inc., 2019.
- Martin L Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, Inc., 1994.
- Martin Riedmiller. Neural fitted q iteration—first experiences with a data efficient neural reinforcement learning method. In *European Conference on Machine Learning*, pp. 317–328. Springer, 2005.
- Axel Ruhe. On the closeness of eigenvalues and singular values for almost normal matrices. *Linear Algebra and its Applications*, 11(1):87–93, 1975.
- David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, et al. Mastering the game of go without human knowledge. *nature*, 550(7676):354–359, 2017.

- Arun Suggala, Adarsh Prasad, and Pradeep K Ravikumar. Connecting optimization and regularization paths. In *Advances in Neural Information Processing Systems*, pp. 10608–10619, 2018.
- Richard S. Sutton, Hamid Reza Maei, Doina Precup, Shalabh Bhatnagar, David Silver, Csaba Szepesvári, and Eric Wiewiora. Fast gradient-descent methods for temporal-difference learning with linear function approximation. In *International Conference on Machine Learning (ICML)*, 2009.
- James Townsend. Differentiating the singular value decomposition. Technical report, Technical Report 2016, <https://j-towns.github.io/papers/svd-derivative...>, 2016.
- Pan Xu and Quanquan Gu. A finite-time analysis of q-learning with neural network function approximation. *arXiv preprint arXiv:1912.04511*, 2019.
- Xin Xu, Tao Xie, Dewen Hu, and Xicheng Lu. Kernel least-squares temporal difference learning. *International Journal of Information Technology*, 11(9):54–63, 2005.
- Xin Xu, Dewen Hu, and Xicheng Lu. Kernel-based least squares policy iteration for reinforcement learning. *IEEE Transactions on Neural Networks*, 18(4):973–992, 2007.
- Yuzhe Yang, Guo Zhang, Zhi Xu, and Dina Katabi. Harnessing structures for value-based planning and reinforcement learning. *arXiv preprint arXiv:1909.12255*, 2019.
- Zhuoran Yang, Yuchen Xie, and Zhaoran Wang. A theoretical analysis of deep q-learning. In *Learning for Dynamics and Control*, pp. 486–489. PMLR, 2020.
- Yufeng Zhang, Qi Cai, Zhuoran Yang, Yongxin Chen, and Zhaoran Wang. Can temporal-difference and q-learning learn representation? a mean-field theory. *arXiv preprint arXiv:2006.04761*, 2020.

Appendices

A ADDITIONAL EVIDENCE FOR IMPLICIT UNDER-PARAMETERIZATION

In this section, we present additional evidence that demonstrates the existence of the implicit under-parameterization phenomenon from Section 3. In all cases, we plot the values of $\text{srnk}_\delta(\Phi)$ computed on a batch size of 2048 i.i.d. sampled transitions from the dataset.

A.1 OFFLINE RL

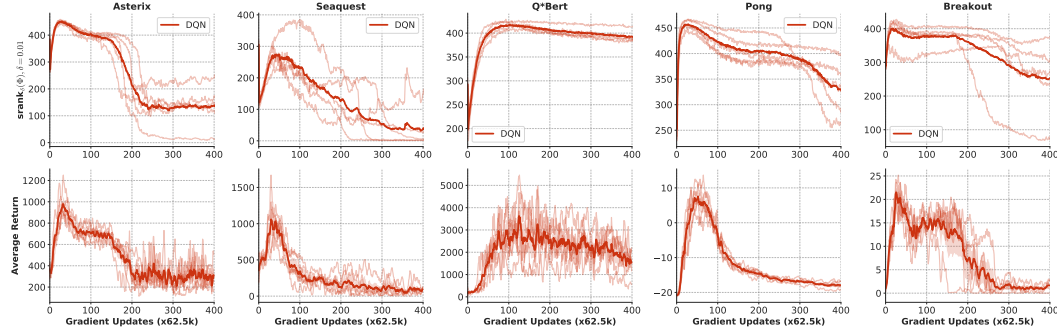


Figure A.1: **Offline DQN on Atari.** $\text{srnk}_\delta(\Phi)$ and performance of DQN on five Atari games in the offline RL setting using the 5% DQN Replay dataset (Agarwal et al., 2020) (marked as **DQN**). Note that low srnk_δ (top row) generally corresponds to worse policy performance (bottom row). Also note that rank collapse begins to generally occur close to the position of peak return. Average across 5 runs is showed for each game with individual runs.

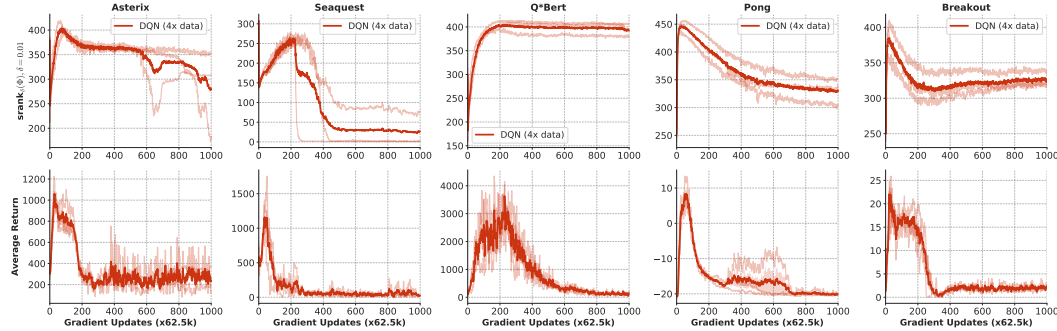


Figure A.2: **Offline DQN on Atari.** $\text{srnk}_\delta(\Phi)$ and performance of DQN on five Atari games in the offline RL setting using the 20% DQN Replay dataset (Agarwal et al., 2020) (marked as **DQN**) trained for 1000 iterations. Note that low srnk_δ (top row) generally corresponds to worse policy performance (bottom row). Average across 5 runs is showed for each game with individual runs.




Figure A.3: **Offline CQL on Atari.** $\text{srnk}_\delta(\Phi)$ and performance of CQL on five Atari games in the offline RL setting using the 5% DQN Replay. Note that low srnk_δ (top row) generally corresponds to worse policy performance (bottom row). Average across 5 runs for each game trained for 1000 iterations.

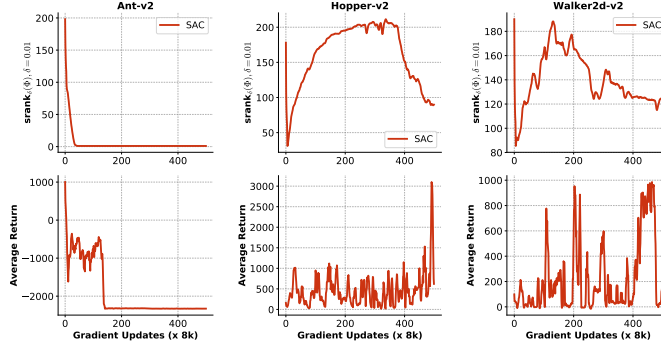


Figure A.4: **Offline Control on MuJoCo.** $\text{rank}_\delta(\Phi)$ and performance of SAC on three Gym environments the offline RL setting. Implicit Under-parameterization is conspicuous from the rank reduction, which highly correlates with performance degradation. We use 20% uniformly sampled data from the entire replay experience of an online SAC agent, similar to the 20% setting from Agarwal et al. (2020).

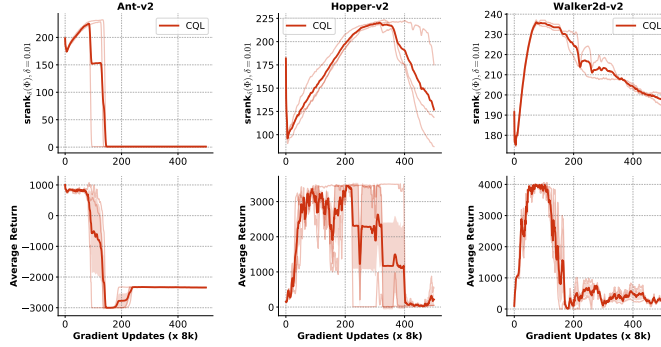


Figure A.5: **Offline Control on MuJoCo.** $\text{rank}_\delta(\Phi)$ and performance of CQL on three Gym environments the offline RL setting. Implicit Under-parameterization is conspicuous from the rank reduction, which highly correlates with performance degradation. We use 20% uniformly sampled data from the entire replay experience of an online SAC agent, similar to the 20% setting from Agarwal et al. (2020).

A.2 DATA EFFICIENT ONLINE RL

In the data-efficient online RL setting, we verify the presence of implicit under-parameterization on both DQN and Rainbow (Hessel et al., 2018) algorithms when larger number of gradient updates are made per environment step. In these settings we find that more gradient updates per environment step lead to a larger decrease in effective rank, whereas effective rank can increase when the amount of data re-use is reduced by taking fewer gradient steps.

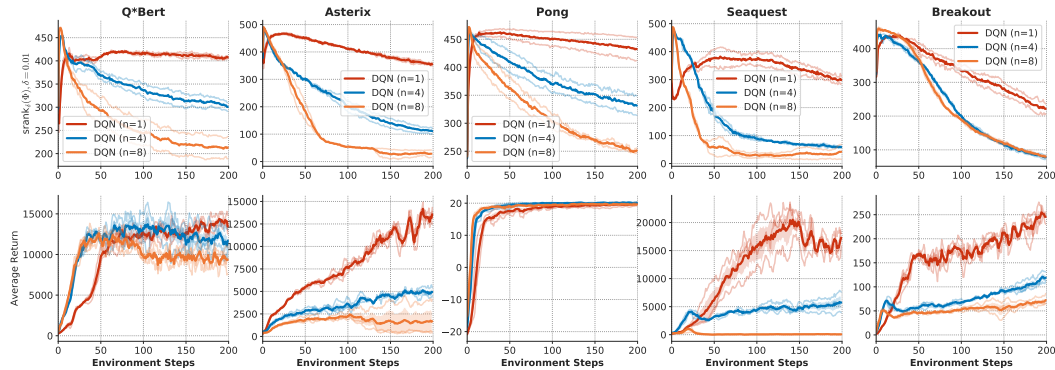


Figure A.6: **Online DQN on Atari.** $\text{rank}_\delta(\Phi)$ and performance of DQN on 5 Atari games in the online RL setting, with varying numbers of gradient steps per environment step (n). Rank collapse happens earlier with more gradient steps, and the corresponding performance is poor. This indicates that implicit under-parameterization aggravates as the rate of data re-use is increased.

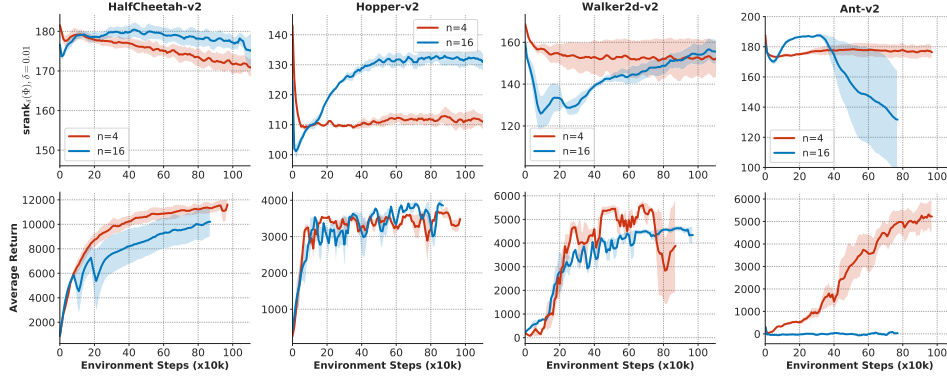


Figure A.7: **Online SAC on MuJoCo.** $\text{srnk}_\delta(\Phi)$ and performance of SAC on three Gym environments the online RL setting, with varying numbers of gradient steps per environment step (n). While in the simpler environments, HalfCheetah-v2, Hopper-v2 and Walker2d-v2 we actually observe an increase in the values of effective rank, which also corresponds to good performance with large n values in these cases, on the more complex Ant-v2 environment rank decreases with larger n , and the corresponding performance is worse with more gradient updates.

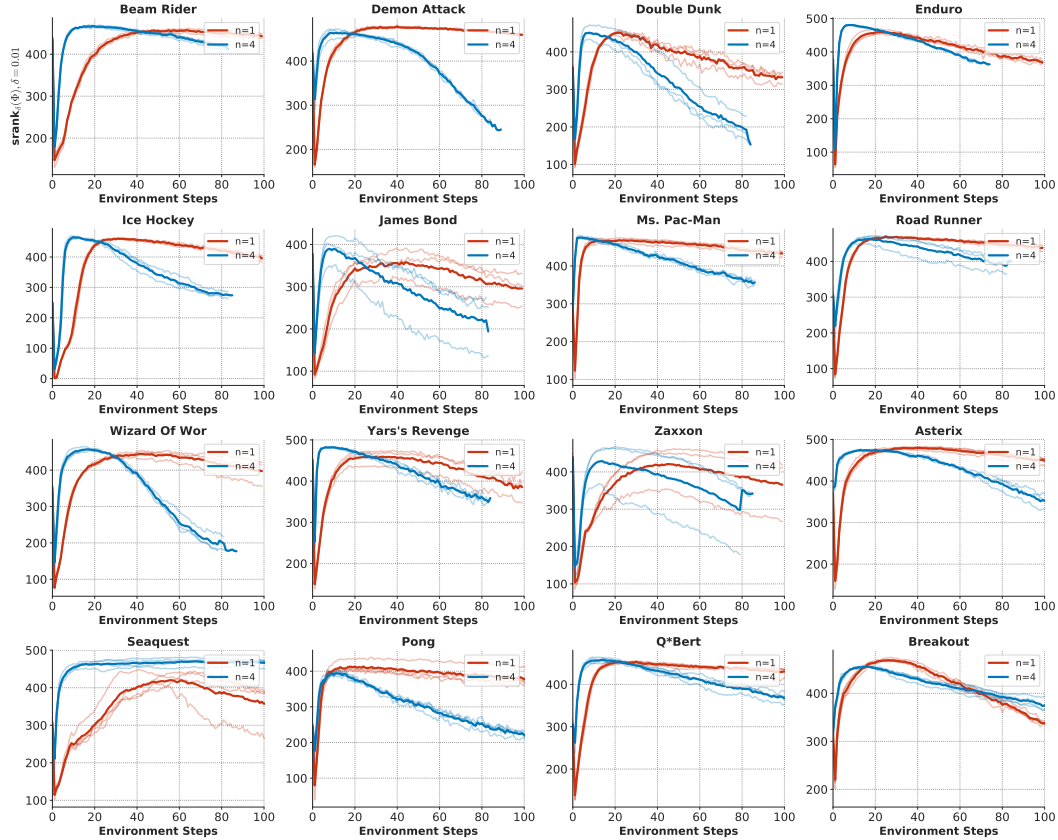


Figure A.8: **Online Rainbow on Atari.** $\text{srnk}_\delta(\Phi)$ Rainbow on 16 Atari games in the data-efficient online setting, with varying numbers of gradient steps per environment step (n). Rank collapse happens earlier with more gradient steps, and the corresponding performance is poor. This plot indicates the multi-step returns, prioritized replay or distributional C51 does not address the implicit under-parameterization issue.

A.3 DOES BOOTSTRAPPING CAUSE IMPLICIT UNDER-PARAMETERIZATION?

In this section, we provide additional evidence to support our claim from Section 3 that suggests that bootstrapping-based updates are a key component behind the existence of implicit under-parameterization. To do so, we empirically demonstrate the following points empirically:

- Implicit under-parameterization occurs even when the form of the bootstrapping update is changed from Q-learning that utilizes a $\max_{a'}$ backup operator to a policy evaluation (fitted Q-evaluation) backup operator, that computes an expectation of the target Q-values under the distributions specified by a different policy. **Thus, with different bootstrapped updates, the phenomenon still appears.**




Figure A.9: **Offline Policy Evaluation on Atari.** $\text{srnk}_\delta(\Phi)$ and performance of offline policy evaluation (FQE) on 5 Atari games in the offline RL setting using the 5% and 20% DQN Replay dataset (Agarwal et al., 2020). The rank degradation shows that under-parameterization is not specific to the Bellman optimality operator but happens even when other bootstrapping-based backup operators are combined with gradient descent. Furthermore, the rank degradation also happens when we increase the dataset size.

- Implicit under-parameterization does not occur when Monte-Carlo regression targets - that compute regression targets for the Q-function by computing a non-parametric estimate of the future trajectory return, i.e., $y_k(s_t, a_t) = \sum_{t'=t}^{\infty} \gamma^{t'-t} r(s_{t'}, a_{t'})$ and do not use bootstrapping. In this setting, we find that the values of effective rank actually increase over time and stabilize, unlike the corresponding case for bootstrapped updates. **Thus, other factors kept identically the same, implicit under-parameterization happens only when bootstrapped updates are used.**




Figure A.10: **Monte Carlo Offline Policy Evaluation.** $\text{srnk}_\delta(\Phi)$ on 5 Atari games in when using Monte Carlo returns for targets and thus removing bootstrapping updates. Rank collapse does not happen in this setting implying that is bootstrapping was essential for under-parameterization. We perform the experiments using 5% and 20% DQN replay dataset from Agarwal et al. (2020).

- For the final point in this section, we verify if the non-stationarity of the policy in the Q-learning (control) setting, i.e., when the Bellman optimality operator is used is not a reason behind the emergence of implicit under-parameterization. The non-stationary policy in a control setting causes the targets to change and, as a consequence, leads to non-zero errors. However, rank drop is primarily caused by bootstrapping rather than non-stationarity of the control objective. To illustrate this, we ran an experiment in the control setting on Gridworld, regressing to the target computed using the true value function Q^π for the current policy π (computed using tabular Q-evaluation) instead of using the bootstrap TD estimate. The results, shown in figure A.11a, show that the srnk_δ doesn't decrease significantly when regressing to true control values and in fact increases with more iterations as compared to Figure 6a where rank drops with bootstrapping. This experiment, alongside with experiments discussed above, ablating bootstrapping in the stationary policy evaluation setting shows that rank-deficiency is caused due to bootstrapping.

A.4 HOW DOES IMPLICIT REGULARIZATION INHIBIT DATA-EFFICIENT RL?

Implicit under-parameterization leads to a trade-off between minimizing the TD error vs. encouraging low rank features as shown in Figure 4b. This trade-off often results in decrease in effective rank, at the expense of increase in TD error, resulting in lower performance. Here we present additional evidence to support this.

Figure A.11b shows a gridworld problem with one-hot features, which naturally leads to reduced state-aliasing. In this setting, we find that the amount of rank drop with respect to the supervised

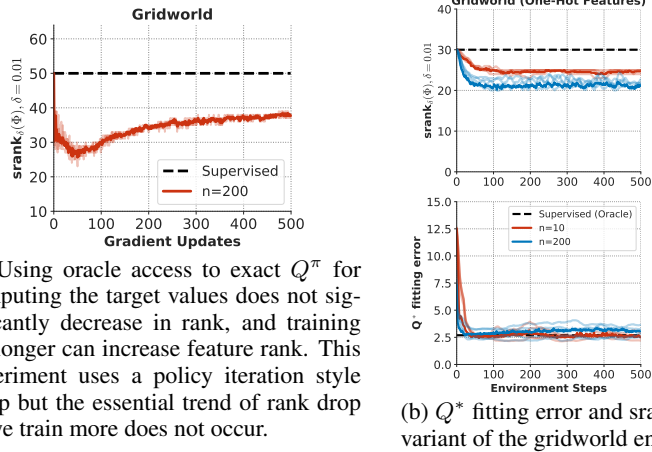


Figure A.11: Gridword: (a) Effective rank change due to non-stationarity (b) Q^* fitting error vs. srnk_{δ} .

projection of oracle computed Q^* values is quite small and the regression error to Q^* actually decreases unlike the case in Figure 4a, where it remains same or even increases. The method is able to learn policies that attain good performance as well. Hence, this justifies that when there's very little rank drop, for example, 5 rank units in the example on the right, FQI methods are generally able to learn Φ that is able to fit Q^* . This provides evidence showing that typical Q-networks learn Φ that can fit the optimal Q-function when rank collapse does not occur.

In Atari, we do not have access to Q^* , and so we instead measure the error in fitting the target value estimates, $\mathbf{R} + \gamma P^{\pi} \mathbf{Q}_k$. As rank decreases, the TD error increases (Figure A.12) and the value function is unable to fit the target values, culminating in a performance plateau (Figure A.6).

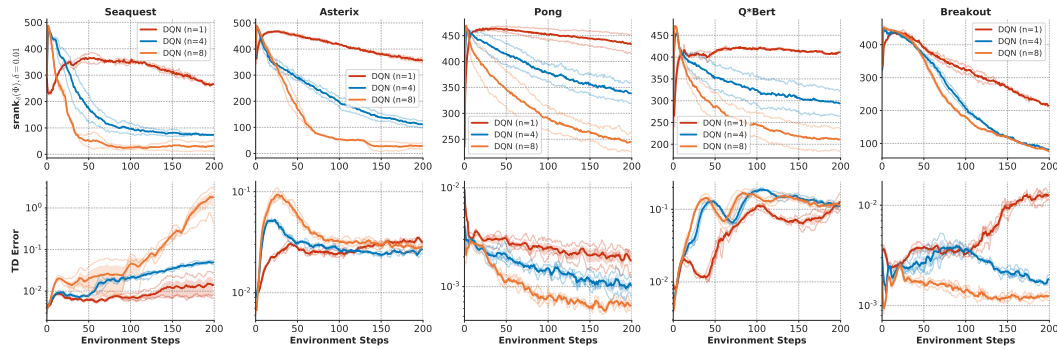


Figure A.12: **TD error vs. Effective rank on Atari.** We observe that Huber-loss TD error is often higher when there is a larger implicit under-parameterization, measured in terms of drop in effective rank. The results are shown for the data-efficient online RL setting.

A.5 TRENDS IN VALUES OF EFFECTIVE RANK WITH PENALTY.

In this section, we present the trend in the values of the effective rank when the penalty $\mathcal{L}_p(\Phi)$ is added. In each plot below, we present the value of $\text{srnk}_{\delta}(\Phi)$ with and without penalty respectively.

A.5.1 OFFLINE RL: DQN

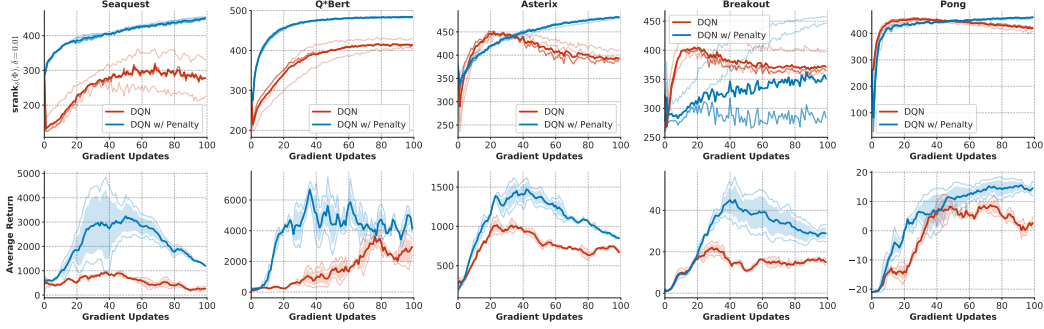


Figure A.13: **Effective rank values with the penalty on DQN.** Trends in effective rank and performance for offline DQN. Note that the performance of DQN with penalty is generally better than DQN and that the penalty (blue) is effective in increasing the values of effective rank. We report performance at the end of 100 epochs, as per the protocol set by Agarwal et al. (2020) in Figure 7.

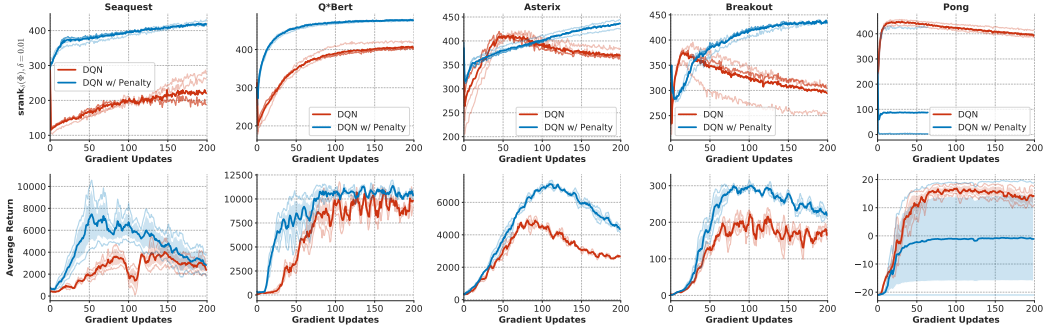


Figure A.14: **Effective rank values with the penalty on DQN on a 4x larger dataset.** Trends in effective rank and performance for offline DQN with a 4x larger dataset, where distribution shift effects are generally removed. Note that the performance of DQN with penalty is generally better than DQN and that the penalty (blue) is effective in increasing the values of effective rank in most cases. Infact in PONG, where the penalty is not effective in increasing rank, we observe suboptimal performance (blue vs. red).

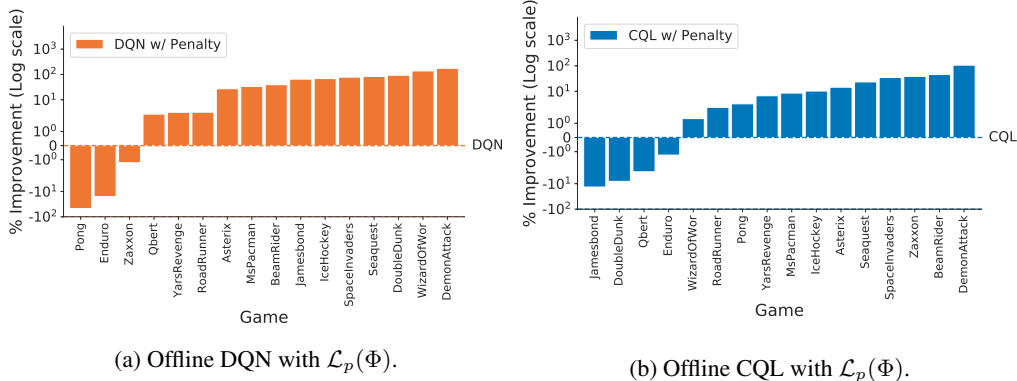


Figure A.15: Performance improvement of (a) offline DQN and (b) offline CQL with the $\mathcal{L}_p(\Phi)$ penalty on 20% Atari dataset, i.e., the dataset referred to as 4x large in Figure 2.

A.5.2 OFFLINE RL: CQL WITH $L_p(\Phi)$ PENALTY

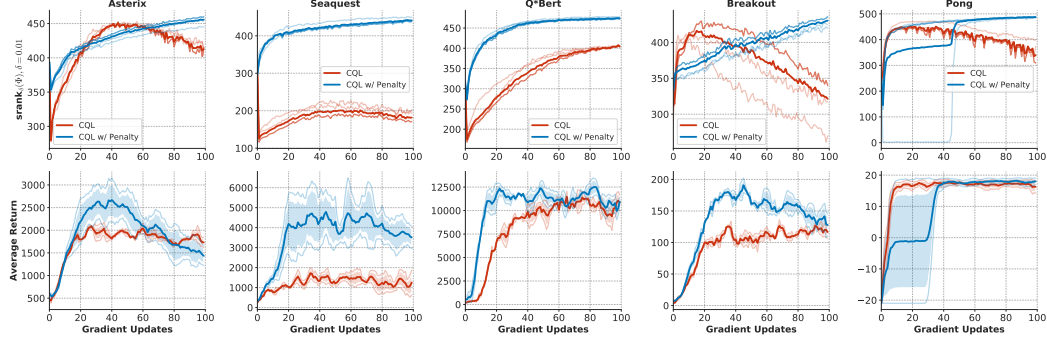


Figure A.16: **Effective rank values with the penalty $L_p(\Phi)$ on CQL.** Trends in effective rank and performance for offline DQN. Note that the performance of CQL with penalty is generally better than vanilla CQL and that the penalty (blue) is effective in increasing the values of effective rank. We report performance at the end of 100 epochs, as per the protocol set by Agarwal et al. (2020) in Figure 7.

A.6 DATA-EFFICIENT ONLINE RL: RAINBOW

A.6.1 RAINBOW WITH $L_p(\Phi)$ PENALTY: RANK PLOTS

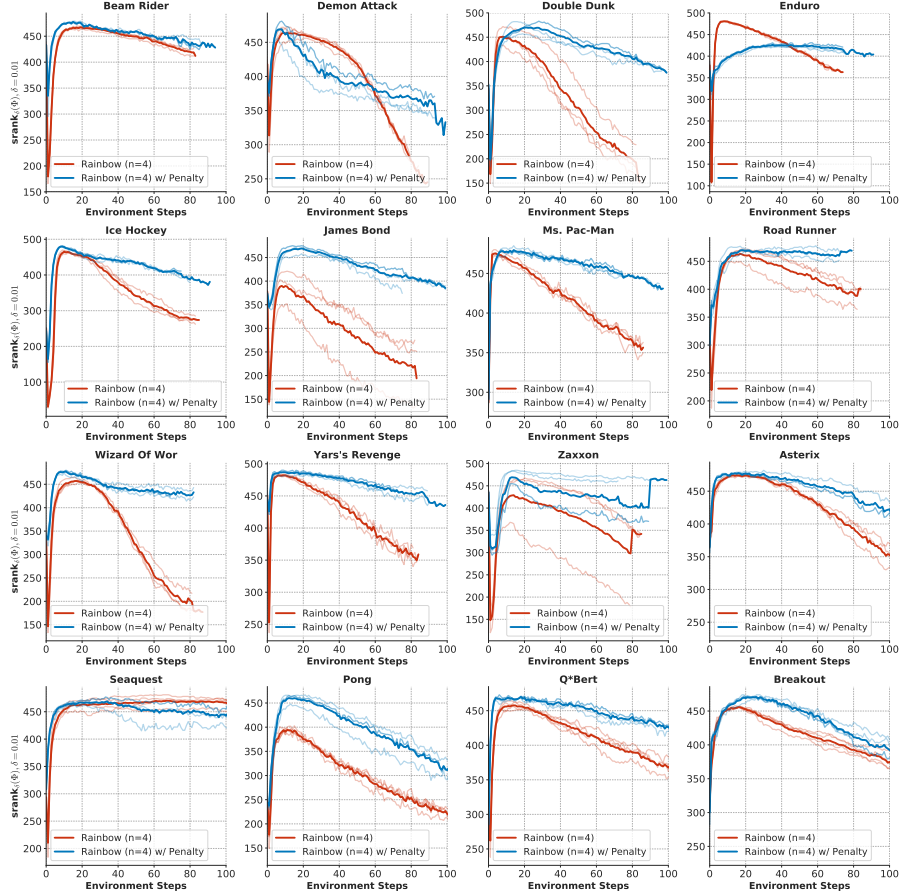



Figure A.17: **Effective rank values with the penalty on Rainbow in the data-efficient online RL setting.** Trends in effective rank and performance for online Rainbow, where distribution shift effects are generally removed. Note that the performance of DQN with penalty is generally better than DQN and that the penalty (blue) is effective in increasing the values of effective rank in most cases. Infact in PONG, where the penalty is not effective in increasing rank, we observe suboptimal performance (blue vs. red).

A.6.2 RAINBOW WITH $L_p(\Phi)$ PENALTY: PERFORMANCE

In this section, we present additional results for supporting the hypothesis that preventing rank-collapse leads to better performance. In the first set of experiments, we apply the proposed \mathcal{L}_p penalty to Rainbow in the data-efficient online RL setting ($n = 4$). In the second set of experiments, we present evidence for prevention of rank collapse by comparing rank values for different runs.



As we will show in the next section, the state-of-the-art Rainbow (Hessel et al., 2018) algorithm also suffers from rank collapse in the data-efficient online RL setting when more updates are performed per gradient step. In this section, we applied our penalty \mathcal{L}_p to Rainbow with $n = 4$, and obtained a median **20.66%** improvement on top of the base method. This result is summarized below.

Figure A.18: Performance of Rainbow ($n = 4$) and Rainbow ($n = 4$) with the $\mathcal{L}_p(\Phi)$ penalty (Equation 6). Note that the penalty improves on the base Rainbow in **12/16** games.

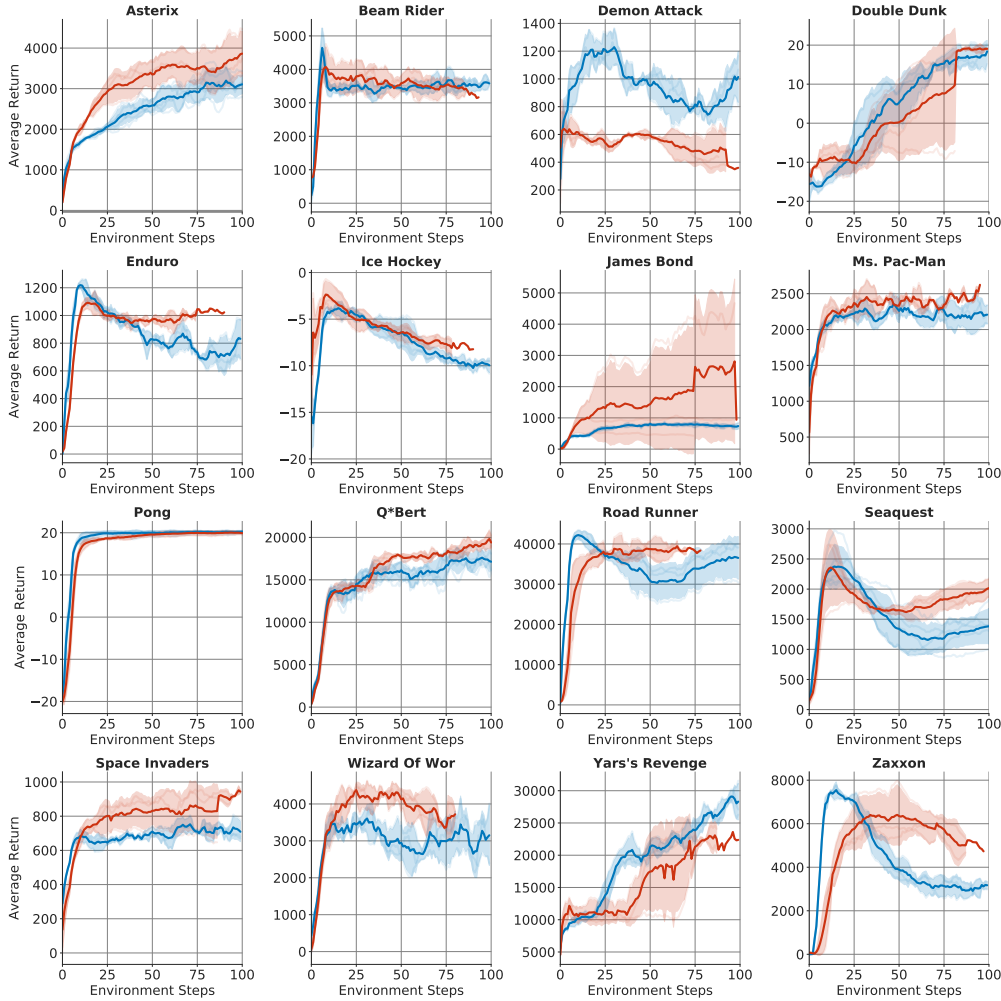


Figure A.19: Learning curves with $n = 4$ gradient updates per environment step for Rainbow(Blue) and Rainbow with the $\mathcal{L}_p(\Phi)$ penalty (Equation 6) (Red) on 16 games, corresponding to the bar plot above. One unit on the x-axis is equivalent to 1M environment steps.

A.7 RELAXING THE NORMALITY ASSUMPTION IN THEOREM 4.1

We can relax the normality assumption on \mathbf{S} in Theorem 4.1. An analogous statement holds for non-normal matrices \mathbf{S} for a slightly different notion of effective rank, denoted as $\text{srank}_{\delta,\lambda}(\mathbf{M}_k)$, that utilizes eigenvalue norms instead of singular values. Formally, let $\lambda_1(\mathbf{M}_k), \dots, \lambda_2(\mathbf{M}_k), \dots$ be the (complex) eigenvalues of \mathbf{M}_k arranged in decreasing order of their norms, *i.e.*, $|\lambda_1(\mathbf{M}_k)| \geq |\lambda_2(\mathbf{M}_k)| \geq \dots$, then,

$$\text{srank}_{\delta,\lambda}(\mathbf{M}_k) = \min \left\{ k : \frac{\sum_{i=1}^k |\lambda_i(\mathbf{M}_k)|}{\sum_{i=1}^d |\lambda_i(\mathbf{M}_k)|} \geq 1 - \delta \right\}.$$

A statement essentially analogous to Theorem 4.1 suggests that in this general case, $\text{srank}_{\delta,\lambda}(\mathbf{M}_k)$ decreases for *all* (complex) diagonalizable matrices \mathbf{S} , which is the set of almost all matrices of size $\dim(\mathbf{S})$. Now, if \mathbf{S} is approximately normal, *i.e.*, when $|\sigma_i(\mathbf{S}) - |\lambda_i(\mathbf{S})||$ is small, then the result in Theorem 4.1 also holds approximately as we discuss at the end of Appendix C.

We now provide empirical evidence showing that the trend in the values of effective rank computed using singular values, $\text{srank}_{\delta}(\Phi)$ is almost identical to the trend in the effective rank computed using normalized eigenvalues, $\text{srank}_{\delta,\lambda}(\Phi)$. Since eigenvalues are only defined for a square matrix Φ , in practice, we use a batch of $d = \dim(\phi(\mathbf{s}, \mathbf{a}))$ state-action pairs for computing the eigenvalue rank and compare to the corresponding singular value rank in Figures A.20 and A.21.




Figure A.20: Comparing different measures of effective rank on a run of offline DQN in the 5% replay setting, previously studied in Figure A.3.

Connection to Theorem 4.1. We computed the effective rank of Φ instead of \mathbf{S} , since \mathbf{S} is a theoretical abstraction that cannot be computed in practice as it depends on the Green’s kernel (Duffy, 2015) obtained by assuming that the neural network behaves as a kernel regressor. Instead, we compare the different notions of ranks of Φ since Φ is the practical counterpart for the matrix, \mathbf{S} , when using neural networks (as also indicated by the analysis in Section 4.2). In fact, on the gridworld (Figure A.21), we experiment with a feature Φ with dimension equal to the number of state-action pairs, *i.e.*, $\dim(\phi(\mathbf{s}, \mathbf{a})) = |S||A|$, with the same number of parameters as a kernel parameterization of the Q-function: $Q(\mathbf{s}, \mathbf{a}) = \sum_{\mathbf{s}', \mathbf{a}'} \mathbf{w}(\mathbf{s}', \mathbf{a}') k(\mathbf{s}, \mathbf{a}, \mathbf{s}', \mathbf{a}')$. This can also be considered as performing gradient descent on a “wide” linear network, and we measure the feature rank while observing similar rank trends.

Since we do not require the assumption that \mathbf{S} is normal in Theorem 4.1 to obtain a decreasing trend in $\text{srank}_{\delta,\lambda}(\Phi)$, and we find that in practical scenarios (Figures A.20 and A.21), $\text{srank}_{\delta}(\Phi) \approx \text{srank}_{\delta,\lambda}(\Phi)$ with an extremely similar qualitative trend we believe that Theorem 4.1 still explains the rank-collapse practically observed in deep Q-learning and is not vacuous.




Figure A.21: Comparison of $\text{srank}_{\delta,\lambda}(\Phi)$ and $\text{srank}_{\delta}(\Phi)$ on the gridworld in the offline setting (left) and the online setting (right), when using a deep linear network with $\dim(\phi(\mathbf{s}, \mathbf{a})) = |S||A|$. Note that both notions of effective rank exhibit a similar decreasing trend and are closely related to each other.

A.8 NORMALIZED PLOTS FOR FIGURE 3/ FIGURE A.6

In this section, we provide a set of normalized srnk and performance trends for Atari games (the corresponding unnormalized plots are found in Figure A.6). In these plots, each unit on the x-axis is equivalent to one gradient update, and so since $n = 8$ prescribes $8\times$ many updates as compared to $n = 1$, it it runs for $8\times$ as long as $n = 1$. These plots are in Figure A.22.

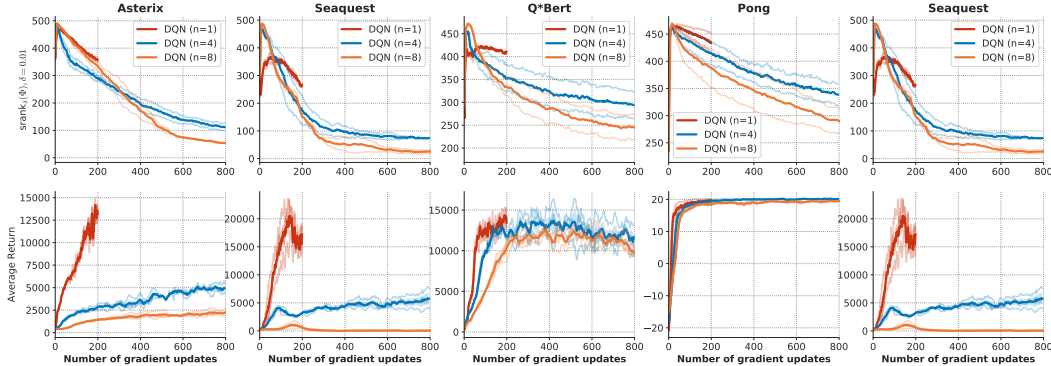


Figure A.22: **Rank collapse in DQN as a function of gradient updates on the x-axis** for five Atari games in the data-efficient online RL setting. This setting was previously studied in Figure A.6. Note that lesser number of updates per unit amount data, *i.e.*, smaller values of n possess larger srnk_δ values.

Note that the trend that effective rank decreases with larger n values also persists when rescaling the x-axis to account for the number of gradient steps, in all but one game. This is expected since it tells us that performing bootstrapping based updates in the data-efficient setting (larger n values) still leads to more aggressive rank drop as updates are being performed on a relatively more static dataset for larger values of n .

B HYPERPARAMETERS & EXPERIMENT DETAILS

B.1 ATARI EXPERIMENTS

We follow the experiment protocol from Agarwal et al. (2020) for all our experiments including hyperparameters and agent architectures provided in Dopamine and report them for completeness and ease of reproducibility in Table B.1. We only use hyperparameter selection over the regularization experiment α_p based on results from 5 Atari games (Asterix, Seaquest, Pong, Breakout and Seaquest). We will also open source our code to further aid in reproducing our results.

Evaluation Protocol. Following Agarwal et al. (2020), the Atari environments used in our experiments are stochastic due to sticky actions, *i.e.*, there is 25% chance at every time step that the environment will execute the agent’s previous action again, instead of the agent’s new action. All agents (online or offline) are compared using the best evaluation score (averaged over 5 runs) achieved during training where the evaluation is done online every training iteration using a ϵ -greedy policy with $\epsilon = 0.001$. We report offline training results with same hyperparameters over 5 random seeds of the DQN replay data collection, game simulator and network initialization.

Offline Dataset. As suggested by Agarwal et al. (2020), we randomly subsample the DQN Replay dataset containing 50 millions transitions to create smaller offline datasets with the same data distribution as the original dataset. We use the 5% DQN replay dataset for most of our experiments. We also report results using the 20% dataset setting (4x larger) to show that our claims are also valid even when we have higher coverage over the state space.

Optimizer related hyperparameters. For existing off-policy agents, step size and optimizer were taken as published. We used the DQN (Adam) algorithm for all our experiments, given its superior performance over the DQN (Nature) which uses RMSProp, as reported by Agarwal et al. (2020).

Atari 2600 games used. For all our experiments in Section 3, we used the same set of 5 games as utilized by Agarwal et al. (2020); Bellemare et al. (2017) to present analytical results. For our

Table B.1: Hyperparameters used by the offline and online RL agents in our experiments.

Hyperparameter	Setting (for both variations)	
Sticky actions	Yes	
Sticky action probability	0.25	
Grey-scaling	True	
Observation down-sampling	(84, 84)	
Frames stacked	4	
Frame skip (Action repetitions)	4	
Reward clipping	[-1, 1]	
Terminal condition	Game Over	
Max frames per episode	108K	
Discount factor	0.99	
Mini-batch size	32	
Target network update period	every 2000 updates	
Training environment steps per iteration	250K	
Update period every	4 environment steps	
Evaluation ϵ	0.001	
Evaluation steps per iteration	125K	
Q-network: channels	32, 64, 64	
Q-network: filter size	$8 \times 8, 4 \times 4, 3 \times 3$	
Q-network: stride	4, 2, 1	
Q-network: hidden units	512	
Hardware	Tesla P100 GPU	
Hyperparameter	Online	Offline
Min replay size for sampling	20,000	-
Training ϵ (for ϵ -greedy exploration)	0.01	-
ϵ -decay schedule	250K steps	-
Fixed Replay Memory	No	Yes
Replay Memory size (Online)	1,000,000 steps	-
Fixed Replay size (5%)	-	2,500,000 steps
Fixed Replay size (20%)	-	10,000,000 steps
Replay Scheme	Uniform	Uniform
Training Iterations	200	500

empirical evaluation in Appendix A.5, we use the set of games employed by Fedus et al. (2020b) which are deemed suitable for offline RL by Gulcehre et al. (2020). Similar in spirit to Gulcehre et al. (2020), we use the set of 5 games used for analysis for hyperparameter tuning for offline RL methods.

5 games subset: ASTERIX, QBERT, PONG, SEAQUEST, BREAKOUT

16 game subset: In addition to 5 games above, the following 11 games: DOUBLE DUNK, JAMES BOND, MS. PACMAN, SPACE INVADERS, ZAXXON, WIZARD OF WOR, YARS’ REVENGE, EN-DURO, ROAD RUNNER, BEAMRIDER, DEMON ATTACK

B.2 GRIDWORLD EXPERIMENTS

We use the gridworld suite from Fu et al. (2019) to obtain gridworlds for our experiments. All of our gridworld results are computed using the 16×16 GRID16SMOOTHOPS environment, which consists of a 256-cell grid, with walls arising randomly with a probability of 0.2. Each state allows 5 different actions (subject to hitting the boundary of the grid): move left, move right, move up, move down and no op. The goal in this environment is to minimize the cumulative discounted distance to a fixed goal location where the discount factor is given by $\gamma = 0.95$. The features for this Q-function are given by randomly chosen vectors which are smoothened spatially in a local neighborhood of a grid cell (x, y) .

We use a deep Q-network with two hidden layers of size (64, 64), and train it using soft Q-learning with entropy coefficient of 0.1, following the code provided by authors of Fu et al. (2019). We use a first-in-first out replay buffer of size 10000 to store past transitions.

C PROOFS FOR SECTION 4.1

In this section, we provide the technical proofs from Section 4.1. We first derive a solution to optimization problem Equation 1 and show that it indeed comes out to have the form described in Equation 2. We first introduce some notation, including definition of the kernel \mathbf{G} which was used for this proof. This proof closely follows the proof from Mobahi et al. (2020).

Definitions. For any universal kernel u , the Green’s function (Duffy, 2015) of the linear kernel operator L given by: $[L\mathbf{Q}](\mathbf{s}, \mathbf{a}) := \sum_{(\mathbf{s}', \mathbf{a}')} u((\mathbf{s}, \mathbf{a}), (\mathbf{s}', \mathbf{a}')) \mathbf{Q}(\mathbf{s}', \mathbf{a}')$ is given by the function $g((\mathbf{s}, \mathbf{a}), (\mathbf{s}', \mathbf{a}'))$ that satisfies:

$$\sum_{(\mathbf{s}, \mathbf{a})} u((\mathbf{s}, \mathbf{a}), (\mathbf{s}', \mathbf{a}')) g((\mathbf{s}', \mathbf{a}'), (\bar{\mathbf{s}}, \bar{\mathbf{a}})) = \delta((\mathbf{s}, \mathbf{a}) - (\bar{\mathbf{s}}, \bar{\mathbf{a}})), \quad (\text{C.1})$$

where δ is the Dirac-delta function. Thus, Green’s function can be understood as a kernel that “inverts” the universal kernel u to the identity (Dirac-delta) matrix. We can then define the matrix \mathbf{G} as the matrix of vectors $\mathbf{g}_{(\mathbf{s}, \mathbf{a})}$ evaluated on the training dataset, \mathcal{D} , however note that the functional $\mathbf{g}_{(\mathbf{s}, \mathbf{a})}$ can be evaluated for other state-action tuples, not present in \mathcal{D} .

$$\mathbf{G}((\mathbf{s}_i, \mathbf{a}_i), (\mathbf{s}_j, \mathbf{a}_j)) := \mathbf{g}((\mathbf{s}_i, \mathbf{a}_i), (\mathbf{s}_j, \mathbf{a}_j)) \quad \text{and} \quad \mathbf{g}_{(\mathbf{s}, \mathbf{a})}[i] = \mathbf{g}((\mathbf{s}, \mathbf{a}), (\mathbf{s}_i, \mathbf{a}_i)) \quad \forall (\mathbf{s}_i, \mathbf{a}_i) \in \mathcal{D}. \quad (\text{C.2})$$

Lemma C.0.1. The solution to Equation 1 is given by Equation 2.

Proof. This proof closely follows the proof of Proposition 1 from (Mobahi et al., 2020). We revisit key aspects the key parts of this proof here.

We restate the optimization problem below, and solve for the optimum \mathbf{Q}_k to this equation by applying the functional derivative principle.

$$\min_{\mathbf{Q} \in \mathcal{Q}} J(\mathbf{Q}) := \sum_{\mathbf{s}_i, \mathbf{a}_i \in \mathcal{D}} (\mathbf{Q}(\mathbf{s}_i, \mathbf{a}_i) - y_k(\mathbf{s}_i, \mathbf{a}_i))^2 + c \sum_{(\mathbf{s}, \mathbf{a})} \sum_{(\mathbf{s}', \mathbf{a}')} u((\mathbf{s}, \mathbf{a}), (\mathbf{s}', \mathbf{a}')) \mathbf{Q}(\mathbf{s}, \mathbf{a}) \mathbf{Q}(\mathbf{s}', \mathbf{a}').$$

The functional derivative principle would say that the optimal \mathbf{Q}_k to this problem would satisfy, for any other function f and for a small enough $\varepsilon \rightarrow 0$,

$$\forall f \in \mathcal{Q} : \left. \frac{\partial J(\mathbf{Q}_k + \varepsilon f)}{\partial \varepsilon} \right|_{\varepsilon=0} = 0. \quad (\text{C.3})$$

By setting the gradient of the above expression to 0, we obtain the following stationarity conditions on \mathbf{Q}_k (also denoting $(\mathbf{s}_i, \mathbf{a}_i) := \mathbf{x}_i$ for brevity:

$$\sum_{\mathbf{x}_i \in \mathcal{D}} \delta(\mathbf{x} - \mathbf{x}_i) (\mathbf{Q}_k(\mathbf{x}_i) - y_k(\mathbf{x}_i)) + c \sum_{\mathbf{x}} u(\mathbf{x}, \mathbf{x}') \mathbf{Q}_k(\mathbf{x}') = 0. \quad (\text{C.4})$$

Now, we invoke the definition of the Green’s function discussed above and utilize the fact that the Dirac-delta function can be expressed in terms of the Green’s function, we obtain a simplified version of the above relation:

$$\sum_{\mathbf{x}} u(\mathbf{x}, \mathbf{x}') \sum_{\mathbf{x}_i \in \mathcal{D}} (\mathbf{Q}_k(\mathbf{x}_i) - y_k(\mathbf{x}_i)) g(\mathbf{x}', \mathbf{x}_i) = -c \sum_{\mathbf{x}} u(\mathbf{x}, \mathbf{x}') \mathbf{Q}_k(\mathbf{x}'). \quad (\text{C.5})$$

Since the kernel $u(\mathbf{x}, \mathbf{x}')$ is universal and positive definite, the optimal solution $\mathbf{Q}_k(\mathbf{x})$ is given by:

$$\mathbf{Q}_k(\mathbf{s}, \mathbf{a}) = -\frac{1}{c} \sum_{(\mathbf{s}_i, \mathbf{a}_i) \in \mathcal{D}} (\mathbf{Q}_k(\mathbf{s}_i, \mathbf{a}_i) - y_k(\mathbf{s}_i, \mathbf{a}_i)) \cdot g((\mathbf{s}, \mathbf{a}), (\mathbf{s}_i, \mathbf{a}_i)). \quad (\text{C.6})$$

Finally we can replace the expression for residual error, $\mathbf{Q}_k(\mathbf{s}_i, \mathbf{a}_i) - y_k(\mathbf{s}_i, \mathbf{a}_i)$ using the green’s kernel on the training data by solving for it in closed form, which gives us the solution in Equation 2.

$$\mathbf{Q}_k(\mathbf{s}, \mathbf{a}) = -\frac{1}{c} \mathbf{g}_{(\mathbf{s}, \mathbf{a})}^T (\mathbf{Q}_k - y_k) = \mathbf{g}_{(\mathbf{s}, \mathbf{a})}^T (c\mathbf{I} + \mathbf{G})^{-1} y_k. \quad (\text{C.7})$$

□

Next, we now state and prove a slightly stronger version of Theorem 4.1 that immediately implies the original theorem.

Theorem C.1. *Let \mathbf{S} be a shorthand for $\mathbf{S} = \gamma P^\pi \mathbf{A}$ and assume \mathbf{S} is a normal matrix. Then there exists an infinite, strictly increasing sequence of fitting iterations, $(k_l)_{l=1}^\infty$ starting from $k_1 = 0$, such that, for any two singular-values $\sigma_i(\mathbf{S})$ and $\sigma_j(\mathbf{S})$ of \mathbf{S} with $\sigma_i(\mathbf{S}) \leq \sigma_j(\mathbf{S})$,*

$$\forall l \in \mathbb{N} \text{ and } l' \geq l, \quad \frac{\sigma_i(\mathbf{M}_{k_{l'}})}{\sigma_j(\mathbf{M}_{k_{l'}})} < \frac{\sigma_i(\mathbf{M}_{k_l})}{\sigma_j(\mathbf{M}_{k_l})} \leq \frac{\sigma_i(\mathbf{S})}{\sigma_j(\mathbf{S})}. \quad (\text{C.8})$$

Therefore, the effective rank of \mathbf{M}_k satisfies: $\text{srnk}_\delta(\mathbf{M}_{k_{l'}}) \leq \text{srnk}_\delta(\mathbf{M}_{k_l})$. Furthermore,

$$\forall l \in \mathbb{N} \text{ and } t \geq k_l, \quad \frac{\sigma_i(\mathbf{M}_t)}{\sigma_j(\mathbf{M}_t)} < \frac{\sigma_i(\mathbf{M}_{k_l})}{\sigma_j(\mathbf{M}_{k_l})} + \mathcal{O}\left(\left(\frac{\sigma_i(\mathbf{S})}{\sigma_j(\mathbf{S})}\right)^{k_l}\right). \quad (\text{C.9})$$

Therefore, the effective rank of \mathbf{M}_t , $\text{srnk}_\delta(\mathbf{M}_t)$, outside the chosen subsequence is also controlled above by the effective rank on the subsequence $(\text{srnk}_\delta(\mathbf{M}_{k_l}))_{l=1}^\infty$.

To prove this theorem, we first show that for any two fitting iterations, $t < t'$, if \mathbf{S}^t and $\mathbf{S}^{t'}$ are positive semi-definite, the ratio of singular values and the effective rank decreases from t to t' . As an immediate consequence, this shows that when \mathbf{S} is positive semi-definite, the effective rank decreases at every iteration, i.e., by setting $k_l = l$ (Corollary C.1.1).

To extend the proof to arbitrary normal matrices, we show that for any \mathbf{S} , a sequence of fitting iterations $(k_l)_{l=1}^\infty$ can be chosen such that \mathbf{S}^{k_l} is (approximately) positive semi-definite. For this subsequence of fitting iterations, the ratio of singular values and effective rank also decreases. Finally, to control the ratio and effective rank on fitting iterations t outside this subsequence, we construct an upper bound on the ratio $f(t): \frac{\sigma_i(\mathbf{M}_t)}{\sigma_j(\mathbf{M}_t)} < f(t)$, and relate this bound to the ratio of singular values on the chosen subsequence.

Lemma C.1.1 ($\text{srnk}_\delta(\mathbf{M}_k)$ decreases when \mathbf{S}^k is PSD.). *Let \mathbf{S} be a shorthand for $\mathbf{S} = \gamma P^\pi \mathbf{A}$ and assume \mathbf{S} is a normal matrix. Choose any $t, t' \in \mathbb{N}$ such that $t < t'$. If \mathbf{S}^t and $\mathbf{S}^{t'}$ are positive semi-definite, then for any two singular-values $\sigma_i(\mathbf{S})$ and $\sigma_j(\mathbf{S})$ of \mathbf{S} , such that $0 < \sigma_i(\mathbf{S}) < \sigma_j(\mathbf{S})$:*

$$\frac{\sigma_i(\mathbf{M}_{t'})}{\sigma_j(\mathbf{M}_{t'})} < \frac{\sigma_i(\mathbf{M}_t)}{\sigma_j(\mathbf{M}_t)} \leq \frac{\sigma_i(\mathbf{S})}{\sigma_j(\mathbf{S})}. \quad (\text{C.10})$$

Hence, the effective rank of \mathbf{M}_k decreases from t to t' : $\text{srnk}_\delta(\mathbf{M}_{t'}) \leq \text{srnk}_\delta(\mathbf{M}_t)$.

Proof. First note that \mathbf{M}_k is given by:

$$\mathbf{M}_k := \sum_{i=1}^k \gamma^{k-i} (P^\pi \mathbf{A})^{k-i} = \sum_{i=1}^k \mathbf{S}^{k-i}. \quad (\text{C.11})$$

From hereon, we omit the leading γ term since it is a constant scaling factor that does not affect ratio or effective rank. Almost every matrix \mathbf{S} admits a complex orthogonal eigendecomposition. Thus, we can write $\mathbf{S} := \mathbf{U} \lambda(\mathbf{S}) \mathbf{U}^H$. And any power of \mathbf{S} , i.e., \mathbf{S}^i can be expressed as: $\mathbf{S}^i = \mathbf{U} \lambda(\mathbf{S})^i \mathbf{U}^H$, and hence, we can express \mathbf{M}_k as:

$$\mathbf{M}_k := \mathbf{U} \left(\sum_{i=0}^{k-1} \lambda(\mathbf{S})^i \right) \mathbf{U}^H = \mathbf{U} \cdot \text{diag} \left(\frac{1 - \lambda(\mathbf{S})^k}{1 - \lambda(\mathbf{S})} \right) \cdot \mathbf{U}^H. \quad (\text{C.12})$$

Since \mathbf{S} is normal, its eigenvalues and singular values are further related as $\sigma_k(\mathbf{S}) = |\lambda_k(\mathbf{S})|$. And this also means that \mathbf{M}_k is normal, indicating that $\sigma_i(\mathbf{M}_k) = |\lambda_i(\mathbf{M}_k)|$. Thus, the singular values of \mathbf{M}_k can be expressed as

$$\sigma_i(\mathbf{M}_k) := \left| \frac{1 - \lambda_i(\mathbf{S})^k}{1 - \lambda_i(\mathbf{S})} \right|, \quad (\text{C.13})$$

When \mathbf{S}^k is positive semi-definite, $\lambda_i(\mathbf{S})^k = \sigma_i(\mathbf{S})^k$, enabling the following simplification:

$$\sigma_i(\mathbf{M}_k) = \frac{|1 - \sigma_i(\mathbf{S})^k|}{|1 - \lambda_i(\mathbf{S})|}. \quad (\text{C.14})$$

To show that the ratio of singular values decreases from t to t' , we need to show that $f(\sigma) = \frac{|1-\sigma^{t'}|}{|1-\sigma^t|}$ is an increasing function of σ when $t' > t$. It can be seen that this is the case, which implies the desired result.

To further show that $\text{srnk}_\delta(\mathbf{M}_t) \geq \text{srnk}_\delta(\mathbf{M}_{t'})$, we can simply show that $\forall i \in [1, \dots, n]$, $h_k(i) := \frac{\sum_{j=1}^i \sigma_j(\mathbf{M}_k)}{\sum_{j=1}^n \sigma_j(\mathbf{M}_k)}$ increases with k , and this would imply that the $\text{srnk}_\delta(\mathbf{M}_k)$ cannot increase from $k = t$ to $k = t'$. We can decompose $h_k(i)$ as:

$$h_k(i) = \sum_{j=1}^i \frac{\sigma_j(\mathbf{M}_k)}{\sum_{l=1}^n \sigma_l(\mathbf{M}_k)} = \frac{1}{1 + \frac{\sum_{j=i+1}^n \sigma_j(\mathbf{M}_k)}{\sum_{j=1}^i \sigma_j(\mathbf{M}_k)}}. \quad (\text{C.15})$$

Since $\sigma_j(\mathbf{M}_k)/\sigma_l(\mathbf{M}_i)$ decreases over time k for all j, l if $\sigma_j(\mathbf{S}) \leq \sigma_l(\mathbf{S})$, the ratio in the denominator of $h_k(i)$ decreases with increasing k implying that $h_k(i)$ increases from t to t' . \square

Corollary C.1.1 ($\text{srnk}_\delta(\mathbf{M}_k)$ decreases for PSD \mathbf{S} matrices.). *Let \mathbf{S} be a shorthand for $\mathbf{S} = \gamma P^\pi \mathbf{A}$. Assuming that \mathbf{S} is positive semi-definite, for any $k, t \in \mathbb{N}$, such that $t > k$ and that for any two singular-values $\sigma_i(\mathbf{S})$ and $\sigma_j(\mathbf{S})$ of \mathbf{S} , such that $\sigma_i(\mathbf{S}) < \sigma_j(\mathbf{S})$,*

$$\frac{\sigma_i(\mathbf{M}_t)}{\sigma_j(\mathbf{M}_t)} < \frac{\sigma_i(\mathbf{M}_k)}{\sigma_j(\mathbf{M}_k)} \leq \frac{\sigma_i(\mathbf{S})}{\sigma_j(\mathbf{S})}. \quad (\text{C.16})$$

Hence, the effective rank of \mathbf{M}_k decreases with more fitting iterations: $\text{srnk}_\delta(\mathbf{M}_t) \leq \text{srnk}_\delta(\mathbf{M}_k)$.

In order to now extend the result to arbitrary normal matrices, we must construct a subsequence of fitting iterations $(k_l)_{l=1}^\infty$ where \mathbf{S}^{k_l} is (approximately) positive semi-definite. To do so, we first prove a technical lemma that shows that rational numbers, i.e., numbers that can be expressed as $r = \frac{p}{q}$, for integers $p, q \in \mathbb{Z}$ are “dense” in the space of real numbers.

Lemma C.1.2 (Rational numbers are dense in the real space.). *For any real number α , there exist infinitely many rational numbers $\frac{p}{q}$ such that α can be approximated by $\frac{p}{q}$ upto $\frac{1}{q^2}$ accuracy.*

$$\left| \alpha - \frac{p}{q} \right| \leq \frac{1}{q^2}. \quad (\text{C.17})$$

Proof. We first use Dirichlet’s approximation theorem (see [Hlawka et al. \(1991\)](#)) for a proof of this result using a pigeonhole argument and extensions) to obtain that for any real numbers α and $N \geq 1$, there exist integers p and q such that $1 \leq q \leq N$ and,

$$|q\alpha - p| \leq \frac{1}{|N| + 1} < \frac{1}{N}. \quad (\text{C.18})$$

Now, since $q \geq 1 > 0$, we can divide both sides by q , to obtain:

$$\left| \alpha - \frac{p}{q} \right| \leq \frac{1}{Nq} \leq \frac{1}{q^2}. \quad (\text{C.19})$$

To obtain infinitely many choices for $\frac{p}{q}$, we observe that Dirichlet’s lemma is valid only for all values of N that satisfy $N \leq \frac{1}{|q\alpha - p|}$. Thus if we choose an N' such that $N' \geq N_{\max}$ where N_{\max} is defined as:

$$N_{\max} = \max \left\{ \frac{1}{|q'\alpha - p'|} \mid p', q' \in \mathbb{Z}, 1 \leq q' \leq q \right\}. \quad (\text{C.20})$$

Equation C.20 essentially finds a new value of N , such that the current choices of p and q , which were valid for the first value of N do not satisfy the approximation error bound. Applying Dirichlet’s lemma to this new value of N' hence gives us a new set of p' and q' which satisfy the $\frac{1}{q'^2}$ approximation error bound. Repeating this process gives us countably many choices of (p, q) pairs that satisfy the approximation error bound. As a result, rational numbers are dense in the space of real numbers, since for any arbitrarily chosen approximation accuracy given by $\frac{1}{q^2}$, we can obtain atleast one rational number, $\frac{p}{q}$ which is closer to α than $\frac{1}{q^2}$. This proof is based on [Johnson \(2016\)](#). \square

Now we utilize Lemmas C.1.1 and C.1.2 to prove Proposition 4.1.

Proof of Proposition 4.1 and Theorem C.1 Recall from the proof of Lemma C.1.1 that the singular values of \mathbf{M}_k are given by:

$$\sigma_i(\mathbf{M}_k) := \left| \frac{1 - \lambda_i(\mathbf{S})^k}{1 - \lambda_i(\mathbf{S})} \right|, \quad (\text{C.21})$$

Bound on Singular Value Ratio: The ratio between $\sigma_i(\mathbf{M}_k)$ and $\sigma_j(\mathbf{M}_k)$ can be expressed as

$$\frac{\sigma_i(\mathbf{M}_k)}{\sigma_j(\mathbf{M}_k)} = \left| \frac{1 - \lambda_i(\mathbf{S})^k}{1 - \lambda_j(\mathbf{S})^k} \right| \left| \frac{1 - \lambda_j(\mathbf{S})}{1 - \lambda_i(\mathbf{S})} \right|. \quad (\text{C.22})$$

For a normal matrix \mathbf{S} , $\sigma_i(\mathbf{S}) = |\lambda_i(\mathbf{S})|$, so this ratio can be bounded above as

$$\frac{\sigma_i(\mathbf{M}_k)}{\sigma_j(\mathbf{M}_k)} \leq \frac{1 + \sigma_i(\mathbf{S})^k}{|1 - \sigma_j(\mathbf{S})^k|} \left| \frac{1 - \lambda_j(\mathbf{S})}{1 - \lambda_i(\mathbf{S})} \right|. \quad (\text{C.23})$$

Defining $f(k)$ to be the right hand side of the equation, we can verify that f is a monotonically decreasing function in k when $\sigma_i < \sigma_j$. This shows that this ratio of singular values is bounded above and in general, must decrease towards some limit $\lim_{k \rightarrow \infty} f(k)$.

Construction of Subsequence: We now show that there exists a subsequence $(k_l)_{l=1}^\infty$ for which \mathbf{S}^{k_l} is approximately positive semi-definite. For ease of notation, let's represent the i -th eigenvalue as $\lambda_i(\mathbf{S}) = |\lambda_i(\mathbf{S})| \cdot e^{i\theta_i}$, where $\theta_i > 0$ is the polar angle of the complex value $\lambda_i(\mathbf{S})$ and $|\lambda_i(\mathbf{S})|$ is its magnitude (norm). Now, using Lemma C.1.2, we can approximate any polar angle, θ_i using a rational approximation, *i.e.*, we apply lemma C.1.2 on $\frac{\theta_i}{2\pi}$

$$\exists p_i, q_i \in \mathbb{N}, \text{ s.t. } \left| \frac{\theta_i}{2\pi} - \frac{p_i}{q_i} \right| \leq \frac{1}{q_i^2}. \quad (\text{C.24})$$

Since the choice of q_i is within our control we can estimate θ_i for all eigenvalues $\lambda_i(\mathbf{S})$ to infinitesimal accuracy. Hence, we can approximate $\theta_i \approx 2\pi \frac{p_i}{q_i}$. We will now use this approximation to construct an infinite sequence $(k_l)_{l=1}^\infty$, shown below:

$$k_l = l \cdot \text{LCM}(q_1, \dots, q_n) \quad \forall j \in \mathbb{N}, \quad (\text{C.25})$$

where LCM is the least-common-multiple of natural numbers q_1, \dots, q_n .

In the absence of any approximation error in θ_i , we note that for any i and for any $l \in \mathbb{N}$ as defined above, $\lambda_i(\mathbf{S})^{k_l} = |\lambda_i(\mathbf{S})|^{k_l} \cdot \exp\left(2i\pi \cdot \frac{p_i}{q_i} \cdot k_l\right) = |\lambda_i(\mathbf{S})|^{k_l}$, since the polar angle for any k_l is going to be a multiple of 2π , and hence it would fall on the real line. As a result, \mathbf{S}^{k_l} will be positive semi-definite, since all eigenvalues are positive and real. Now by using the proof for Lemma C.1.1, we obtain the ratio of i and j singular values are increasing over the sequence of iterations $(k_j)_{j=1}^\infty$. Since the approximation error in θ_i can be controlled to be infinitesimally small to prevent any increase in the value of srnk_δ due to it (this can be done given the discrete form of srnk_δ), we note that the above argument applies even with the approximation, proving the required result on the subsequence.

Controlling All Fitting Iterations using Subsequence:

We now relate the ratio of singular values within this chosen subsequence to the ratio of singular values elsewhere. Choose $t, l \in \mathbb{N}$ such that $t > k_l$. Earlier in this proof, we showed that the ratio between singular values is bounded above by a monotonically decreasing function $f(t)$, so

$$\frac{\sigma_i(\mathbf{M}_t)}{\sigma_j(\mathbf{M}_t)} \leq f(t) < f(k_l). \quad (\text{C.26})$$

Now, we show that that $f(k_l)$ is in fact very close to the ratio of singular values:

$$f(k_l) = \left| \frac{1 - \sigma_i(\mathbf{S})^{k_l}}{1 - \sigma_j(\mathbf{S})^{k_l}} \right| \left| \frac{1 - \lambda_j(\mathbf{S})}{1 - \lambda_i(\mathbf{S})} \right| \leq \frac{\sigma_i(\mathbf{M}_t)}{\sigma_j(\mathbf{M}_t)} + \frac{2\sigma_i(\mathbf{S})^{k_l}}{|1 - \sigma_j(\mathbf{S})^{k_l}|} \left| \frac{1 - \lambda_j(\mathbf{S})}{1 - \lambda_i(\mathbf{S})} \right|. \quad (\text{C.27})$$

The second term goes to zero as k_l increases; algebraic manipulation shows that this gap be bounded by

$$f(k_l) \leq \frac{\sigma_i(\mathbf{M}_{k_l})}{\sigma_j(\mathbf{M}_{k_l})} + \underbrace{\left(\frac{\sigma_i(\mathbf{S})}{\sigma_j(\mathbf{S})} \right)^{k_l} \frac{2\sigma_j(\mathbf{S})}{|1 - \sigma_j(\mathbf{S})|} \left| \frac{1 - \lambda_j(\mathbf{S})}{1 - \lambda_i(\mathbf{S})} \right|}_{\text{constant}}. \quad (\text{C.28})$$

Putting these inequalities together proves the final statement,

$$\frac{\sigma_i(\mathbf{M}_t)}{\sigma_j(\mathbf{M}_t)} \leq \frac{\sigma_i(\mathbf{M}_{k_l})}{\sigma_j(\mathbf{M}_{k_l})} + \mathcal{O}\left(\left(\frac{\sigma_i(S)}{\sigma_j(S)}\right)^{k_l}\right). \quad (\text{C.29})$$

Extension to approximately-normal \mathbf{S} . We can extend the result in Theorem C.1 (and hence also Theorem 4.1) to approximately-normal \mathbf{S} . Note that the main requirement for normality of \mathbf{S} (i.e., $\sigma_i(\mathbf{S}) = |\lambda_i(\mathbf{S})|$) is because it is straightforward to relate the eigenvalue of \mathbf{S} to \mathbf{M} as shown below.

$$|\lambda_i(\mathbf{M}_k)| := \left| \frac{1 - \lambda_i(\mathbf{S})^k}{1 - \lambda_i(\mathbf{S})} \right|, \quad (\text{C.30})$$

Now, since the matrix \mathbf{S} is approximately normal, we can express it using its Schur’s triangular form as, $\mathbf{S} = \mathbf{U} \cdot (\Lambda + \mathbf{N}) \cdot \mathbf{U}^H$, where Λ is a diagonal matrix and \mathbf{N} is an “offset” matrix. The *departure from normality* of \mathbf{S} is defined as: $\Delta(\mathbf{S}) := \inf_{\mathbf{N}} \|\mathbf{N}\|_2$, where the infimum is computed over all matrices \mathbf{N} that can appear in the Schur triangular form for \mathbf{S} . For a normal \mathbf{S} only a single value of $\mathbf{N} = 0$ satisfies the Schur’s triangular form. For an approximately normal matrix \mathbf{S} , $\|\mathbf{N}\|_2 \leq \Delta(\mathbf{S}) \leq \varepsilon$, for a small ε .

Furthermore note that from Equation 6 in Ruhe (1975), we obtain that

$$|\sigma_i(\mathbf{S}) - |\lambda_i(\mathbf{S})|| \leq \Delta(\mathbf{S}) \leq \varepsilon, \quad (\text{C.31})$$

implying that singular values and norm-eigenvalues are close to each other for \mathbf{S} .

Next, let us evaluate the departure from normality of \mathbf{M}_k . First note that, $\mathbf{S}^j = \mathbf{U} \cdot (\Lambda + \mathbf{N})^j \cdot \mathbf{U}^H$, and so, $\mathbf{M}_k = \mathbf{U} \cdot \left(\sum_{j=1}^k (\Lambda + \mathbf{N})^j\right) \cdot \mathbf{U}^H$ and if $\|\mathbf{N}\|_2 \leq \varepsilon$, for a small epsilon (i.e., considering only terms that are linear in \mathbf{N} for $(\Lambda + \mathbf{N})^j$), we note that:

$$|\sigma_i(\mathbf{M}_k) - |\lambda_i(\mathbf{M}_k)|| \leq \sum_{j=1}^k j \cdot |\lambda_1(\mathbf{S})|^{j-1} \Delta(\mathbf{S}) \leq \frac{1}{(1 - |\lambda_1(\mathbf{S})|)^2} \cdot \varepsilon. \quad (\text{C.32})$$

Thus, the matrix \mathbf{M}_k is also approximately normal provided that the max eigenvalue norm of \mathbf{S} is less than 1. This is true, since $\mathbf{S} = \gamma P^\pi \mathbf{A}$ (see Theorem 4.1, where both P^π and \mathbf{A} have eigenvalues less than 1, and $\gamma < 1$).

Given that we have shown that \mathbf{M}_k is approximately normal, we can show that $\text{srnk}_\delta(\mathbf{M}_k)$ only differs from $\text{srnk}_{\delta, \lambda}(\mathbf{M}_k)$, i.e., the effective rank of eigenvalues, in a bounded amount. If the value of ε is then small enough, we still retain the conclusion that $\text{srnk}_\delta(\mathbf{M}_k)$ generally decreases with more training by following the proof of Theorem C.1.

D PROOFS FOR SECTION 4.2

In this section, we provide technical proofs from Section 4.2. We start by deriving properties of optimization trajectories of the weight matrices of the deep linear network similar to Arora et al. (2018) but customized to our set of assumptions, then prove Proposition 4.1, and finally discuss how to extend these results to the fitted Q-iteration setting and some extensions not discussed in the main paper. Similar to Section 4.1, we assume access to a dataset of transitions, $\mathcal{D} = \{(\mathbf{s}_i, \mathbf{a}_i, r(\mathbf{s}_i, \mathbf{a}_i), \mathbf{s}'_i)\}$ in this section, and assume that the same data is used to re-train the function.

Notation and Definitions. The Q-function is represented using a deep linear network with at least 3 layers, such that

$$Q(\mathbf{s}, \mathbf{a}) = \mathbf{W}_N \mathbf{W}_{N-1} \cdots \mathbf{W}_1 [\mathbf{s}; \mathbf{a}], \text{ where } N \geq 3, \mathbf{W}_N \in \mathbb{R}^{1 \times d_{N-1}}, \quad (\text{D.1})$$

and $\mathbf{W}_i \in \mathbb{R}^{d_i \times d_{i-1}}$ for $i = 1, \dots, N-1$. We index the weight matrices by a tuple (k, t) : $\mathbf{W}_j(k, t)$ denotes the weight matrix \mathbf{W}_j at the t -th step of gradient descent during the k -th fitting iteration (Algorithm 1). Let the end-to-end weight matrix $\mathbf{W}_N \mathbf{W}_{N-1} \cdots \mathbf{W}_1$ be denoted shorthand as $\mathbf{W}_{N:1}$, and let the features of the penultimate layer of the network, be denoted as $\mathbf{W}_\phi(k, t) :=$

$\mathbf{W}_{N-1}(k, t) \cdots \mathbf{W}_1(k, t)$. $\mathbf{W}_\phi(k, t)$ is the matrix that maps an input $[\mathbf{s}; \mathbf{a}]$ to corresponding features $\Phi(\mathbf{s}, \mathbf{a})$. In our analysis, it is sufficient to consider the effective rank of $\mathbf{W}_\phi(k, t)$ since the features Φ are given by: $\Phi(k, t) = \mathbf{W}_\phi(k, t)[\mathcal{S}; \mathcal{A}]$, which indicates that:

$$\text{rank}(\Phi(k, t)) = \text{rank}(\mathbf{W}_\phi(k, t)[\mathcal{S}; \mathcal{A}]) \leq \min(\text{rank}(\mathbf{W}_\phi(k, t)), \text{rank}([\mathcal{S}; \mathcal{A}])).$$

Assuming the state-action space has full rank, we are only concerned about $\text{rank}(\mathbf{W}_\phi(k, t))$ which justifies our choice for analyzing $\text{srnk}_\delta(\mathbf{W}_\phi(k, t))$.

Let $L_{k+1}(\mathbf{W}_{N:1}(k, t))$ denote the mean squared Bellman error optimization objective in the k -th fitting iteration.

$$L_{k+1}(\mathbf{W}_{N:1}(k, t)) = \sum_{i=1}^{|\mathcal{D}|} (\mathbf{W}_N(k, t) \mathbf{W}_\phi(k, t) [\mathbf{s}_i; \mathbf{a}_i] - \mathbf{y}_k(\mathbf{s}_i, \mathbf{a}_i))^2, \text{ where } \mathbf{y}_k = \mathbf{R} + \gamma P^\pi \mathbf{Q}_k.$$

When gradient descent is used to update the weight matrix, the updates to $\mathbf{W}_i(k, t)$ are given by:

$$\mathbf{W}_j(k, t+1) \leftarrow \mathbf{W}_j(k, t) - \eta \frac{\partial L_{k+1}(\mathbf{W}_{N:1}(k, t))}{\partial \mathbf{W}_j(k, t)}.$$

If the learning rate η is small, we can approximate this discrete time process with a continuous-time differential equation, which we will use for our analysis. We use $\dot{\mathbf{W}}(k, t)$ to denote the derivative of $\mathbf{W}(k, t)$ with respect to t , for a given k .

$$\dot{\mathbf{W}}_j(k, t) = -\eta \frac{\partial L_{k+1}(\mathbf{W}_{N:1}(k, t))}{\partial \mathbf{W}_j(k, t)} \quad (\text{D.2})$$

In order to quantify the evolution of singular values of the weight matrix, $\mathbf{W}_\phi(k, t)$, we start by quantifying the evolution of the weight matrix $\mathbf{W}_\phi(k, t)$ using a more interpretable differential equation. In order to do so, we make an assumption similar to but not identical as [Arora et al. \(2018\)](#), that assumes that all except the last weight matrix are “balanced” at initialization $t = 0, k = 0$. i.e.

$$\forall i \in [0, \dots, N-2] : \mathbf{W}_{i+1}^T(0, 0) \mathbf{W}_{i+1}(0, 0) = \mathbf{W}_i(0, 0) \mathbf{W}_i(0, 0)^T. \quad (\text{D.3})$$

Note the distinction from [Arora et al. \(2018\)](#), the last layer is not assumed to be balanced. As a result, we may not be able to comment about the learning dynamics of the end-to-end weight matrix, but we prevent the vacuous case where all the weight matrices are rank 1. Now we are ready to derive the evolution of the feature matrix, $\mathbf{W}_\phi(k, t)$.

Lemma D.0.1 (Adaptation of balanced weights ([Arora et al., 2018](#)) across FQI iterations). *Assume the weight matrices evolve according to Equation D.2, with respect to L_k for all fitting iterations k . Assume balanced initialization only for the first $N-1$ layers, i.e., $\mathbf{W}_{j+1}(0, 0)^T \mathbf{W}_{j+1}(0, 0) = \mathbf{W}_j(0, 0) \mathbf{W}_j(0, 0)^T, \forall j \in 1, \dots, N-2$. Then the weights remain balanced throughout, i.e.*

$$\forall k, t \quad \mathbf{W}_{j+1}(k, t)^T \mathbf{W}_{j+1}(k, t) = \mathbf{W}_j(k, t) \mathbf{W}_j(k, t)^T, \forall j \in 1, \dots, N-2. \quad (\text{D.4})$$

Proof. First consider the special case of $k = 0$. To begin with, in order to show that weights remain balanced throughout training in $k = 0$ iteration, we will follow the proof technique in [Arora et al. \(2018\)](#), with some modifications. First note that the expression for $\frac{\partial L_{k+1}(\mathbf{W}_{N:1}(k, t))}{\partial \mathbf{W}_j(k, t)}$ can be expressed as:

$$\begin{aligned} \frac{\partial L_{k+1}(\mathbf{W}_{N:1}(k, t))}{\partial \mathbf{W}_j(k, t)} &= \left(\prod_{i=j+1}^N \mathbf{W}_i^T \right) \cdot \frac{dL_k(\mathbf{W}_{N:1})}{d\mathbf{W}_{N:1}} \cdot \left(\prod_{i=1}^{j-1} \mathbf{W}_i^T \right) \\ &= (\mathbf{W}_{j+1}^T \mathbf{W}_{j+2}^T \cdots \mathbf{W}_N^T) \cdot \frac{dL_k(\mathbf{W}_{N:1})}{d\mathbf{W}_{N:1}} \cdot \left(\prod_{i=1}^{j-1} \mathbf{W}_i^T \right). \end{aligned}$$

Now, since the weight matrices evolve as per Equation D.2, by multiplying the similar differential equation for \mathbf{W}_j with $\mathbf{W}_j^T(k, t)$ on the right and multiplying evolution of \mathbf{W}_{j+1} with $\mathbf{W}_{j+1}^T(k, t)$ from the left, and adding the two equations, we obtain:

$$\forall j \in [0, \dots, N-2] : \mathbf{W}_{j+1}^T(0, t) \dot{\mathbf{W}}_{j+1}(0, t) = \dot{\mathbf{W}}_j(0, t) \mathbf{W}_j^T(0, t). \quad (\text{D.5})$$

We can then take transpose of the equation above, and add it to itself, to obtain an easily integrable expression:

$$\begin{aligned} \frac{d \mathbf{W}_{j+1}(0, t) \mathbf{W}_{j+1}(0, t)^T}{dt} &= \mathbf{W}_{j+1}^T(0, t) \dot{\mathbf{W}}_{j+1}(0, t) + \dot{\mathbf{W}}_{j+1}(0, t) \mathbf{W}_{j+1}^T(0, t) = \\ &\dot{\mathbf{W}}_j(0, t) \mathbf{W}_j^T(0, t) + \mathbf{W}_j(0, t) \dot{\mathbf{W}}_j^T(0, t) = \frac{d \mathbf{W}_j^T(0, t) \mathbf{W}_j(0, t)}{dt}. \end{aligned} \quad (\text{D.6})$$

Since we have assumed the balancedness condition at the initial timestep 0, and the derivatives of the two quantities are equal, their integral will also be the same, hence we obtain:

$$\mathbf{W}_{j+1}^T(0, t) \mathbf{W}_{j+1}(0, t) = \mathbf{W}_j(0, t) \mathbf{W}_j(0, t)^T. \quad (\text{D.7})$$

Now, since the weights after T iterations in fitting iteration $k = 0$ are still balanced, the initialization for $k = 1$ is balanced. Note that since the balancedness property does not depend on which objective gradient is used to optimize the weights, as long as \mathbf{W}_j and \mathbf{W}_{j+1} utilize the same gradient of the loss function. Formalizing this, we can show inductively that the weights will remain balanced across all fitting iterations k and at all steps t within each fitting iteration. Thus, we have shown the result in Equation D.4. \square

Our next result aims at deriving the evolution of the feature-matrix that under the balancedness condition. We will show that the feature matrix, $\mathbf{W}_\phi(k, t)$ evolves according to a similar, but distinct differential equation as the end-to-end weight matrix, $\mathbf{W}_{N:1}(k, t)$, which still allows us to appeal to techniques and results from Arora et al. (2019) to study properties of the singular value evolution and hence, discuss properties related to the effective rank of the matrix, $\mathbf{W}_\phi(k, t)$.

Lemma D.0.2 ((Extension of Theorem 1 from Arora et al. (2018)). *Under conditions specified in Lemma D.0.1, the feature matrix, $\mathbf{W}_\phi(k, t)$ evolves as per the following continuous-time differential equation, for all fitting iterations k :*

$$\dot{\mathbf{W}}_\phi(k, t) = -\eta \sum_{j=1}^{N-1} [\mathbf{W}_\phi(k, t) \mathbf{W}_\phi(k, t)^T]^{\frac{N-j}{N-1}} \cdot \mathbf{W}_N(k, t)^T \frac{dL_k(\mathbf{W}_{N:1}(k, t))}{d\mathbf{W}_{N:1}} \cdot [\mathbf{W}_\phi(k, t)^T \mathbf{W}_\phi(k, t)]^{\frac{j-1}{N-1}}.$$

Proof. In order to prove this statement, we utilize the fact that the weights upto layer $N - 2$ are balanced throughout training. Now consider the singular value decomposition of any weight \mathbf{w}_j (unless otherwise states, we use \mathbf{W}_j to refer to $\mathbf{W}_j(k, t)$ in this section, for ease of notation. $\mathbf{W}_j = \mathbf{U}_j \Sigma_j \mathbf{V}_j^T$). The balancedness condition re-written using SVD of the weight matrices is equivalent to

$$\mathbf{V}_{j+1} \Sigma_{j+1}^T \Sigma_{j+1} \mathbf{V}_{j+1}^T = \mathbf{U}_j \Sigma_j \Sigma_j^T \mathbf{U}_j^T. \quad (\text{D.8})$$

Thus for all j on which the balancedness condition is valid, it must hold that $\Sigma_{j+1}^T \Sigma_{j+1} = \Sigma_j \Sigma_j^T$, since these are both the eigendecompositions of the same matrix (as they are equal). As a result, the weight matrices \mathbf{W}_j and \mathbf{W}_{j+1} share the same singular value space which can be written as $\rho_1 \geq \rho_2 \geq \dots \geq \rho_m$. The ordering of eigenvalues can be different, and the matrices \mathbf{U} and \mathbf{V} can also be different (and be rotations of one other) but the unique values that the singular values would take are the same. Note the distinction from Arora et al. (2018), where they apply balancedness on all matrices, and that in our case would trivially give a rank-1 matrix.

Now this implies, that we can express the feature matrix, also in terms of the common singular values, $(\rho_1, \rho_2, \dots, \rho_m)$, for example, as $\mathbf{W}_j(k, t) = \mathbf{U}_{j+1} \text{Diag}(\sqrt{\rho_1}, \dots, \sqrt{\rho_m}) \mathbf{V}_j^T$, where $\mathbf{U} \mathbf{U}^T = \mathbf{V} \mathbf{V}^T = \mathbf{O}_j$, where \mathbf{O}_j is an orthonormal matrix. Using this relationship, we can say the following:

$$\begin{aligned} \prod_{i=j}^{N-1} \mathbf{w}_i(k, t) \prod_{i=j}^{N-1} \mathbf{w}_i(k, t)^T &= [\mathbf{W}_\phi(k, t) \mathbf{W}_\phi(k, t)^T]^{\frac{N-j}{N-1}} \\ \prod_{i=1}^j \mathbf{w}_i(k, t)^T \prod_{i=1}^j \mathbf{w}_i(k, t) &= [\mathbf{W}_\phi(k, t)^T \mathbf{W}_\phi(k, t)]^{\frac{j}{N-1}}. \end{aligned}$$

Now, we can use these expressions to obtain the desired result, by taking the differential equations governing the evolution of $\mathbf{W}_i(k, t)$, for $i \in [1, \dots, N-1]$, multiplying the i -th equation by $\prod_{j=1}^{N-1} \mathbf{W}_j(k, t)$ from the left, and $\prod_{j=1}^{i-1} \mathbf{W}_j(k, t)$ to the right, and then summing over i .

$$\begin{aligned} \dot{\mathbf{W}}_\phi(k, t) &= \sum_{i=1}^{N-1} \left(\prod_{j=i+1}^{N-1} \mathbf{W}_j(k, t) \right) \dot{\mathbf{W}}_i(k, t) \left(\prod_{j=1}^{i-1} \mathbf{W}_j(k, t) \right) \\ &= -\eta \sum_{i=1}^{N-1} \left(\prod_{j=i+1}^{N-1} \mathbf{W}_j(k, t) \prod_{j=1}^N \mathbf{W}_j(k, t)^T \right) \frac{dL_k(\mathbf{W}_{N:1})}{d\mathbf{W}_{N:1}} \left(\prod_{j=1}^{i-1} \mathbf{W}_j(k, t)^T \prod_{j=1}^{i-1} \mathbf{W}_j(k, t) \right) \end{aligned}$$

The above equation simplifies to the desired result by taking out $\mathbf{W}_N(k, t)$ from the first summation, and using the identities above for each of the terms. \square

Comparing the previous result with Theorem 1 in [Arora et al. \(2018\)](#), we note that the resulting differential equation for weights holds true for arbitrary representations or features in the network provided that the layers from the input to the feature layer are balanced. A direct application of [Arora et al. \(2018\)](#) restricts the model class to only fully balanced networks for convergence analysis and the resulting solutions to the feature matrix will then only have one active singular value, leading to less-expressive neural network configurations.

Proof of Proposition 4.1. Finally, we are ready to use Lemma D.0.2 to prove the relationship with evolution on singular values. This proof can be shown via a direct application of Theorem 3 in [Arora et al. \(2019\)](#). Given that the feature matrix, $\mathbf{W}_\phi(k, t)$ satisfies a very similar differential equation as the end-to-end matrix, with the exception that the gradient of the loss with respect to the end-to-end matrix is pre-multiplied by $\mathbf{W}_N(k, t)^T$. As a result, we can directly invoke [Arora et al. \(2019\)](#)'s result and hence, we have $\forall r \in [1, \dots, \dim(W)]$ that:

$$\dot{\sigma}_r(k, t) = -N \cdot (\sigma_r^2(k, t))^{1 - \frac{1}{N-1}} \cdot \left\langle \mathbf{W}_N(k, t)^T \frac{dL_{N,k}(\mathbf{W}_{K,t})}{d\mathbf{W}}, \mathbf{u}_r(k, t) \mathbf{v}_r(k, t)^T \right\rangle. \quad (\text{D.9})$$

Further, as another consequence of the result describing the evolution of weight matrices, we can also obtain a result similar to [Arora et al. \(2019\)](#) that suggests that the goal of the gradient update on the singular vectors $\mathbf{U}(k, t)$ and $\mathbf{V}(k, t)$ of the features $\mathbf{W}_\phi(k, t) = \mathbf{U}(k, t) \mathbf{S}(k, t) \mathbf{V}(K, t)^T$, is to align these spaces with $\mathbf{W}_N(k, t)^T \frac{dL_{N,k}(\mathbf{W}_{K,t})}{d\mathbf{W}}$.

D.1 EXPLAINING RANK DECREASE BASED ON SINGULAR VALUE EVOLUTION




Figure D.1: **Evolution of singular values of Φ on Atari.** The larger singular values of the feature matrix Φ grow at a disproportionately higher rate than other smaller singular values as described by equation 4.

In this section, we discuss why the evolution of singular values discussed in Equation 4 indicates a decrease in the rank of the feature matrix *within* a fitting iteration k . To see this, let's consider the case when gradient descent has been run long enough (*i.e.*, the data-efficient RL case) for the singular vectors to stabilize, and consider the evolution of singular values post timestep $t \geq t_0$ in training. First of all, when the singular vectors stabilize, we obtain that $\mathbf{u}_r(k, t)^T \mathbf{W}_N(k, t)^T \frac{dL_{N,k}(\mathbf{W}_{K,t})}{d\mathbf{W}} \mathbf{v}_r(k, t)$ is diagonal (extending result of Corollary 1 from [Arora et al. \(2019\)](#)). Thus, we can assume that

$$\mathbf{u}_r^T(k, t) \mathbf{W}_N(k, t)^T \frac{dL_{N,k}(\mathbf{W}_{K,t})}{d\mathbf{W}} \mathbf{v}_r(k, t) = \mathbf{f}(k, t) \cdot \mathbf{e}_r \cdot \mathbf{d}_r,$$

where \mathbf{e}_r is given by the unit basis vector for the standard orthonormal basis, \mathbf{f} is a shorthand for the gradient norm of the loss function pre-multiplied by the transpose of the last layer weight matrix,

and \mathbf{d}_r denotes the singular values of the state-action input matrix, $[\mathcal{S}, \mathcal{A}]$. In this case, we can re-write Equation 4 as:

$$\dot{\sigma}_r(k, t) = -N (\sigma_r^2(k, t))^{1-1/N} \cdot \mathbf{f}(k, t) \cdot \mathbf{e}_r \cdot \mathbf{d}_r. \quad (\text{D.10})$$

Note again that unlike Arora et al. (2019), the expression for $\mathbf{f}(k, t)$ is different from the gradient of the loss, since the weight matrix $\mathbf{W}_N(k, t)$ is multiplied to the expression of the gradient in our case. By using the fact that the expression for $\mathbf{f}(k, t)$ is shared across all singular values, we can obtain differential equations that are equal across different $\sigma_r(k, t)$ and $\sigma_{r'}(k, t)$. Integrating, we can show that depending upon the ratio $\frac{\mathbf{e}_r \cdot \mathbf{d}_r}{\mathbf{e}_{r'} \cdot \mathbf{d}_{r'}}$, and the value of N , the singular values $\sigma_r(k, t)$ and $\sigma_{r'}(k, t)$ will take on different magnitude values, in particular, they will grow at disproportionate rates. By then appealing to the result from Proposition 4.1 in Section 4.1 (kernel regression argument), we can say that disproportionately growing singular values would imply that the value of $\text{srnk}_\delta(\mathbf{W}_\phi(k, t))$ decreases within each iteration k as we will discuss next.

Interpretation of the abstract rank-regularized objective. We now discuss the form for the abstract rank-regularized objective in Equation 5 that we utilize in Section 4.2. Intuitively, the justification for Equation 5 is that larger singular values grow much faster than small singular values, which dictates how the ratio of singular values evolves through time, thereby reducing effective rank.

The discussion in the previous subsection shows that the effective rank, $\text{srnk}_\delta(\mathbf{W}_\phi(k, t))$ decreases within each fitting iteration (*i.e.*, decreases over the variable t). Now, we will argue that for a given value of effective rank at $t = 0$ at a fitting iteration k (denoted $\text{srnk}_\delta(\mathbf{W}_\phi(k, 0))$), the effective rank after T steps of gradient descent, $\text{srnk}_\delta(\mathbf{W}_\phi(k, T))$, is constrained to be equal to some constant $\varepsilon_k < \text{srnk}_\delta(\mathbf{W}_\phi(k, 0))$ due to the implicit regularization effect of gradient descent. To see this, note that if $\sigma_i(k, 0) > \sigma_j(k, 0)$, then $\sigma_i(k, t)$ increases at a much faster rate than $\sigma_j(k, t)$ with increasing t (see Equation D.10 and Arora et al. (2019), Equation 11) and hence, the value of $h_{k,t}(i)$ is increasing over t :

$$h_{k,t}(i) = \frac{\sum_{j=1}^i \sigma_j(\mathbf{W}_\phi(k, t))}{\sum_{l=1}^i \sigma_l(\mathbf{W}_\phi(k, t))} = \frac{1}{1 + \frac{\sum_{j=i+1}^n \sigma_j(\mathbf{W}_\phi(k, t))}{\sum_{j=1}^i \sigma_j(\mathbf{W}_\phi(k, t))}} \geq \frac{1}{1 + \sum_{j=i+1}^n \frac{\sigma_j(\mathbf{W}_\phi(k, t))}{\sigma_1(\mathbf{W}_\phi(k, t))}}, \quad (\text{D.11})$$

as the ratio between larger and smaller singular values increases. Using Equation D.11, we note that the value of $h_{k,t}(i)$ approaches 1 as the ratio σ_j/σ_1 decreases (in the limit, the ratio is zero). An increase in the value of $h_k(i)$ implies a decrease in srnk_δ , which can be expressed as:

$$\text{srnk}_\delta(\mathbf{W}_\phi(k, t)) = \min_j \{ h_k(j) \geq 1 - \delta \}.$$

Thus, by running gradient descent for $T_0 = T(\varepsilon_k)$ (a function of ε_k) steps within a fitting iteration, k , the effective rank of the resulting $\mathbf{W}_\phi(k, T_0)$ satisfies $\text{srnk}_\delta(\mathbf{W}_\phi(k, T_0)) = \varepsilon_k$. Now, since the solution $\mathbf{W}_\phi(k, T_0)$ satisfies the constraint $\text{srnk}_\delta(\mathbf{W}_\phi(k, T_0)) = \varepsilon_k < \text{srnk}_\delta(\mathbf{W}_\phi(k, 0))$, we express this solution using the solution of a penalized optimization problem (akin to penalty methods) that penalizes $\text{srnk}_\delta(\mathbf{W}_\phi)$ for a suitable coefficient λ_k , which is a function of ε_k . We assume $\lambda_k > 0$ since when running gradient descent long enough (*i.e.*, for T_0 steps in the above discussion), we can reduce $\text{srnk}_\delta(\mathbf{W}_\phi)$ from its initial value (note that $\varepsilon_k < \text{srnk}_\delta(\mathbf{W}_\phi(k, 0))$), indicating a non-zero value of regularization.

D.2 PROOF OF PROPOSITION 4.2: COMPOUNDING RANK DROP IN SECTION 4.2

In this section, we illustrate that the rank drop effect compounds due to bootstrapping and prove Proposition 4.2 formally. Our goal is to demonstrate the compounding effect: a change in the rank at a given iteration gets propagated into the next iteration. We first illustrate this compounding effect and then discuss Theorem 4.2 as a special case of this phenomenon.

Assumptions. We require two assumptions pertaining to closure of the function class and the change in effective rank due to reward and dynamics transformations, to be able for this analysis. We assume that the following hold for any fitting iteration k of the algorithm.

Assumption D.1 (Closure). *The chosen function class $\mathbf{W}_1 \mathbf{W}_2 \cdots \mathbf{W}_N$ is closed under the Bellman evaluation backup for policy π . That is, if $\mathbf{Q}_{k-1}^T = \mathbf{W}_N(k-1) \mathbf{W}_\phi(k-1) [\mathcal{S}; \mathcal{A}]^T$, then there exists an assignment of weights to the deep linear net such that the corresponding target value $\mathbf{R} + \gamma P^\pi \mathbf{Q}_k$ can be written as $(\mathbf{R} + \gamma P^\pi \mathbf{Q}_{k-1})^T = \mathbf{W}_N^P(k) \mathbf{W}_\phi^P(k) [\mathcal{S}; \mathcal{A}]^T$. This assumption is commonly used in analyses of approximate dynamic programming with function approximation.*

Assumption D.2 (Change in srank_δ). For every fitting iteration k , we assume that the difference between $\text{srank}_\delta(\mathbf{W}_\phi^P(k))$, equal to the effective rank to feature matrix obtained when $\mathbf{R} + \gamma P^\pi \mathbf{Q}_{k-1}$ is expressed in the function class, and $\text{srank}_\delta(\mathbf{W}_\phi(k-1))$ by an iteration specific threshold c_k , i.e.,

$$\text{srank}_\delta(\mathbf{W}_\phi^P(k)) \leq \text{srank}_\delta(\mathbf{W}_\phi(k-1)) + c_k.$$

We will characterize the properties of c_k in special cases by showing how Theorem 4.2 is a special case of our next argument, where we can analyze the trend in c_k .

To first see how a change in rank indicated by c_k in Assumption D.2 propagates through bootstrapping, note that we can use Assumption D.2 and D.1 to observe the following relation:

$$\begin{aligned} \text{srank}_\delta(\mathbf{W}_\phi(k)) &\leq \text{srank}_\delta(\mathbf{W}_\phi^P(k)) - \frac{\|\mathbf{Q}_k - \mathbf{y}_k\|}{\lambda_k} \quad (\text{since we can copy over the weights}) \\ &\leq \text{srank}_\delta(\mathbf{W}_\phi(k-1)) + c_k - \frac{\|\mathbf{Q}_k - \mathbf{y}_k\|}{\lambda_k} \quad (\text{using Assumption D.2}) \\ &\leq \text{srank}_\delta(\mathbf{W}_\phi(k-2)) + c_{k-1} - \frac{\|\mathbf{Q}_{k-1} - \mathbf{y}_{k-1}\|}{\lambda_{k-1}} + c_k - \frac{\|\mathbf{Q}_k - \mathbf{y}_k\|}{\lambda_k} \\ \text{srank}_\delta(\mathbf{W}_\phi(k)) &\leq \boxed{\text{srank}_\delta(\mathbf{W}_\phi(0)) + \sum_{j=1}^k c_j - \sum_{j=1}^k \frac{\|\mathbf{Q}_j - \mathbf{y}_j\|}{\lambda_j}} \end{aligned} \quad (\text{D.12})$$

Derivation of the steps: The first inequality holds via the weight-copying argument: since the weights for the target value $\mathbf{R} + \gamma P^\pi \mathbf{Q}_{k-1}$ can be directly copied to obtain a zero TD error feasible point with an equal srank , we will obtain a better solution at the optimum with a smaller srank in Equation 5. We then use Assumption D.2 to relate $\text{srank}_\delta(\mathbf{W}_\phi^P(k))$ to the effective rank of the previous Q-function, $\text{srank}_\delta(\mathbf{W}_\phi(k-1))$, which gives us a recursive inequality. Finally, the last two steps follow from a repeated application of the recursive inequality in the second step. This proves the result shown in Proposition 4.2.

Now if the decrease in effective rank due to accumulating TD errors is larger than the possible cumulative increase in rank c_k , then we observe a rank drop. Also note that c_k need not be positive, c_k can be negative, in which case both terms contribute to a drop in rank. But in the most general case, c_k can also be positive for some iterations. This equation indicates how a change in rank in one iteration gets compounded on the next iteration due to bootstrapping.

What can we comment about the general value of c_k ? In the general case, c_k can be positive, for example when the addition of the reward and dynamics projection increases srank . Our current set of assumptions are insufficient to guarantee something about c_k in the general case and more assumptions are needed. For instance, as shown in Theorem 4.2, under Assumption D.1, we obtain that if we are in a small neighborhood around the fixed point, we will find that $c_k \leq 0$.

Empirical verification: While our analysis does not dictate the value of c_k , we performed an experiment on Seaquest to approximately visualize $\text{srank}_\delta(\mathbf{W}_\phi^P(k))$ by computing the effective rank for the feature matrix Φ obtained when $P^\pi \mathbf{Q}_{k-1}$ is expressed in the Q-function class. We do not include the reward function in this experiment, but we believe that a “simple” reward function, i.e., it takes only three possible values in Atari, and hence it may not contribute too much to $\text{srank}_\delta(\mathbf{W}_\phi^P(k))$. As shown in Figure D.2, the value of the target value feature effective rank decreases, indicating that the coefficient c_k is expected to be bounded in practice.




Figure D.2: Two pairs of runs depicting the trends of target feature (approximately equal to $\text{srank}_\delta(\mathbf{W}_\phi^P(k))$) (denoted as “Target”) and $\text{srank}_\delta(\mathbf{W}_\phi(k-1))$ (denoted as “Main”) on Seaquest with in the data-efficient online RL setting. Note that the contribution of dynamics transformation in the value of c_k is generally negative for the Pair 1 and is small/positive for Pair 2.

D.3 PROOF FOR THEOREM 4.2: RANK COLLAPSE NEAR A FIXED POINT

Now, we will prove Theorem 4.2 by showing that when the current Q-function iterate \mathbf{Q}_k is close to a fixed point but have not converged yet, i.e., when $\|\mathbf{Q}_k - (\mathbf{R} + \gamma P^\pi \mathbf{Q}_k)\| \leq \varepsilon$, then rank-decrease

happens. We will prove this as a special case of the discussion in Section D.2, by showing that for any iteration k , when the current Q-function iterate \mathbf{Q}_k is close to a fixed point, the value of c_k can be made close to 0. While the compounding argument (Equation D.12) may not directly hold here since we cannot consistently comment on whether making an update on the Q-function pushes it out of the ε -ball near the fixed point, so we instead prove a “one-step” result here.

To prove Theorem 4.2, we evaluate the (infinitesimal) change in the singular values of the features of $\mathbf{W}_\phi(k, t)$ as a result of an addition in the value of ε . In this case, the change (or differential) in the singular value matrix, $\mathbf{S}(k, T)$, ($\mathbf{W}_\phi(k, T) = \mathbf{U}(k, T)\mathbf{S}(k, T)\mathbf{V}(k, T)^T$) is given by:

$$d\mathbf{S}(k, T) = \mathbf{I}_d \circ \left[\mathbf{U}(k, T)^T \cdot d\mathbf{W}_\phi(k, t) \cdot \mathbf{V}(k, T) \right], \quad (\text{D.13})$$

using results on computing the derivative of the singular value decomposition (Townsend, 2016).

Proof strategy: Our high-level strategy in this proof is to design a form for ε such that the value of effective rank for feature matrix obtained when the resulting target value $\mathbf{y}_k = \mathbf{Q}_{k-1} + \varepsilon$ is expressed in the deep linear network function class, let’s say with feature matrix, $\mathbf{W}'_\phi(k, T)$ and the last layer weights, $\mathbf{W}_N(k, T) = \mathbf{W}_N(k-1, T)$, then the value of $\text{srnk}_\delta(\mathbf{W}'_\phi(k, T))$ can be larger than $\text{srnk}_\delta(\mathbf{W}_\phi(k-1, T))$ by only a bounded limited amount α that depends on ε . More formally, we desire a result of the form

$$\begin{aligned} \text{srnk}_\delta(\mathbf{W}'_\phi(k, T)) &\leq \text{srnk}_\delta(\mathbf{W}_\phi(k-1, T)) + \alpha, \\ \text{where } \|\varepsilon\| &\ll \|\mathbf{Q}_{k-1}\|, \text{ and } \mathbf{y}_k(\mathbf{s}, \mathbf{a}) = \mathbf{W}_N(k-1, T)\mathbf{W}'_\phi(k, T)[\mathbf{s}; \mathbf{a}]. \end{aligned} \quad (\text{D.14})$$

Once we obtain such a parameterization of \mathbf{y}_k in the linear network function class, using the argument in Section 4.2 about self-training, we can say that just copying over the weights $\mathbf{W}'_\phi(k, T)$ is sufficient to obtain a bounded increase in $\text{srnk}_\delta(\mathbf{W}_\phi(k, T))$ at the cost of incurring zero TD error (since the targets can be exactly fit). As a result, the optimal solution found by Equation 5 will attain lower $\text{srnk}_\delta(\mathbf{W}_\phi(k, T))$ value if the TD error is non-zero. Specifically,

$$\text{srnk}_\delta(\mathbf{W}'_\phi(k, T)) + \frac{\|\mathbf{Q}_k - \mathbf{y}_k\|}{\lambda_k} \leq \text{srnk}_\delta(\mathbf{W}_\phi(k-1, T)) + \alpha \quad (\text{D.15})$$

As a result we need to examine the factors that affect $\text{srnk}_\delta(\mathbf{W}_\phi(k, T)) - (1)$ an increase in $\text{srnk}_\delta(\mathbf{W}_\phi(k, T))$ due to an addition of α to the rank, and (2) a decrease in $\text{srnk}_\delta(\mathbf{W}_\phi(k, T))$ due to the implicit behavior of gradient descent.

Proof: In order to obtain the desired result (Equation D.14), we can express $\varepsilon(\mathbf{s}, \mathbf{a})$ as a function of the last layer weight matrix of the *current* Q-function, $\mathbf{W}_N(k-1, T)$ and the state-action inputs, $[\mathbf{s}; \mathbf{a}]$. Formally, we assume that $\varepsilon(\mathbf{s}, \mathbf{a}) = \mathbf{W}_N(k-1, T)\zeta[\mathbf{s}; \mathbf{a}]$, where ζ is a matrix of the same size as \mathbf{W}_ϕ such that $\|\zeta\|_\infty \ll \tau = \mathcal{O}(\|\mathbf{W}_\phi(k, T)\|_\infty)$, i.e., ζ has entries with ignorable magnitude compared the actual feature matrix, $\mathbf{W}_\phi(k, T)$. This can be computed out of the closure assumption (Assumption D.1) from Section D.2.

Using this form of $\varepsilon(\mathbf{s}, \mathbf{a})$, we note that the targets \mathbf{y}_k used for the backup can be written as

$$\begin{aligned} \mathbf{y}_k &= \mathbf{Q}_{k-1} + \varepsilon = \mathbf{W}_N(k-1, T)\mathbf{W}_\phi(k-1, T)[\mathbf{S}; \mathcal{A}] + \mathbf{W}_N(k-1, T)\zeta[\mathbf{S}; \mathcal{A}] \\ &= \mathbf{W}_N(k-1, T) \cdot \underbrace{(\mathbf{W}_\phi(k-1, T) + \zeta)}_{\mathbf{W}'_\phi(k, T)}[\mathbf{S}; \mathcal{A}] \end{aligned}$$

Using the equation for sensitivity of singular values due to a change in matrix entries, we obtain the maximum change in the singular values of the resulting “effective” feature matrix of the targets \mathbf{y}_k in the linear function class, denoted as $\mathbf{W}'_\phi(k, T)$, is bounded:

$$d\mathbf{S}'(k, T) = \mathbf{I}_d \circ \left[\mathbf{U}(k, T) \cdot d\mathbf{W}'_\phi(k, T) \cdot \mathbf{V}(k, T)^T \right] \implies \|d\mathbf{S}'(k, T)\|_\infty \leq \zeta. \quad (\text{D.16})$$

Now, let's use this inequality to find out the maximum change in the function, $h_k(i)$ used to compute $\text{srnk}_\delta(\mathbf{W}_\phi)$: $h_k(i) = \frac{\sum_{j=1}^i \sigma_j(\mathbf{W}_\phi(k, T))}{\sum_{j=1}^N \sigma_j(\mathbf{W}_\phi(k, T))}$ ($\text{srnk}_\delta(\mathbf{W}_\phi(k, T)) = \min \{j : h_k(j) \geq 1 - \delta\}$).

$$\begin{aligned} |dh_k(i)| &= \left| \frac{\sum_{j=1}^i d\sigma_j(\mathbf{W}_\phi)}{\sum_{j=1}^N \sigma_j(\mathbf{W}_\phi)} - \left(\frac{\sum_{j=1}^i \sigma_j(\mathbf{W}_\phi)}{\sum_{j=1}^N \sigma_j(\mathbf{W}_\phi)} \right) \left(\frac{\sum_{j=1}^N d\sigma_j(\mathbf{W}_\phi)}{\sum_{j=1}^N \sigma_j(\mathbf{W}_\phi)} \right) \right| \\ &\leq \frac{i\zeta}{\sum_{j=1}^N \sigma_j(\mathbf{W}_\phi)} + \underbrace{\left(\frac{\sum_{j=1}^i \sigma_j(\mathbf{W}_\phi)}{\sum_{j=1}^N \sigma_j(\mathbf{W}_\phi)} \right)}_{\leq 1} \frac{N\zeta}{\sum_{j=1}^N \sigma_j(\mathbf{W}_\phi)} \\ &\leq \frac{(i+N)\zeta}{\sum_{j=1}^N \sigma_j(\mathbf{W}_\phi)}. \end{aligned}$$

The above equation implies that the maximal change in the effective rank of the feature matrix generated by the targets, \mathbf{y}_k (denoted as $\mathbf{W}'_\phi(k, T)$) and the effective rank of the features of the current Q-function \mathbf{Q}_k (denoted as $\mathbf{W}_\phi(k, T)$) are given by:

$$\text{srnk}_\delta(\mathbf{W}'_\phi(k, T)) - \text{srnk}_\delta(\mathbf{W}_\phi(k, T)) \leq \alpha$$

where α can be formally written by the cardinality of the set:

$$\alpha = \left| \left\{ j : h_k(j) - \frac{(j+N)\zeta}{\sum_{j=1}^N \sigma_j(\mathbf{W}_\phi)} \geq (1-\delta), \quad h_k(j) \leq (1-\delta) \right\} \right|. \quad (\text{D.17})$$

Note that by choosing $\varepsilon(\mathbf{s}, \mathbf{a})$ and thus ζ to be small enough, we can obtain $\mathbf{W}'_\phi(k, T)$ such that $\alpha = 0$. Now, the self-training argument discussed above applies and gradient descent in the next iteration will give us solutions that reduce rank.

Assuming $r > 1$ and $\text{srnk}_\delta(\mathbf{W}_\phi(k, T)) = r$, we know that $h_k(r-1) < 1-\delta$, while $h_k(r) \geq 1-\delta$. Thus, $\text{srnk}_\delta(\mathbf{W}'_\phi(k, T))$ to be equal to r , it is sufficient to show that $h_k(r) - |dh_k(r)| \geq 1-\delta$ since both $h_k(i)$ and $dh_k(i)$ are increasing for $i = 0, 1, \dots, r$. Thus, $\text{srnk}_\delta(\mathbf{W}'_\phi(k, T)) = r \forall \zeta$, whenever

$$\begin{aligned} \zeta &\leq \frac{\sum_{j=1}^N \sigma_j(\mathbf{W}_\phi)}{r+N} (h_k(r) - 1 - \delta) \\ &= \frac{\sum_{j=1}^r \sigma_j(\mathbf{W}_\phi) - (1-\delta) \sum_{j=1}^N \sigma_j(\mathbf{W}_\phi)}{r+N} \end{aligned}$$

This implies that $\varepsilon \leq \|W_N(k, T)\|_\infty \frac{\sum_{j=1}^r \sigma_j(\mathbf{W}_\phi) - (1-\delta) \sum_{j=1}^N \sigma_j(\mathbf{W}_\phi)}{r+N}$.

Proof summary: We have thus shown that there exists a neighborhood around the optimal fixed point of the Bellman equation, parameterized by $\varepsilon(\mathbf{s}, \mathbf{a})$ where bootstrapping behaves like self-training. In this case, it is possible to reduce srnk while the TD error is non-zero. And of course, this would give rise to a rank reduction close to the optimal solution.