

445 **A Related Works**

446 **Value Divergence with Neural Network.** In online reinforcement learning (RL), off-policy algo-
 447 rithms that employ value function approximation and bootstrapping can experience value divergence,
 448 a phenomenon known as the deadly triad [35, 4, 36, 37, 11]. Deep Q-Networks (DQN) typify this
 449 issue. As they employ neural networks for function approximation, they are particularly susceptible
 450 to Q-value divergence [18, 13, 37]. Past research has sought to empirically address this divergence
 451 problem through various methods, such as the use of separate target networks [29] and Double-Q
 452 learning [18, 13]. Achiam *et al.* [1] analyze a linear approximation of Q-network to characterizes
 453 the diverge, while CrossNorm [6] uses a elaborated version of BatchNorm [19] to achieve stable
 454 learning. Value divergence becomes even more salient in offline RL, where the algorithm learns
 455 purely from a fixed dataset without additional environment interaction [14, 24]. Much of the fo-
 456 cus in the field of offline RL has been on controlling the extent of off-policy learning, *i.e.*, policy
 457 constraint [12, 30, 23, 14, 40, 38, 8, 32]. Several previous studies [31, 5] have empirically utilized
 458 LayerNorm to enhance performance in online and offline RL. These empirical results partially align
 459 with the experimental section of our work. However, our study makes a theoretical explanation for
 460 how LayerNorm mitigates divergence through the NTK analysis. Specifically, we empirically and
 461 theoretically illustrate how LayerNorm reduces SEEM. In addition to LayerNorm, our contribution
 462 extends to explaining divergence and proposing promising solutions from the perspective of reducing
 463 SEEM. Specially, we discover that WeightNorm can also be an effective tool and explain why other
 464 regularization techniques fall short. Finally, we perform comprehensive experiments to empirically
 465 verify the effectiveness of LayerNorm on the %X dataset, a practical setting not explored in previous
 466 work. Thus, our contributions are multifaceted and extend beyond the mere application of LayerNorm.

467 **Offline RL.** Offline RL presents significant challenges due to severe off-policy issues and extrapola-
 468 tion errors. Some existing methods focuses on designs explicit or implicit policy regularizations to
 469 minimize the discrepancy between the learned and behavior policies. For example, TD3+BC [13, 12]
 470 directly adds a behavior cloning loss to mimic the behavior policy, Diffusion-QL [38] further replace
 471 the BC loss with a diffusion loss and using diffusion models as the policy. CRR [39] and AWR [32]
 472 impose an implicit policy regularization by performing policy gradient-style policy updates. Mean-
 473 while, some other works try to alleviate the extrapolation errors by modifying the policy evaluation
 474 procedure. Specifically, CQL [25] penalizes out-of-distribution actions for having higher Q-values,
 475 while IQL [23] and OneStep RL [7] only uses in-distribution data for policy evaluation, thus avoiding
 476 querying unseen actions. Alternatively, decision transformer (DT) [9] and trajectory transformer [21]
 477 cast offline RL as a sequence generation problem, which are beyond the scope of this paper. Despite
 478 the effectiveness of the above methods, they usually neglect the effect of function approximator and
 479 are thus sensitive to hyperparameters for trading off performance and training stability. Exploration
 480 into the function approximation aspect of the deadly triad is lacking in offline RL. Moreover, a
 481 theoretical analysis of divergence in offline RL that does not consider the function approximator
 482 would be inherently incomplete. We instead focus on this orthogonal perspective and provide both
 483 theoretical understanding and empirical solution to the offline RL problem.

484 **B Proof of Main Theorems**

485 Before proving our main theorem, we first state an important lemma.

486 **Lemma 1.** *For any L -layer ReLU-activate MLP and any fixed input \mathbf{x}, \mathbf{x}' . If we scale up every*
 487 *parameter of f_{θ} to λ times, namely $\theta' = \lambda\theta$ where λ is a large number such that the bias term is*
 488 *negligible, then we have following equations hold*

$$\begin{aligned} f_{\theta'}(\mathbf{x}) &= \lambda^L f_{\theta}(\mathbf{x}), \\ \nabla_{\theta} f_{\theta'}(\mathbf{x}) &\approx \lambda^{L-1} \nabla_{\theta} f_{\theta}(\mathbf{x}), \\ \mathbf{G}_{\theta'}(\mathbf{x}, \mathbf{x}') &\approx \lambda^{2(L-1)} \mathbf{G}_{\theta}(\mathbf{x}, \mathbf{x}'). \end{aligned}$$

Proof. Recursively define

$$\begin{aligned} \mathbf{z}_{\ell+1} &= \mathbf{W}_{\ell} \tilde{\mathbf{z}}_{\ell} + \mathbf{b}_{\ell}, \quad \mathbf{z}_0 = \mathbf{x} \\ \tilde{\mathbf{z}}_{\ell} &= \sigma(\mathbf{z}_{\ell}). \end{aligned}$$

489 Then it is easy to see that if we multiply each \mathbf{W}_ℓ and \mathbf{b}_ℓ by λ , denote the new corresponding value
490 to be \mathbf{z}'_ℓ we have

$$\begin{aligned} \mathbf{z}'_1 &= \lambda \mathbf{z}_1 \\ \mathbf{z}'_2 &= \lambda^2 \mathbf{z}_2 \\ &\dots \\ \mathbf{z}'_L &= \lambda^L \mathbf{z}_L \end{aligned}$$

491 Hence we know $f_{\theta'}(\mathbf{x}) = \lambda^L f_\theta(\mathbf{x})$.

Taking gradient backwards, we know that $\left\| \frac{\partial f}{\partial \mathbf{W}_\ell} \right\|$ is proportional to both $\|\tilde{\mathbf{z}}_\ell\|$ and $\left\| \frac{\partial f}{\partial \mathbf{z}_{\ell+1}} \right\|$. Therefore we know

$$\frac{\partial f_{\theta'}}{\partial \mathbf{W}'_\ell} = \tilde{\mathbf{z}}'_\ell \frac{\partial f_{\theta'}}{\partial \mathbf{z}'_{\ell+1}} = \lambda^\ell \tilde{\mathbf{z}}_\ell \cdot \lambda^{L-\ell-1} \frac{\partial f_\theta}{\partial \mathbf{z}_{\ell+1}} = \lambda^{L-1} \frac{\partial f_\theta}{\partial \mathbf{W}_\ell}.$$

492 This suggests that all gradients with respect to the weights become scaled by a factor of λ^{L-1} . The
493 gradients with respect to the biases are proportional to λ^{L-l} . When λ is large enough to render
494 the gradient of the bias term negligible, it follows that $\nabla_{\theta} f_{\theta'}(\mathbf{x}) \approx \lambda^{L-1} \nabla_{\theta} f_\theta(\mathbf{x})$. This equation
495 implies that the gradient updates for the model parameters are dominated by the weights, with
496 negligible contribution from the bias terms. And since NTK is the inner product between gradients,
497 we know $\mathbf{G}_{\theta'}(\mathbf{x}, \mathbf{x}') \approx \lambda^{2(L-1)} \mathbf{G}_\theta(\mathbf{x}, \mathbf{x}')$. \square

498 **Theorem 1** Suppose that the network's parameter at iteration t is θ_t . For each transition
499 (s_i, a_i, s_{i+1}, r_i) in dataset, denote $\mathbf{r} = [r_1, \dots, r_M]^\top \in \mathbb{R}^M$, $\hat{\pi}_{\theta_t}(s) = \arg \max_a \hat{Q}_{\theta_t}(s, a)$. Denote
500 $\mathbf{x}_{i,t}^* = (s_{i+1}, \hat{\pi}_{\theta_t}(s_{i+1}))$. Concatenate all $\mathbf{x}_{i,t}^*$ to be \mathbf{X}_t^* . Denote $\mathbf{u}_t = f_{\theta_t}(\mathbf{X}) - (\mathbf{r} + \gamma \cdot f_{\theta_t}(\mathbf{X}_t^*))$
501 to be TD error vector at iteration t . The learning rate η is infinitesimal. We have the following
502 evolving equation for \mathbf{u}_{t+1}

$$\mathbf{u}_{t+1} = (\mathbf{I} + \eta \mathbf{A}_t) \mathbf{u}_t. \quad (2)$$

503 where $\mathbf{A} = (\gamma \phi_{\theta_t}(\mathbf{X}_t^*) - \phi_{\theta_t}(\mathbf{X}))^\top \phi_{\theta_t}(\mathbf{X}) = \gamma \mathbf{G}_{\theta_t}(\mathbf{X}_t^*, \mathbf{X}) - \mathbf{G}_{\theta_t}(\mathbf{X}, \mathbf{X})$.

504 *Proof.* For the sake of simplicity, denote $\mathbf{Z}_t = \nabla_{\theta} f_{\theta}(\mathbf{X}) \Big|_{\theta_t}$, $\mathbf{Z}_t^* = \nabla_{\theta} f_{\theta}(\mathbf{X}_t^*) \Big|_{\theta_t}$. The Q-value
505 iteration minimizes loss function \mathcal{L} defined by $\mathcal{L}(\theta) = \frac{1}{2} \|f_{\theta}(\mathbf{X}) - (\mathbf{r} + \gamma \cdot f_{\theta_t}(\mathbf{X}_t^*))\|_2^2$. Therefore
506 we have the gradient as

$$\frac{\partial \mathcal{L}(\theta)}{\partial \theta} = \mathbf{Z}_t (f_{\theta}(\mathbf{X}) - (\mathbf{r} + \gamma \cdot f_{\theta_t}(\mathbf{X}_t^*))) = \mathbf{Z}_t \mathbf{u}_t. \quad (3)$$

507 According to gradient descent, we know $\theta_{t+1} = \theta_t - \eta \mathbf{Z}_t \mathbf{u}_t$. Since η is very small, we know θ_{t+1}
508 stays within the neighborhood of θ_t . We can Taylor-expand function $f_{\theta}(\cdot)$ near θ_t as

$$f_{\theta}(\mathbf{X}) \approx \nabla_{\theta}^\top f_{\theta}(\mathbf{X}) \Big|_{\theta_t} (\theta - \theta_t) + f_{\theta_t}(\mathbf{X}) = \mathbf{Z}_t^\top (\theta - \theta_t) + f_{\theta_t}(\mathbf{X}). \quad (4)$$

$$f_{\theta}(\mathbf{X}_t^*) \approx (\mathbf{Z}_t^*)^\top (\theta - \theta_t) + f_{\theta_t}(\mathbf{X}_t^*). \quad (5)$$

509 When η is infinitesimally small, the equation holds. Plug in θ_{t+1} , we know

$$f_{\theta_{t+1}}(\mathbf{X}) - f_{\theta_t}(\mathbf{X}) = -\eta \mathbf{Z}_t^\top \mathbf{Z}_t \mathbf{u}_t = -\eta \cdot \mathbf{G}_{\theta_t}(\mathbf{X}, \mathbf{X}). \quad (6)$$

$$f_{\theta_{t+1}}(\mathbf{X}_t^*) - f_{\theta_t}(\mathbf{X}_t^*) = -\eta (\mathbf{Z}_t^*)^\top \mathbf{Z}_t \mathbf{u}_t = -\eta \cdot \mathbf{G}_{\theta_t}(\mathbf{X}_t^*, \mathbf{X}). \quad (7)$$

510 Since the change of θ is small, we know $\mathbf{X}_{t+1}^* \approx \mathbf{X}_t^*$. So \mathbf{u}_{t+1} boils down to

$$\mathbf{u}_{t+1} = f_{\theta_{t+1}}(\mathbf{X}) - \mathbf{r} - \gamma f_{\theta_{t+1}}(\mathbf{X}_{t+1}^*) \quad (8)$$

$$= f_{\theta_t}(\mathbf{X}) - \eta \cdot \mathbf{G}_{\theta_t}(\mathbf{X}, \mathbf{X}) \mathbf{u}_t - \mathbf{r} - \gamma (f_{\theta_t}(\mathbf{X}_t^*) - \eta \cdot \mathbf{G}_{\theta_t}(\mathbf{X}_t^*, \mathbf{X})) \mathbf{u}_t \quad (9)$$

$$= \underbrace{f_{\theta_t}(\mathbf{X}) - \mathbf{r} - \gamma f_{\theta_t}(\mathbf{X}_t^*)}_{\mathbf{u}_t} + \eta \cdot (\gamma \mathbf{G}_{\theta_t}(\mathbf{X}_t^*, \mathbf{X}) - \mathbf{G}_{\theta_t}(\mathbf{X}, \mathbf{X})) \mathbf{u}_t \quad (10)$$

$$= (\mathbf{I} + \eta \mathbf{A}_t) \mathbf{u}_t. \quad (11)$$

511 where $\mathbf{A} = \gamma \mathbf{G}_{\theta_t}(\mathbf{X}_t^*, \mathbf{X}) - \mathbf{G}_{\theta_t}(\mathbf{X}, \mathbf{X})$. \square

512

513 **Theorem 3** Given iteration $t > t_0$, and $\mathbf{A} = \gamma \mathbf{G}_{\theta_{t_0}}(\mathbf{X}_{t_0}^*, \mathbf{X}) - \mathbf{G}_{\theta_{t_0}}(\mathbf{X}, \mathbf{X})$. The divergence of
514 \mathbf{u}_t is equivalent to whether there exists an eigenvalue λ of \mathbf{A} such that $\text{Re}(\lambda) > 0$. If converge, we
515 have $\mathbf{u}_t = (\mathbf{I} + \eta \mathbf{A})^{t-t_0} \cdot \mathbf{u}_{t_0}$. Otherwise, \mathbf{u}_t becomes parallel to the eigenvector of the largest
516 eigenvalue λ of \mathbf{A} , and its norm diverges to infinity at following order.

$$\|\mathbf{u}_t\|_2 = O\left(\frac{1}{(1 - C' \lambda \eta t)^{L/(2L-2)}}\right). \quad (12)$$

517 for some constant C' to be determined and L is the number of layers of MLP. Specially, when $L = 2$,
518 it reduces to $O\left(\frac{1}{1 - C' \lambda \eta t}\right)$.

519 *Proof.* According to Assumption 2, max action becomes stable after t_0 . It implies $\mathbf{X}_t^* = \mathbf{X}_{t_0}^* := \mathbf{X}^*$.
520 The stability of the NTK direction implies that for some scalar k_t and the specific input \mathbf{X}^*, \mathbf{X} ,
521 we have $\mathbf{G}_{\theta_t}(\mathbf{X}^*, \mathbf{X}) = k_t \mathbf{G}_{\theta_{t_0}}(\mathbf{X}^*, \mathbf{X})$ and $\mathbf{G}_{\theta_t}(\mathbf{X}, \mathbf{X}) = k_t \mathbf{G}_{\theta_{t_0}}(\mathbf{X}, \mathbf{X})$. Further, we have
522 $\mathbf{A}_t = k_t \mathbf{A}$. It equals 1 if the training is convergent, but will float up if the model's predicted Q-value
523 blows up.

524 we know all the eigenvalues of $\mathbf{I} + \eta \mathbf{A}$ have form $1 + \eta \lambda_i$. Considering η is small enough, we
525 have $|1 + \eta \lambda_i|^2 \approx 1 + 2\eta \text{Re}(\lambda)$. Now suppose if there does not exists eigenvalue λ of \mathbf{A} satisfies
526 $\text{Re}(\lambda) > 0$, we have $|1 + \eta \lambda_i| \leq 1$. Therefore, the NTK will become perfectly stable so $k_t = 1$ for
527 $t > t_0$, and we have

$$\mathbf{u}_t = (\mathbf{I} + \eta \mathbf{A}_{t-1}) \mathbf{u}_{t-1} = (\mathbf{I} + \eta \mathbf{A}_{t-1})(\mathbf{I} + \eta \mathbf{A}_{t-2}) \mathbf{u}_{t-2} = \dots = \prod_{s=t_0}^{t-1} (\mathbf{I} + \eta \mathbf{A}_s) \mathbf{u}_{t_0} \quad (13)$$

$$= \prod_{s=t_0}^{t-1} (\mathbf{I} + \eta \mathbf{A}) \mathbf{u}_{t_0} = (\mathbf{I} + \eta \mathbf{A})^{t-t_0} \cdot \mathbf{u}_{t_0}. \quad (14)$$

528 Otherwise, there exists an eigenvalue for \mathbf{A} satisfying $\text{Re}(\lambda) > 0$. Denote the one with the largest
529 real part as λ , and \mathbf{v} to be the corresponding eigenvector. We know matrix $\mathbf{I} + \eta \mathbf{A}$ also has left
530 eigenvector \mathbf{v} , whose eigenvalue is $1 + \eta \lambda$. In this situation, we know after each iteration, $\|\mathbf{u}_{t+1}\|$
531 will become larger than $\|\mathbf{u}_t\|$. Moreover, to achieve larger and larger prediction values, the model's
532 parameter's norm $\|\theta_t\|$ also starts to explode. We know \mathbf{u}_t is homogeneous with respect θ_t for
533 ReLU networks. The output $f_{\theta_t}(\mathbf{X})$ enlarges p^L times when θ_t enlarges p times. When the reward
534 values is small with respect to the divergent Q-value, TD error $\mathbf{u}_t = O(f_{\theta_t}(\mathbf{X})) = O(\theta_t^L)$. Besides,
535 according to lemma1, we know $k_t = O(\|\theta_t\|^{2(L-1)}) = O(\|\mathbf{u}_t\|^{2(L-1)/L}) = O(\|\mathbf{u}_t\|^{2-2/L})$.

536 Denote $g(\eta t) = \mathbf{v}^\top \mathbf{u}_t$, left multiply \mathbf{v} to equation $\mathbf{u}_{t+1} = (\mathbf{I} + \eta k_t \mathbf{A}) \mathbf{u}_t$. we have $g(\eta t + \eta) =$
537 $(1 + \eta \lambda k_t) g(\eta t)$. Since we know such iteration will let \mathbf{u}_t to be dominated by \mathbf{v} and align with \mathbf{v} , we
538 know $g(\eta t) = O(\|\mathbf{u}_t\|)$ for large t . Therefore $k_t = O(\|\mathbf{u}_t\|^{2(L-1)/L}) = C \cdot g(\eta t)^{2-2/L}$. This boils
539 down to $g(\eta t + \eta) = g(\eta t) + C \eta \lambda g(\eta t)^2$, which further becomes

$$\frac{g(\eta t + \eta) - g(\eta t)}{\eta} = C \lambda g(\eta t)^{3-2/L} \quad (15)$$

540 Let $\eta \rightarrow 0$, we have an differential equation $\frac{dg}{dt} = C \lambda g(t)^{3-2/L}$. When $L = 1$, the MLP network
541 degenerates to a linear function. The solution of ODE is

$$\|\mathbf{u}_t\| = g(\eta t) = C' e^{\lambda t}, \quad (16)$$

542 reflecting the exponential growth under linear function that has been studied in previous works [36].
543 When $L > 2$, Solving this ODE gives

$$g(t) = \frac{1}{(1 - C' \lambda t)^{L/(2L-2)}}. \quad (17)$$

544 So at an infinite limit, we know $\|\mathbf{u}_t\| = g(\eta t) = O\left(\frac{1}{(1 - C' \lambda \eta t)^{L/(2L-2)}}\right)$. Specially, for the exper-
545 imental case we study in Figure 3 where $L = 2$, it reduces to $O\left(\frac{1}{1 - C' \lambda \eta t}\right)$. We conduct more
546 experiments with $L = 3$ in Appendix C.2 to verify our theoretical findings. \square

547 **C More Observations and Deduction**

548 **C.1 Model Alignment**

549 In addition to the findings presented in Theorem 1 and Theorem 3, we have noticed several intriguing
 550 phenomena. Notably, beyond the critical point, gradients tend to align along a particular direction,
 551 leading to an infinite growth of the model’s parameters in that same direction. This phenomenon is
 552 supported by the observations presented in Figure 12, Figure 13, and Figure 14, where the cosine
 553 similarity between the current model parameters and the ones at the ending of training remains close
 554 to 1 after reaching a critical point, even as the norm of the parameters continually increases.

555 **C.2 Terminal Time**

556 Theorem 3 claims $\|\mathbf{u}_t\| = O(f_{\theta_t}(\mathbf{X})) = O\left(\frac{1}{(1-C'\lambda\eta t)^{L/(2L-2)}}\right)$, implying the relation

$$1/q^{(2L-2)/L} \propto 1 - C'\lambda\eta t. \quad (18)$$

557 Besides, it implies the existence of a “terminal time” $\frac{1}{C'\eta\lambda}$ that the model must crash at a singular
 558 point. When the training approaches this singular point, the estimation value and the model’s norm
 559 explode rapidly in very few steps. We have run an experiment with $L = 2$ in Figure 3, from which we
 560 can see that Q-value’s inverse proves to decay linearly and eventually becomes Nan at the designated
 561 time step. When $L = 3$, from our theoretical analysis, we have $1/q^{4/3} \propto 1 - C'\lambda\eta t$. The experimental
 562 results in Figure 10 corroborate this theoretical prediction, where the inverse Q-value raised to the
 563 power of 4/3 is proportional to $1 - C'\lambda\eta t$ after a critical point and it eventually reaches a NAN value
 564 at the terminal time step.

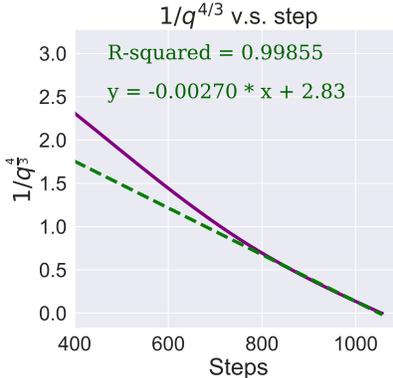


Figure 10: Linear decay with SGD and L=3.

565 **C.3 Adam Case**

566 In this section, we will prove that if the algorithm employs Adam as the optimizer, the model still
 567 suffers divergence. Moreover, we demonstrate that the norm of the network increase linearly, of
 568 which the slope is $\eta\sqrt{P}$, where P is the number of parameters and η is the learning rate. Also, the
 569 Q-value prediction will increase at L_{th} -polynomial’s rate, where L is the number of layers of model
 570 f_{θ} . Experimental results in Figure 4 verified our findings. Besides, we show that all runnings across
 571 D4RL environments represents the linear growth of the norm of the Q-network in Figure 12, Figure 13,
 572 and Figure 14.

573 **Theorem 4.** *Suppose we use Adam optimizer for Q-value iteration and all other settings are the same*
 574 *as Theorem 3. After $t > t_0$, the model will diverge if and only if $\lambda_{\max}(\mathbf{A}) > 0$. If it diverges, we*
 575 *have $\|\theta_t\| = \eta\sqrt{P}t + o(t)$ and $\|\mathbf{u}_t\| = \Theta(t^L)$ where P and L are the number of parameters and the*
 576 *number of layers for network f_{θ} , respectively.*

577 *Proof.* We only focus on the asymptotic behavior of Adam. So we only care about the dynamics for
 578 $t > T$ for some large T . Also, at this regime, we know that the gradient has greatly aligned with the

579 model parameters. So we assume that

$$\nabla L(\theta_t) = -C \cdot \theta_t. \quad C > 0 \quad (19)$$

580 Recall that each iteration of the Adam algorithm has the following steps.

$$g_t = \nabla L(\theta_t), \quad (20)$$

$$m_t = \beta_1 m_{t-1} + (1 - \beta_1) g_t, \quad (21)$$

$$v_t = \beta_2 v_{t-1} + (1 - \beta_2) g_t^2, \quad (22)$$

$$\hat{m}_t = \frac{m_t}{1 - \beta_1^t}, \quad (23)$$

$$\hat{v}_t = \frac{v_t}{1 - \beta_2^t}, \quad (24)$$

$$\theta_{t+1} = \theta_t - \frac{\eta}{\sqrt{\hat{v}_t + \epsilon}} \hat{m}_t. \quad (25)$$

581 Instead of exactly solving this series, we can verify linear growth is indeed the terminal behavior of
 582 θ_t since we only care about asymptotic order. Assume that $\theta_t = kt$ for $t > T$, we can calculate m_t
 583 by dividing both sides of the definition of m_t by β_1^t , which gives

$$\frac{m_t}{\beta_1^t} = \frac{m_{t-1}}{\beta_1^{t-1}} + \frac{1 - \beta_1}{\beta_1^t} g_t. \quad (26)$$

$$\frac{m_t}{\beta_1^t} = \sum_{s=0}^{t-1} \frac{1 - \beta_1}{\beta_1^s} g_s. \quad (27)$$

$$m_t = -C \sum_{s=0}^{t-1} (1 - \beta_1) \beta_1^{t-s} k s = -kCt + o(t) \quad (28)$$

584 , where g_t is given in Equation (19). Similarly, we have

$$v_t = kC^2 t^2 + o(t^2) \quad (29)$$

Hence we verify that

$$\theta_{t+1} - \theta_t = -\eta \cdot \frac{m_t}{1 - \beta_1^t} \cdot \sqrt{\frac{1 - \beta_2^t}{v_t}} \rightarrow \eta \cdot \frac{kCt}{\sqrt{k^2 C^2 t^2}} = \eta$$

585 therefore we know each iteration will increase each parameter by exactly constant η . This in turn
 586 verified our assumption that parameter θ_t grows linearly. The slope for the overall parameter is thus
 587 $\eta\sqrt{P}$. This can also be verified in Figure 4. When we have $\theta_t = \eta\sqrt{P}\bar{\theta}$, where $\bar{\theta}$ is the normalized
 588 parameter, we can further deduce the increasing order of the model's estimation. According to lemma
 589 1, the Q-value estimation (also the training error) increase at speed $O(t^L)$. \square

590 D LayerNorm's Effect on NTK

591 In this section, we demonstrate the effect of LayerNorm on SEEM. Our demonstration is just an
 592 intuitive explanation rather than a rigorous proof. We show that adding a LayerNorm can effectively
 593 reduce the NTK between any x_0 and extreme input x down from linear to constant. Since each
 594 entry of Gram matrix \mathbf{G} is an individual NTK value, we can informally expect that $\mathbf{G}(\mathbf{X}_t^*, \mathbf{X})$'s
 595 eigenvalue are greatly reduced when every individual NTK value between any x_0 and extreme input
 596 x is reduced.

We consider a two-layer MLP. The input is $\mathbf{x} \in \mathbb{R}^{d_{in}}$, and the hidden dimension is d . The parameters include $\mathbf{W} = [\mathbf{w}_1, \dots, \mathbf{w}_d]^\top \in \mathbb{R}^{d \times d_{in}}$, $\mathbf{b} \in \mathbb{R}^d$ and $\mathbf{a} \in \mathbb{R}^d$. Since for the NTK value, the last layer's bias term has a constant gradient, we do not need to consider it. The forward function of the network is

$$f_{\theta}(\mathbf{x}) = \sum_{i=1}^d a_i \sigma(\mathbf{w}_i^\top \mathbf{x} + b_i).$$

597 **Proposition 1.** For any input \mathbf{x} and network parameter $\boldsymbol{\theta}$, if $\nabla_{\boldsymbol{\theta}} f_{\boldsymbol{\theta}}(\mathbf{x}) \neq \mathbf{0}$, then we have

$$\lim_{\lambda \rightarrow \infty} k_{\text{NTK}}(\mathbf{x}, \lambda \mathbf{x}) = \Omega(\lambda) \rightarrow \infty. \quad (30)$$

598 *Proof.* Denote $z_i = \mathbf{w}_i^\top \mathbf{x} + b_i$, according to condition $\nabla_{\boldsymbol{\theta}} f_{\boldsymbol{\theta}}(\mathbf{x}) \neq \mathbf{0}$, we know there must exist at
 599 least one i such that $z_i > 0$, denote this set as P . Now consider all the $i \in [d]$ that satisfy $z_i > 0$ and
 600 $\mathbf{w}_i^\top \mathbf{x} > 0$ (otherwise take opposite sign of λ), we have

$$\left. \frac{\partial f}{\partial a_i} \right|_{\mathbf{x}} = \sigma(\mathbf{w}_i^\top \mathbf{x} + b_i) = \mathbf{w}_i^\top \mathbf{x} + b_i, \quad (31)$$

$$\left. \frac{\partial f}{\partial \mathbf{w}_i} \right|_{\mathbf{x}} = a_i \mathbf{x}, \quad (32)$$

$$\left. \frac{\partial f}{\partial b_i} \right|_{\mathbf{x}} = a_i. \quad (33)$$

601 Similarly, we have

$$\left. \frac{\partial f}{\partial a_i} \right|_{\lambda \mathbf{x}} = \sigma(\lambda \mathbf{w}_i^\top \mathbf{x} + b_i) = \lambda \mathbf{w}_i^\top \mathbf{x} + b_i, \quad (34)$$

$$\left. \frac{\partial f}{\partial \mathbf{w}_i} \right|_{\lambda \mathbf{x}} = \lambda a_i \mathbf{x}, \quad (35)$$

$$\left. \frac{\partial f}{\partial b_i} \right|_{\lambda \mathbf{x}} = a_i. \quad (36)$$

So we have

$$\sum_{i \in P} \left\langle \frac{\partial f(\mathbf{x})}{\partial \boldsymbol{\theta}_i}, \frac{\partial f(\lambda \mathbf{x})}{\partial \boldsymbol{\theta}_i} \right\rangle = \lambda ((\mathbf{w}_i^\top \mathbf{x})^2 + b_i \mathbf{w}_i^\top \mathbf{x} + a_i^2 \|\mathbf{x}\|^2) + O(1) = \Theta(\lambda).$$

Denote $N = \{1, \dots, d\} \setminus P$. We know for every $j \in N$ either $\frac{\partial f(\mathbf{x})}{\partial a_j} = \frac{\partial f(\mathbf{x})}{\partial b_j} = \frac{\partial f(\mathbf{x})}{\partial \mathbf{w}_j} = 0$, or
 $\mathbf{w}_j^\top \mathbf{x} < 0$. For the latter case, we know $\lim_{\lambda \rightarrow \infty} \frac{\partial f(\lambda \mathbf{x})}{\partial a_j} = \frac{\partial f(\lambda \mathbf{x})}{\partial b_j} = \frac{\partial f(\lambda \mathbf{x})}{\partial \mathbf{w}_j} = 0$. In both cases, we
 have

$$\lim_{\lambda \rightarrow \infty} \sum_{j \in N} \left\langle \frac{\partial f(\mathbf{x})}{\partial \boldsymbol{\theta}_j}, \frac{\partial f(\lambda \mathbf{x})}{\partial \boldsymbol{\theta}_j} \right\rangle = 0.$$

Therefore, according to the definition of NTK, we have

$$\lim_{\lambda \rightarrow \infty} k_{\text{NTK}}(\mathbf{x}, \lambda \mathbf{x}) = \lim_{\lambda \rightarrow \infty} \left\langle \frac{\partial f(\mathbf{x})}{\partial \boldsymbol{\theta}_i}, \frac{\partial f(\lambda \mathbf{x})}{\partial \boldsymbol{\theta}_i} \right\rangle = \Theta(\lambda) \rightarrow \infty.$$

602

□

For the model equipped with LayerNorm, the forward function becomes

$$f_{\boldsymbol{\theta}}(\mathbf{x}) = \mathbf{a}^\top \sigma(\psi(\mathbf{W}\mathbf{x} + \mathbf{b})),$$

where $\psi(\cdot)$ is the layer normalization function defined as

$$\psi(\mathbf{x}) = \sqrt{d} \cdot \frac{\mathbf{x} - \mathbf{1}\mathbf{1}^\top \mathbf{x}/d}{\|\mathbf{x} - \mathbf{1}\mathbf{1}^\top \mathbf{x}/d\|}.$$

603 Denote $\mathbf{P} = \mathbf{I} - \mathbf{1}\mathbf{1}^\top/d$, note that the derivative of $\psi(\cdot)$ is

$$\dot{\psi}(\mathbf{x}) = \frac{\partial \psi(\mathbf{x})}{\partial \mathbf{x}} = \sqrt{d} \cdot \left(\frac{\mathbf{I}}{\|\mathbf{P}\mathbf{x}\|} - \frac{\mathbf{P}\mathbf{x}\mathbf{x}^\top \mathbf{P}}{\|\mathbf{P}\mathbf{x}\|^3} \right) \mathbf{P}. \quad (37)$$

604 Specially, we have

$$\psi(\lambda \mathbf{x}) = \sqrt{d} \cdot \frac{\lambda \mathbf{x} - \lambda \mathbf{1}\mathbf{1}^\top \mathbf{x}/d}{\lambda \|\mathbf{x} - \mathbf{1}\mathbf{1}^\top \mathbf{x}/d\|} = \psi(\mathbf{x}). \quad (38)$$

605 Now we state the second proposition.

606 **Proposition 2.** For any input \mathbf{x} and network parameter $\boldsymbol{\theta}$ and any direction $\mathbf{v} \in \mathbb{R}^{d_{in}}$, if the network
607 has LayerNorm, then we know there exists a universal constant C , such that for any $\lambda \geq 0$, we have

$$k_{\text{NTK}}(\mathbf{x}, \mathbf{x} + \lambda \mathbf{v}) \leq C. \quad (39)$$

608 *Proof.* Since for finite range, there always exists a constant upper bound, we just need to analyze the
609 case for $\lambda \rightarrow +\infty$ and shows that it is constant bounded. First compute $\nabla_{\boldsymbol{\theta}} f_{\boldsymbol{\theta}}(\mathbf{x})$ and get

$$\left. \frac{\partial f}{\partial \mathbf{a}} \right|_{\mathbf{x}} = \sigma(\psi(\mathbf{W}\mathbf{x} + \mathbf{b})), \quad (40)$$

$$\left. \frac{\partial f}{\partial \mathbf{W}} \right|_{\mathbf{x}} = \mathbf{a}^{\top} \sigma'(\psi(\mathbf{W}\mathbf{x} + \mathbf{b})) \dot{\psi}(\mathbf{W}\mathbf{x} + \mathbf{b}) \mathbf{x}, \quad (41)$$

$$\left. \frac{\partial f}{\partial \mathbf{b}} \right|_{\mathbf{x}} = \mathbf{a}^{\top} \sigma'(\psi(\mathbf{W}\mathbf{x} + \mathbf{b})) \dot{\psi}(\mathbf{W}\mathbf{x} + \mathbf{b}). \quad (42)$$

610 These quantities are all constant bounded. Next we compute $\lim_{\lambda \rightarrow \infty} \nabla_{\boldsymbol{\theta}} f_{\boldsymbol{\theta}}(\mathbf{x} + \lambda \mathbf{v})$

$$\left. \frac{\partial f}{\partial \mathbf{a}} \right|_{\mathbf{x} + \lambda \mathbf{v}} = \sigma(\psi(\mathbf{W}(\mathbf{x} + \lambda \mathbf{v}) + \mathbf{b})), \quad (43)$$

$$\left. \frac{\partial f}{\partial \mathbf{W}} \right|_{\mathbf{x} + \lambda \mathbf{v}} = \mathbf{a}^{\top} \sigma'(\psi(\mathbf{W}(\mathbf{x} + \lambda \mathbf{v}) + \mathbf{b})) \dot{\psi}(\mathbf{W}(\mathbf{x} + \lambda \mathbf{v}) + \mathbf{b})(\mathbf{x} + \lambda \mathbf{v}), \quad (44)$$

$$\left. \frac{\partial f}{\partial \mathbf{b}} \right|_{\mathbf{x} + \lambda \mathbf{v}} = \mathbf{a}^{\top} \sigma'(\psi(\mathbf{W}\mathbf{x} + \mathbf{b})) \dot{\psi}(\mathbf{W}(\mathbf{x} + \lambda \mathbf{v}) + \mathbf{b}). \quad (45)$$

611 According to the property of LayerNorm in Equation (38), we have

$$\overline{\lim}_{\lambda \rightarrow \infty} \left. \frac{\partial f}{\partial \mathbf{a}} \right|_{\mathbf{x} + \lambda \mathbf{v}} = \overline{\lim}_{\lambda \rightarrow \infty} \sigma(\psi(\mathbf{W}(\mathbf{x} + \lambda \mathbf{v}) + \mathbf{b})) \quad (46)$$

$$= \sigma(\psi(\mathbf{W}(\lambda \mathbf{v}))) \quad (47)$$

$$= \sigma(\psi(\mathbf{W}\mathbf{v})) = \text{Constant} \quad (48)$$

$$\overline{\lim}_{\lambda \rightarrow \infty} \left. \frac{\partial f}{\partial \mathbf{W}} \right|_{\mathbf{x} + \lambda \mathbf{v}} = \overline{\lim}_{\lambda \rightarrow \infty} \mathbf{a}^{\top} \sigma'(\psi(\mathbf{W}(\mathbf{x} + \lambda \mathbf{v}) + \mathbf{b})) \dot{\psi}(\mathbf{W}(\mathbf{x} + \lambda \mathbf{v}) + \mathbf{b})(\mathbf{x} + \lambda \mathbf{v}) \quad (49)$$

$$= \overline{\lim}_{\lambda \rightarrow \infty} \mathbf{a}^{\top} \sigma'(\psi(\mathbf{W}\mathbf{v})) \dot{\psi}(\mathbf{W}(\mathbf{x} + \lambda \mathbf{v}) + \mathbf{b})(\mathbf{x} + \lambda \mathbf{v}) \quad (50)$$

$$= \overline{\lim}_{\lambda \rightarrow \infty} \mathbf{a}^{\top} \sigma'(\psi(\mathbf{W}\mathbf{v})) \sqrt{d} \cdot \left(\frac{\mathbf{I}}{\|\mathbf{P}\lambda \mathbf{W}\mathbf{v}\|} - \frac{\mathbf{P}(\lambda \mathbf{W}\mathbf{v})(\lambda \mathbf{W}\mathbf{v})^{\top} \mathbf{P}}{\|\mathbf{P}(\lambda \mathbf{W}\mathbf{v})\|^3} \right) \mathbf{P}(\mathbf{x} + \lambda \mathbf{v}) \quad (51)$$

$$= \overline{\lim}_{\lambda \rightarrow \infty} \mathbf{a}^{\top} \sigma'(\psi(\mathbf{W}\mathbf{v})) \sqrt{d} \cdot \left(\frac{\mathbf{P}(\mathbf{x} + \lambda \mathbf{v})}{\lambda \|\mathbf{P}\mathbf{W}\mathbf{v}\|} - \frac{\mathbf{P}\mathbf{W}\mathbf{v}\mathbf{v}^{\top} \mathbf{W}^{\top} \mathbf{P}(\mathbf{x} + \lambda \mathbf{v})}{\lambda \|\mathbf{P}\mathbf{W}\mathbf{v}\|^3} \right) \quad (52)$$

$$= \mathbf{a}^{\top} \sigma'(\psi(\mathbf{W}\mathbf{v})) \sqrt{d} \cdot \left(\frac{\mathbf{P}\mathbf{v}}{\|\mathbf{P}\mathbf{W}\mathbf{v}\|} - \frac{\mathbf{P}\mathbf{W}\mathbf{v}\mathbf{v}^{\top} \mathbf{W}^{\top} \mathbf{P}\mathbf{v}}{\|\mathbf{P}\mathbf{W}\mathbf{v}\|^3} \right) \quad (53)$$

$$= \text{Constant}. \quad (54)$$

$$\overline{\lim}_{\lambda \rightarrow \infty} \left. \frac{\partial f}{\partial \mathbf{b}} \right|_{\mathbf{x} + \lambda \mathbf{v}} = \overline{\lim}_{\lambda \rightarrow \infty} \mathbf{a}^{\top} \sigma'(\psi(\mathbf{W}\mathbf{v})) \sqrt{d} \cdot \left(\frac{\mathbf{I}}{\lambda \|\mathbf{P}\mathbf{W}\mathbf{v}\|} - \frac{\mathbf{P}\mathbf{W}\mathbf{v}\mathbf{W}\mathbf{v}^{\top} \mathbf{P}}{\lambda \|\mathbf{P}(\mathbf{W}\mathbf{v})\|^3} \right) \mathbf{P} \quad (55)$$

$$= 0. \quad (56)$$

612 Therefore we know its limit is also constant bounded. So we know there exists a universal constant
613 with respect to $\boldsymbol{\theta}, \mathbf{x}, \mathbf{v}$ such that $k_{\text{NTK}}(\mathbf{x}, \mathbf{x} + \lambda \mathbf{v}) = \left\langle \frac{\partial f(\mathbf{x})}{\partial \boldsymbol{\theta}_i}, \frac{\partial f(\mathbf{x} + \lambda \mathbf{v})}{\partial \boldsymbol{\theta}_i} \right\rangle \leq C$.

614 E Experiment Setup

615 **SEEM Experiments** For the experiments presented in Section Section 3.1, we adopted TD3 as our
616 baseline, but with a modification: instead of using an exponential moving average (EMA), we directly

617 copied the current Q-network as the target network. The Adam optimizer was used with a learning rate
 618 of 0.0003, $\beta_1 = 0.9$, and $\beta_2 = 0.999$. The discount factor, γ , was set to 0.99. Our code builds upon
 619 the existing TD3+BC framework, which can be found at https://github.com/sfujim/TD3_BC.

620 **SEEM Reduction Experiments** For the experiments discussed in Section Section 4, we maintained
 621 the same configuration as in the SEEM experiments, with the exception of adding regularizations and
 622 normalizations. LayerNorm was implemented between the linear and activation layers with learnable
 623 affine parameters, applied to all hidden layers excluding the output layer. WeightNorm was applied
 624 to the output layer weights.

625 **Offline RL Algorithm Experiments** For the experiments presented in Section Section 5, we used
 626 true offline RL algorithms including TD3+BC, IQL, Diff-QL, and CQL as baselines. We implement
 627 our method on the top of official implementations of TD3+BC and IQL; for CQL and Diff-QL, we
 628 use reliable JAX implementations. LayerNorm was directly added to the critic network in these
 629 experiments.

630 **Linear Decay of Inverse Q-value with SGD** Given that the explosion in D4RL environments
 631 occurs very quickly in the order of $\frac{1}{1-C'\lambda\eta t}$ and is difficult to capture, we opted to use a simple
 632 toy task for these experiments. The task includes a continuous two-dimensional state space $s =$
 633 $(x_1, x_2) \in \mathcal{S} = \mathbb{R}^2$, where the agent can freely navigate the plane. The action space is discrete, with
 634 8 possible actions representing combinations of forward or backward movement in two directions.
 635 Each action changes the state by a value of 0.01. All rewards are set to zero, meaning that the true
 636 Q-value should be zero for all state-action pairs. For this task, we randomly sampled 100 state-action
 637 pairs as our offline dataset. The Q-network was implemented as a two-layer MLP with a hidden size
 638 of 200. We used SGD with a learning rate of 0.01, and the discount factor, γ was set to 0.99.

639 F More Experiments

640 **Benchmarking Normalizations.** Previously, we have demonstrated that LayerNorm, BatchNorm,
 641 and WeightNorm can effectively maintain a low SEEM and stabilize Q convergence in Section 4. Our
 642 next goal is to identify the most suitable regularization method for the value network in offline RL.
 643 Prior research has shown that divergence is correlated with poor control performance[37, 18]. In this
 644 context, we evaluate the effectiveness of various regularization techniques based on their performance
 645 in two distinct settings - the Antmaze task and the X% Mujoco dataset we mentioned above. Previous
 646 offline RL algorithms have not performed particularly well in these challenging scenarios. As
 647 displayed in Figure 11, TD3+BC, when coupled with layer normalization or batch normalization,
 648 yields significant performance enhancement on the 10% Mujoco datasets. The inability of batch
 649 normalization to improve the performance might be attributed to the oscillation issue previously
 650 discussed in Section 4. In the case of Antmaze tasks, which contain numerous suboptimal trajectories,
 651 we select TD3 with a diffusion policy, namely Diff-QL [38], as our baseline. The diffusion policy is
 652 capable of capturing multi-modal behavior. As demonstrated in Figure 11 and Table 2, LayerNorm
 653 can markedly enhance performance on challenging Antmaze tasks. In summary, we empirically find
 654 LayerNorm to be a suitable normalization for the critic in offline RL.

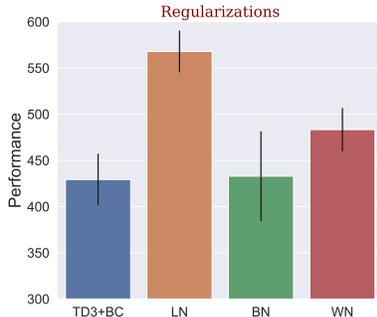


Figure 11: Normalizations effect on 10% Mujoco Locomotion Datasets.

Table 2: Normalizations effect on two challenging Antmaze tasks.

Dataset	diff-QL	LN	BN	WN
antmaze-large-play-v0	1.6	72.7	1.0	35.0
antmaze-large-diverse-v0	4.4	66.5	2.1	42.5

655 **How LayerNorm should be added.** The inclusion of LayerNorm is situated between the linear
656 and activation layers. However, the ideal configuration for adding LayerNorm can vary and may
657 depend on factors such as 1) the specific layers to which LayerNorm should be added, and 2) whether
658 or not to apply learnable per-element affine parameters. To explore these variables, we conducted
659 an assessment of their impacts on performance in the two most challenging Antmaze environments.
660 Our experimental setup mirrored that of the Antmaze experiments mentioned above, utilizing a three-
661 layer MLP critic with a hidden size configuration of (256,256,256). We evaluated variants where
662 LayerNorm was only applied to a portion of hidden layers and where learnable affine parameters were
663 disabled. As seen in Table 3, the performances with LayerNorm applied solely to the initial layers
664 LN (0), LN (0,1) are considerably lower compared to the other setups in the ‘antmaze-large-play-v0’
665 task, while applying LayerNorm to all layers LN(0,1,2) seems to yield the best performance. For the
666 ‘antmaze-large-diverse-v0’ task, performances seem to be more consistent across different LayerNorm
667 applications. Overall, this analysis suggests that applying LayerNorm to all layers tends to yield
668 the best performance in these tasks. Also, the utilization of learnable affine parameters appears less
669 critical in this context.

Table 3: The effect of LayerNorm implementations on two challenging Antmaze tasks.

Dataset	w.o. LN	LN (0)	LN (0,1,)	LN (1,2)	LN (2)	LN (0,1,2)	LN (no learnable)
antmaze-large-play-v0	1.6	0	0	8.3	17.8	72.7	72.8
antmaze-large-diverse-v0	4.4	60.2	68	77.1	65.5	66.5	66.7

670 G Discussion

671 **SEEM and Deadly Triad.** Deadly Triad is a term that refers to a problematic interaction observed
672 in reinforcement learning algorithms, where off-policy learning, function approximation, and boot-
673 strapping converge, leading to divergence during training. Existing studies primarily analyze linear
674 functions as Q-values, which tend to limit the analysis to specific toy examples. In contrast, our work
675 uses NTK theory to provide an in-depth understanding of the divergence of Q-values in non-linear
676 neural networks in realistic settings, and introduces SEEM as a tool to depict such divergence. SEEM
677 can be used to understand the Deadly Triad as follows: If a policy is nearly on-policy, \mathbf{X}_t^* is merely a
678 perturbation of \mathbf{X} . Consequently, $\mathbf{A}_t = \gamma \mathbf{G}_{\theta_t}(\mathbf{X}_t^*, \mathbf{X}) - \mathbf{G}_{\theta_t}(\mathbf{X}, \mathbf{X}) \approx (\gamma - 1) \mathbf{G}_{\theta_t}(\mathbf{X}, \mathbf{X})$, with
679 \mathbf{G} tending to be negative-definite. Without function approximation, the update of $Q(\mathbf{X})$ will not
680 influence $Q(\mathbf{X}_t^*)$, and the first term in \mathbf{A}_t becomes zero. $\mathbf{A}_t = -\mathbf{G}_{\theta_t}(\mathbf{X}, \mathbf{X})$ ensures that SEEM
681 is non-positive and Q-value iteration remains non-expansive. If we avoid bootstrapping, the value
682 iteration transforms into a supervised learning problem with well-understood convergence properties.
683 However, when all three components in Deadly Triad are present, the NTK analysis gives rise to the
684 form $\mathbf{A}_t = \gamma \mathbf{G}_{\theta_t}(\mathbf{X}_t^*, \mathbf{X}) - \mathbf{G}_{\theta_t}(\mathbf{X}, \mathbf{X})$, which may result in divergence if the SEEM is positive.

685 **Policy Constraint and LayerNorm.** We have established a connection between SEEM and value
686 divergence. As shown in Figure 6, policy constraint alone can also control SEEM and prevent
687 divergence. In effect, policy constraint addresses an aspect of the Deadly Triad by managing the
688 degree of off-policy learning. However, an overemphasis on policy constraint, leading to excessive
689 bias, can be detrimental to the policy and impair performance, as depicted in Figure 7. Building on
690 this insight, we focus on an orthogonal perspective in deadly triad - regularizing the generalization
691 capacity of the critic network. Specifically, we propose the use of LayerNorm in the critic network to
692 inhibit value divergence and enhance agent performance. Policy constraint introduces an explicit bias
693 into the policy, while LayerNorm does not. Learning useful information often requires some degree
694 of prior bias towards offline dataset, but too much can hinder performance. LayerNorm, offering an
695 orthogonal perspective to policy constraint, aids in striking a better balance.

696 **H More Visualization Results**

697 In Assumption 2, we posit that the direction of NTK and the policy remains stable following a
 698 certain period of training. We validate this assumption through experimental studies. We observe
 699 the convergence of the NTK trajectory and policy in all D4RL Mujoco Locomotion and Antmaze
 700 tasks, as depicted in the first two columns of Figures Figure 12, Figure 13, and Figure 14. We also
 701 illustrate the linear growth characteristic of Adam optimization (as outlined in Theorem Theorem 4)
 702 in the fourth column. As a consequence, the model parameter vectors maintain a parallel trajectory,
 703 keeping the cosine similarity near 1 as shown in the third column. Figure 15 and Figure 16 showcase
 704 how SEEM serves as a "divergence detector" in Mujoco and Antmaze tasks. The surge in the SEEM
 705 value is consistently synchronized with an increase in the estimated Q-value.

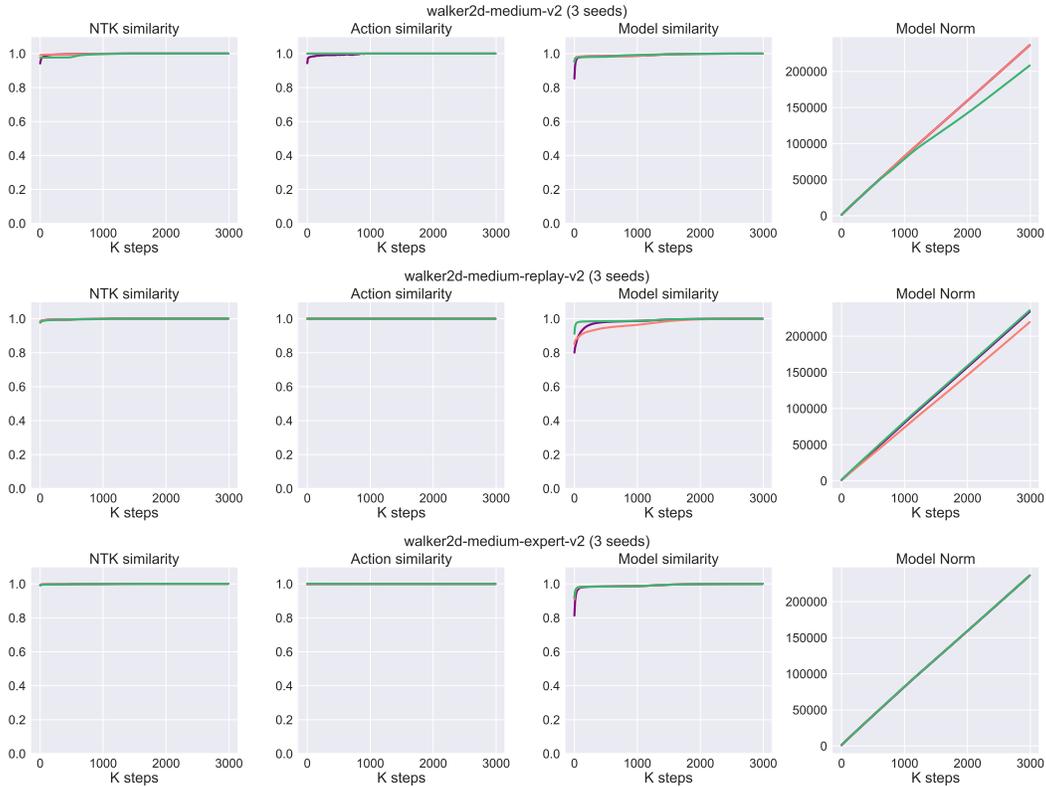


Figure 12: NTK similarity, action similarity, model parameter similarity, and model parameter norm curves in D4RL Mujoco Walker2d tasks.

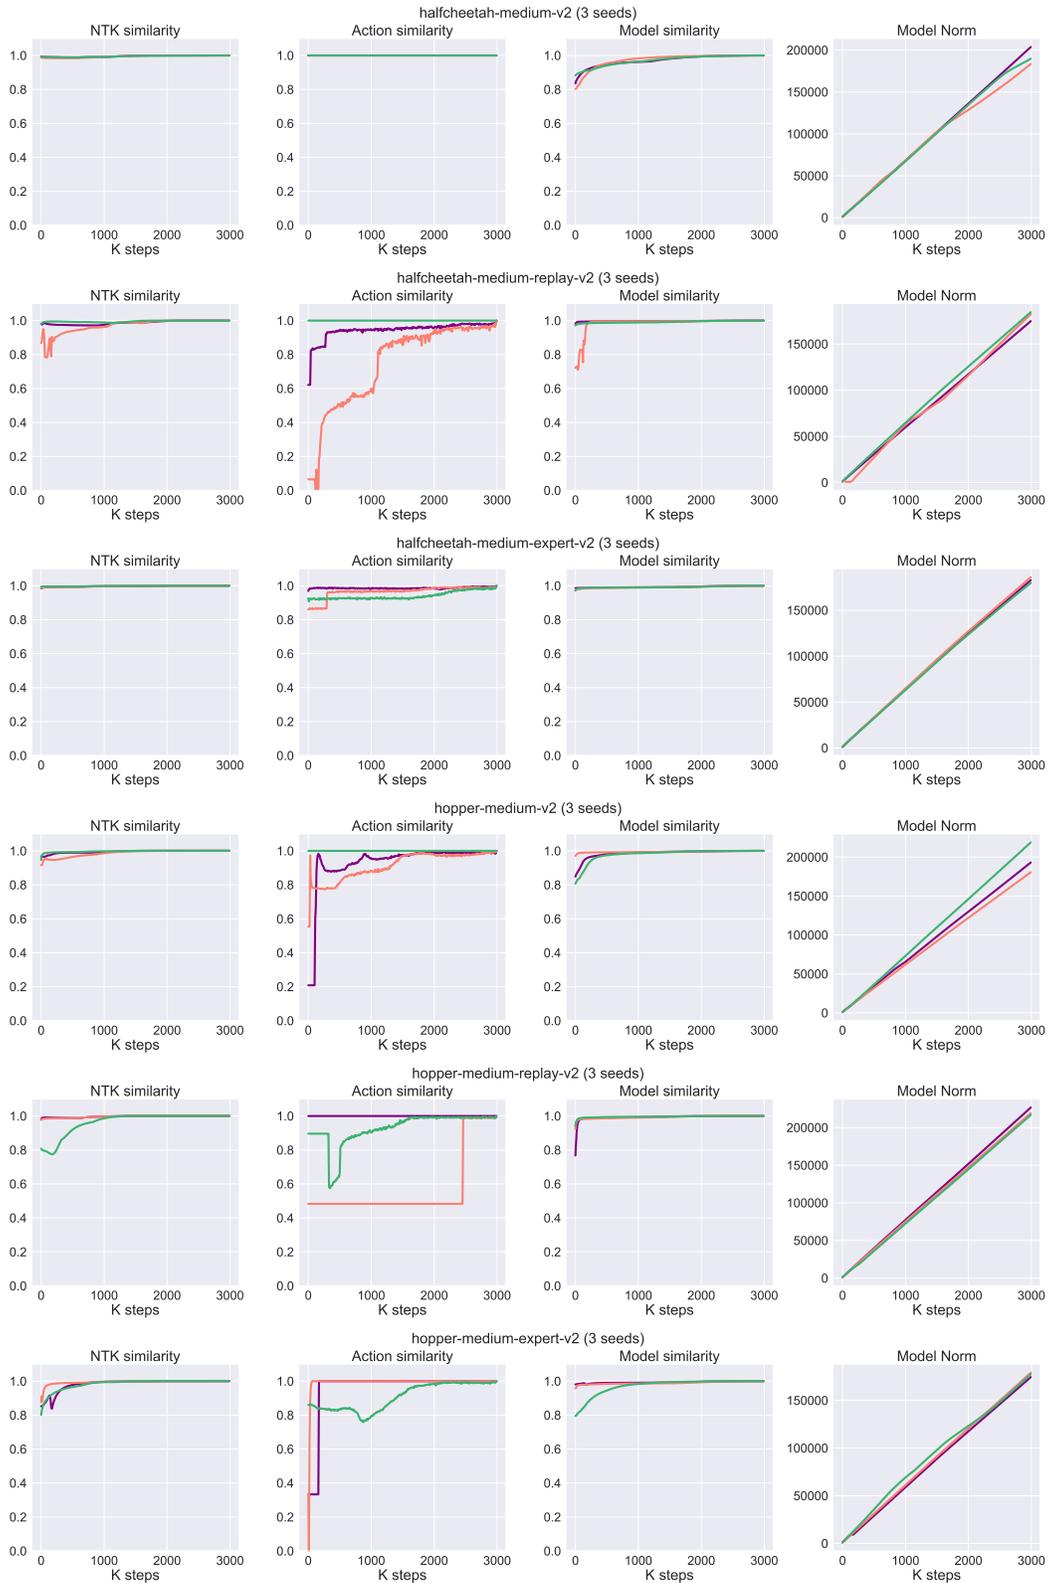


Figure 13: NTK similarity, action similarity, model parameter similarity, and model parameter norm curves in D4RL Mujoco Halfcheetah and Hopper tasks.

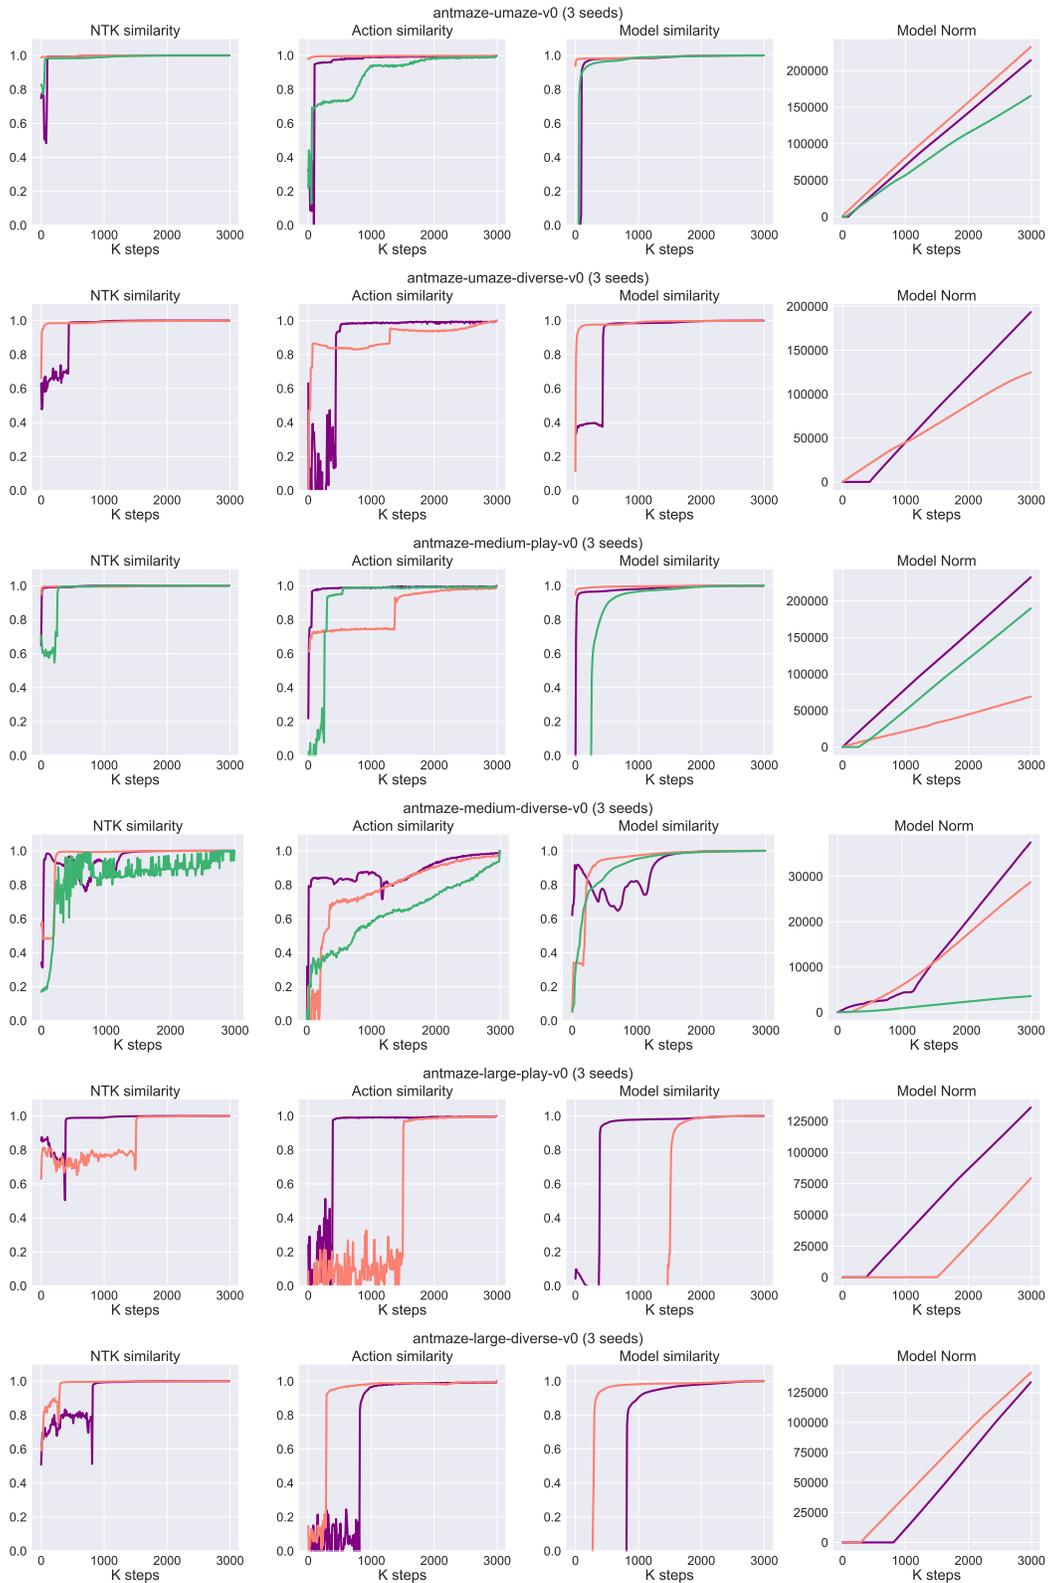


Figure 14: NTK similarity, action similarity, model parameter similarity, and model parameter norm curves in Antmaze tasks.

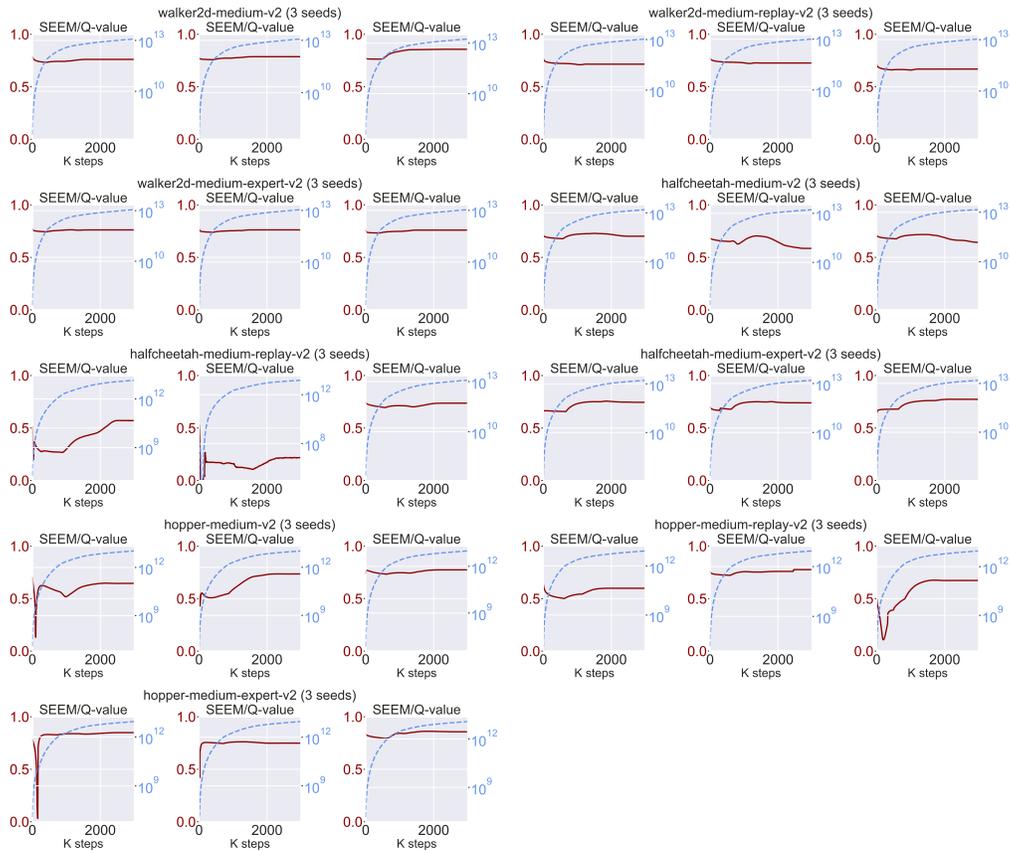
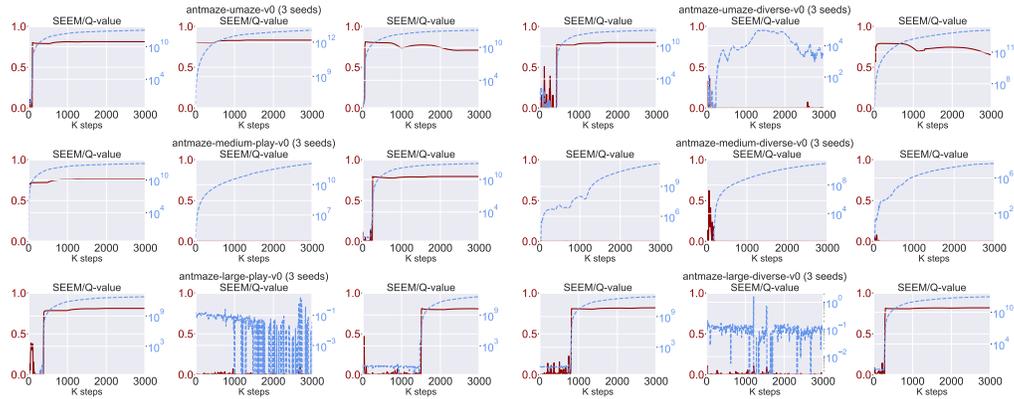
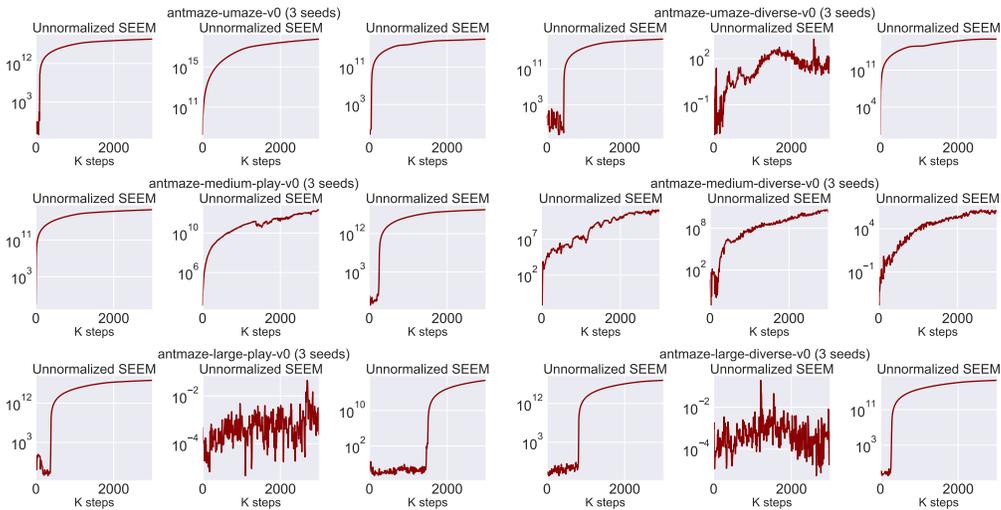


Figure 15: The normalized kernel matrix's **SEEM** (in red) and the estimated **Q-value** (in blue) in D4RL Mujoco tasks. For each environment, results from three distinct seeds are reported.



(a) The normalized kernel matrix's SEEM (in red) and the estimated Q-value (in blue) in D4RL Mujoco tasks.



(b) The unnormalized kernel matrix's SEEM. The three curves in each environment correspond directly to those presented in Figure (a)

Figure 16: In Figure (a), an inflation in the estimated Q-value coincides with a surge in the normalized SEEM. However, there are some anomalies, such as the second running in the 'umaze-diverse' environment, where the Q-value rises while the unnormalized SEEM remains low. However, the corresponding normalized SEEM in Figure (b) suggests an actual inflation of SEEM. Furthermore, for scenarios where the Q-value converges, as seen in the second running in 'large-diverse', the unnormalized SEEM maintains an approximate zero value.