

000 FIRE: FROBENIUS-ISOMETRY REINITIALIZATION FOR 001 BALANCING THE STABILITY-PLASTICITY TRADEOFF 002

003 **Anonymous authors**

004 Paper under double-blind review

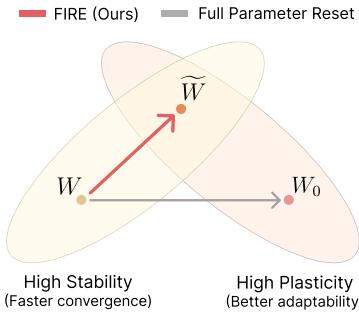
005 ABSTRACT

006 Deep neural networks trained on nonstationary data must balance stability (i.e.,
007 retaining prior knowledge) and plasticity (i.e., adapting to new tasks). Standard
008 reinitialization methods, which reinitialize weights toward their original values,
009 are widely used but difficult to tune: conservative reinitializations fail to restore
010 plasticity, while aggressive ones erase useful knowledge. We propose FIRE, a
011 principled reinitialization method that explicitly balances the stability–plasticity
012 tradeoff. FIRE quantifies stability through Squared Frobenius Error (SFE), mea-
013 suring proximity to past weights, and plasticity through Deviation from Isometry
014 (DfI), reflecting weight isotropy. The reinitialization point is obtained by solving
015 a constrained optimization problem, minimizing SFE subject to DfI being zero,
016 which is efficiently approximated by Newton–Schulz iteration. FIRE is eval-
017 uated on continual visual learning (CIFAR-10 with ResNet-18), language modeling
018 (OpenWebText with GPT-0.1B), and reinforcement learning (HumanoidBench
019 with SAC and Atari games with DQN). Across all domains, FIRE consistently
020 outperforms both naive training without intervention and standard reinitialization
021 methods, demonstrating effective balancing of the stability–plasticity tradeoff.

022 1 INTRODUCTION

023 Deep neural networks are typically trained under a fixed,
024 stationary data distribution (Brown et al., 2020; Podell
025 et al., 2024). However, many real-world applications
026 require models to adapt continually as new data and
027 shifting distributions emerge. In computer vision, au-
028 tonomous driving systems must recognize unseen traffic
029 signs, road layouts, or weather conditions that were
030 absent during training (Verwimp et al., 2023). Large lan-
031 guage models, trained once and deployed with a fixed
032 knowledge cutoff date, quickly become outdated unless
033 continually updated (Ke et al., 2023). Likewise, robots
034 deployed in dynamic physical environments must adjust
035 their perception and control policies as the environment
036 changes (Wołczyk et al., 2021). In all of these domains,
037 a central challenge is reliable adaptation to nonstationary
038 data while preserving prior knowledge.

039 This challenge is often framed as a balance between two competing properties: *stability*, the retention
040 of learned knowledge, and *plasticity*, the ability to incorporate new information (Mermilliod et al.,
041 2013). Different research communities emphasize these properties to varying degrees. Conventional
042 continual learning assumes limited access to past data, prioritizing stability to mitigate catastrophic
043 forgetting (Kirkpatrick et al., 2017; Rebuffi et al., 2017; Rusu et al., 2016). In contrast, most
044 foundation models and robotic agents are trained on expanding datasets where past data remain
045 accessible (Achiam et al., 2023; Team et al., 2025), making plasticity loss a central challenge (Lyle
046 et al., 2023; Berariu et al., 2021). In this regime, stability is less about preserving past knowledge and
047 more about accelerating adaptation to new tasks by leveraging prior representations. Motivated by
048 these real world scenarios, we study the stability–plasticity tradeoff under the assumption of access to
049 past data during continual learning.



050 Figure 1: **Illustration of FIRE.** Solving
051 a constrained optimization problem, FIRE
052 places weights at the intersection of high-
053 stability and high-plasticity manifolds.

Existing approaches to mitigating plasticity loss fall broadly into two categories: regularization-based and reinitialization-based. Regularization-based methods constrain parameters or features near their initialization (Kumar et al., 2025b; Lyle et al., 2022), or enforce weight orthogonality (Chung et al., 2024). While these methods can preserve a favorable geometry for future learning, overly strong constraints slow convergence, and overly weak ones fail to prevent plasticity degradation. Reinitialization-based methods instead reset weights to earlier checkpoints when new data arrive (Ash & Adams, 2020; Nikishin et al., 2022; Lee et al., 2024a; Shin et al., 2024). Their advantage lies in avoiding interference with current optimization, often yielding faster adaptation with lower overhead. However, they also suffer from a tuning dilemma: aggressive resets erase useful knowledge, while conservative ones provide little plasticity gain.

We aim to resolve this dilemma by treating reinitialization as a principled constrained optimization problem. Our approach relies on two complementary measures that capture the core dimensions of the stability–plasticity tradeoff. First, we define stability using the Squared Frobenius Error (SFE) between current and past weights, which measures the sum of squared differences across all weight entries. A smaller SFE indicates greater similarity, meaning the model remains closer to its previous representations. For plasticity, prior work has linked plasticity loss to sharp loss curvature (Lyle et al., 2023), dormant neurons (Sokar et al., 2023), and low rank features (Kumar et al., 2021a), but these metrics depend on incoming data and are non differentiable, limiting their use for optimization. We instead propose the Deviation from Isometry (DfI) (Pennington et al., 2017; Xiao et al., 2018), which measures how close weight matrices are to orthonormal. We show that reducing DfI simultaneously decreases curvature, prevents neuron dormancy, increases feature rank, and remains differentiable, making it a practical measure of plasticity to optimize. A formal proof is provided in Section 3.

We propose FIRE (Frobenius–Isometry REinitialization), which minimizes the SFE subject to the DfI being 0. As illustrated in Figure 1, FIRE avoids the pitfalls of either overly conservative or aggressive reinitialization by projecting weights onto the isotropy manifold while remaining close to their previous subspace. While directly solving this constrained optimization is costly, it can be implemented efficiently with the Newton–Schulz iteration, adding less than 1% to training time.

We evaluate FIRE on continual learning benchmarks in vision, language, and reinforcement learning, assuming access to past data. For vision, we split CIFAR-10, CIFAR-100, and Tiny-ImageNet into chunks under random or class-incremental protocols using ResNet He et al. (2016) and Vision Transformer Dosovitskiy et al. (2020). For language, we use a warm-start setup where GPT-0.1B Karpathy (2023) is pretrained on WikiText-103 and then continually trained on a mixture of OpenWebText and WikiText-103. For reinforcement learning, we test continuous control with SAC Haarnoja et al. (2018) on HumanoidBench Sferrazza et al. (2024) and discrete control with DQN Mnih et al. (2015) on Atari Bellemare et al. (2013). For vision and language tasks, reinitialization is applied whenever new data arrive, while in reinforcement learning it is applied once at the midpoint of training. Across all domains, FIRE consistently outperforms naive training and standard reinitialization, showing its effectiveness as a unified solution to the stability plasticity tradeoff.

2 RELATED WORK

2.1 STABILITY-PLASTICITY TRADEOFF

The stability–plasticity tradeoff (Mermilliod et al., 2013; Kim & Han, 2023) is a fundamental challenge in continual learning. Stability refers to the ability of a model to preserve previously acquired knowledge and avoid catastrophic forgetting when exposed to new data. Plasticity refers to the ability of a model to adapt flexibly and effectively to novel tasks. These two properties often conflict with each other since strong stability can make the model rigid and resistant to new learning while excessive plasticity can lead to the loss of past knowledge.

Research in continual learning has therefore focused on methods that balance these competing requirements such as constraining parameter updates through regularization (Kirkpatrick et al., 2017), revisiting earlier data through replay (Rebuffi et al., 2017; Rolnick et al., 2019; Lopez-Paz & Ranzato, 2017; Chaudhry et al., 2018; Aljundi et al., 2019), or designing architectures that separate parameters across tasks (Rusu et al., 2016; Mallya & Lazebnik, 2018; Mallya et al., 2018; Yoon et al., 2017; Wortsman et al., 2020). Their aim is to develop models that can maintain previously learned skills while remaining adaptive to new experiences.

108
109

2.2 LOSS OF PLASTICITY

110
111
112
113
114
115
116

Deep learning has traditionally been studied under stationary datasets (Glorot & Bengio, 2010; He et al., 2015), yet real-world applications often involve non-stationary streams (Shen et al., 2024; Kumar et al., 2025a). Training in such environments leads to a loss of plasticity (Lyle et al., 2023; Dohare et al., 2024; Kumar et al., 2025b), where models fail to adapt to new distributions. Prior work has identified potential indicators of this phenomenon, including dormant neurons (Sokar et al., 2023; Xu et al., 2023), shifts in pre-activations (Lyle et al., 2024), feature rank collapse (Kumar et al., 2021a), and diverging weight magnitudes (Lyle et al., 2024).

117
118
119
120
121

Loss of plasticity hinders not only the ability to fit the training data, but also the ability to generalize to unseen data. Models trained incrementally often generalize worse than those trained from scratch (Ash & Adams, 2020; Berariu et al., 2021; Lyle et al., 2025), due to factors such as diminished gradient norms (Ash & Adams, 2020), weak feature changes (Lyle et al., 2025), and the compounding effects of small pretraining datasets or noisy labels (Lee et al., 2024a).

122
123
124
125
126
127
128
129

To counteract plasticity loss, reinitialization-based strategies such as S&P (Ash & Adams, 2020), DASH (Shin et al., 2024) reinitialize weights into an intermediate checkpoint, weight regularizers constrain parameters to initialization or specific subspaces (Kumar et al., 2025b; Elsayed et al., 2024; Lewandowski et al., 2024a), and spectral or rank-based approaches explicitly maintain representation quality (Kumar et al., 2021a;b; He et al., 2024). Another approach proposed reinitializing at the neuronal level, based on the utility of each neuron (Sokar et al., 2023; Dohare et al., 2024; Elsayed & Mahmood, 2024). Recent work further leverages the fact that linear networks do not suffer from plasticity loss (Dohare et al., 2024; Lewandowski et al., 2024b; Park et al., 2025).

130
131

3 METHOD

132
133
134
135
136
137

In this section we explain how we frame the stability-plasticity tradeoff as a constrained optimization problem. To do this we first need two metrics: one for stability and one for plasticity loss. With these two pieces in place the optimization problem naturally emerges, and from this formulation we introduce our method FIRE, which provides an efficient approximation to the solution.

138
139

3.1 MEASURE FOR STABILITY

140
141
142
143

To measure stability, we propose a simple yet effective metric, the *Squared Frobenius Error (SFE)*. SFE provides a natural way to quantify the preservation of learned information by comparing an original weight matrix W with its modified counterpart \widetilde{W} . Formally,

144
145

$$\text{SFE}(W, \widetilde{W}) = \|W - \widetilde{W}\|_F^2, \quad (1)$$

146
147
148
149
150
151
152
153

which captures the element-wise squared deviation between the two weight configurations. However, it remains unclear whether SFE can be used as a metric that can meaningfully capture similarity of feature representations. To clarify this point, we establish a theoretical link between SFE and the normalized feature covariance, a metric widely used in prior work to measure representation similarity (Lyle et al., 2025; Yang et al., 2022). In particular, we show that SFE provides an upper bound on the discrepancy between the normalized feature covariances of two distinct neural networks' output features (Theorem 1).

154
155
156

Theorem 1 (SFE bounds output feature covariance between two deep neural networks). *Let $\Theta = \{W^1, \dots, W^L\}$ and $\widetilde{\Theta} = \{\widetilde{W}^1, \dots, \widetilde{W}^L\}$ be the parameters of two depth- L feedforward networks with elementwise activations σ_ℓ (Lipschitz constants L_{σ_ℓ}).*

157
158
159
160
161

For an input batch $Z \in \mathbb{R}^{n \times d_0}$, we denote the layer outputs recursively by $H_\Theta^0(Z) = Z$ and $H_\Theta^\ell(Z) = \sigma_\ell(H_\Theta^{\ell-1}(Z)W^\ell)$. Let $B_\ell = \max\{\|W^\ell\|_2, \|\widetilde{W}^\ell\|_2\}$ be the maximum spectral norm of the weights in layer ℓ and $B_\Pi^\ell = \prod_{k=1}^\ell B_k$ the product across all layers. We further define $m_\ell = \min\{\|H_\Theta^\ell(Z)\|_F, \|H_{\widetilde{\Theta}}^\ell(Z)\|_F\} > 0$. The normalized feature covariances of the two networks are given by $C_\Theta^\ell(Z) = H_\Theta^\ell(Z)H_\Theta^\ell(Z)^\top / \|H_\Theta^\ell(Z)\|_F^2$, $C_{\widetilde{\Theta}}^\ell(Z) = H_{\widetilde{\Theta}}^\ell(Z)H_{\widetilde{\Theta}}^\ell(Z)^\top / \|H_{\widetilde{\Theta}}^\ell(Z)\|_F^2$.

162 Then the difference between the output feature covariances of the two networks is bounded as follows:
 163

$$164 \|C_{\Theta}^{\ell}(Z) - C_{\tilde{\Theta}}^{\ell}(Z)\|_F^2 \leq \frac{16 \|Z\|_F^2}{m_{\ell}^2} \left(\prod_{k=1}^{\ell} L_{\sigma_k} \right)^2 B_{\Pi}^2 \left(\sum_{j=1}^{\ell} \frac{1}{B_j^2} \right) \text{SFE}(\Theta, \tilde{\Theta}).$$

167 In particular, if each activation is 1-Lipschitz ($L_{\sigma_k} \leq 1$) and each spectral norm is bounded
 168 ($B_j \leq S$), then

$$169 \|C_{\Theta}^{\ell}(Z) - C_{\tilde{\Theta}}^{\ell}(Z)\|_F \leq \frac{4 \|Z\|_F}{m_{\ell}} \sqrt{\ell} S^{\ell-1} \sqrt{\text{SFE}(\Theta, \tilde{\Theta})}.$$

172 Theorem 1 shows that the discrepancy between normalized feature covariances is bounded by four
 173 factors: the input norm ($\|Z\|_F^2$), the Lipschitz constants of activation functions (L_{σ_k}), the spectral
 174 norms of the weight matrices (B_{Π} and B_j), and SFE. In practice, inputs are typically normalized and
 175 have a fixed upper bound, and commonly used activation functions have small Lipschitz constants
 176 (e.g., 1 for ReLU and tanh). Thus, the contributions from $\|Z\|_F^2$ and L_{σ_k} can be ignored. The
 177 remaining dominant factors are the spectral norms of the weight matrices (B_{Π} and B_j) and SFE.
 178 This has two implications. First, for any fixed architecture and weight scale, minimizing SFE
 179 monotonically tightens the upper bound, and thus is an effective way to preserve feature similarity
 180 between two networks. Second, the slope of the bound with respect to SFE is proportional to the
 181 spectral norms: larger spectral norms make the worst-case discrepancy more sensitive to changes in
 182 SFE. In such regimes, reducing SFE becomes even more important for stability, since even a small
 183 increase in SFE can, in principle, lead to a substantial change in the resulting feature representations.

184 3.2 MEASURE FOR PLASTICITY LOSS

185 In this work, we view plasticity loss as the situation where weights learned from previous data fail to
 186 serve as a favorable initialization point for new data. This perspective is particularly useful when
 187 designing reinitialization strategies that partially reinitialize weights before the arrival of new data.
 188 Prior analyses of plasticity loss suggest that one of its main causes is the increasing sharpness of
 189 the loss landscape curvature with respect to new data during training on the current task, which
 190 destabilizes optimization (Lyle et al., 2023). In addition, an increase in dormant neurons (Sokar
 191 et al., 2023) and a collapse in effective rank (Kumar et al., 2021a) have also been known to indicate
 192 plasticity loss.

193 However, as these measures are data-dependent and non-differentiable, they are inappropriate for the
 194 optimization objective. Instead, we propose Deviation from Isometry (DfI) (Pennington et al., 2017;
 195 Xiao et al., 2018), an optimizable metric that closely connected to previous plasticity measure.

$$196 \text{DfI}(W) = \|W^T W - I\|_F^2. \quad (2)$$

197 Our theoretical analysis reveals that minimizing DfI results in a smoother loss landscape curvature
 198 (Theorem 2), a smaller number of dormant neurons (Theorem 4), and a higher effective rank
 199 (Theorem 3). This supports the use of DfI as a suitable measure for plasticity measure. In addition,
 200 minimizing DfI also makes the weights close to isotropy, which is known as a key property of
 201 favorable neural network initializations, termed as dynamical isometry (Xiao et al., 2018).

202 **Theorem 2** (Hessian spectral norm bounded by layerwise DfIs). *We assume that*

- 204 • the inputs are whitened, so that the empirical covariance $\Sigma_Z = \frac{1}{n} Z^T Z$ is approximately
 205 the identity: $\Sigma_Z := \frac{1}{n} Z^T Z \approx I$
- 207 • for every sample i and relevant parameter vector u , the Hessian norm satisfies $|\nabla_u^2 \ell_i(u)|_2 \leq$
 208 β , while the gradient norm is bounded by $|\nabla_u \ell_i(u)|_2 \leq \gamma$.
- 209 • finally, we focus on a fixed ReLU activation pattern at the point of interest, so that the
 210 network is piecewise linear in that region. In particular, each diagonal gating matrix arising
 211 from the ReLU has operator norm ≤ 1 .

212 Let $\nu_k = 1 + \sqrt{\text{DfI}(W_k)}$, then the Hessian spectral norm can be bounded as follows:

$$214 215 \|\nabla_{\theta}^2 \mathcal{L}(W_{1:L})\|_2 \leq \beta \sum_{k=1}^L \prod_{j \neq k} \nu_j + 2\gamma \sum_{1 \leq k < \ell \leq L} \prod_{j \notin \{k, \ell\}} \nu_j.$$

Theorem 3 (DfI controls effective rank). *Let $Z \in \mathbb{R}^{n \times a}$ and $W \in \mathbb{R}^{a \times b}$ be an input and weight matrix, respectively, $\Phi = ZW$ be a feature matrix. Let $\Sigma_Z = \frac{1}{n}Z^\top Z$ be the empirical covariance of the inputs and let $W = QS$ be the right polar decomposition of W , where $Q \in \mathbb{R}^{a \times b}$ has orthonormal columns. Then $\eta_1 \geq \dots \geq \eta_d > 0$ denote the positive eigenvalues of $Q^\top \Sigma_Z Q$, with $d = \text{rank}(Q^\top \Sigma_Z Q)$, and $\text{srank}_\delta(\Phi)$ be defined from the nonzero singular values $\{\sigma_i(\Phi)\}_{i=1}^d$ by $\text{srank}_\delta(\Phi) = \min\left\{k : \left(\sum_{i=1}^k \sigma_i(\Phi) / \sum_{i=1}^d \sigma_i(\Phi)\right) \geq 1 - \delta\right\}$.*

Then $\varepsilon = \sqrt{\text{DfI}(W)} < 1$ gives a lower bound on the srank as left inequality below. If additionally $\Sigma_Z = I$, the bound simplifies to the right inequality below.

$$\text{srank}_\delta(\Phi) \geq \left\lceil \frac{(1 - \delta) d}{\delta \sqrt{\frac{1+\varepsilon}{1-\varepsilon}} \sqrt{\frac{\eta_1}{\eta_d}} + (1 - \delta)} \right\rceil, \quad \text{srank}_\delta(\Phi) \geq \left\lceil \frac{(1 - \delta) d}{\delta \sqrt{\frac{1+\varepsilon}{1-\varepsilon}} + (1 - \delta)} \right\rceil. \quad (3)$$

Theorem 4 (Minimizing DfI increases neuron activity score). *Let $W \in \mathbb{R}^{a \times b}$ and let σ be positive-homogeneous ($\sigma(\alpha t) = \alpha\sigma(t)$ for $\alpha \geq 0$). Assume the input vector $z \sim \mathcal{N}(0, I_a)$ (isotropic Gaussian). For neuron j with column w_j of W , define activity score of neuron j as $s_j = \frac{\mathbb{E}_z[|\sigma(\langle z, w_j \rangle)|]}{\frac{1}{b} \sum_{k=1}^b \mathbb{E}_z[|\sigma(\langle z, w_k \rangle)|]}$. If $\varepsilon := \sqrt{\text{DfI}(W)} < 1$, then for all j ,*

$$\sqrt{\frac{1 - \varepsilon}{1 + \varepsilon}} \leq s_j \leq \sqrt{\frac{1 + \varepsilon}{1 - \varepsilon}}.$$

Sokar et al. (2023) classified neurons with $s_j < \tau$ as dormant neurons. Hence, reducing the number of dormant neurons requires increasing the activity scores s_j . Theorem 4 states that minimizing DfI increases the lower bound $\sqrt{\frac{1 - \varepsilon}{1 + \varepsilon}} < s_j$. Note that minimizing DfI also decreases the corresponding upper bound. This is particularly meaningful because the neuron activity score s_j is a relative measure, defined as an activation normalized by the mean activation within the same layer. Consequently, it is impossible to increase all s_j simultaneously, since they are normalized by their average. To reduce dormant neurons, it is therefore critical to reduce the discrepancy in activations across neurons rather than uniformly scaling them. Theorem 4 supports this perspective: minimizing DfI tightens both the lower and upper bounds on s_j , thereby limiting the score discrepancy between neurons and effectively reduces dormant neurons.

We provided detailed proofs in Appendix A. Theorem 2, 3, 4 demonstrates that reducing DfI is directly associated with maintaining a smoother loss landscape, a high effective rank, and large number of active units.

3.3 BALANCING BETWEEN STABILITY AND PLASTICITY

To achieve low DfI while minimizing the loss of information, we formulate the problem as a constrained optimization. Specifically, we minimize the SFE between the original weights W and their orthogonalized counterpart \widetilde{W} , subject to the orthogonality constraint $\widetilde{W}^\top \widetilde{W} = I$, which is equivalent to requiring $\text{DfI}(\widetilde{W}) = \|\widetilde{W}^\top \widetilde{W} - I\|_F^2 = 0$.

$$\min_{\widetilde{W}} \|\widetilde{W} - W\|_F^2 \quad \text{s.t.} \quad \widetilde{W}^\top \widetilde{W} = I. \quad (4)$$

This formulation is mathematically equivalent to the well-studied *Orthogonal Procrustes Problem* (Schönemann, 1966), whose solution can be expressed in closed form via the polar decomposition:

$$\widetilde{W}^* = W(W^\top W)^{-\frac{1}{2}}. \quad (5)$$

While the optimization itself is classical, our contribution lies in leveraging this operation as a principled mechanism to balance stability and plasticity in neural networks. In particular, we reinterpret equation 5 as a projection that simultaneously drives the spectrum of W toward isotropy (low DfI) while maintaining stability (low SFE). We provide a derivation of Equation 5 in the Appendix A.

270 3.4 APPROXIMATING THE SOLUTION
271272 Directly computing \tilde{W}^* exactly can be expensive, we efficiently approximate it using the Newton–
273 Schulz iteration, making the approach scalable to large networks.
274275 Our method, FIRE, orthogonalizes neural network
276 weights after training on the current dataset but be-
277 fore learning on new data, using the Newton–Schulz
278 iteration (Shown in Algorithm 1.) to efficiently approx-
279 imate the above solution. Specifically, given a weight
280 matrix $W \in \mathbb{R}^{m \times n}$, we apply the Newton–Schulz up-
281 date defined as $X_{k+1} = aX_k + bX_k(X_k^\top X_k)$ with con-
282 stants $a = 1.5$ and $b = -0.5$, where $X_0 = W/\|W\|_F$.
283 This iterative process progressively drives the singular
284 values of W toward 1, thereby enforcing approximate
285 orthonormality.
286287 In convolutional layers, the update is applied kernel-wise along the spatial dimensions, ensuring that
288 each convolutional filter is orthogonalized independently.
289290 4 EXPERIMENTS
291292 To demonstrate the effectiveness of FIRE, we evaluated it on three settings: continual visual learning,
293 continual pretraining of LLMs, and reinforcement learning.
294295 4.1 CONTINUAL VISUAL LEARNING
296297 In continual visual learning experiments, we evaluated FIRE on various dataset–architecture pairs:
298 the CIFAR-10 dataset with ResNet-18, the CIFAR-100 dataset with ViT-Tiny, and the Tiny ImageNet
dataset with VGG-16.
299300 To evaluate both the ability of FIRE to recover plasticity and its capacity to restore performance after
301 a reset, we compare against two representative reset-based baselines: S&P (Ash & Adams, 2020),
302 and DASH (Shin et al., 2024). Also, to evaluate the impact of regularization on the convergence
303 speed, we adopt Parseval regularization (Chung et al., 2024), which constrains the weights to remain
304 close to orthogonal, as a baseline. We also used L2init (Kumar et al., 2025b) as a baseline, which is a
305 representative regularization-based method. Node-resetting methods such as Continual Backprop
306 (CBP) (Dohare et al., 2024), Self-Normalized Resets (SNR) (Farias & Jozefiak, 2024), and Recycling
307 Dormant Neurons (ReDo) (Sokar et al., 2023) are employed as baselines. In addition, we use the
308 Muon optimizer (Jordan et al.) as a baseline, as it also employs the Newton–Schulz iteration. While
309 FIRE minimizes the Deviation from Isometry (DfI) when resetting, Parseval regularization enforces
310 the same constraint continuously throughout training. Thus, both FIRE and Parseval regularization
311 share the same underlying optimization objective. Detailed experiment settings for continual visual
312 learning are provided in Appendix E.1.
313314 Following Ash & Adams (2020) and Lee et al. (2024a), we evaluate our approach in the warm-start
315 setting, where the model is first trained on a subset of the dataset and then on the full dataset. Since
316 plasticity loss is most severe when the subset ratio is small (Lee et al., 2024a), we warm-start with
317 only 10% of the data before continuing on the entire dataset. As shown in Figure 2 (a), FIRE provides
318 consistent performance gains across all three benchmarks. In particular, on CIFAR-10 with ResNet-18
319 and Tiny ImageNet with VGG-16, FIRE outperforms all baselines, demonstrating a strong ability
320 to recover plasticity. On CIFAR-100 with ViT-Tiny, the improvement is less pronounced. However
321 FIRE still outperforms all other baselines except DASH, and competitive to S&P. In this setting,
322 DASH, which employs a data-dependent shrinking strategy, proves especially effective, suggesting
323 that guidance from data can be beneficial when reinitializing transformer architectures. Notably,
Parseval regularization and L2init converges more slowly than FIRE, particularly on CIFAR-10
with ResNet-18 and Tiny ImageNet with VGG-16. This suggests that continuously enforcing the
constraint during training can hinder convergence, leading to longer training and incurring additional
computational cost.
3245 Algorithm 1: Newton–Schulz Iteration
(Pytorch-like)

```
# X: a two-dimensional matrix
# N: number of iteration
a, b = (1.5, -0.5)
X = X / X.norm()
for _ in range(N):
    A = X.T @ X
    X = a * X + b * (X @ A)
return X
```

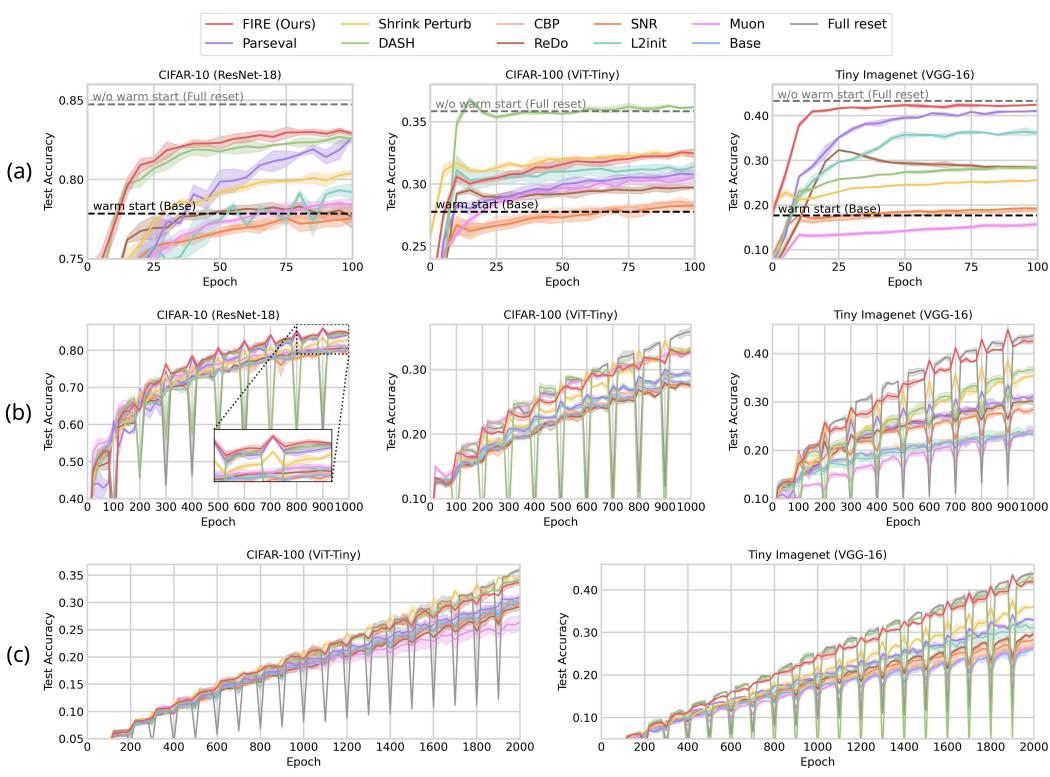


Figure 2: **Continual visual learning results.** Warm-start setting (a): training begins with only 10% of the data before continuing on the full dataset. Continual setting (b): the dataset is revealed in ten stages, expanding from 10% to 100% in 10% increments. Class-incremental setting (c): new classes are introduced over 20 phases, with an equal number of classes added at each phase.

To examine whether these findings hold in a setting where data are continuously added, which is a more realistic and natural setting, we evaluate FIRE in the continual setting (Lee et al., 2024a). Here, training is divided into ten stages, starting with 10% of the dataset and adding an additional 10% at each stage. In this way, data gradually expand from 10% to the full 100%. As shown in Figure 2 (b), FIRE delivers consistent gains across all datasets. The improvements are particularly pronounced on CIFAR-10 with ResNet-18 and Tiny ImageNet with VGG-16, while on CIFAR-100 with ViT-Tiny it achieves performance comparable to the best alternatives. In contrast, full reset and DASH suffer a sharp drop immediately after each reset, and although S&P avoids such drops, its performance remains suboptimal compared to FIRE. In contrast, FIRE incurs only a slight or negligible drop, suggesting that it successfully balances stability and plasticity, thereby achieving high performance with minimal drop in performance.

To assess the effectiveness of FIRE under large distribution shifts, we conducted experiments in a class-incremental learning scenario, which is widely used setup in the continual learning literature (Rebuffi et al., 2017; Dohare et al., 2024; Lewandowski et al., 2024b). New classes were gradually introduced at regular intervals. The training process was divided into 20 phases, with an equal number of classes added in each phase. Since CIFAR-10 does not contain a sufficient number of classes for this setting, we exclude it in this experiment. Figure 2 (c) reports the results in the class-incremental setting. Consistent with our earlier findings, FIRE shows strong performance without exhibiting a performance drop after resets by effectively balancing stability and plasticity, and full reset and DASH show sharp drop after reset while S&P show suboptimal performance.

Node-resetting methods such as CBP, SNR, and ReDo show poor overall performance. This is consistent with Lee et al. (2024a), which found that methods aiming to improve plasticity by maintaining trainability provide only limited gains in generalization. The Muon optimizer also

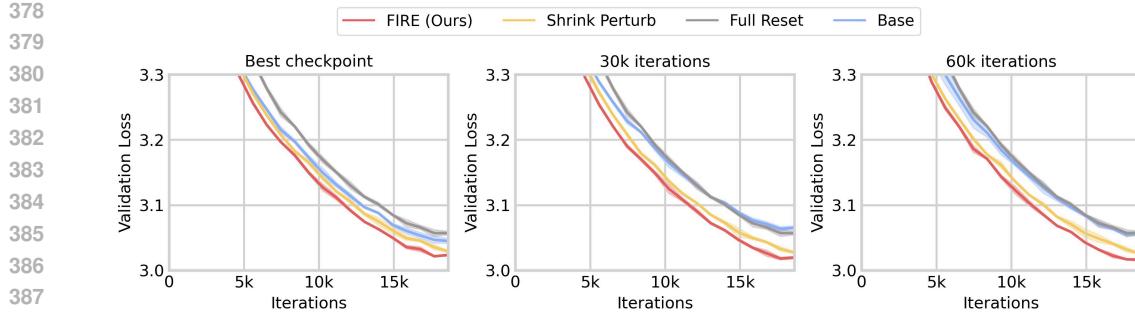


Figure 3: **Continual pretraining of GPT-0.1B.** Models are first pretrained on WikiText-103 and then continually trained on a new dataset consisting of a mixture of OpenWebText and WikiText-103. From left to right, results correspond to models initialized from the best checkpoint during pretraining, from 30k pretraining iterations, and from 60k pretraining iterations.

performs poorly overall, suggesting that periodically reinitializing weights using the Newton–Schulz iteration (FIRE) is substantially more effective than applying this iteration to the gradients (Muon).

4.2 CONTINUAL PRETRAINING OF LLMs

Setup. We also tested FIRE in the continual pretraining of LLMs. We first pretrained a GPT-0.1B model on WikiText-103 and then trained on a combination of OpenWebText and WikiText-103. For the second phase, we used the best, 30k, and 60k checkpoints from initial pretraining to examine how plasticity loss worsens beyond the best checkpoint and how effectively FIRE mitigates this degradation at different stages. We present the detailed settings for the LLM experiments in Appendix E.2.

Results. As shown in Figure 3, the gap between the base model and full reset narrows as pretraining progresses, since the base model’s validation loss increases with longer training. This aligns with prior findings that plasticity loss becomes more severe as pretraining duration grows (Ash & Adams, 2020). While S&P improves performance by moving parameters toward intermediate trade-off points between stability and plasticity, it remains suboptimal compared to FIRE, which achieves a more principled balance. Notably, FIRE was applied without any tuning, using a fixed 5 iterations, whereas S&P was carefully tuned over varying reinitialization degrees. Moreover, while the performance of the base model deteriorates with longer pretraining, FIRE maintains strong performance even when initialized from the 60k checkpoint. This demonstrates that FIRE can effectively balance the stability–plasticity trade-off even under severe plasticity loss.

In addition, unlike in continual visual learning (Section 4.1), full reset performs poorly in this setting. The main reason is the lack of stability inherent to full reset. Consequently, the full reseted model cannot outperform the base model, even though the base model itself already suffers from plasticity loss. In other words, the instability introduced by erasing all past information outweighs the potential benefit of restoring plasticity. These findings indicate that full parameter resetting is not an effective strategy for mitigating plasticity loss in continual pretraining of LLMs. Instead of providing a stable improvement, it wastes useful prior knowledge and leads to extreme inefficiency, making it an impractical choice in this setting.

4.3 REINFORCEMENT LEARNING

Setup. Finally, we evaluated FIRE in reinforcement learning. We evaluated the effectiveness of FIRE in a high Replay Ratio (RR) setting (Nikishin et al., 2022; Sokar et al., 2023), where loss of plasticity is severe and acts as a critical bottleneck for sample efficiency. For a comprehensive evaluation, we consider both continuous and discrete control environments. For discrete control, we focus on three Atari 2600 (Bellemare et al., 2013) games (Asterix, BeamRider, and DemonAttack), which have been reported to suffer from severe plasticity loss (Sokar et al., 2023). We use standard nature CNN with DQN algorithm (Mnih et al., 2015). For continuous control, we choose three primary tasks from HumanoidBench (Mnih et al., 2015): balance, walk, and run. We use SimBa (Lee et al., 2024b) with SAC algorithm (Haarnoja et al., 2018) as our baseline, whose replay ratio has failed to scale beyond

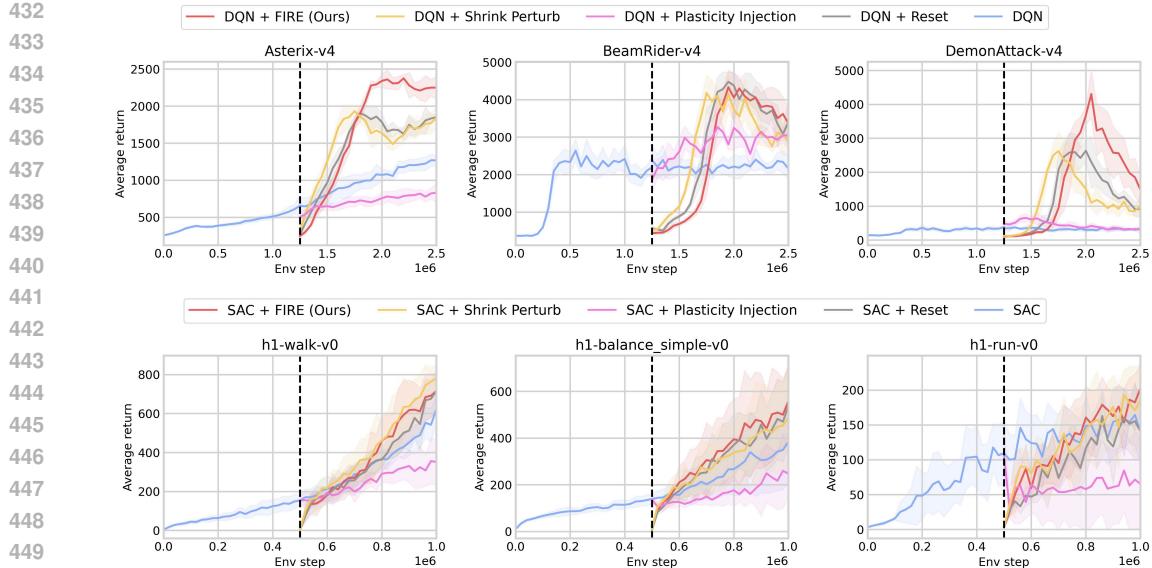


Figure 4: **Reinforcement learning results.** Discrete control with DQN on three Atari environments that suffer from severe plasticity loss (a) and continuous control with SAC on three HumanoidBench tasks (b). The black dashed line indicates the point at which reinitialization is applied.

1 without resets (Lee et al., 2025). We considered three baselines: full reset, Shrink and Perturb (S&P) (Ash & Adams, 2020; D’Oro et al., 2022), and Plasticity Injection (Nikishin et al., 2023). To eliminate performance differences caused by randomness before reinitialization, we reinitialized the network using the same checkpoint and replay buffer.

Results. As shown in Figure 4, FIRE achieves superior or competitive performance across environments compared to S&P, Plasticity Injection, and full reset. In DQN, FIRE consistently outperforms S&P, surpasses full reset in Asterix, and remains competitive in other environments. Although S&P provides a slight improvement in convergence speed, it is still suboptimal relative to both full reset and FIRE. In continuous control tasks, S&P performs competitively, but it falls short of FIRE in all Atari environments. Plasticity Injection, which introduces additional parameters to balance stability and plasticity, shows poor performance across discrete and control tasks. These results suggest that manually tuning hyperparameters to balance stability and plasticity is less effective in visual reinforcement learning—where plasticity loss is particularly severe—than our principled approach, FIRE, which explicitly balances the two.

4.4 ABLATION STUDY

To better understand the underlying factor of FIRE’s strong performance, we conducted an ablation study. To verify whether FIRE indeed effectively balances stability and plasticity, we evaluated the stability metric (SFE) and the plasticity metric (DfI), and compared FIRE against reinitialization baselines. In addition, we measured the loss landscape curvature with respect to upcoming data immediately after a reset, to examine whether our theoretical findings are also reflected in practice.

As shown in Figure 5 (b), FIRE achieves the lowest DfI while maintaining the lowest SFE, which suggests that FIRE successfully balances stability and plasticity in practice. Moreover, FIRE produces a smoother loss landscape compared to S&P, while still preserving a lower SFE. This indicates that our theoretical insights on DfI and loss curvature are indeed manifested in practice. Although DASH is particularly effective in smoothing the loss landscape, it also exhibits the highest SFE, which may contribute to an erasure of useful learned knowledge, thereby leading to instability after reset.

In addition, we evaluated FIRE with various hyperparameters in the warm-start setting to assess its sensitivity. The only hyperparameter in FIRE is the number of iterations for the Newton–Schulz iteration. As the number of iterations increases, we obtain a more accurate estimate of the solution to the constrained optimization problem discussed in Section 3.3. Therefore, our interest is to identify the minimum number of iterations that provides a sufficiently accurate estimate of the solution to yield

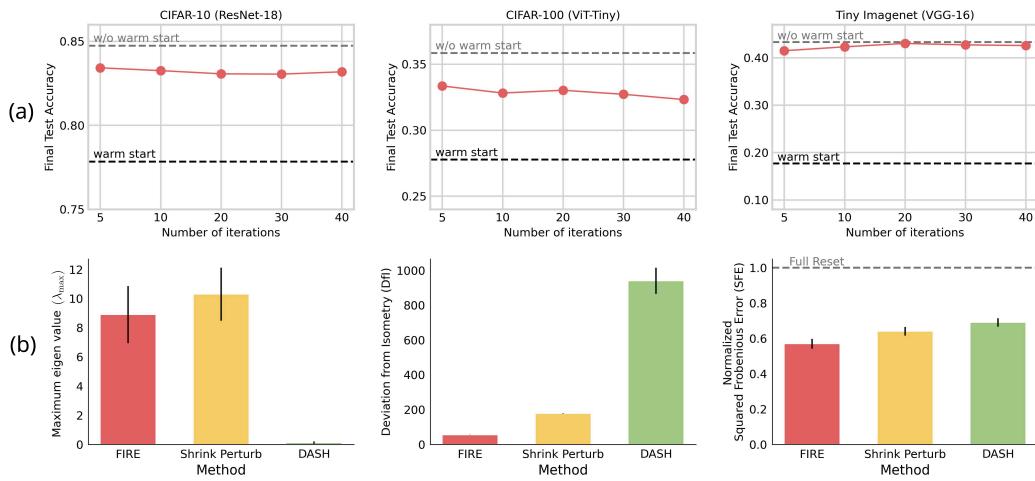


Figure 5: **Ablation study results.** Final performance of FIRE with varying numbers of iterations for Netwon-Schulz algorithm (a). Comparison of FIRE and baselines in terms of loss curvature (maximum eigenvalue of the Hessian), plasticity (DfI), and stability (normalized SFE) (b).

performance benefits. As shown in Figure 5 (a), FIRE is highly robust to the number of iterations and already provides strong performance gains even with as few as five iterations.

5 CONCLUSION

In this work, we addressed stability–plasticity trade-off, which is the long-standing problem in continual learning, by introducing FIRE. By approaching stability–plasticity tradeoff as a constrained optimization problem, FIRE enables a principled reinitialization without heavy hyperparameter tuning. Across continual visual learning, reinforcement learning, and language learning benchmarks, FIRE achieved superior or competitive performance, underscoring the importance of effective stability–plasticity management for advancing continual learning.

The main limitation of our work is the assumption of access to past data. Since our focus is on balancing stability and plasticity when such access is available, we did not evaluate FIRE under restricted data scenarios. Future work should, therefore, examine FIRE under restricted access to past data. In addition, we only used relatively small models for continual pretraining of LLMs. Evaluating FIRE on the larger models and applying FIRE not only pretraining, but also continual fine-tuning of LLMs can be a promising direction for future works.

REPRODUCIBILITY STATEMENT

We provide hyperparameter configurations and implementation details of our experiments in Appendix E. The core algorithmic part of our method is described in Algorithm 1. The proofs and assumption of theoretical works provided in this paper are described in Appendix A.

REFERENCES

Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.

Rahaf Aljundi, Eugene Belilovsky, Tinne Tuytelaars, Laurent Charlin, Massimo Caccia, Min Lin, and Lucas Page-Caccia. Online continual learning with maximal interfered retrieval. *Advances in neural information processing systems*, 32, 2019.

Jordan Ash and Ryan P Adams. On warm-starting neural network training. *Advances in neural information processing systems*, 33:3884–3894, 2020.

540 Marc G Bellemare, Yavar Naddaf, Joel Veness, and Michael Bowling. The arcade learning environment:
 541 An evaluation platform for general agents. *Journal of artificial intelligence research*, 47:
 542 253–279, 2013.

543 Tudor Berariu, Wojciech Czarnecki, Soham De, Jorg Bornschein, Samuel Smith, Razvan Pascanu, and
 544 Claudia Clopath. A study on the plasticity of neural networks. *arXiv preprint arXiv:2106.00042*,
 545 2021.

546 Jeremy Bernstein and Laker Newhouse. Modular duality in deep learning. In *Forty-second Interna-*
 547 *tional Conference on Machine Learning*, 2025. URL <https://openreview.net/forum?id=hErdffTsLu>.

548 Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal,
 549 Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are
 550 few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.

551 Arslan Chaudhry, Marc’Aurelio Ranzato, Marcus Rohrbach, and Mohamed Elhoseiny. Efficient
 552 lifelong learning with a-gem. *arXiv preprint arXiv:1812.00420*, 2018.

553 Wesley Chung, Lynn Cherif, David Meger, and Doina Precup. Parseval regularization for continual
 554 reinforcement learning. *Advances in Neural Information Processing Systems*, 37:127937–127967,
 555 2024.

556 Shibhansh Dohare, J Fernando Hernandez-Garcia, Qingfeng Lan, Parash Rahman, A Rupam Mah-
 557 mood, and Richard S Sutton. Loss of plasticity in deep continual learning. *Nature*, 632(8026):
 558 768–774, 2024.

559 Pierluca D’Oro, Max Schwarzer, Evgenii Nikishin, Pierre-Luc Bacon, Marc G Bellemare, and Aaron
 560 Courville. Sample-efficient reinforcement learning by breaking the replay ratio barrier. In *Deep*
 561 *Reinforcement Learning Workshop NeurIPS 2022*, 2022.

562 Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas
 563 Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An
 564 image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint*
 565 *arXiv:2010.11929*, 2020.

566 Mohamed Elsayed and A Rupam Mahmood. Addressing loss of plasticity and catastrophic forgetting
 567 in continual learning. *arXiv preprint arXiv:2404.00781*, 2024.

568 Mohamed Elsayed, Qingfeng Lan, Clare Lyle, and A Rupam Mahmood. Weight clipping for deep
 569 continual and reinforcement learning. *arXiv preprint arXiv:2407.01704*, 2024.

570 Vivek F Farias and Adam D Jozefiak. Self-normalized resets for plasticity in continual learning.
 571 *arXiv preprint arXiv:2410.20098*, 2024.

572 Scott Fujimoto, Herke Hoof, and David Meger. Addressing function approximation error in actor-
 573 critic methods. In *International conference on machine learning*, pp. 1587–1596. PMLR, 2018.

574 Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feedforward neural
 575 networks. In *Proceedings of the thirteenth international conference on artificial intelligence and*
 576 *statistics*, pp. 249–256. JMLR Workshop and Conference Proceedings, 2010.

577 Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy
 578 maximum entropy deep reinforcement learning with a stochastic actor. In *International conference*
 579 *on machine learning*, pp. 1861–1870. PMLR, 2018.

580 Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing
 581 human-level performance on imagenet classification. In *Proceedings of the IEEE international*
 582 *conference on computer vision*, pp. 1026–1034, 2015.

583 Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image
 584 recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*,
 585 pp. 770–778, 2016.

594 Qiang He, Tianyi Zhou, Meng Fang, and Setareh Maghsudi. Adaptive regularization of representation
 595 rank as an implicit constraint of bellman equation. *arXiv preprint arXiv:2404.12754*, 2024.
 596

597 Peter J Huber. Robust estimation of a location parameter. In *Breakthroughs in statistics: Methodology*
 598 and distribution

599 pp. 492–518. Springer, 1992.

600 Keller Jordan, Yuchen Jin, Vlado Boza, You Jiacheng, Franz Cecista, Laker Newhouse, and
 601 Jeremy Bernstein. Muon: An optimizer for hidden layers in neural networks, 2024. URL
 602 <https://kellerjordan.github.io/posts/muon>, 6.

603 Andrej Karpathy. NanoGPT. <https://github.com/karpathy/nanoGPT>, 2023.

604 Zixuan Ke, Yijia Shao, Haowei Lin, Tatsuya Konishi, Gyuhak Kim, and Bing Liu. Continual
 605 pre-training of language models. In *The Eleventh International Conference on Learning Represen-*
 606 *tations*, 2023.

607 Dongwan Kim and Bohyung Han. On the stability-plasticity dilemma of class-incremental learning.
 608 In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp.
 609 20196–20204, 2023.

610 Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint*
 611 *arXiv:1412.6980*, 2014.

612 James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A
 613 Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. Overcoming
 614 catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*, 114
 615 (13):3521–3526, 2017.

616 Aviral Kumar, Rishabh Agarwal, Dibya Ghosh, and Sergey Levine. Implicit under-parameterization
 617 inhibits data-efficient deep reinforcement learning. In *International Conference on Learning
 618 Representations*, 2021a.

619 Aviral Kumar, Rishabh Agarwal, Tengyu Ma, Aaron Courville, George Tucker, and Sergey Levine.
 620 Dr3: Value-based deep reinforcement learning requires explicit regularization. *arXiv preprint*
 621 *arXiv:2112.04716*, 2021b.

622 Saurabh Kumar, Henrik Marklund, Ashish Rao, Yifan Zhu, Hong Jun Jeon, Yueyang Liu, Ben-
 623 jamin Van Roy, et al. Continual learning as computationally constrained reinforcement learning.
 624 *Foundations and Trends® in Machine Learning*, 18(5):913–1053, 2025a.

625 Saurabh Kumar, Henrik Marklund, and Benjamin Van Roy. Maintaining plasticity in continual
 626 learning via regenerative regularization. In *Conference on Lifelong Learning Agents*, pp. 410–430.
 627 PMLR, 2025b.

628 Hojoon Lee, Hyeonseo Cho, Hyunseung Kim, Donghu Kim, Dugki Min, Jaegul Choo, and Clare
 629 Lyle. Slow and steady wins the race: Maintaining plasticity with hare and tortoise networks. In
 630 *International Conference on Machine Learning*, pp. 26416–26438. PMLR, 2024a.

631 Hojoon Lee, Dongyoong Hwang, Donghu Kim, Hyunseung Kim, Jun Jet Tai, Kaushik Subramanian,
 632 Peter R. Wurman, Jaegul Choo, Peter Stone, and Takuma Seno. Simba: Simplicity bias for scaling
 633 up parameters in deep reinforcement learning. *arXiv preprint arXiv:2410.09754*, 2024b.

634 Hojoon Lee, Youngdo Lee, Takuma Seno, Donghu Kim, Peter Stone, and Jaegul Choo. Hyperspherical
 635 normalization for scalable deep reinforcement learning. *arXiv preprint arXiv:2502.15280*, 2025.

636 Alex Lewandowski, Michał Bortkiewicz, Saurabh Kumar, András György, Dale Schuurmans, Mateusz
 637 Ostaszewski, and Marlos C Machado. Learning continually by spectral regularization. *arXiv*
 638 *preprint arXiv:2406.06811*, 2024a.

639 Alex Lewandowski, Dale Schuurmans, and Marlos C Machado. Plastic learning with deep fourier
 640 features. *arXiv preprint arXiv:2410.20634*, 2024b.

641 David Lopez-Paz and Marc’Aurelio Ranzato. Gradient episodic memory for continual learning.
 642 *Advances in neural information processing systems*, 30, 2017.

648 Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint*
 649 *arXiv:1711.05101*, 2017.

650

651 Clare Lyle, Mark Rowland, and Will Dabney. Understanding and preventing capacity loss in
 652 reinforcement learning. In *International Conference on Learning Representations*, 2022.

653

654 Clare Lyle, Zeyu Zheng, Evgenii Nikishin, Bernardo Avila Pires, Razvan Pascanu, and Will Dabney.
 655 Understanding plasticity in neural networks. In *International Conference on Machine Learning*,
 656 pp. 23190–23211. PMLR, 2023.

657

658 Clare Lyle, Zeyu Zheng, Khimya Khetarpal, Hado van Hasselt, Razvan Pascanu, James Martens,
 659 and Will Dabney. Disentangling the causes of plasticity loss in neural networks. *arXiv preprint*
 660 *arXiv:2402.18762*, 2024.

661

662 Clare Lyle, Gharda Sokar, Razvan Pascanu, and Andras Gyorgy. What can grokking teach us about
 663 learning under nonstationarity? *arXiv preprint arXiv:2507.20057*, 2025.

664

665 Guozheng Ma, Lu Li, Sen Zhang, Zixuan Liu, Zhen Wang, Yixin Chen, Li Shen, Xueqian Wang, and
 666 Dacheng Tao. Revisiting plasticity in visual reinforcement learning: Data, modules and training
 667 stages. *arXiv preprint arXiv:2310.07418*, 2023.

668

669 Arun Mallya and Svetlana Lazebnik. Packnet: Adding multiple tasks to a single network by iterative
 670 pruning. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pp.
 671 7765–7773, 2018.

672

673 Arun Mallya, Dillon Davis, and Svetlana Lazebnik. Piggyback: Adapting a single network to multiple
 674 tasks by learning to mask weights. In *Proceedings of the European conference on computer vision*
 675 (*ECCV*), pp. 67–82, 2018.

676

677 Martial Mermilliod, Aurélia Bugaiska, and Patrick Bonin. The stability-plasticity dilemma: Investi-
 678 gating the continuum from catastrophic forgetting to age-limited learning effects, 2013.

679

680 Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare,
 681 Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control
 682 through deep reinforcement learning. *nature*, 518(7540):529–533, 2015.

683

684 Evgenii Nikishin, Max Schwarzer, Pierluca D’Oro, Pierre-Luc Bacon, and Aaron Courville. The
 685 primacy bias in deep reinforcement learning. In *International conference on machine learning*, pp.
 686 16828–16847. PMLR, 2022.

687

688 Evgenii Nikishin, Junhyuk Oh, Georg Ostrovski, Clare Lyle, Razvan Pascanu, Will Dabney, and André
 689 Barreto. Deep reinforcement learning with plasticity injection. *Advances in Neural Information
 690 Processing Systems*, 36:37142–37159, 2023.

691

692 Sangyeon Park, Isaac Han, Seungwon Oh, and KyungJoong Kim. Activation by interval-wise dropout:
 693 A simple way to prevent neural networks from plasticity loss. In *Forty-second International
 694 Conference on Machine Learning*, 2025.

695

696 Jeffrey Pennington, Samuel Schoenholz, and Surya Ganguli. Resurrecting the sigmoid in deep
 697 learning through dynamical isometry: theory and practice. *Advances in neural information
 698 processing systems*, 30, 2017.

699

700 Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe
 701 Penna, and Robin Rombach. Sdxl: Improving latent diffusion models for high-resolution image
 702 synthesis. In *The Twelfth International Conference on Learning Representations*, 2024.

703

704 Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H Lampert. icarl:
 705 Incremental classifier and representation learning. In *Proceedings of the IEEE conference on*
 706 *Computer Vision and Pattern Recognition*, pp. 2001–2010, 2017.

707

708 David Rolnick, Arun Ahuja, Jonathan Schwarz, Timothy Lillicrap, and Gregory Wayne. Experience
 709 replay for continual learning. *Advances in neural information processing systems*, 32, 2019.

702 Andrei A Rusu, Neil C Rabinowitz, Guillaume Desjardins, Hubert Soyer, James Kirkpatrick, Koray
 703 Kavukcuoglu, Razvan Pascanu, and Raia Hadsell. Progressive neural networks. *arXiv preprint*
 704 *arXiv:1606.04671*, 2016.

705

706 Peter H Schönemann. A generalized solution of the orthogonal procrustes problem. *Psychometrika*,
 707 31(1):1–10, 1966.

708 Carmelo Sferrazza, Dun-Ming Huang, Xingyu Lin, Youngwoon Lee, and Pieter Abbeel. Humanoid-
 709 bench: Simulated humanoid benchmark for whole-body locomotion and manipulation, 2024.

710

711 Maying Shen, Hongxu Yin, Pavlo Molchanov, Lei Mao, and Jose M Alvarez. Step out and seek
 712 around: On warm-start training with incremental data. *arXiv preprint arXiv:2406.04484*, 2024.

713 Baekrok Shin, Junsoo Oh, Hanseul Cho, and Chulhee Yun. Dash: Warm-starting neural network
 714 training in stationary settings without loss of plasticity. *Advances in Neural Information Processing*
 715 *Systems*, 37:43300–43340, 2024.

716

717 Ghada Sokar, Rishabh Agarwal, Pablo Samuel Castro, and Utku Evci. The dormant neuron phe-
 718 nomenon in deep reinforcement learning. In *International Conference on Machine Learning*, pp.
 719 32145–32168. PMLR, 2023.

720 Gemini Robotics Team, Saminda Abeyruwan, Joshua Ainslie, Jean-Baptiste Alayrac, Montser-
 721 rat Gonzalez Arenas, Travis Armstrong, Ashwin Balakrishna, Robert Baruch, Maria Bauza,
 722 Michiel Blokzijl, et al. Gemini robotics: Bringing ai into the physical world. *arXiv preprint*
 723 *arXiv:2503.20020*, 2025.

724

725 Eli Verwimp, Kuo Yang, Sarah Parisot, Lanqing Hong, Steven McDonagh, Eduardo Pérez-Pellitero,
 726 Matthias De Lange, and Tinne Tuytelaars. Clad: A realistic continual learning benchmark for
 727 autonomous driving. *Neural Networks*, 161:659–669, 2023.

728

729 Maciej Wołczyk, Michał Zajac, Razvan Pascanu, Łukasz Kuciński, and Piotr Miłoś. Continual
 730 world: A robotic benchmark for continual reinforcement learning. *Advances in Neural Information*
Processing Systems, 34:28496–28510, 2021.

731

732 Mitchell Wortsman, Vivek Ramanujan, Rosanne Liu, Aniruddha Kembhavi, Mohammad Rastegari,
 733 Jason Yosinski, and Ali Farhadi. Supermasks in superposition. *Advances in neural information*
processing systems, 33:15173–15184, 2020.

734

735 Lechao Xiao, Yasaman Bahri, Jascha Sohl-Dickstein, Samuel Schoenholz, and Jeffrey Pennington.
 736 Dynamical isometry and a mean field theory of cnns: How to train 10,000-layer vanilla convolu-
 737 tional neural networks. In *International conference on machine learning*, pp. 5393–5402. PMLR,
 738 2018.

739

740 Guowei Xu, Ruijie Zheng, Yongyuan Liang, Xiyao Wang, Zhecheng Yuan, Tianying Ji, Yu Luo,
 741 Xiaoyu Liu, Jiaxin Yuan, Pu Hua, et al. Drm: Mastering visual reinforcement learning through
 742 dormant ratio minimization. *arXiv preprint arXiv:2310.19668*, 2023.

743

744 Greg Yang, Edward J Hu, Igor Babuschkin, Szymon Sidor, Xiaodong Liu, David Farhi, Nick Ryder,
 745 Jakub Pachocki, Weizhu Chen, and Jianfeng Gao. Tensor programs v: Tuning large neural networks
 746 via zero-shot hyperparameter transfer. *arXiv preprint arXiv:2203.03466*, 2022.

747

748 Jaehong Yoon, Eunho Yang, Jeongtae Lee, and Sung Ju Hwang. Lifelong learning with dynamically
 749 expandable networks. *arXiv preprint arXiv:1708.01547*, 2017.

750

751

752

753

754

755

A PROOF OF THEOREMS

Proof of Theorem 1

Proof. Unless explicitly subscripted, $\|\cdot\|$ denotes the Frobenius norm $\|\cdot\|_F$, and $\|\cdot\|_2$ denotes the spectral (operator) norm. We also recall

$$\text{SFE}(\Theta, \tilde{\Theta}) := \sum_{j=1}^L \|W^j - \tilde{W}^j\|_F^2.$$

We will use two elementary facts.

For any F, \tilde{F} with $m := \min\{\|F\|, \|\tilde{F}\|\} > 0$,

$$\left\| \frac{F}{\|F\|} - \frac{\tilde{F}}{\|\tilde{F}\|} \right\| \leq \frac{2}{m} \|F - \tilde{F}\|.$$

If arbitrary two matrix U and V satisfies $\|U\| = \|V\| = 1$ then

$$\|UU^\top - VV^\top\| \leq 2\|U - V\|.$$

Combining these two facts with $U := F/\|F\|$ and $V := \tilde{F}/\|\tilde{F}\|$ yields

$$\left\| \frac{FF^\top}{\|F\|^2} - \frac{\tilde{F}\tilde{F}^\top}{\|\tilde{F}\|^2} \right\| \leq \frac{4}{m} \|F - \tilde{F}\|. \quad (6)$$

Applying equation 6 with $F = H_\Theta^\ell(Z)$ and $\tilde{F} = H_{\tilde{\Theta}}^\ell(Z)$, and $m_\ell := \min\{\|H_\Theta^\ell(Z)\|, \|H_{\tilde{\Theta}}^\ell(Z)\|\}$ gives

$$\|C_\Theta^\ell(Z) - C_{\tilde{\Theta}}^\ell(Z)\| \leq \frac{4}{m_\ell} \|H_\Theta^\ell(Z) - H_{\tilde{\Theta}}^\ell(Z)\|. \quad (7)$$

Introduce the *hybrid* outputs $H_{(\leq j)}^\ell(Z)$: the first j layers use Θ and the remaining $j+1, \dots, \ell$ layers use $\tilde{\Theta}$. Note that $H_{(\leq \ell)}^\ell(Z) = H_\Theta^\ell(Z)$ and $H_{(\leq 0)}^\ell(Z) = H_{\tilde{\Theta}}^\ell(Z)$. Then

$$H_\Theta^\ell(Z) - H_{\tilde{\Theta}}^\ell(Z) = \sum_{j=1}^{\ell} \left(H_{(\leq j)}^\ell(Z) - H_{(\leq j-1)}^\ell(Z) \right),$$

so by the triangle inequality,

$$\|H_\Theta^\ell(Z) - H_{\tilde{\Theta}}^\ell(Z)\| \leq \sum_{j=1}^{\ell} \|H_{(\leq j)}^\ell(Z) - H_{(\leq j-1)}^\ell(Z)\|. \quad (8)$$

Each summand differs in *only the j -th layer weights*. Let $X_{j-1} := H_{\tilde{\Theta}}^{j-1}(Z)$ denote the shared input fed to layer j in both hybrids. Consider the *backend subnetwork*

$$\mathcal{T}_{j \rightarrow \ell}^{(\tilde{\Theta})}(Y) := \sigma_\ell(\dots \sigma_{j+1}(Y \tilde{W}^{j+1}) \dots \tilde{W}^\ell).$$

For arbitrary Y_1 and Y_2 , its input-Lipschitz constant is

$$\|\mathcal{T}_{j \rightarrow \ell}^{(\tilde{\Theta})}(Y_1) - \mathcal{T}_{j \rightarrow \ell}^{(\tilde{\Theta})}(Y_2)\| \leq \left(\prod_{k=j+1}^{\ell} L_{\sigma_k} \|\tilde{W}^k\|_2 \right) \|Y_1 - Y_2\|. \quad (9)$$

Hence

$$\|H_{(\leq j)}^\ell(Z) - H_{(\leq j-1)}^\ell(Z)\| \leq \left(\prod_{k=j+1}^{\ell} L_{\sigma_k} \|\tilde{W}^k\|_2 \right) L_{\sigma_j} \|X_{j-1}\| \|W^j - \tilde{W}^j\|. \quad (10)$$

810 Meanwhile,

811

$$812 \|X_{j-1}\| = \|H_{\Theta}^{j-1}(Z)\| \leq \|Z\| \prod_{k=1}^{j-1} L_{\sigma_k} \|W^k\|_2. \quad (11)$$

813

814

815 Combining equation 8, equation 10, and equation 11 yields

816

$$817 \|H_{\Theta}^{\ell}(Z) - H_{\tilde{\Theta}}^{\ell}(Z)\| \leq \|Z\| \left(\prod_{k=1}^{\ell} L_{\sigma_k} \right) \sum_{j=1}^{\ell} \left(\prod_{k \neq j} B_k \right) \|W^j - \tilde{W}^j\|. \quad (12)$$

818

819 By Cauchy–Schwarz,

820

821

$$822 \sum_{j=1}^{\ell} \left(\prod_{k \neq j} B_k \right) \|W^j - \tilde{W}^j\| \leq B_{\Pi}^{\ell} \left(\sum_{j=1}^{\ell} \frac{1}{B_j^2} \right)^{1/2} \sqrt{\text{SFE}(\Theta, \tilde{\Theta})}.$$

823

824 Therefore,

825

826

$$827 \|H_{\Theta}^{\ell}(Z) - H_{\tilde{\Theta}}^{\ell}(Z)\| \leq \|Z\| \left(\prod_{k=1}^{\ell} L_{\sigma_k} \right) B_{\Pi}^{\ell} \left(\sum_{j=1}^{\ell} \frac{1}{B_j^2} \right)^{1/2} \sqrt{\text{SFE}(\Theta, \tilde{\Theta})}. \quad (12)$$

828

829 Substitute equation 12 into equation 7:

830

831

$$832 \|C_{\Theta}^{\ell}(Z) - C_{\tilde{\Theta}}^{\ell}(Z)\| \leq \frac{4\|Z\|}{m_{\ell}} \left(\prod_{k=1}^{\ell} L_{\sigma_k} \right) B_{\Pi}^{\ell} \left(\sum_{j=1}^{\ell} \frac{1}{B_j^2} \right)^{1/2} \sqrt{\text{SFE}(\Theta, \tilde{\Theta})}.$$

833

834 Squaring both sides gives

835

836

$$837 \|C_{\Theta}^{\ell} - C_{\tilde{\Theta}}^{\ell}\|^2 \leq \frac{16\|Z\|^2}{m_{\ell}^2} \left(\prod_{k=1}^{\ell} L_{\sigma_k} \right)^2 (B_{\Pi}^{\ell})^2 \left(\sum_{j=1}^{\ell} \frac{1}{B_j^2} \right) \text{SFE}(\Theta, \tilde{\Theta}).$$

838

839 If $L_{\sigma_k} \leq 1$ and $B_j \leq S$ for all j , then $(B_{\Pi}^{\ell})^2 \sum_{j=1}^{\ell} B_j^{-2} \leq \ell S^{2\ell-2}$. Therefore

840

841

$$842 \|C_{\Theta}^{\ell} - C_{\tilde{\Theta}}^{\ell}\| \leq \frac{4\|Z\|}{m_{\ell}} \sqrt{\ell} S^{\ell-1} \sqrt{\text{SFE}(\Theta, \tilde{\Theta})}.$$

843

844 \square

845 **Network, Loss, and Notation.** Let $Z \in \mathbb{R}^{n \times d_0}$ be the input matrix and $W_k \in \mathbb{R}^{d_{k-1} \times d_k}$ ($k = 1, \dots, L$) be the weight matrices. Define

846

847

$$848 A_k = H_{k-1} W_k \in \mathbb{R}^{n \times d_k}, \quad H_k = \rho(A_k) \quad (\rho = \text{ReLU}), \quad U := A_L \in \mathbb{R}^{n \times d_L},$$

849

850 with $H_0 := Z$. The empirical risk is

851

852

$$853 \mathcal{L}(W_{1:L}) = \frac{1}{n} \sum_{i=1}^n \ell_i(u_i), \quad u_i \in \mathbb{R}^{d_L}.$$

854

855 Let $\theta = \text{vec}(W_1, \dots, W_L) \in \mathbb{R}^p$ and $H_{\theta} := \nabla_{\theta}^2 \mathcal{L} \in \mathbb{R}^{p \times p}$. We use $\|\cdot\|_2$ for spectral norm and

856 $\|\cdot\|_F$ for the Frobenius norm.

857

858 We denote the maximum eigen values as λ_{\max}

864 **Deviation From Isometry (Dfl).** For a matrix W , $\text{Dfl}(W) := \|W^\top W - I\|_F^2$. For each layer, set

$$865 \quad \nu_k := 1 + \sqrt{\text{Dfl}(W_k)}, \quad \alpha_k := \sqrt{\nu_k}.$$

868 First, let us examine the lemmas required for the proof of Theorem 2.

869 **Lemma 1** (Dfl controls the spectral norm). *For each layer k , $\|W_k\|_2^2 \leq \nu_k$ and $\|W_k\|_2 \leq \alpha_k$.*

871 *Proof.* $\|W\|_2^2 = \lambda_{\max}(W^\top W) \leq \lambda_{\max}(W^\top W - I) + 1 \leq \|W^\top W - I\|_2 + 1 \leq \|W^\top W - I\|_F +$
 872 $1 = 1 + \sqrt{\text{Dfl}(W)}$. \square

873 **Lemma 2** (Covariance/spectral growth through layers). *Let $\Sigma_{H_k} = \frac{1}{n} H_k^\top H_k$. Then*

$$875 \quad \lambda_{\max}(\Sigma_{H_k}) \leq \lambda_{\max}(\Sigma_{H_{k-1}}) \|W_k\|_2^2 \leq \lambda_{\max}(\Sigma_{H_{k-1}}) \nu_k.$$

877 *Consequently, with $\Sigma_{H_0} = \Sigma_Z$,*

$$878 \quad \lambda_{\max}(\Sigma_{H_k}) \leq \lambda_{\max}(\Sigma_Z) \prod_{j=1}^k \nu_j.$$

881 *In particular, under $\Sigma_Z \approx I$ we have $\lambda_{\max}(\Sigma_{H_k}) \leq \prod_{j=1}^k \nu_j$.*

884 *Proof.* ReLU is 1-Lipschitz (applied elementwise), hence $\|H_k\|_2 \leq \|A_k\|_2 \leq \|H_{k-1}\|_2 \|W_k\|_2$.
 885 Therefore $\frac{1}{n} \|H_k\|_2^2 \leq \frac{1}{n} \|H_{k-1}\|_2^2 \|W_k\|_2^2$, i.e., $\lambda_{\max}(\Sigma_{H_k}) \leq \lambda_{\max}(\Sigma_{H_{k-1}}) \|W_k\|_2^2$. Apply
 886 Lemma 1. \square

887 **Lemma 3** (Block-Jacobian bound). *Let $J \in \mathbb{R}^{(nd_L) \times p}$ be the Jacobian of $\text{vec}(U)$ w.r.t. θ , and
 888 $J = [J_1 \ J_2 \ \dots \ J_L]$ the block-columns corresponding to $\text{vec}(W_k)$. Then*

$$889 \quad \frac{1}{n} \|J_k\|_2^2 \leq \left(\prod_{j=1}^{k-1} \nu_j \right) \left(\prod_{j=k+1}^L \nu_j \right) = \prod_{j \neq k} \nu_j.$$

892 *Consequently,*

$$894 \quad \frac{1}{n} \|J\|_2^2 \leq \sum_{k=1}^L \frac{1}{n} \|J_k\|_2^2 \leq \sum_{k=1}^L \prod_{j \neq k} \nu_j.$$

897 *Proof.* Consider a perturbation ΔW_k . Without loss of generality, we may assume $\|\Delta W_k\|_F = 1$,
 898 since the operator norm is defined by the supremum over unit perturbations. With fixed ReLU gates
 899 (op. norm ≤ 1), the output perturbation over all n samples satisfies

$$900 \quad \Delta U = H_{k-1} \Delta W_k B_{k+1:L}, \quad B_{k+1:L} := \underbrace{D_k W_{k+1} D_{k+1} \cdots D_{L-1}}_{\text{diag gates, } \|\cdot\|_2 \leq 1} W_L.$$

903 Thus $\|\Delta U\|_F \leq \|H_{k-1}\|_2 \|\Delta W_k\|_F \|B_{k+1:L}\|_2$, so the operator norm of the linear map $\Delta W_k \mapsto$
 904 ΔU is at most $\|H_{k-1}\|_2 \|B_{k+1:L}\|_2$. Hence $\|J_k\|_2 \leq \|H_{k-1}\|_2 \|B_{k+1:L}\|_2$ and

$$906 \quad \frac{1}{n} \|J_k\|_2^2 \leq \frac{1}{n} \|H_{k-1}\|_2^2 \|B_{k+1:L}\|_2^2 = \lambda_{\max}(\Sigma_{H_{k-1}}) \|B_{k+1:L}\|_2^2.$$

908 Using Lemma 2, $\lambda_{\max}(\Sigma_{H_{k-1}}) \leq \prod_{j=1}^{k-1} \nu_j$. Also $\|B_{k+1:L}\|_2 \leq \prod_{j=k+1}^L \|W_j\|_2 \leq \prod_{j=k+1}^L \alpha_j$,
 909 thus $\|B_{k+1:L}\|_2^2 \leq \prod_{j=k+1}^L \nu_j$. Multiplying the two bounds yields the claim. Finally, since J is a
 910 horizontal concatenation of blocks, $\|J\|_2^2 \leq \sum_k \|J_k\|_2^2$. \square

912 **Lemma 4** (Gauss–Newton part). *For each sample, $\nabla_\theta^2 \ell_i = J_i^\top (\nabla_u^2 \ell_i) J_i + R_i$ with some remainder
 913 R_i . By (A2), $\|\nabla_u^2 \ell_i\|_2 \leq \beta$, hence*

$$915 \quad \left\| \frac{1}{n} \sum_{i=1}^n J_i^\top (\nabla_u^2 \ell_i) J_i \right\|_2 \leq \frac{\beta}{n} \|J\|_2^2 \leq \beta \sum_{k=1}^L \prod_{j \neq k} \nu_j,$$

917 *where the last inequality uses Lemma 3.*

918 **Lemma 5** (Remainder term). *Let $R := \frac{1}{n} \sum_{i=1}^n R_i$. Under (A2) and (A3),*

$$919 \quad \|R\|_2 \leq 2\gamma \sum_{1 \leq k < \ell \leq L} \prod_{j \notin \{k, \ell\}} \nu_j.$$

920 *Proof.* Fix the ReLU gates locally (piecewise linear region) and $\sum_k \|\Delta W_k\|_F = 1$, since the
921 operator norm is defined by the supremum over unit perturbations. Then the network output U is
922 *multilinear* in $\{W_k\}_{k=1}^L$. For $k < \ell$, the mixed second derivative block maps $(\Delta W_k, \Delta W_\ell)$ to
923

$$924 \quad H_{k-1} \Delta W_k C_{k+1:\ell-1} \Delta W_\ell B_{\ell+1:L},$$

925 where $C_{k+1:\ell-1}$ is the product of intermediate gated weights, and $B_{\ell+1:L}$ the tail product as in
926 Lemma 3. By submultiplicativity,
927

$$928 \quad \|H_{k-1} \Delta W_k C_{k+1:\ell-1} \Delta W_\ell B_{\ell+1:L}\|_F \leq \|H_{k-1}\|_2 \|\Delta W_k\|_F \|C_{k+1:\ell-1}\|_2 \|\Delta W_\ell\|_F \|B_{\ell+1:L}\|_2.$$

929 Using Lemma 2 and Lemma 1,

$$930 \quad \|H_{k-1}\|_2 \leq \sqrt{n} \prod_{j=1}^{k-1} \alpha_j, \quad \|C_{k+1:\ell-1}\|_2 \leq \prod_{j=k+1}^{\ell-1} \alpha_j, \quad \|B_{\ell+1:L}\|_2 \leq \prod_{j=\ell+1}^L \alpha_j.$$

931 Therefore, after dividing by n (from the prefactor $1/n$ in \mathcal{L}) and summing the symmetric contribution
932 (ℓ, k) , the bilinear remainder contributes at most
933

$$934 \quad \frac{2}{n} \sum_{k < \ell} \|H_{k-1}\|_2 \|C_{k+1:\ell-1}\|_2 \|B_{\ell+1:L}\|_2 \|\Delta W_k\|_F \|\Delta W_\ell\|_F \leq 2 \sum_{k < \ell} \left(\prod_{j \notin \{k, \ell\}} \alpha_j^2 \right) \|\Delta W_k\|_F \|\Delta W_\ell\|_F.$$

935 Finally, by $2ab \leq a^2 + b^2$ and $\sum_k \|\Delta W_k\|_F^2 = 1$ (unit parameter direction), the operator norm
936 of the second-derivative map is bounded by $\sum_{k < \ell} \prod_{j \notin \{k, \ell\}} \nu_j$. Multiplying by $\|\nabla_u \ell_i\|_2 \leq \gamma$ and
937 averaging over i gives the claim. \square

944 Proof of Theorem 2

945 *Proof.* Combine Lemma 4 and Lemma 5 and use $\|A + B\|_2 \leq \|A\|_2 + \|B\|_2$. \square

946 **Corollary 1** (Near-interpolation or small-gradient regime). *If the training gradients at the outputs
947 are small so that $\gamma \approx 0$, then*

$$948 \quad \|\nabla_\theta^2 \mathcal{L}(W_{1:L})\|_2 \lesssim \beta \sum_{k=1}^L \prod_{j \neq k} (1 + \sqrt{\text{Dfl}(W_j)}).$$

949 Next, we present the lemma required for the proof of Theorem 3.

950 **Lemma 6** (A basic spectral lemma from Dfl). *Let $\varepsilon = \sqrt{\text{Dfl}(W)}$. Then*

$$951 \quad \|W^\top W - I\|_2 \leq \|W^\top W - I\|_F = \varepsilon,$$

952 *hence every eigenvalue μ of $W^\top W = S^2$ satisfies $1 - \varepsilon \leq \mu \leq 1 + \varepsilon$. Equivalently,*

$$953 \quad \sqrt{1 - \varepsilon} I \preceq S \preceq \sqrt{1 + \varepsilon} I.$$

954 *Proof.* By the definition of ε , we have

$$955 \quad \|W^\top W - I\|_2 \leq \|W^\top W - I\|_F = \varepsilon.$$

956 Therefore, all eigenvalues μ of $W^\top W$ lie within the interval

$$957 \quad 1 - \varepsilon \leq \mu \leq 1 + \varepsilon.$$

958 Since $W^\top W = S^2$ with $S \succeq 0$, this is equivalent to the spectral bound

$$959 \quad \sqrt{1 - \varepsilon} I \preceq S \preceq \sqrt{1 + \varepsilon} I.$$

960 \square

961 Using this Lemma, we provide proof of Theorem 3 below:

972 **Proof of Theorem 3**
973974 *Proof.* Let $Z \in \mathbb{R}^{n \times a}$ be the input matrix and $W \in \mathbb{R}^{a \times b}$ a weight matrix. The resulting feature
975 matrix is $\Phi = ZW \in \mathbb{R}^{n \times b}$ and the empirical covariances are
976

977
$$\Sigma_Z = \frac{1}{n} Z^\top Z \in \mathbb{R}^{a \times a}, \quad \Sigma_\Phi = \frac{1}{n} \Phi^\top \Phi = W^\top \Sigma_Z W \in \mathbb{R}^{b \times b}.$$

978

979 Let $W = QS$ denote the right polar decomposition of W , where $Q \in \mathbb{R}^{a \times a}$ has orthonormal columns
980 ($Q^\top Q = I_b$) and $S = (W^\top W)^{1/2} \in \mathbb{R}^{b \times b}$ is positive definite. Then
981

982
$$\Sigma_\Phi = W^\top \Sigma_Z W = S (Q^\top \Sigma_Z Q) S.$$

983

984 Write $M := Q^\top \Sigma_Z Q \succeq 0$, and let its positive eigenvalues be $\eta_1 \geq \dots \geq \eta_d > 0$, where
985 $d = \text{rank}(M) \leq \min\{b, \text{rank}(\Sigma_Z)\}$. Let $\sigma_1(\Phi) \geq \dots \geq \sigma_d(\Phi) > 0$ denote the nonzero singular
986 values of Φ .
987988 For any $x \in \mathbb{R}^b$ with $\|x\| = 1$,

989
$$x^\top \Sigma_\Phi x = x^\top S M S x = (Sx)^\top M (Sx).$$

990

991 Let $y = Sx$. Then $x^\top \Sigma_\Phi x = y^\top My$ and, by Lemma 6,
992

993
$$\|y\|_2^2 = \|Sx\|_2^2 \in [1 - \varepsilon, 1 + \varepsilon].$$

994

995 Therefore

996
$$(1 - \varepsilon) \lambda_{\min}^+(M) \leq x^\top \Sigma_\Phi x \leq (1 + \varepsilon) \lambda_{\max}(M),$$

997

998 where $\lambda_{\min}^+(M) = \eta_d$ denotes the smallest positive eigenvalue of M (the lower bound is interpreted
999 on the subspace where $My \neq 0$). Taking the maximum over unit x yields
1000

1001
$$\lambda_{\max}(\Sigma_\Phi) \leq (1 + \varepsilon) \eta_1,$$

1002

1003 and taking the minimum Rayleigh quotient over the orthogonal complement of $\ker(\Sigma_\Phi)$ yields
1004

1005
$$\lambda_{\min}^+(\Sigma_\Phi) \geq (1 - \varepsilon) \eta_d.$$

1006

1007 Since $\sigma_{\max}(\Phi)^2 = n \lambda_{\max}(\Sigma_\Phi)$ and $(\sigma_{\min}^+(\Phi))^2 = n \lambda_{\min}^+(\Sigma_\Phi)$, below inequality holds with
1008 $d = \text{rank}(M)$.
1009

1010
$$\sqrt{n} \sqrt{1 - \varepsilon} \sqrt{\eta_d} \leq \sigma_{\min}^+(\Phi) \leq \sigma_{\max}(\Phi) \leq \sqrt{n} \sqrt{1 + \varepsilon} \sqrt{\eta_1}.$$

1011

1012 Here $\sigma_{\min}^+(\Phi)$ denotes the smallest positive singular value of Φ (defined only when $d \geq 1$).
10131014 Therefore, if $d \geq 1$, then
1015

1016
$$\rho_\Phi := \frac{\sigma_{\max}(\Phi)}{\sigma_{\min}^+(\Phi)} \leq \sqrt{\frac{1 + \varepsilon}{1 - \varepsilon}} \cdot \sqrt{\frac{\eta_1}{\eta_d}}. \quad (13)$$

1017

1018 Consider the worst-case allocation of the nonzero singular values that maximizes the cumulative ratio
1019 $\sum_{i=1}^k \sigma_i / \sum_{i=1}^d \sigma_i$ given a fixed condition number bound ρ_Φ : the top k singular values all equal
1020 σ_{\max} and the remaining $d - k$ equal σ_{\min}^+ . Then
1021

1022
$$\frac{\sum_{i=1}^k \sigma_i(\Phi)}{\sum_{i=1}^d \sigma_i(\Phi)} \leq \frac{k \sigma_{\max}}{k \sigma_{\max} + (d - k) \sigma_{\min}^+} = \frac{k \rho_\Phi}{k \rho_\Phi + (d - k)}. \quad (14)$$

1023

1024 To achieve a coverage level of $1 - \delta$ with k singular values, it is necessary that
1025

1026
$$\frac{k \rho_\Phi}{k \rho_\Phi + (d - k)} \geq 1 - \delta \implies k \geq \frac{(1 - \delta) d}{\delta \rho_\Phi + (1 - \delta)}.$$

1027

1028 Taking the ceiling and substituting the bound on ρ_Φ from (13), establishes the left inequality in (3)
10291030 When $\Sigma_Z = I$, we have $M = Q^\top I Q = I$, so $\eta_1 = \dots = \eta_b = 1$ and $d = b$. Then,
1031

1032
$$\sqrt{n} \sqrt{1 - \varepsilon} \leq \sigma_i(\Phi) \leq \sqrt{n} \sqrt{1 + \varepsilon} \quad (\forall i),$$

1033

1026 Which leads to

$$1027 \quad 1028 \quad \rho_\Phi \leq \sqrt{\frac{1+\epsilon}{1-\epsilon}}$$

1029 Substituting ρ_Φ in first inequality of (3) leads to the right inequality (3). Without whitening, the
1030 achievable flattening is limited by the compressed input spectrum $M = Q^\top \Sigma_Z Q$.

□

1031
1032
1033
1034 Next, we present the proof for Theorem 4.

1035
1036 **Proof of Theorem 4**

1037
1038 *Proof.* By isotropy and positive homogeneity, there exists a constant $c_\sigma > 0$ such that
1039 $\mathbb{E}_z[\sigma(\langle z, w_j \rangle)] = c_\sigma \|w_j\|$ for all j . Hence $s_j = \|w_j\| / (\frac{1}{b} \sum_k \|w_k\|)$. Let $u_j = \|w_j\|^2 =$
1040 $[W^\top W]_{jj}$. Since

$$1041 \quad 1042 \quad \text{DfI}(W) = \|W^\top W - I\|_F^2 = \sum_{j=1}^b (u_j - 1)^2 + 2 \sum_{i < j} \langle w_i, w_j \rangle^2 \geq \sum_{j=1}^b (u_j - 1)^2,$$

1044
1045 we obtain $|u_j - 1| \leq \sqrt{\sum_k (u_k - 1)^2} \leq \epsilon$, i.e., $1 - \epsilon \leq \|w_j\|^2 \leq 1 + \epsilon$ for all j . The same bounds
1046 imply $\sqrt{1 - \epsilon} \leq \bar{r} := \frac{1}{b} \sum_k \|w_k\| \leq \sqrt{1 + \epsilon}$. Therefore

$$1047 \quad 1048 \quad \frac{\sqrt{1 - \epsilon}}{\sqrt{1 + \epsilon}} \leq s_j = \frac{\|w_j\|}{\bar{r}} \leq \frac{\sqrt{1 + \epsilon}}{\sqrt{1 - \epsilon}}.$$

□

1049
1050
1051 **Corollary 2** (Absence of τ -dormant neurons). *Fix $\tau \in (0, 1)$. If $\text{DfI}(W) \leq \left(\frac{1-\tau^2}{1+\tau^2}\right)^2$, then $s_j \geq \tau$ for all j .*

1052
1053 Note that Sokar et al. (2023) measured neurons with a dormancy score of 0.025 or lower as dormant.
1054 In this threshold, based on our theoretical analysis, $\text{DfI}(W) < 0.9975$ can eliminate dormant neurons
1055 from the network.

1056
1057 **Derivation of Equation 5**

1058
1059 Here we provide derivation of Equation 5.

1060
1061 First expand the norm:

$$1062 \quad 1063 \quad \|W - \widetilde{W}\|_F^2 = \|W\|_F^2 + \|\widetilde{W}\|_F^2 - 2 \text{tr}(\widetilde{W}^\top W).$$

1064 From the constraint $\widetilde{W}^\top \widetilde{W} = I$, we have $\|\widetilde{W}\|_F^2 = \text{tr}(I) = n$, so

$$1066 \quad 1067 \quad \min_{\widetilde{W}^\top \widetilde{W} = I} \|W - \widetilde{W}\|_F^2 \iff \max_{\widetilde{W}^\top \widetilde{W} = I} \text{tr}(\widetilde{W}^\top W).$$

1068
1069 Let $S := W^\top W \succ 0$, and define $Q := WS^{-1/2}$. Then

$$1070 \quad 1071 \quad Q^\top Q = S^{-1/2} W^\top W S^{-1/2} = S^{-1/2} S S^{-1/2} = I,$$

1072 so Q is feasible, and

$$1073 \quad W = Q S^{1/2}$$

1074
1075 is the (column) polar decomposition of W .

1076 Now take any feasible \widetilde{W} and set

$$1077 \quad 1078 \quad Z := \widetilde{W}^\top Q.$$

1079 Then

$$\text{tr}(\widetilde{W}^\top W) = \text{tr}(\widetilde{W}^\top Q S^{1/2}) = \text{tr}(Z S^{1/2}).$$

1080 Because \widetilde{W} and Q have orthonormal columns, one can show $Z^\top Z \leq I$, so all singular values $\sigma_i(Z)$
 1081 satisfy $0 \leq \sigma_i(Z) \leq 1$. By von Neumann's trace inequality,
 1082

$$1083 \text{tr}(ZS^{1/2}) \leq \sum_{i=1}^n \sigma_i(Z) \sigma_i(S^{1/2}) \leq \sum_{i=1}^n \sigma_i(S^{1/2}) = \text{tr}(Q^\top W),$$

$$1084$$

$$1085$$

1086 with equality when $Z = I$, i.e., when $\widetilde{W} = Q$.
 1087

1088 Therefore the solution is:
 1089

$$\widetilde{W}^* = Q = W(W^\top W)^{-\frac{1}{2}}.$$

$$1090$$

$$1091$$

$$1092$$

$$1093$$

$$1094$$

$$1095$$

$$1096$$

$$1097$$

$$1098$$

$$1099$$

$$1100$$

$$1101$$

$$1102$$

$$1103$$

$$1104$$

$$1105$$

$$1106$$

$$1107$$

$$1108$$

$$1109$$

$$1110$$

$$1111$$

$$1112$$

$$1113$$

$$1114$$

$$1115$$

$$1116$$

$$1117$$

$$1118$$

$$1119$$

$$1120$$

$$1121$$

$$1122$$

$$1123$$

$$1124$$

$$1125$$

$$1126$$

$$1127$$

$$1128$$

$$1129$$

$$1130$$

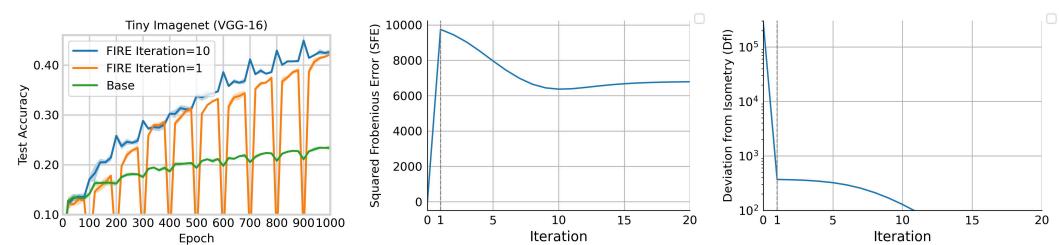
$$1131$$

$$1132$$

$$1133$$

1134 **B ADDITIONAL RESULTS**
11351136 **B.1 COMPUTATIONAL EFFICIENCY**
11371138 To prove computational efficiency of FIRE, we provide wall-clock time and GPU memory usage in
1139 Table 1.1140 Table 1: Wall-Clock Time and GPU memory footprint of FIRE and baseline methods
1141

Method	Wall-Clock Time	GPU Memory
Shrink Perturb	0.002 sec	27 MB
FIRE	0.06 sec	55 MB
DASH	69 sec	2834 MB

1142
1143
1144
1145
1146
1147
1148
1149 As shown in the table, FIRE introduces negligible computational cost and memory usage similar to
1150 Shrink Perturb, while significantly efficient compared to DASH.
11511152 The result is averaged across 10 trials, on VGG16 architecture with TinyImageNet dataset. We used
1153 a machine consisting of TITAN RTX 24GB GPU and AMD Ryzen 7 5800X 8-Core Processor, with
1154 64GB RAM.1155 **B.2 NUMBER OF ITERATIONS FOR NEWTON-SCHULZ ITERATION**
11561157 In this section, we provide a more detailed analysis which illustrates how SFE and DfI evolve across
1158 FIRE iterations.
11591160
1161 Figure 6: **Effect of number of FIRE iterations.** Test accuracy of FIRE with single iteration and
1162 10 iterations (left). Change of SFE during FIRE iterations (middle). Change of DfI during FIRE
1163 iterations (right).
11641165 As shown in Figure 6 (right), DfI decreases substantially after only a single iteration. This suggests
1166 that using a small number of iterations (< 5) is sufficient to bring performance benefits. However, as
1167 shown in Figure 6 (middle), SFE reaches its peak at the first iteration and then decreases as the number
1168 of iterations increases, indicating that using only a few iterations (< 5) can introduce instability and
1169 ultimately lead to performance degradation.
11701171 We validate this result in continual visual learning (VGG-16 with Tiny-ImageNet). Figure 6 (left)
1172 shows comparison between FIRE with 10 iterations and with single iteration. The result shows that
1173 even with single iteration still can achieve comparable performance with 10 iterations, but show
1174 significant drop after reinitialization, which supports aforementioned findings.
11751176 **B.3 COEFFICIENTS FOR NEWTON-SCHULZ ITERATION**
11771178 Our orthogonalization strategy builds on the Newton–Schulz (NS) iteration, which has also been
1179 adopted in recent works such as Muon. However, the exact recurrence used in Muon differs from
1180 ours. Muon employs a tuned quintic polynomial of the form
1181

1182
$$\varphi(x) = ax + bx^3 + cx^5,$$

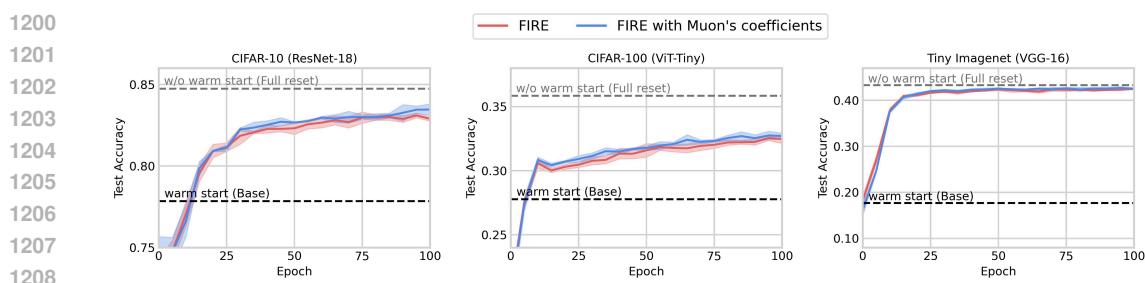
1183

1188 with optimized coefficients such as $(a, b, c) = (3.4445, -4.7750, 2.0315)$, chosen to accelerate
 1189 convergence so that only a few iterations are needed in practice. However, this sacrifices accuracy
 1190 for speed, since the singular values do not converge to 1, but oscillate near it. Since our interest
 1191 is accuracy rather than speed, we adopt the standard coefficients $(a, b, c) = (2, -1.5, 0.5)$, which
 1192 correspond to a well-known rectangular variant of NS:

$$X_{k+1} = 2X_k - 1.5X_k(X_k^\top X_k) + 0.5X_k(X_k^\top X_k)^2.$$

1193 Although this more slowly increases small singular values than Muon’s tuned version, it accurately
 1194 converges to orthogonal matrix.

1195 Empirically, we did not observe significant difference in performance when we tested both coeffi-
 1196 cients in the warm-start setting (results are shown in Figure 7).



1200
 1201
 1202
 1203
 1204
 1205
 1206
 1207
 1208
 1209
 1210
 1211
 1212
 1213
 1214
 1215
 1216
 1217
 1218
 1219
 1220
 1221
 1222
 1223
 1224
 1225
 1226
 1227
 1228
 1229
 1230
 1231
 1232
 1233
 1234
 1235
 1236
 1237
 1238
 1239
 1240
 1241

Figure 7: Effect of Newton Schulz iteration coefficients on FIRE. FIRE and FIRE with Muon’s coefficients are evaluated on warm-start setting under CIFAR-10 with ResNet-18 (left), CIFAR-100 with ViT-Tiny (middle), and Tiny ImageNet with VGG-16 (right).

B.4 TRAIN ACCURACY

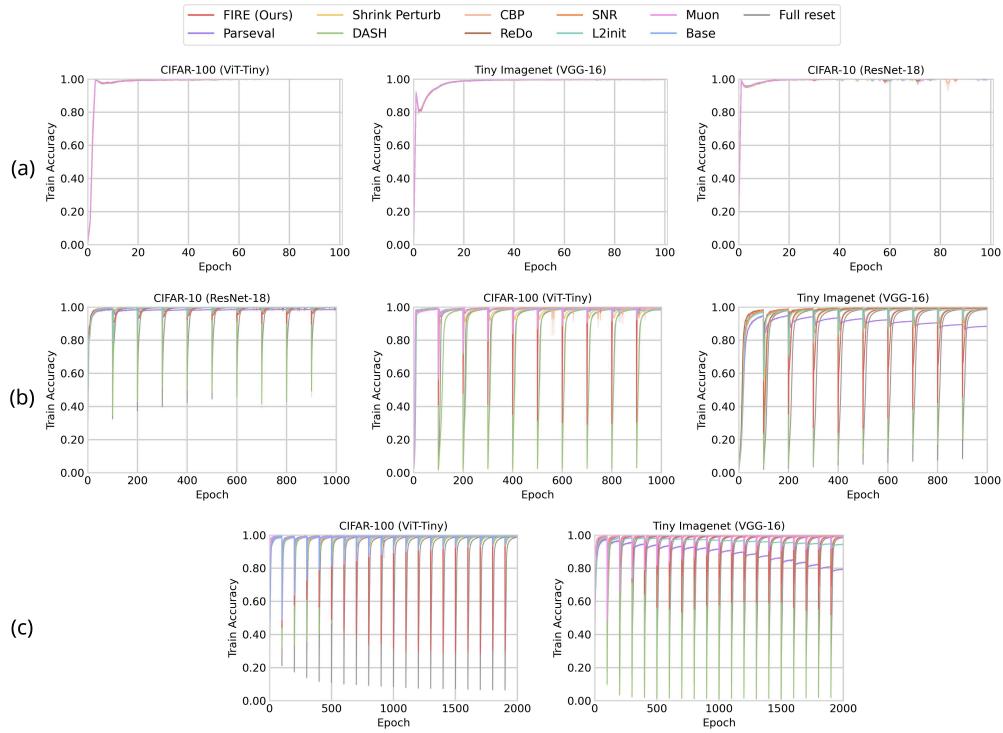


Figure 8: Train accuracy of continual visual learning experiment. Warm-start setting (a), Continual setting (b), and Class-incremental setting (c).

1242 **C IMPLEMENTATION DETAILS OF FIRE**
1243

1244 Here we describe how FIRE is applied in practice to different modules of the network.
1245

1246 **Linear layers.** For fully-connected weights $W \in \mathbb{R}^{d_{\text{out}} \times d_{\text{in}}}$, we first normalize and then apply NS
1247 iteration to approximate an orthogonal matrix. Since orthogonalization alone changes the scale of the
1248 outputs, we multiply the result by

1249
$$\text{scale} = \sqrt{\frac{d_{\text{out}}}{d_{\text{in}}}}.$$

1250

1251 This factor is motivated by the Modular Duality framework (Bernstein & Newhouse, 2025), which
1252 shows that taking the ratio of output to input dimension is sufficient to preserve stable signal variance.
1253 In short, the orthogonalization ensures the weights are well-conditioned, and the scaling factor
1254 restores the right magnitude.

1255 **Convolutional layers.** For convolutional filters $W \in \mathbb{R}^{C_{\text{out}} \times C_{\text{in}} \times k_h \times k_w}$, we apply the same proce-
1256 dure slice by slice over the spatial indices. Here the scaling factor additionally accounts for the size
1257 of the kernel:

1258
$$\text{scale} = \frac{\sqrt{C_{\text{out}}/C_{\text{in}}}}{k_h k_w}.$$

1259

1260 Intuitively, the larger the kernel, the more input values contribute to each output, so we divide by the
1261 kernel area to prevent the output variance from exploding.

1262 Note that for each spatial location (i, j) , the slice $W[:, :, i, j] \in \mathbb{R}^{C_{\text{out}} \times C_{\text{in}}}$ is orthogonalized indepen-
1263 dently by applying the Newton–Schulz iteration.

1264 **Attention modules.** In Vision Transformers (ViTs), we restrict orthogonalization to the query
1265 (Q) and key (K) projections. Empirically, applying it to the feedforward MLP layers or the output
1266 projections does not provide clear benefits and may even reduce performance. Because the dot-
1267 product QK^{\top} is the part most sensitive to poor conditioning, orthogonalizing Q and K helps improve
1268 the stability of similarity scores while leaving the value (V), output (O), and MLP weights unchanged.

1269
1270
1271
1272
1273
1274
1275
1276
1277
1278
1279
1280
1281
1282
1283
1284
1285
1286
1287
1288
1289
1290
1291
1292
1293
1294
1295

1296 **D BASELINE METHODS**
 1297
 1298

1299 **Shrink & Perturb.** Shrink & Perturb (S&P) is a Reset-based method that shrinks weight parameters
 1300 and injects noise (Ash & Adams, 2020). This method has proven particularly beneficial for warm-start
 1301 training. Following the setup in prior work Lee et al. (2024a), we control both the noise level and the
 1302 shrinkage strength using a single hyperparameter. Formally, letting θ denote the learnable parameters,
 1303 θ_0 the initial parameters, and λ the S&P coefficient, the update rule is: $\theta \leftarrow (1 - \lambda)\theta + \lambda\theta_0$.

1304 **DASH.** Direction-Aware SHrinking (DASH) (Shin et al., 2024) is a Reinitialization-based method
 1305 that selectively shrinks network weights according to their directional alignment with the loss
 1306 gradient, measured by cosine similarity. This method suppresses parameters that contribute to noise
 1307 memorization while preserving weights that encode task-relevant features. This method enhances
 1308 training efficiency and preserves model plasticity, leading to improved generalization under stationary
 1309 data distributions. In our experiments, we applied DASH when new data was added.

1310 **Parseval Regularization.** Parseval Reg. introduces a regularization term that enforces the weight
 1311 matrices of neural network layers to remain approximately orthogonal (Chung et al., 2024). Formally,
 1312 letting W denote a weight matrix, I identity matrix, and $\|\cdot\|_F$ the Frobenius norm. The loss term is
 1313 $\lambda\|WW^\top - sI\|_F^2$, where $s > 0$ is a scaling factor and λ is the regularization strength. It penalizes
 1314 the deviation of WW^\top from a scaled identity matrix, encouraging the rows of each weight matrix W
 1315 to be orthogonal and have controlled norms. This constraint keeps the singular values of W close
 1316 to a constant, preventing gradient explosion or vanishing and leading to more stable and efficient
 1317 optimization. We used $s = 1$ in all experiments and only swept λ .

1318 **Continual Backpropagation.** Continual Backpropagation (CBP) selectively reinitializes low-utility
 1319 hidden units using a contribution-utility measure (Dohare et al., 2024). Contribution-utility scores
 1320 are computed as an exponential moving average of the unit’s activation magnitude multiplied by
 1321 the summed magnitude of its outgoing weights. Units with persistently low contribution utility are
 1322 considered uninformative and are periodically reset. CBP is controlled by two hyperparameters: the
 1323 maturity threshold m , which protects units from reinitialization for at least m update steps to allow
 1324 stable utility estimation, and the replacement rate ρ , which determines the expected fraction of units
 1325 to reset at each update step via fractional accumulation.

1326 **Recycling Dormant neurons.** Recycling Dormant neurons (ReDo) (Sokar et al., 2023) is another
 1327 unit-reinitialization method that assigns a neuron score to each hidden unit in every layer, and resets
 1328 units whose scores fall below the hyperparameter τ . The neuron score s is computed as the ratio
 1329 between a unit’s average activation magnitude and the average activation magnitude of all units in
 1330 the same layer, formally defined as $s_i^\ell = \frac{\mathbb{E}_{x \in D} |h_i^\ell(x)|}{\frac{1}{H^\ell} \sum_{k \in h} \mathbb{E}_{x \in D} |h_k^\ell(x)|}$, where $h_i^\ell(x)$ denotes the activation of
 1331 neuron i in layer ℓ for input $x \in D$, and H^ℓ is the number of neurons in layer ℓ .

1332 **L2 Init.** L2 Init (Kumar et al., 2025b), as known as Regen (Regenerative regularization), is a weight
 1333 regularization method designed to mitigate plasticity loss by leveraging the property that the initial
 1334 network exhibits the highest plasticity. L2 Init regularizes the weights to stay close to the initial
 1335 weights by adding a term $\lambda\|W - W_0\|_F^2$ to the loss function, where λ is the regularization strength,
 1336 W is the current weight matrix and W_0 is initial weight matrix.

1337 **Self-Normalized Resets.** Self-Normalized Resets (SNR) (Farias & Jozefiak, 2024) is a reset-based
 1338 method that detects inactive neurons by monitoring their firing statistics and statistically testing
 1339 whether a neuron’s activity is consistent with its past behavior. For each neuron, SNR maintains an
 1340 empirical distribution of inter-firing times (the number of consecutive updates with zero activation).
 1341 If the computed probability falls below a threshold $1 - \tau$, the neuron is classified as inactive and
 1342 reset. This procedure adaptively replaces neurons whose activity has effectively vanished, mitigating
 1343 plasticity loss without relying on a fixed, hand-tuned inactivity window.

1344 **Muon.** Muon (Jordan et al.) is an optimizer that augments SGD with momentum by orthogonalizing
 1345 its update matrices. Concretely, Muon first forms the usual SGD-momentum update G for each
 1346 weight matrix and then applies a small fixed number of Newton–Schulz iterations to approximate
 1347 the closest semi-orthogonal matrix $Ortho(G)$, effectively replacing G by a matrix with singular
 1348 values near one while staying close in Frobenius norm. Following the reference implementation, in
 1349 our experiments we apply Muon only to the middle weight matrices of hidden layers, while scalar

1350 and vector parameters, as well as input and output layers, are optimized with AdamW. We set the
1351 momentum to 0.95 as recommended by the authors.
1352

1353 **Plasticity Injection.** Plasticity injection (Nikishin et al., 2023) restores neural network’s plasticity by
1354 adding a fresh, randomly initialized copy of the prediction head while leaving current predictions
1355 unchanged at the moment of the change. The original prediction head is frozen, and two identical
1356 new heads are created, one that is allowed to learn and one that always stays fixed. At the start, the
1357 learned and fixed new heads cancel each other out, so the overall output of the agent stays exactly the
1358 same. As training continues, the learnable new head adapts to new data, giving the agent renewed
1359 flexibility, while the original and the fixed new head act as a stable reference. For DQN, we applied
1360 plasticity injection to MLP layers, and for SAC, we applied it to whole critic network which is known
1361 to suffer from severe plasticity loss (Ma et al., 2023).
1362
1363
1364
1365
1366
1367
1368
1369
1370
1371
1372
1373
1374
1375
1376
1377
1378
1379
1380
1381
1382
1383
1384
1385
1386
1387
1388
1389
1390
1391
1392
1393
1394
1395
1396
1397
1398
1399
1400
1401
1402
1403

1404 **E DETAILED EXPERIMENT SETTINGS**
14051406 **E.1 CONTINUAL VISUAL LEARNING**
14071408 For continual visual learning, we report the results with 3 seeds.
14091410 Table 2: Detailed settings in continual visual learning.
1411

1412 Parameter	1413 Value
1414 Optimizer	1415 Adam (Kingma & Ba, 2014)
1416 Learning Rate	1417 $1e-3$
1417 Learning Rate Scheduler	1418 Warmup
1418 Gradient norm clipping	1419 0.5
1419 Batch Size	256
1420 Epochs per Chunk	100
1421 Data Augmentation	1422 False

1423 In this section, we describe the detailed settings for conducting continual visual learning. We note
1424 that most of the hyperparameters we used were adopted from Lee et al. (2024a).
14251426 For the Warmup scheduler, the learning rate is increased linearly from 0 to the target learning rate
1427 during the first 10% of training on each dataset. In other words, in the case of Table 2, the learning
1428 rate is gradually raised from 0 to $1e-3$ over the first 10 epochs of each data chunk.
14291430 In the warm-start scenario described in Section 4.1, we trained for 1000 epochs before new data
1431 arrived and for 100 epochs after its arrival, in order to balance the total number of gradient updates
1432 before and after the introduction of new data.
14331434 **E.2 CONTINUAL PRETRAINING OF LLMs**
14351436 For continual pretraining of LLMs, we report the results with 3 seeds.
14371438 Table 3: Detailed settings in continual pretraining of LLMs.
1439

1440 Parameter	1441 Value
1441 Optimizer	1442 AdamW (Loshchilov & Hutter, 2017)
1442 Weight Decay	1443 $1e-1$
1443 Learning Rate	1444 $6e-4$
1444 Minimum Learning Rate	1445 $6e-5$
1445 Learning Rate Scheduler	1446 Warmup + Linearly Decaying
1446 Gradient norm clipping	1447 1.0
1447 Batch Size	1448 480

1449 We used implementation and hyperparameters of nanoGPT from Karpathy (2023).
14501451 During first phase, the learning rate linearly increases from 0 to the target learning rate ($6e-4$) during
1452 2,000 steps, then annealed to minimum learning rate ($6e-5$) until the end of the phase. In the second
1453 phase, the learning rate linearly increases from 0 to the target learning rate ($6e-4$) during 10% of
1454 training iterations of second phase. Then, it decreases linearly to minimum learning rate ($6e-5$) until
1455 the end of the phase.
14561457 **E.3 REINFORCEMENT LEARNING**
14581459 For reinforcement learning, we report the results with 5 seeds.
14601461 For S&P method, we apply S&P to the encoder and Reset to the fully connected layers (D’Oro et al.,
1462 2022) for discrete control, and S&P with $\lambda = 0.8$ to whole parameters for continuous control tasks.
1463 We perform a single intervention (Full Reset, S&P, FIRE) at the midpoint of learning. We followed
1464 the hyperparameter configurations used in prior work (Sokar et al., 2023).
1465

1458 Table 4: Hyperparameters used in the ALE environment with DQN algorithm.
1459

1460 Parameter	1461 Value
1462 Optimizer	1463 Adam (Kingma & Ba, 2014)
1464 Optimizer: ϵ	1465 1.5e - 4
1466 Optimizer: Learning rate	1467 6.25e - 5
1468 Minimum ϵ for training	1469 0.01
1470 Evaluation ϵ	1471 0.001
1472 Discount factor γ	1473 0.99
1474 Replay buffer size	1475 10^6
1476 Minibatch size	1477 32
1478 Initial collect steps	1479 20000
1480 Training iterations	1481 10
1482 Training environment steps per iteration	1483 $250K$
1484 Updates per environment step (Replay Ratio)	1485 1
1486 Target network update period	1487 2000
1488 Loss function	1489 Huber Loss (Huber, 1992)

1490 For continuous tasks, the hyperparameter setting is followed by Lee et al. (2024b).
14911492 Table 5: Hyperparameters used in HumanoidBench environments with SimBa.
1493

1494 Parameter	1495 Value
1496 Optimizer	1497 AdamW (Loshchilov & Hutter, 2017)
1498 Optimizer: Learning rate	1499 1e - 4
1500 Optimizer: Weight decay	1501 0.01
1502 Actor hidden dim	1503 128
1504 Actor num blocks	1505 1
1506 Critic hidden dim	1507 512
1508 Critic num blocks	1509 2
1510 Discount factor γ	1511 0.99
1512 Clipped Double-Q (Fujimoto et al., 2018)	1513 True
1514 Replay buffer size	1515 10^6
1516 Minibatch size	1517 256
1518 Initial collect steps	1519 5000
1520 Updates per environment step (Replay Ratio)	1521 2
1522 Soft target update factor τ	1523 0.005

1524 **E.4 HYPERPARAMETER SEARCH SPACE**1525 Table 6 presents the hyperparameter search space, and Tables 7-10 present their optimal values.
1526

Table 6: Hyperparameter search space for all experiments.

Experiment	Method	Hyperparameters	Search Space
Warm-Starting	S&P	λ	0.2, 0.4, 0.6, 0.8
	DASH	α	0.1, 0.3
		λ	0.05, 0.1, 0.3
	Parseval Reg.	λ	1e-3, 1e-4, 1e-5
		CBP	ρ
			1e-4, 1e-5
		m	100, 1000
		ReDo	τ
		L2 Init	λ
		SNR	τ
Continual Learning	S&P	λ	0.2, 0.4, 0.6, 0.8
	DASH	α	0.1, 0.3
		λ	0.05, 0.1, 0.3
	Parseval Reg.	λ	1e-3, 1e-4, 1e-5
		CBP	ρ
			1e-4, 1e-5
		m	100, 1000
		ReDo	τ
		L2 Init	λ
		SNR	τ
Class-Incremental Learning	S&P	λ	0.2, 0.4, 0.6, 0.8
	DASH	α	0.1, 0.3
		λ	0.05, 0.1, 0.3
	Parseval Reg.	λ	1e-3, 1e-4, 1e-5
		CBP	ρ
			1e-4, 1e-5
		m	100, 1000
		ReDo	τ
		L2 Init	λ
		SNR	τ
Continual pretraining of GPT-0.1B	S&P	λ	0.2, 0.5, 0.8

Table 7: Hyperparameters for Warm-Start setting.

Dataset	Method	Value
CIFAR-10 (ResNet-18)	S&P	$\lambda = 0.8$
	DASH	$\alpha = 0.3, \lambda = 0.05$
	Parseval Reg.	$\lambda = 1e-3$
	CBP	$\tau = 1e-4, m = 1000$
	ReDo	$\tau = 0.5$
	L2 Init	$\lambda = 1e-3$
	SNR	$\tau = 0.01$
CIFAR-100 (ViT-Tiny)	FIRE	iter = 10
	S&P	$\lambda = 0.8$
	DASH	$\alpha = 0.1, \lambda = 0.05$
	Parseval Reg.	$\lambda = 1e-5$
	CBP	$\tau = 1e-4, m = 100$
	ReDo	$\tau = 0.5$
	L2 Init	$\lambda = 1e-3$
Tiny ImageNet (VGG-16)	SNR	$\tau = 0.01$
	FIRE	iter = 10
	S&P	$\lambda = 0.8$
	DASH	$\alpha = 0.1, \lambda = 0.05$
	Parseval Reg.	$\lambda = 1e-3$
	CBP	$\tau = 1e-4, m = 1000$
	ReDo	$\tau = 0.5$
1590	L2 Init	$\lambda = 1e-3$
	SNR	$\tau = 0.08$
	FIRE	iter = 10

1566
1567
1568
1569
1570
1571
1572
1573
1574
1575
1576
1577
1578
1579
1580
1581
1582
1583
1584
1585
1586
1587
1588
1589
1590
1591
1592
1593
1594
1595
1596
1597
1598
1599
1600
1601
1602
1603
1604
1605
1606
1607
1608
1609
1610
1611
1612
1613
1614
1615
1616
1617
1618
1619

Table 8: Hyperparameters for Continual Setting.

Dataset	Method	Value
CIFAR-10 (ResNet-18)	S&P	$\lambda = 0.8$
	DASH	$\alpha = 0.3, \lambda = 0.05$
	Parseval Reg.	$\lambda = 1e-3$
	CBP	$\tau = 1e-5, m = 1000$
	ReDo	$\tau = 0.05$
	L2 Init	$\lambda = 1e-5$
	SNR	$\tau = 0.08$
CIFAR-100 (ViT-Tiny)	FIRE	iter = 10
	S&P	$\lambda = 0.8$
	DASH	$\alpha = 0.1, \lambda = 0.05$
	Parseval Reg.	$\lambda = 1e-5$
	CBP	$\tau = 1e-5, m = 1000$
	ReDo	$\tau = 0.01$
	L2 Init	$\lambda = 1e-4$
Tiny ImageNet (VGG-16)	SNR	$\tau = 0.04$
	FIRE	iter = 10
	S&P	$\lambda = 0.8$
	DASH	$\alpha = 0.1, \lambda = 0.1$
	Parseval Reg.	$\lambda = 1e-4$
	CBP	$\tau = 1e-5, m = 1000$
	ReDo	$\tau = 0.01$
Tiny ImageNet (VGG-16)	L2 Init	$\lambda = 1e-5$
	SNR	$\tau = 0.01$
	FIRE	iter = 10

Table 9: Hyperparameters for Class-Incremental Setting.

Dataset	Method	Value
CIFAR-100 (ViT-Tiny)	S&P	$\lambda = 0.8$
	DASH	$\alpha = 0.3, \lambda = 0.3$
	Parseval Reg.	$\lambda = 1e-5$
	CBP	$\tau = 1e-5, m = 1000$
	ReDo	$\tau = 0.01$
	L2 Init	$\lambda = 1e-5$
	SNR	$\tau = 0.08$
Tiny ImageNet (VGG-16)	FIRE	iter = 10
	S&P	$\lambda = 0.8$
	DASH	$\alpha = 0.3, \lambda = 0.1$
	Parseval Reg.	$\lambda = 1e-3$
	CBP	$\tau = 1e-4, m = 1000$
	ReDo	$\tau = 0.05$
	L2 Init	$\lambda = 1e-4$
Tiny ImageNet (VGG-16)	SNR	$\tau = 0.02$
	FIRE	iter = 10

Table 10: Hyperparameters for Continual pretraining of LLMs.

Method	CKPT	Value
S&P	Best	$\lambda = 0.5$
	30k	$\lambda = 0.8$
	60k	$\lambda = 0.5$
FIRE	Best	iter = 5
	30k	iter = 5
	60k	iter = 5