

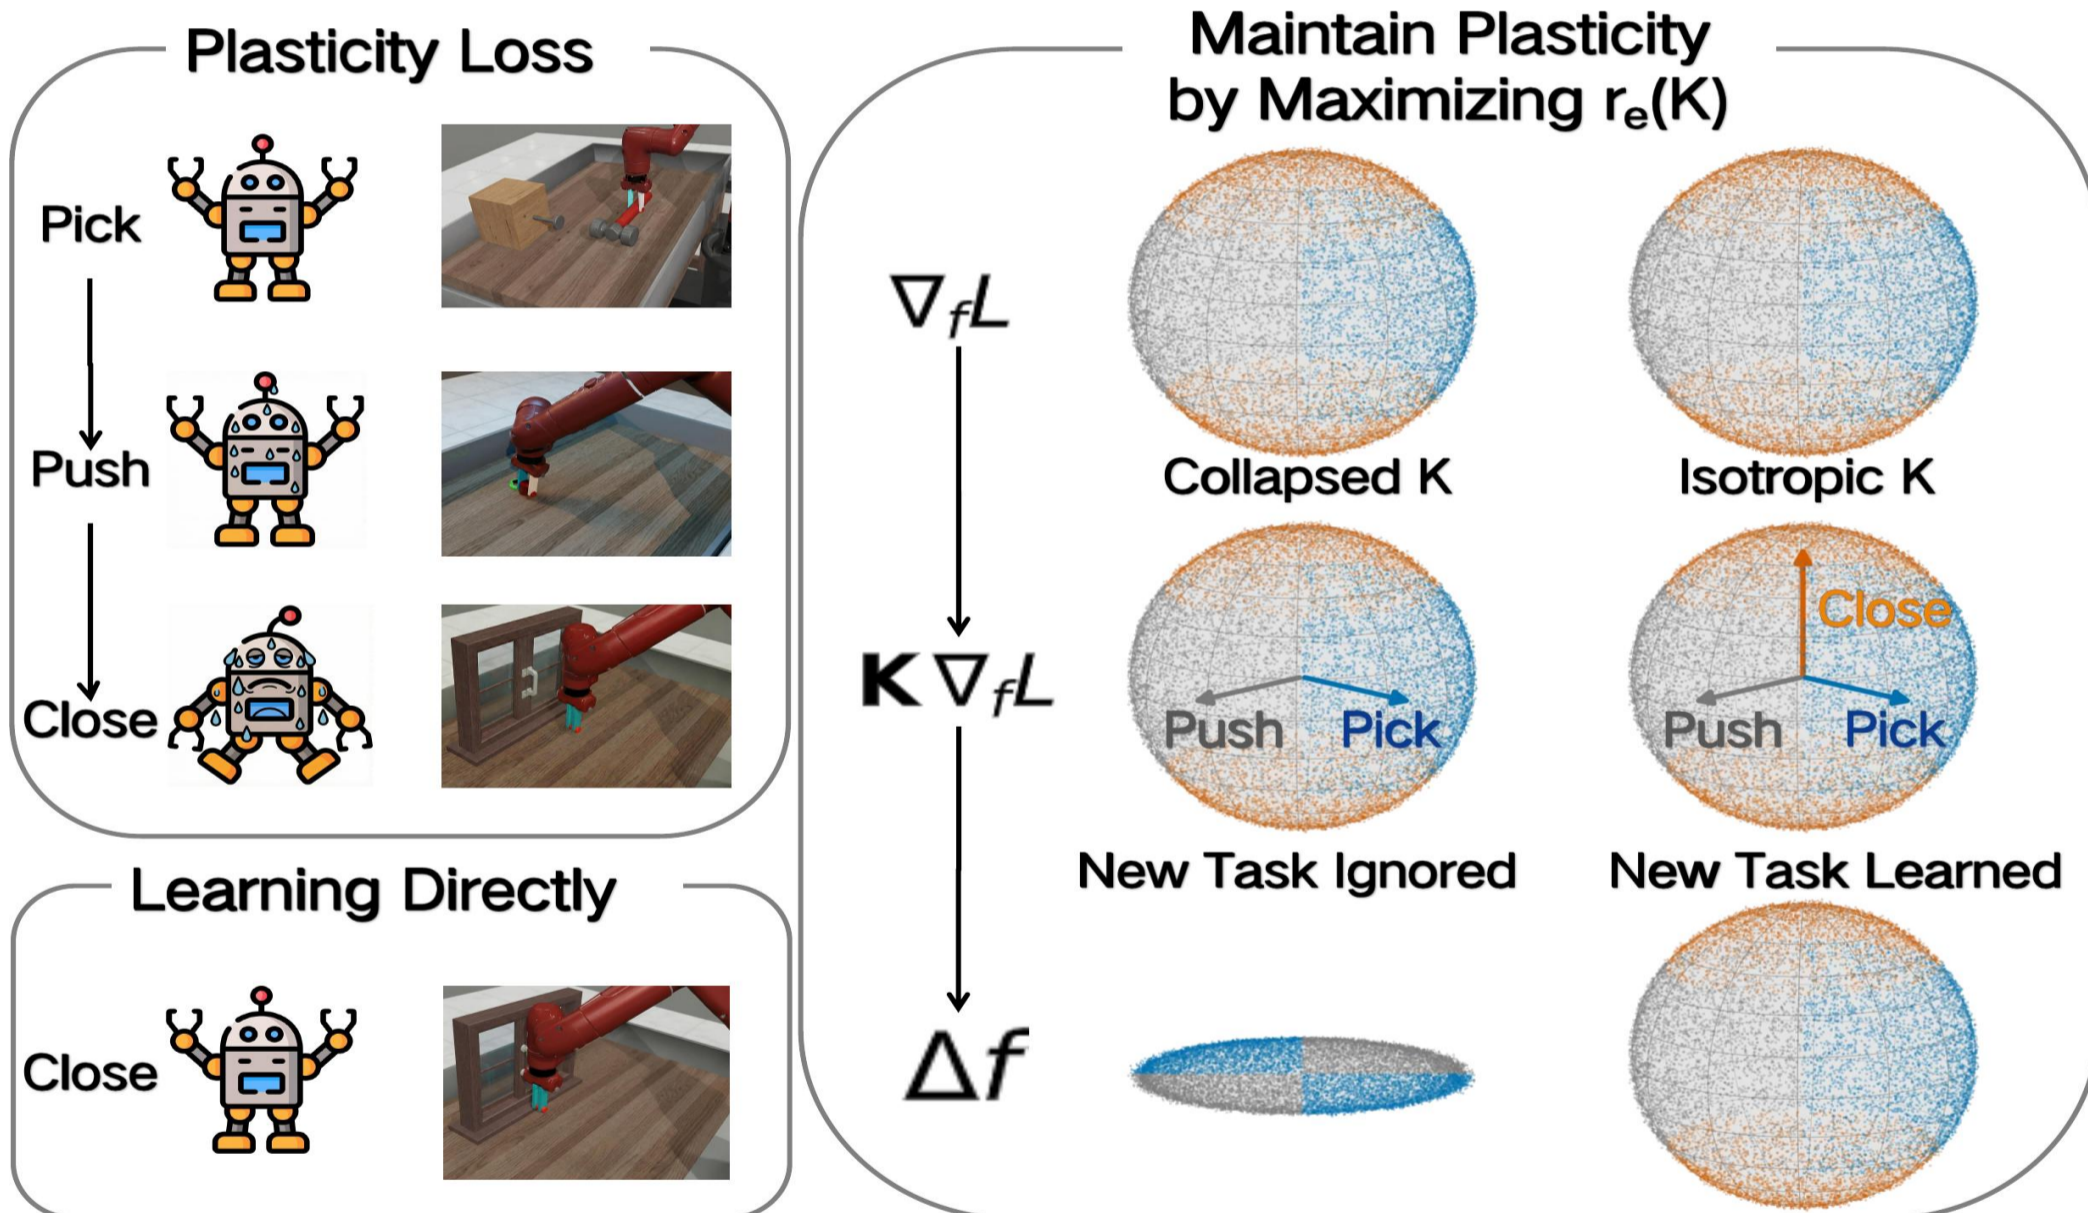
## Introduction

In continual reinforcement learning, Mixture-of-Experts (MoE) networks experience "plasticity loss": their ability to learn new skills diminishes as training proceeds.

### Our Contribution:

1. We formalize plasticity loss as the loss of spectral plasticity and derive a tractable proxy building on Neural Tangent Kernel (NTK) theory.
2. We introduce SPHERE, a principled and practical regularizer to counter the loss of spectral plasticity.
3. We evaluate SPHERE on MetaWorld and HumanoidBench, showing consistent gains over MoE baselines and prior plasticity methods, and sustained spectral plasticity throughout training.

## Plasticity loss as loss of spectral plasticity



Left: Continual RL can fail on later tasks despite isolated learnability.

Right: Low effective rank of the empirical Neural Tangent Kernel (eNTK) restricts updates to a few directions (collapsed spectrum), while high eNTK effective rank enables diverse directions (isotropic spectrum).

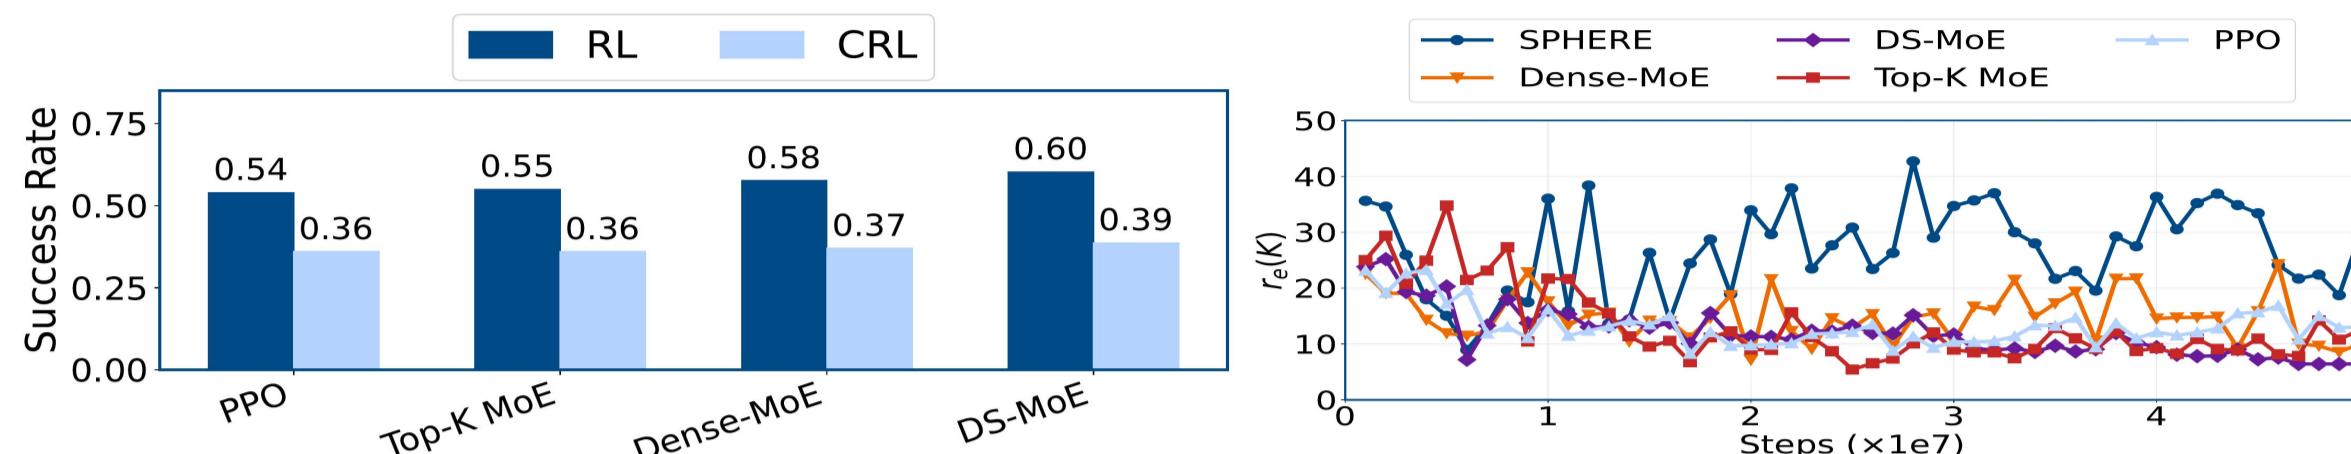
$$\Delta f = -\eta \mathbf{K} \nabla_f L$$

Spectral plasticity is defined as the effective rank of an NTK spectrum, i.e., the uniformity of the NTK spectrum. The larger the effective rank, the higher the spectral plasticity.

$$p_i = \sigma_i(M) / \sum_j \sigma_j(M) \quad r_e(M) = \exp\left(-\sum_{\{i: \sigma_i(M) > 0\}} p_i \log p_i\right)$$

## Evidence of Loss of Spectral Plasticity

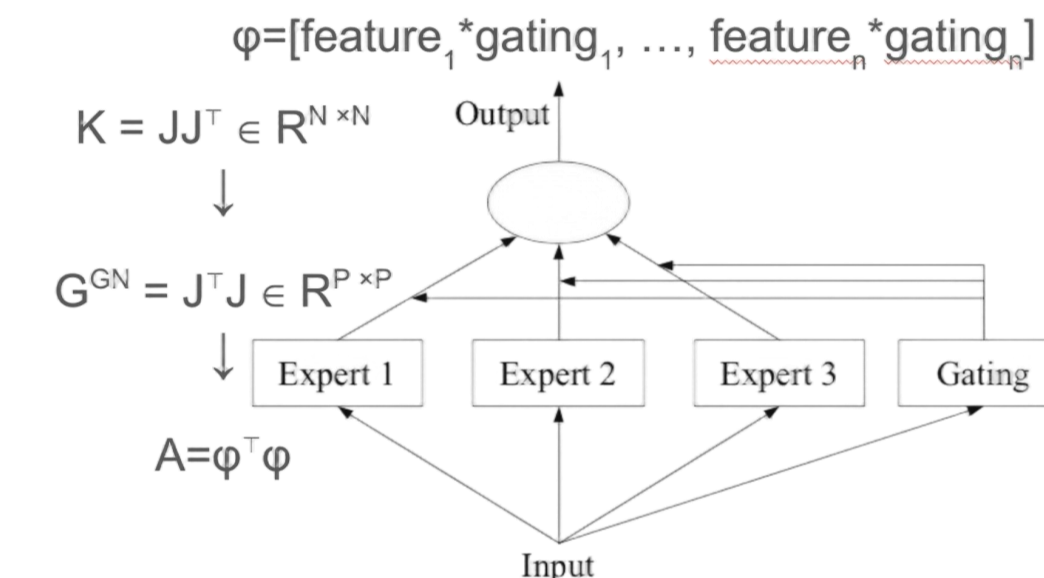
All methods show degraded performance under CRL. During training, the spectral plasticity decreased for all methods.



## Approach

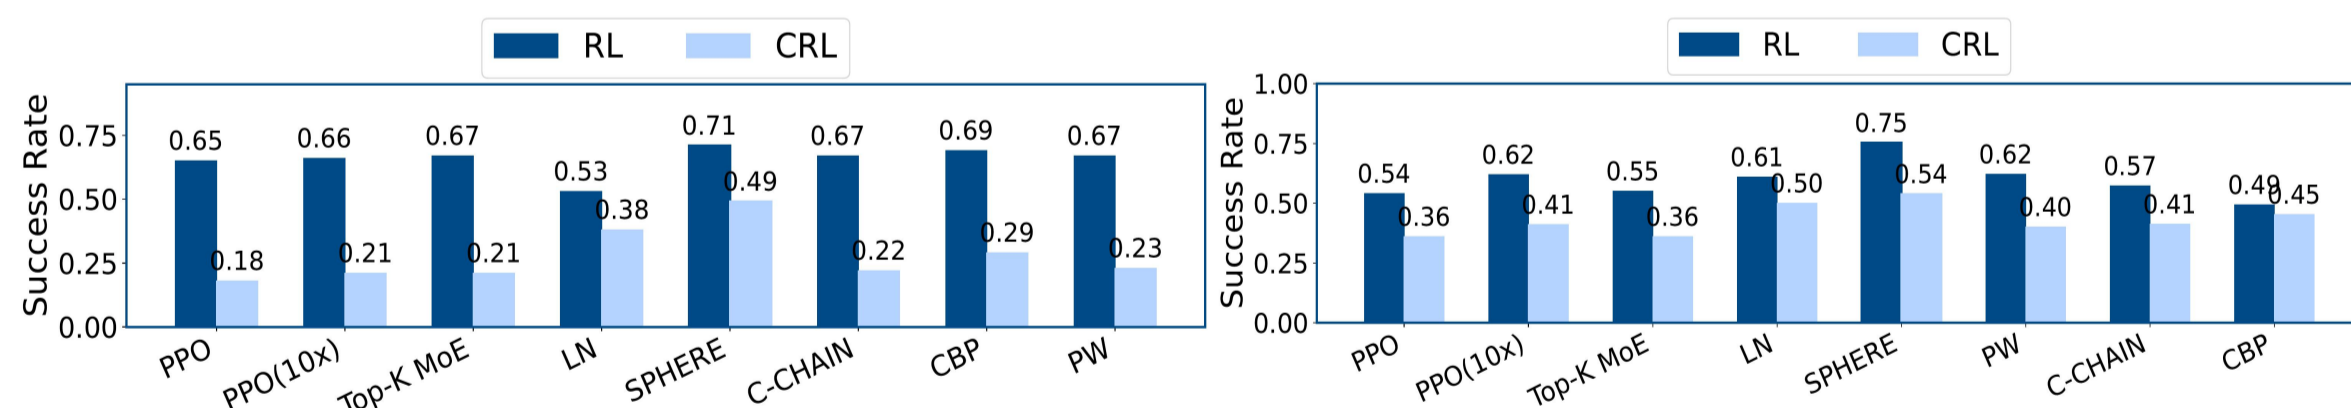
SPHERE achieves best rank on Humanoidbench and Metaworld tasks.

$$\begin{aligned} \mathcal{L}_{\text{SPHERE}}(\mathbf{A}_{\text{last}}^{\text{exp}}) &= \left\| \mathbf{A}_{\text{last}}^{\text{exp}} - \frac{\text{Tr}(\mathbf{A}_{\text{last}}^{\text{exp}})}{m} \mathbf{I}_m \right\|_F^2 \\ &= \left\| \mathbf{A}_{\text{last}}^{\text{exp}} \right\|_F^2 - \frac{\text{Tr}(\mathbf{A}_{\text{last}}^{\text{exp}})^2}{m} \end{aligned}$$



## Results

SPHERE achieves best rank on Humanoidbench and Metaworld tasks.



SPHERE avoids spectral collapse to a single component.

