ADIS-Law: Unified Scaling Laws for Annealing-Phase Domain Injection in Large Language Models

Lyuxin Xue Hui Cai Xiaoyun Feng Xuanwei Hu Wei Wu Jiangming Yang Xin Zhang*

xuelyuxin.xlx@ant-intl.com *Corresponding author

ADIS-Law predicts loss trajectories and ranks domain-injection policies at larger model scales and training budgets using small-scale experiments.

Jointly modeling injection timing, domain-data mixture ratio and re-warmup peak learning rate.

Error: <5%

Mean relative domain-trajectory error

D-Up: 2.78%; M-Up: 4.31%

Accuracy: 88.1–95.2%

Pairwise policy-ranking accuracy

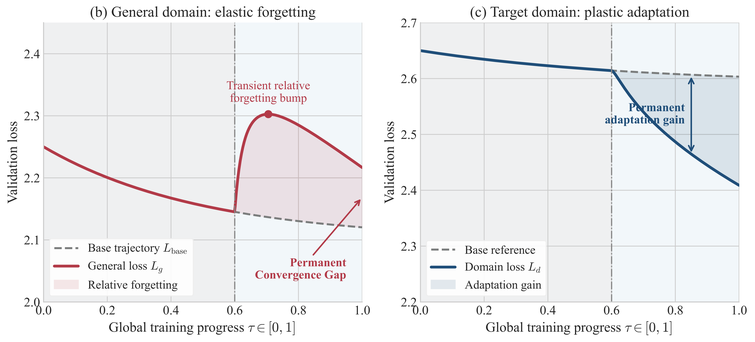
D-Up: 88.1%; M-Up: 95.2%

01 Motivation

Annealing-phase continual pre-training (CPT) introduces domain data within a fixed training budget. **Three coupled hyperparameters** govern the trade-off between general-domain retention and target-domain adaptation.

- λ **Injection fraction:** CPT budget share; injection starts at $\tau = 1 - \lambda$.
- r **Mixture ratio:** domain-data fraction after injection.
- η **Re-warmup peak LR:** peak LR during post-injection re-warmup.

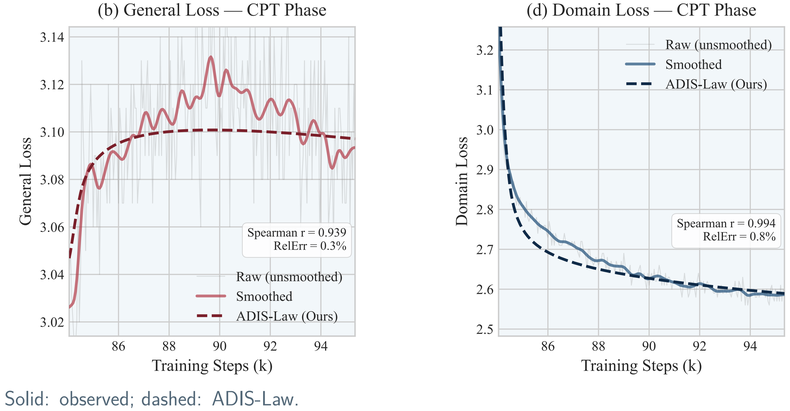
Asymmetric loss dynamics



General domain: transient forgetting followed by partial recovery.

Target domain: adaptation reduces loss below the baseline.

Observed trajectories and model fit



Solid: observed; dashed: ADIS-Law.

Research question. Can a unified scaling law predict and target domain loss trajectories across injection settings, model sizes and training budgets?

02 Methodology

ADIS-Law uses an additive loss decomposition to separate the base pre-training trajectory from domain-injection effects.

Pre-training Trajectory

$$L_{\text{base}} = \frac{A}{N^{\alpha}} + \frac{B}{S_1^{\beta}} - C_1 S_2^{\text{pt}} - C_2 S_2^{\text{cpt}} + L_{\infty}$$

N is model size; S_1 is the cumulative learning rate.
S_2^{pt} and S_2^{cpt} track LR decay before and after injection.
General and target loss branches are fitted independently.

Injection Intensity & Saturation

$$I = r\,\frac{\eta}{\eta_0}, \qquad \phi = 1 - \left(1 + E S_1^{\text{cpt}}\right)^{-\beta_\phi}$$

Injection intensity combines the mixture ratio r with the re-warmup peak learning rate η, normalized by the reference rate η_0.

The saturation gate grows from 0 toward 1 as the cumulative post-injection learning rate S_1^{cpt} increases.

Global progress is $\tau = t/T$; injection begins at $\tau = 1 - \lambda$.

General-domain Forgetting & Recovery

$$L_g = L_{\text{base}} + \phi\,(aI - b\tau)$$

Forgetting $(+\phi a I)$ raises general-domain loss with injection intensity, coupling mixture ratio r and peak learning rate η.

Recovery $(-\phi b \tau)$ offsets this increase as training progresses. A residual loss gap can remain at $\tau = 1$ under a finite budget.

Target-domain Adaptation

$$L_d = L_{\text{base}} - \phi\,\Psi\,r^{\mu}\tau^{\nu}$$

Adaptation gain $(\phi\Psi r^{\mu}\tau^{\nu})$ reduces target-domain loss below the baseline as the gate activates after injection.

Diminishing returns. Exponents $0 < \mu, \nu \leq 1$ govern the response to mixture ratio r and global training progress τ.

03 Experiments

Cross-Scale Extrapolation

D-Up: 25B/50B → 100B tokens. **M-Up:** ≤895M → 1.67B parameters.
Target scales are held out during fitting.

General-domain loss

Protocol	Relative error ↓	Pairwise accuracy ↑
D-Up	**0.73%**	**85.7%**
M-Up	**0.83%**	**81.0%**

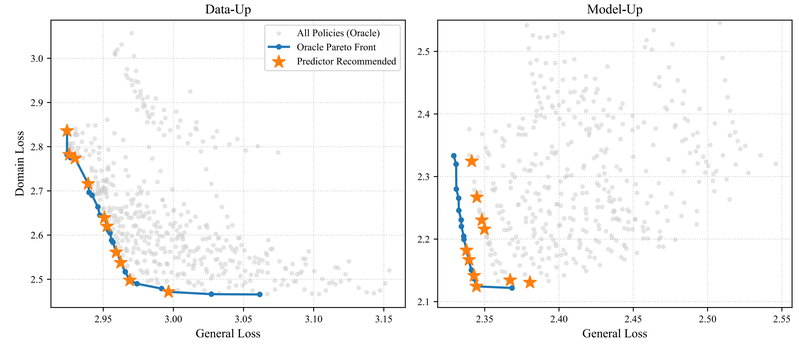
Target-domain loss

Protocol	Relative error ↓	Pairwise accuracy ↑
D-Up	**2.78%**	**88.1%**
M-Up	**4.31%**	**95.2%**

Four-domain means. Error: relative trajectory error.
Accuracy: pairwise final-loss ranking.

Pareto-based Policy Selection

486 synthetic BioMed policies, scored using proxy-scale fits.
Recommendations are evaluated by an oracle fitted at the target scale.



Orange stars: predictor recommendations. **Blue curves:** oracle frontiers.

Takeaway Messages

Coupled injection effects. Injection timing, mixture ratio and peak LR jointly shape general-domain forgetting and target-domain adaptation.

Cross-scale prediction. Proxy-scale fits predict loss trajectories and policy rankings at unseen model scales and token budgets.

Policy selection. Proxy-recommended policies lie near the target-scale oracle frontier in the synthetic-policy evaluation.