




FreezeAsGuard: Mitigating Illegal Adaptation of Diffusion Models via Selective Tensor Freezing

Kai Huang, Haoming Wang and Wei Gao
University of Pittsburgh




Background and Motivation

❖ Illegal Adaptation of Diffusion Models



Forging public figures




Original artwork
by Hollie Mengert

Mimicked artwork
in Hollie's style

Copying artistic styles

- Open-source text-to-image diffusion models can be efficiently adapted to new identities, artistic styles, and application domains using small custom datasets.
- The same flexibility can enable unauthorized uses, including forging public figures' portraits, duplicating copyrighted artistic styles, and generating explicit content.

❖ Limitation of Existing works




Only detects watermarked data;
cannot prevent adaptation.

Easy to bypass;
high false-positive rate.

Requires control of training data;
fails on unprotected data.

Knowledge can be relearned;
may harm legal adaptation.

Method Details



❖ Mask Learning via Multi-objective Bi-level Optimization

$$\mathbf{m}^* = \arg \min_{\mathbf{m}} \left(-\mathcal{L}_{\theta^*(\mathbf{m})}(x_{\text{illegal}}), \mathcal{L}_{\theta^*(\mathbf{m})}(x_{\text{legal}}) \right) \quad \text{s.t.} \quad \theta^*(\mathbf{m}) = \arg \min_{\theta(m)} \left(\mathcal{L}_{\theta(m)}(x_{\text{illegal}}), \mathcal{L}_{\theta(m)}(x_{\text{legal}}) \right)$$

linear scalarization

$$-\lambda_1 \mathcal{L}_{\theta^*(\mathbf{m})}(x_{\text{illegal}}) + \lambda_2 \mathcal{L}_{\theta^*(\mathbf{m})}(x_{\text{legal}})$$

- Estimate $\theta(m)$ to apply gradient solver

$$\theta(m) = m \odot \theta_{\text{pre}} + (1 - m) \odot \theta_{\text{ft}}$$

- Controlling mitigation power via mask ratio

$$\mathcal{L}_{\text{sparcity}} = \|\mathbf{1}^\top \mathbf{m} / N - \rho\|_2^2 \rightarrow \text{Added to the loss of mask learning loop}$$

❖ Efficient Optimization

Early Stop strategy



- fine-tuning loss typically drops fast in the first few iterations and then violently fluctuates
- do not wait for the loss to converge to reduce cost


Reuse model states across iterations

Stores only the masked weights $\theta(m)$ and the difference $\theta_d = \theta_{\text{pre}} - \theta_{\text{ft}}$, instead of three full model copies.


Method Overview

- FreezeAsGuard constrains what the model can change by selectively freezing tensors that are critical to illegal adaptations while leaving tensors needed for legal adaptations trainable.

Selectively Freezing Tensors
before model adaptation




Bilevel optimization to identify
which tensor to freeze



Experiment results

❖ Comparison with baselines



❖ Different Frozen ratio

