

# 000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 WARP: WEIGHT TELEPORTATION FOR ATTACK- RESILIENT UNLEARNING PROTOCOLS

Anonymous authors

Paper under double-blind review

## ABSTRACT

Approximate machine unlearning aims to efficiently remove the influence of specific data points from a trained model, offering a practical alternative to full retraining. However, it introduces privacy risks: an adversary with access to pre- and post-unlearning models can exploit their differences for membership inference or data reconstruction. We show these vulnerabilities arise from two factors: large gradient norms of *forget-set* samples and the close proximity of unlearned parameters to the original model. To demonstrate their severity, we propose unlearning-specific membership inference and reconstruction attacks, showing that several state-of-the-art methods (e.g., NGP, SCRUB) remain vulnerable. To mitigate this leakage, we introduce WARP, a *plug-and-play teleportation defense* that leverages neural network symmetries to reduce *forget-set* gradient energy and increase parameter dispersion while preserving predictions. This reparameterization obfuscates the signal of forgotten data, making it harder for attackers to distinguish forgotten samples from non-members or recover them via reconstruction. Across six unlearning algorithms, our approach achieves consistent privacy gains, reducing adversarial advantage (AUC) by up to 64% in black-box and 92% in white-box settings, while maintaining accuracy on retained data. These results highlight teleportation as a general tool for reducing attack success in approximate unlearning.

## 1 INTRODUCTION

Machine unlearning (MU) aims to enforce the “right to be forgotten” by updating a trained model so that a designated *forget-set* has no influence Bourtoule et al. (2021); Zhao et al. (2024). The ideal outcome matches retraining from scratch on the remaining *retain-set*, with both the model’s parameters and predictions unaffected by the forgotten data, and without degrading generalization. A primary motivation for machine unlearning is to ensure privacy compliance for sensitive information Wang et al. (2025a). Once personal data is used for training, models may memorize specific details Ravikumar et al. (2024a), creating risks of privacy breaches Bourtoule et al. (2021); Carlini et al. (2022b). Unlearning addresses this by eliminating such traces, preventing exposure. The most direct solution is retraining from scratch without the *forget set*, but this is computationally prohibitive. *Exact Unlearning* methods such as SISA Bourtoule et al. (2021) reduce cost by modifying training to allow provable deletion, but they require proactive deployment and add overhead. To avoid full retraining, *Approximate Unlearning* methods finetune the original model to forget the target data while preserving utility Kurmanji et al. (2023); Chundawat et al. (2023a); Golatkar et al. (2020); Thudi et al. (2022), trading computational efficiency against formal guarantees.

At the same time, ML models are vulnerable to privacy attacks Rigaki & Garcia (2023). In Membership Inference Attacks (MIA), an adversary determines whether a given sample was part of the training set Shokri et al. (2017). In Data Reconstruction Attacks (DRA), the adversary seeks to recover raw data (or a close approximation) from model outputs or parameters Yin et al. (2021); Li et al. (2022); Jeon et al. (2021); Fang et al. (2023). These attacks have been demonstrated in both black-box (access to outputs) and white-box (access to weights) settings Nasr et al. (2019).

Ironically, MU itself can leak the very data it aims to erase. Given access to both the original and unlearned models, an adversary can mount differencing attacks Hu et al. (2024); Bertran et al. (2024), which substantially improve reconstruction success. Even models previously resistant to MIAs can become vulnerable once deletion is performed Bertran et al. (2024); Chen et al. (2021). The key

observation is that the parameter difference between the two models approximates the gradient of the forgotten sample (up to second-order terms), effectively releasing it to the adversary. Gradient inversion techniques, as in federated learning Geiping et al. (2020), can then reconstruct the forgotten data. Thus, approximate unlearning methods, especially gradient-ascent variants Kurmanji et al. (2023), can inadvertently compromise privacy instead of ensuring it.

In this work, we aim to strengthen MU against privacy attacks by characterizing two key factors driving leakage. The first, illustrated in Figure 1, is that a forgotten sample’s privacy risk correlates with its gradient norm in the original model. Intuitively, samples with large gradient magnitudes during training or finetuning induce stronger parameter changes when removed, making them more detectable via MIA and more exploitable for reconstruction Ye et al. (2023).

Second, as shown in prior work Thudi et al. (2022); Kurmanji et al. (2023), most approximate unlearning methods make minor parameter updates, typically by maximizing the *forget-set* loss while keeping retain-set accuracy stable. This keeps the unlearned model close to the original, so the parameter difference encodes information about the forgotten data. In gradient-ascent-based methods Kurmanji et al. (2023); Chundawat et al. (2023a), this difference is essentially the *forget-set* gradient. Recent studies confirm that such updates expose information equivalent to a single gradient step on the forgotten sample Bertran et al. (2024), which attackers can invert to reconstruct it.

To mitigate these risks, we propose WARP, a plug-and-play defense that integrates into existing unlearning algorithms without training-time statistics. Our method leverages neural network teleportation Armenta et al. (2023), exploiting parameter-space symmetries (e.g., rescaling or permutation) that preserve predictions. By applying selective teleportation steps before or during unlearning, we reduce *forget-set* gradient norms while injecting symmetry-preserving randomness. This yields unlearned models that retain accuracy yet are displaced in parameter space, making it harder for an attacker to disentangle forgetting from teleportation. Consequently, membership inference and reconstruction attacks are significantly weakened, as shown in Sections 4.2, 4.3, and 4.4.

Our **contributions** are summarized as follows:

- **Tailored privacy attacks.** We design MIA and DRA for the unlearning setting, where the adversary compares pre- and post-unlearning models. These attacks show that leading methods remain vulnerable, as parameter updates still expose information about the *forget-set*.
- **Symmetry-based defense.** We propose WARP, a plug-and-play defense that, building on existing teleportation and symmetry constructions, applies loss-preserving transformations to reduce *forget-set* gradient norms and increase parameter dispersion, thereby obscuring the signal exploited in reconstruction and inference, while remaining agnostic to the particular symmetry mechanism used to realize these transformations. WARP integrates into gradient-based post-hoc unlearning algorithms without requiring training-time statistics.
- **Comprehensive evaluation.** We evaluate our attacks and defense across three datasets—CIFAR-10, Tiny-ImageNet, and ImageNet-1K—using ResNet-18 and ViT-B/16 models under both black-box and white-box settings. Results across multiple unlearning algorithms show that teleportation consistently reduces privacy leakage while preserving accuracy on the retain set.

Overall, our work reframes unlearning privacy risk through the lens of *gradient norm reduction* and connects it to neural network symmetry, an underexplored optimization principle that lays a conceptual foundation for more privacy attack-resilient unlearning algorithms. Related works to ours are discussed in more detail in Appendix A. The code is available at [this link](#).

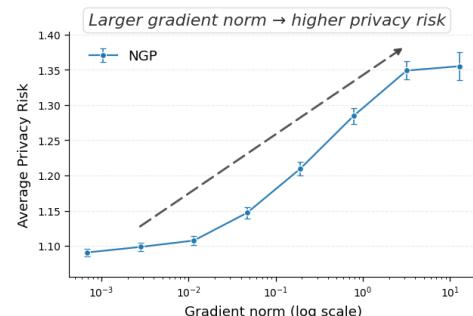


Figure 1: Privacy risk vs. gradient norms of *forget-set* samples, measured with U-LiRA.

108  
109  

## 2 THREAT MODEL

110  
111 We consider a strong adversary performing *sample-wise membership inference*, distinguishing whether  
112 a sample belongs to the *forget-set*  $\mathcal{D}_f$  or the *test set*  $\mathcal{D}_{\text{test}}$ . The attacker has access to both the pre- and  
113 post-unlearning models.114  
115 **Attacker Capabilities.** The attacker has full access to both the original  $\theta^{\text{org}}$  and unlearned model  
116  $\theta^u$ , as well as complete knowledge of the unlearning algorithm  $\mathcal{A}_{\text{unlearn}}$  and its hyperparameters  
117  $\mathcal{H}_{\text{unlearn}}$  (e.g., optimizer, learning rate, update steps, retain-set size).118 We consider two settings: **Black-box** — the attacker queries outputs  $f(x; \theta^u)$ . **White-box** — the  
119 attacker additionally accesses full internals of both models  $(\theta^{\text{org}}, \theta^u)$ , including weights.120  
121 **Attack Objective** Given a sample  $(x, y)$  from either the *forget-set*  $\mathcal{D}_{\text{forg}}$  or the held-out test set  
122  $\mathcal{D}_{\text{test}}$ , the attacker computes a score  $A'(x, y)$  and predicts membership as  $A(x, y) = \mathbb{I}[A'(x, y) > \tau]$ ,  
123 where  $\mathbb{I}[\cdot]$  is the indicator function and  $\tau$  is a decision threshold. The attacker seeks a high true  
124 positive rate (TPR) on forgotten samples while maintaining a low false positive rate (FPR) on test  
125 samples. This directly measures privacy risk: if membership can be reliably inferred, incomplete  
126 unlearning is exposed and the forgotten samples identified. Unlike prior work, our goal is to audit  
127 unlearning algorithms from a *privacy perspective*, rather than evaluating indistinguishability between  
128 approximate and exact unlearning outcomes.129  
130  

## 3 METHODOLOGY

131  
132  

### 3.1 PRIVACY ATTACKS

133 To systematically evaluate privacy leakage in unlearning, we consider two complementary classes of  
134 attacks: *membership inference* and *data reconstruction*.135  
136 **Black-box (U-LiRA).** For the black-box setting, we adopt U-LiRA (Hayes et al., 2025), an  
137 adaptation of LiRA (Carlini et al., 2022a) to unlearning. U-LiRA leverages shadow models trained  
138 and unlearned with the same algorithm as the target, yielding a strong adaptive baseline for auditing  
139 privacy. We defer full algorithmic details to Appendix B.140  
141 **White-box (Gaussian Gradient-Difference).** In the white-box setting, we extend the Gaussian  
142 gradient-difference framework of Leemann et al. (2023) to the unlearning case by contrasting  
143 gradients computed on both the original and unlearned models. This contrast provides a powerful  
144 signal of residual membership leakage when both model versions are available to attacker. The  
145 detailed proposed formulation and test statistic are presented in Appendix C.146  
147 **Reconstruction Attack in Unlearning.** We develop a *white-box* reconstruction attack tailored to  
148 approximate unlearning with retain-set updates. Let  $\Delta\theta = \theta^u - \theta^{\text{org}}$  be the observed parameter  
149 change after one unlearning stage (possibly aggregating multiple optimizer steps). As in gradient  
150 inversion, we seek an input whose parameter-gradient aligns with a target vector; here the natural  
151 target is  $\Delta\theta$ . Our baseline (single-sample) objective is:

152  
153 
$$\hat{x}, \hat{y} \in \arg \min_{x, y} \mathcal{D}(\nabla_{\theta} \ell(f(x; \theta^{\text{org}}), y), \Delta\theta), \quad (1)$$

154 where  $\ell$  is the training loss,  $f(\cdot; \theta)$  the network, and  $\mathcal{D}$  a distance (e.g.,  $\ell_2$  or negative cosine).155 With approximate unlearning, the update  $\Delta\theta$  mixes retain and forget gradients. For a forget example  
156  $(x_f, y_f)$  and a retain minibatch  $\mathcal{B}_r$ ,

157  
158 
$$\Delta\theta \approx -\eta(g_r - \alpha g_f), \quad g_r = \frac{1}{|\mathcal{B}_r|} \sum_{(x_r, y_r) \in \mathcal{B}_r} \nabla_{\theta} \ell(f(x_r; \theta^{\text{org}}), y_r), \quad g_f = \nabla_{\theta} \ell(f(x_f; \theta^{\text{org}}), y_f), \quad (2)$$
  
159  
160

161 with effective step size  $\eta$  and ascent weight  $\alpha > 0$ . Directly targeting  $\Delta\theta$  in equation 1 is therefore  
162 confounded by  $g_r$ . Even when equation 1 is instantiated with state-of-the-art gradient inversion

162 methods, naively inverting the unfiltered update  $\Delta\theta$  remains ineffective, producing low accuracy of  
 163 the reconstruction (see Section 4, Table 2).

164 Let  $G_{\text{org}} = [g(b_i; \theta^{\text{org}})]_{i=1}^m$  and  $G_u = [g(b_i; \theta^u)]_{i=1}^m$  be gradient snapshots on a small probe set drawn  
 165 from the training distribution. We compute thin SVDs,  $G_{\text{org}} = U_{\text{org}} \Sigma_{\text{org}} V_{\text{org}}^\top$  and  $G_u = U_u \Sigma_u V_u^\top$ ,  
 166 and keep the top- $k$  left singular vectors to obtain orthonormal bases (columns) for the dominant  
 167 gradient subspaces. Define the *orthogonal projectors*

$$\Pi_{\text{org}} = U_{\text{org}} U_{\text{org}}^\top, \quad \Pi_u = U_u U_u^\top, \quad \Pi_u^\perp = I - \Pi_u.$$

171 Unlearning attenuates the forget component, so retain gradients are expected to persist in both models,  
 172 whereas the forget component is prominent in  $\theta^{\text{org}}$  but suppressed in  $\theta^u$ . We therefore *orthogonalize*  
 173 the update against the unlearned subspace and keep only directions supported by the original model:

$$\tilde{g}_f = \Pi_{\text{org}} \Pi_u^\perp \left( -\frac{1}{\eta} \Delta\theta \right). \quad (3)$$

174 Intuitively,  $\Pi_u^\perp$  removes directions consistent with retain gradients that remain after unlearning, while  
 175  $\Pi_{\text{org}}$  preserves directions active before unlearning where the forget signal resides. If the retain  
 176 subspace is well captured, then  $\Pi_u^\perp g_r \approx 0$  and  $\Pi_{\text{org}} \Pi_u^\perp (\alpha g_f) \approx \alpha g_f$ , yielding a high-SNR estimate of  
 177 the forget gradient.

178 We reconstruct the forgotten sample by solving the filtered inversion:

$$\hat{x}_f, \hat{y}_f \in \arg \min_{x, y} \mathcal{D}(\nabla_\theta \ell(f(x; \theta^{\text{org}}), y), \tilde{g}_f), \quad (4)$$

179 with optional priors or constraints on  $(x, y)$ . In practice, we choose  $k$  to retain a fixed fraction  
 180 of gradient energy (e.g., 90–95%), which stabilizes the projectors and reliably isolates the forget  
 181 component via orthogonalization. We empirically validate that orthogonal subspace filtering boosts  
 182 reconstruction success across models and datasets; see Section 4.4 and Appendix Table 3.

### 183 3.2 WARP (TELEPORTATION-BASED DEFENSE)

184 **Motivation I: Parameter closeness increases privacy leakage.** We formulate post-hoc unlearning  
 185 as minimizing a composite objective that balances forgetting on  $\mathcal{D}_f$  with utility on  $\mathcal{D}_r$ :

$$\min_{\theta} \underbrace{\ell_f(\theta | \mathcal{D}_f)}_{\text{Forget}} + \lambda \underbrace{\ell_r(\theta | \mathcal{D}_r)}_{\text{Retain}}, \quad \lambda \geq 0, \quad (5)$$

186 where  $\theta$  denotes model parameters;  $\ell_f$  is any differentiable *forgetting surrogate* that penalizes high  
 187 confidence or reduces fidelity on  $\mathcal{D}_f$  (e.g., loss-inflation, uniform/soft labels, margin expansion); and  
 188  $\ell_r$  is the standard training/consistency loss on  $\mathcal{D}_r$  to preserve performance. The trade-off coefficient  $\lambda$   
 189 controls how strongly the unlearning step remains anchored to the retain-set: larger  $\lambda$  keeps  $\theta^u$  closer  
 190 to  $\theta^{\text{org}}$ , preserving accuracy but reducing the parameter shift introduced by forgetting. A first-order  
 191 optimizer with mini-batches  $\mathcal{B}_f \subset \mathcal{D}_f$  and  $\mathcal{B}_r \subset \mathcal{D}_r$  yields the iterative update

$$\theta_{t+1} = \theta_t - \eta_t \left( \nabla_\theta \ell_f(\theta_t | \mathcal{B}_f) + \lambda \nabla_\theta \ell_r(\theta_t | \mathcal{B}_r) \right), \quad (6)$$

192 which encompasses common post-training approximate unlearning schemes; for instance, “negative-  
 193 gradient” methods are recovered by taking  $\ell_f(\cdot) = -\ell_{\text{train}}(\cdot)$  (i.e., ascent on the standard training  
 194 loss over  $\mathcal{D}_f$ ), whereas rehearsal/consistency-based approaches instantiate  $\ell_r$  with supervised loss or  
 195 distillation on  $\mathcal{D}_r$  Thudi et al. (2022); Kurmanji et al. (2023); Chundawat et al. (2023a).

196 Because equation 5 explicitly regularizes utility on  $\mathcal{D}_r$  and is optimized with small steps and early  
 197 stopping on  $\mathcal{D}_f$ , the resulting unlearned parameters  $\theta^u$  typically remain *close* to the original  $\theta^{\text{org}}$   
 198 in parameter space. The displacement  $\Delta\theta = \theta^u - \theta^{\text{org}}$  is well-approximated (to first order) by a  
 199 weighted combination of gradients on the *forget-set*, mildly contaminated by retain gradients Thudi  
 200 et al. (2022); Kurmanji et al. (2023); Huang et al. (2024). This proximity creates a privacy attack  
 201 surface: An adversary with access to  $(\theta^{\text{org}}, \theta^u)$  can leverage  $\Delta\theta$  to perform membership inference or  
 202 gradient-based reconstruction of  $\mathcal{D}_f$  Hu et al. (2024); Bertran et al. (2024), motivating the defenses  
 203 applied over unlearning algorithms.

**Motivation II: Gradient norm and curvature amplify leakage.** Recent evidence suggests that the per-sample gradient trajectory is a strong predictor of privacy vulnerability. Tobaben et al. (2024) show that training examples that accumulate larger gradient norms during optimization are significantly more prone to MIA, reflecting the intuition from differential privacy that each update’s privacy loss scales with gradient magnitude. Complementing this, Ravikumar et al. (2024b) demonstrate that curvature around training samples—captured via local sharpness of the loss—serves as a reliable discriminator between members and non-members, with sharper regions implying higher membership exposure. These findings aligns with theoretical analyses such as Ye et al. (2023), who prove that large per-sample gradients at initialization inflate the KL divergence between neighboring training trajectories, directly increasing the sample’s privacy risk. Motivated by this, we hypothesize that approximate unlearning inherits the same vulnerability: samples with higher gradient norms tend to push parameters towards sharper local extrema during both training and unlearning, thereby overshooting the target update and leaving a stronger privacy footprint. Our experiments (Fig.1) confirm this intuition, revealing a clear correlation between a sample’s gradient norm in the original model and its susceptibility to membership inference after unlearning.

To simultaneously address (i) the parameter–space proximity that enables differencing and (ii) the gradient–norm driver of leakage, we leverage *loss-invariant symmetries* of deep networks.

**Symmetry framework.** Let  $\mathcal{G}$  denote a set of symmetry transformations acting on parameters  $\theta$  (and, when needed, internal representations) such that the task loss is invariant:  $\mathcal{L}(X, \theta) = \mathcal{L}(g(X, \theta))$  for all  $g \in \mathcal{G}$  Zhao et al. (2022; 2023); Armenta et al. (2023); Simsek et al. (2021). A *teleportation* step chooses  $g$  and updates  $\theta \leftarrow g \cdot \theta$ , moving within the loss level set. In our defense, we select  $g$  to reduce the gradient norm of the *forget-set* while preserving utility on the *retain-set*:

$$g^* \in \arg \min_{g \in \mathcal{G}} \left\{ \underbrace{\sum_{(x,y) \in \mathcal{D}_f} \|\nabla_{\theta} \ell(f(x; g \cdot \theta), y)\|_2^2}_{\text{shrink forget-set gradients}} - \beta \underbrace{\|g \cdot \theta - \theta\|_2^2}_{\text{increase parameter dispersion}} \right\} \quad (7)$$

s.t.  $\ell_r(g \cdot \theta | \mathcal{D}_r) \leq \ell_r(\theta | \mathcal{D}_r) + \varepsilon$ .

with trade-off  $\beta \geq 0$  and tolerance  $\varepsilon \geq 0$ . The first term reduces squared gradient norms of forget examples (Motivation II); the dispersion term adds symmetry-preserving randomness, displacing parameters from  $\theta^{\text{org}}$  (Motivation I); the constraint preserves retain performance.

WARP operates on an abstract prediction-preserving symmetry map  $T_\phi$ , and any such symmetry family can instantiate the framework. In practice, we use two concrete realizations—the retain–null-space projection introduced in the next paragraph, and the change-of-basis teleportation detailed in Appendix D—to illustrate this generality. To complement this algorithmic view, Appendix O develops teleportation-aware information-theoretic bounds on gradient-based reconstruction, showing how injecting symmetry-induced noise via  $T_\phi$  expands the symmetry orbit and provably increases the expected reconstruction error for attackers observing  $(\theta^{\text{org}}, \theta^u)$ .

**Primary instantiation: teleportation with retain null-space projection.** We first describe one convenient way to instantiate  $T_\phi$  using retain–null-space projections Wu et al. (2025). To optimize equation 7 efficiently on modern architectures without explicit group actions, we adopt *teleportation with input null-space gradient projection* Wu et al. (2025) and instantiate it using the recent projector formulation that keeps updates on the loss-invariant level set by per-layer projections onto the input null space (thus leaving the task loss unchanged up to numerical error). Concretely, define the *teleportation loss*

$$\mathcal{L}_{\text{tel}}(\theta) = \sum_{(x,y) \in \mathcal{B}_f} \|\nabla_{\theta} \ell(f(x; \theta), y)\|_2^2 - \beta \|\theta - \theta^{\text{org}}\|_2^2,$$

where  $\mathcal{B}_f$  is a minibatch from  $\mathcal{D}_f$ . Let  $R_\ell$  be the per-layer representation matrix from a *retain* minibatch (layer- $\ell$  inputs), with thin SVD  $R_\ell = U_\ell \Sigma_\ell V_\ell^\top$ . We keep the top- $k$  left singular vectors  $B_\ell = U_{\ell,1:k}$  to span the retain subspace and define the orthogonal projector onto its complement  $\Pi_\ell^\perp = I - B_\ell B_\ell^\top$ . A teleportation step then applies the layer-wise update

$$W_\ell^{t+1} \leftarrow W_\ell^t - \eta_{\text{tel}} \Pi_\ell^\perp (\nabla_{W_\ell} \mathcal{L}_{\text{tel}}(\theta^t)) \quad (8)$$

270 which (i) *reduces* the forget-set gradient norms by descending on  $\mathcal{L}_{\text{tel}}$ , (ii) *preserves* the function  
 271 on the retain-set by restricting motion to the retain-orthogonal subspace. The projection operator  
 272 in equation 8 corresponds to the input-null-space projector. This is implemented by subtracting the  
 273 component in the subspace of the core gradient, leaving only the residual for the teleport step.

274 To align the invariance with utility preservation, we compute  $B_\ell$  *only from retain data*. Let  $R_\ell(\mathcal{D}_r) =$   
 275  $[\phi_\ell(x)]_{x \in \mathcal{B}_r}$  denote the matrix formed by stacking the layer- $\ell$  inputs for a retain minibatch  $\mathcal{B}_r$ . Then:

$$277 \quad R_\ell(\mathcal{D}_r) = U_\ell \Sigma_\ell V_\ell^\top, \quad B_\ell = U_{\ell,1:k}, \quad \Pi_\ell^\perp = I - B_\ell B_\ell^\top. \quad (9)$$

279 We set  $k$  to capture a fixed fraction of retain variance (typically 95%–99%) and apply the resulting  
 280 projectors in equation 8. This confines each teleport step to the retain-orthogonal subspace, stabilizing  
 281 predictions on  $\mathcal{D}_r$  while suppressing gradient energy on  $\mathcal{D}_f$ . Since  $\Pi_\ell^\perp$  removes directions spanned  
 282 by retain representations, suitable choices of rank  $k$  and step size  $\eta_{\text{tel}}$  ensure that

$$283 \quad |\ell_r(g \cdot \theta | \mathcal{D}_r) - \ell_r(\theta | \mathcal{D}_r)| \leq \varepsilon,$$

285 which matches the constraint below equation 7; in practice, prediction drift on  $\mathcal{D}_r$  remains within  
 286 numerical tolerance (see Appendix P for hyperparameter sensitivity). To underline that WARP is not  
 287 tied to retain-null-space projections, Appendix D instantiates  $T_\phi$  using the SVD-free change-of-basis  
 288 symmetries introduce in Armenta et al. (2023).

289 **Plug-and-play scope.** Teleportation is interleaved with the standard unlearning update equation 6,  
 290 requiring no training-time per-sample gradients or stored statistics. The update equation 8 is applied  
 291 at intervals  $t \in K \subset 0, \dots, T-1$  (e.g., every  $S$  steps), keeping *forget-set* gradient norms low while  
 292 preserving retention performance. The full algorithm appears in Appendix K.

## 295 4 EXPERIMENTS

297 We now empirically evaluate the proposed teleportation-based defense across multiple unlearning  
 298 algorithms, datasets, and model architectures. Our experiments are designed to answer the following  
 299 research questions: (i) How vulnerable are state-of-the-art unlearning algorithms to privacy attacks  
 300 under both black-box and white-box threat models? (ii) To what extent does teleportation reduce  
 301 membership and reconstruction leakage without sacrificing utility on the retain-set?

302 **Experimental Setup.** We conduct experiments on CIFAR-10, Tiny-ImageNet, and ImageNet-  
 303 1K. On CIFAR-10 we use ResNet-18, while on ImageNet we evaluate ViT-B/16, covering both  
 304 convolutional and transformer models. All models are trained with SGD and standard augmentation.  
 305 Following prior work Kurmanji et al. (2023); Chundawat et al. (2023a), forget sets  $\mathcal{D}_f$  are sampled  
 306 as roughly 1% of training data per class, with retain sets  $\mathcal{D}_r$  comprising the rest.

308 **Baselines.** We benchmark six representative unlearning algorithms—NEGGRAD+ Kurmanji  
 309 et al. (2023), SCRUB Kurmanji et al. (2023), SALUN Fan et al. (2023), PGU Hoang et al. (2024),  
 310 BADTEACHER Chundawat et al. (2023a), and SRF-ON Huang et al. (2024)—covering paradigms of  
 311 gradient ascent, regularization, saliency, projection, and distillation. Full details are in Appendix E.

### 313 4.1 OVERVIEW EFFECTIVENESS OF WARP

315 Figure 2 summarizes privacy and utility across six unlearning methods with and without our plug-in  
 316 defense. Each radar chart reports black-box membership inference risk (AUC and TPR at low FPR),  
 317 accuracy on the most-memorized subset, white-box membership inference risk (AUC and TPR at  
 318 low FPR), and standard test accuracy. The most-memorized subset is selected following our U-LiRA  
 319 protocol in Sec. 4.2, motivated by prior findings that highly memorized samples carry elevated  
 320 unlearning risk Naderloui et al. (2025). For visualization, all metrics are min–max normalized across  
 321 methods. Privacy metrics in which lower is better are inverted by plotting  $1 - \text{metric}$ , so that larger  
 322 polygons correspond to stronger privacy, while higher test accuracy remains preferable.

323 Three key observations emerge. First, no unlearning algorithm dominates across all axes. For  
 instance, SF performs well under black-box auditing but is weaker under white-box auditing and in

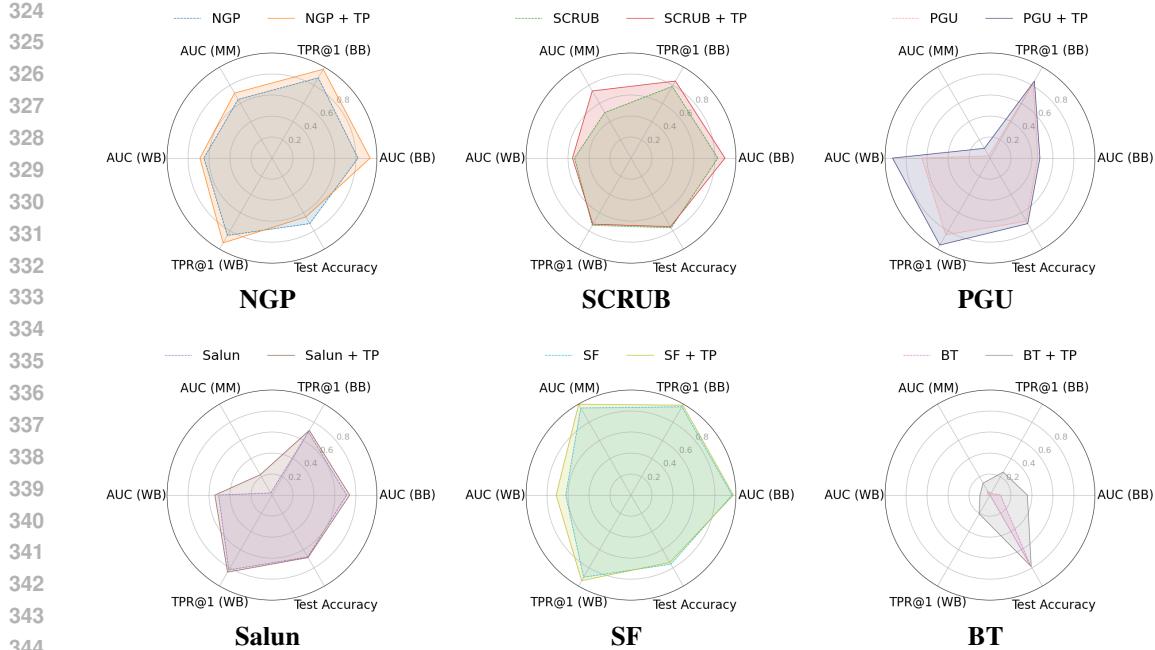


Figure 2: Comparison of unlearning vs. teleportation across six unlearning methods.

test accuracy, illustrating the necessity of evaluating under both threat models. Second, algorithms that appear robust under black-box evaluation such as NGP and SF still exhibit substantial leakage under our white-box test, underscoring the importance of auditing with gradient- or weight-based evidence. Third, adding our symmetry-based teleportation module, instantiated via retain null-space projection, consistently improves privacy across both black-box and white-box metrics while maintaining utility. In some cases, such as BT and SF, teleportation even improves test accuracy. The only noticeable accuracy drop occurs for NGP (about one percentage point), for which we provide a detailed privacy–utility trade-off analysis in Appendix I. The runtime overhead of teleportation is analyzed separately in Appendix J, and Appendix P presents ablations showing that WARP’s performance does not hinge on fragile choices of teleportation hyperparameters. Overall, these results demonstrate that the proposed defense empirically reduces attack success consistently and effectively across a diverse set of unlearning algorithms and threat models. For completeness, we also compare WARP against the strongest noise-based alternative, namely projected DP–Langevin unlearning Chien et al. (2024b), using its formally calibrated update rule; the full comparison is provided in Appendix M.

#### 4.2 U-LiRA (BLACK-BOX)

We evaluate our teleportation defense with U-LiRA Hayes et al. (2025), a state-of-the-art black-box unlearning auditor. Following Deep Unlearn Cadet et al. (2024a), we train  $T = 64$  shadow models with 10 random forget sets each. To model a strong adaptive adversary, shadows use the same unlearning algorithm, teleportation, and hyperparameters as the target, reducing proxy–target miscalibration Cretu et al. (2023). Details of U-LiRA appear in Appendix B.

As emphasized in prior work Carlini et al. (2022a), the most informative regime is low false-positive rates (FPR), where practical attacks must operate. We therefore report AUC as well as TPR@0.1, TPR@1, and TPR@5, which capture attacker success in this stringent regime. In addition, following RULI Naderlou et al. (2025), we stratify the *forget-set* by *memorization* (ranked by training confidence) and evaluate U-LiRA on the most–memorized slice. These points carry elevated privacy risk, so we report low-FPR TPR on this subset alongside aggregate metrics.

Table 1 shows that adding our teleportation plug-in reduces black-box membership leakage across all methods, on both the full *forget-set* and the most–memorized slice, with the largest gains at low FPR. For example, NGP’s TPR@1 nearly halves (0.030 → 0.014), SCRUB’s memorized-slice AUC

378 Table 1: **Privacy (Black-box) with and without WARP.** Reported are risks on *all forget samples*  
 379 and the *most-memorized* 1% (AUC, TPR@0.1/1/5%), plus test accuracy. Each row shows baseline,  
 380 WARP, and relative improvement (%).

381

382

383

384

385

386

387

388

389

390

391

392

393

394

395

396

397

398

399

400

401

402

403

404

405

406

407

408

409

410

411

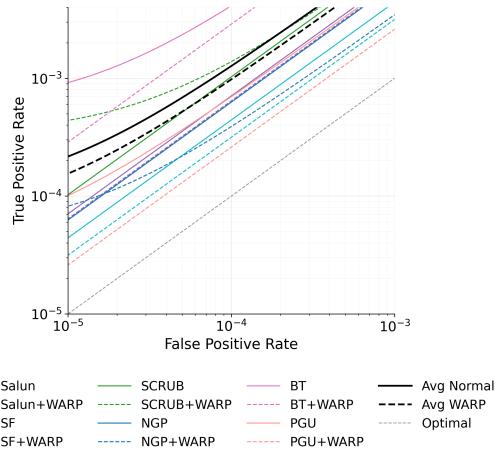
412

413

414

415

Method	All samples (BB)					Most-memorized (top 1%)					Acc.
	AUC	TPR@.1	TPR@1	TPR@5		AUC	TPR@.1	TPR@1	TPR@5	Test	
NGP (base)	0.545	0.012	0.030	0.077	0.649	0.058	0.157	0.277	0.277	<b>0.808</b>	
+ WARP	<b>0.516</b>	<b>0.003</b>	<b>0.014</b>	<b>0.055</b>	<b>0.598</b>	<b>0.015</b>	<b>0.082</b>	<b>0.206</b>	<b>0.206</b>	0.797	
Improvement (%)	64.4	81.8	80.0	81.5	34.2	75.4	51.0	31.3	31.3	-5.7	
SCRUB (base)	0.543	0.020	0.047	0.092	0.710	0.086	0.227	0.397	0.397	<b>0.815</b>	
+ WARP	<b>0.526</b>	<b>0.015</b>	<b>0.036</b>	<b>0.078</b>	<b>0.610</b>	<b>0.041</b>	<b>0.119</b>	<b>0.213</b>	<b>0.213</b>	0.813	
Improvement (%)	39.5	26.3	29.7	33.3	47.6	52.9	49.8	53.0	53.0	-1.1	
PGU (base)	0.636	0.024	0.040	<b>0.098</b>	0.910	0.201	0.511	0.706	0.706	0.804	
+ WARP	<b>0.631</b>	<b>0.018</b>	<b>0.036</b>	0.104	<b>0.875</b>	<b>0.160</b>	<b>0.431</b>	<b>0.663</b>	<b>0.663</b>	<b>0.808</b>	
Improvement (%)	3.7	26.1	13.3	-12.5	8.5	20.5	16.0	6.6	6.6	+2.0	
Salun (base)	0.572	0.020	0.062	0.121	0.910	0.129	0.321	0.520	0.520	0.802	
+ WARP	<b>0.565</b>	<b>0.019</b>	<b>0.059</b>	<b>0.113</b>	<b>0.826</b>	<b>0.107</b>	<b>0.264</b>	<b>0.487</b>	<b>0.487</b>	<b>0.803</b>	
Improvement (%)	9.7	5.3	5.8	11.3	20.5	17.2	18.3	7.0	7.0	+0.5	
SF (base)	0.509	0.004	0.015	0.056	0.518	0.089	0.034	0.079	0.079	<b>0.814</b>	
+ WARP	<b>0.506</b>	<b>0.002</b>	<b>0.012</b>	<b>0.051</b>	<b>0.501</b>	<b>0.006</b>	<b>0.026</b>	<b>0.068</b>	<b>0.068</b>	0.811	
Improvement (%)	33.3	66.7	60.0	83.3	94.4	94.3	33.3	37.9	37.9	-1.6	
BT (base)	0.725	0.000	0.177	0.287	0.902	0.119	0.295	0.582	0.582	0.816	
+ WARP	<b>0.661</b>	0.000	<b>0.137</b>	<b>0.219</b>	<b>0.865</b>	<b>0.113</b>	<b>0.275</b>	<b>0.537</b>	<b>0.537</b>	<b>0.818</b>	
Improvement (%)	28.4	—	24.0	28.7	9.2	5.1	7.0	8.5	8.5	+1.1	



Method	AUC	TPR@.1	TPR@1	TPR@5
NGP (base)	0.642	0.004	0.034	0.139
+ WARP	<b>0.614</b>	<b>0.002</b>	<b>0.021</b>	<b>0.097</b>
Improvement (%)	17.0	50.0	40.6	34.2
SCRUB (base)	0.700	0.011	0.102	0.287
+ WARP	<b>0.657</b>	<b>0.006</b>	<b>0.061</b>	<b>0.193</b>
Improvement (%)	14.3	54.5	42.5	33.5
PGU (base)	0.659	0.007	0.064	0.215
+ WARP	<b>0.533</b>	<b>0.002</b>	<b>0.025</b>	<b>0.085</b>
Improvement (%)	92.9	83.3	64.5	65.5
Salun (base)	0.721	0.008	0.069	0.230
+ WARP	<b>0.705</b>	<b>0.006</b>	<b>0.062</b>	<b>0.214</b>
Improvement (%)	9.5	33.3	10.1	7.0
SF (base)	0.670	0.005	0.043	0.161
+ WARP	<b>0.629</b>	<b>0.003</b>	<b>0.030</b>	<b>0.124</b>
Improvement (%)	29.2	50.0	34.9	23.2
BT (base)	0.938	0.037	0.346	0.809
+ WARP	<b>0.907</b>	<b>0.028</b>	<b>0.279</b>	<b>0.684</b>
Improvement (%)	49.2	25.7	19.4	18.4

416 Figure 3: **White-box privacy with and without WARP.** Gaussian gradient-diff test on 640 unlearned  
 417 models. ROC curves (left) and AUC/TPRs (right); full ROC plots are in Appendix F.

418

419

420

421

422

423

424

425

426

427

428

429

430

431

420 drops by 0.10 (0.710→0.610), and SF’s AUC falls to near-random (0.501). Low-FPR TPR gains  
 421 are often large even when aggregate AUC shifts are modest, showing that teleportation suppresses  
 422 the high-confidence tails attacks exploit. Some methods remain leaky on memorized points, but  
 423 teleportation frequently drives this slice close to random without hurting accuracy. Its impact is  
 424 strongest on TPR@0.1 and TPR@1, as retain-null-space projection reduces forget gradients and  
 425 shrinks extreme margins, weakening the rare signals enabling low-FPR success.

426

427

428

### 4.3 WHITE-BOX MIA

429 We evaluate the Gaussian gradient-difference test of Section C under the setup of Section 4, using  
 430 ResNet-18 on CIFAR-10 and ViT-B/16 on Tiny-ImageNet (full ViT in Appendix H). For the null  
 431 background we draw  $m=1000$  non-members from  $\mathcal{D}_{\text{test}}$ , estimate  $(\hat{\mu}, \hat{\Sigma})$  with ridge  $\lambda=10^{-3}$ , and  
 432 restrict the test to the top-10% most-variant  $\Delta(b)$  coordinates. Figure 3 shows ROC curves with and

432 Table 2: **Effect of teleportation defense** on reconstruction (ImageNet-1K, ResNet-18, NGP).  
433

434 <b>Variant</b>	435 <b>PSNR (dB) <math>\uparrow</math></b>	436 <b>LPIPS (VGG) <math>\downarrow</math></b>	437 <b>LPIPS (Alex) <math>\downarrow</math></b>	438 <b>SSIM <math>\uparrow</math></b>	439 <b>Test MSE <math>\downarrow</math></b>	440 <b>Feat MSE <math>\downarrow</math></b>
Ours (normal unlearning)	$10.74 \pm 0.31$	$0.56 \pm 0.013$	$0.34 \pm 0.015$	$0.12 \pm 0.008$	$0.10 \pm 0.007$	$5.39 \pm 0.50$
Ours + WARP	$7.38 \pm 0.40$	$0.68 \pm 0.01$	$0.46 \pm 0.02$	$0.08 \pm 0.006$	$0.21 \pm 0.02$	$11.28 \pm 1.89$
<i>Improvement of Defense (%)</i>	+45.5	+21.2	+26.1	+31.6	+52.4	+52.2

441 Figure 4: **Reconstructions under NGP vs. NGP+WARP.**  
442

443 without teleportation (log-log for low-FPR). Across methods, teleported variants shift toward chance  
444 ( $TPR = FPR$ ) and flatten between  $10^{-5}$ – $10^{-2}$  FPR, suppressing high-confidence tails. The strongest  
445 effect appears for BT and PGU, which show the largest AUC drops, while NGP, SF, and SALUN  
446 show smaller but consistent shifts. An exception is SCRUB, where teleportation lowers ROC above  
447  $10^{-3}$  FPR but slightly raises TPR at  $< 10^{-3}$ , due to knowledge distillation interacting with symmetry  
448 moves that amplify high-leverage directions. Overall, null-space teleportation reduces white-box  
449 evidence at low FPR, with a narrow corner case for SCRUB.

461 

#### 4.4 RECONSTRUCTION ATTACK RESULTS

462 We evaluate the white-box reconstruction attack of Section 3.1 on **ImageNet-1K** with **ResNet-18**,  
463 focusing on **NGP**. We reconstruct a *single* forgotten example and average over **100** uniformly sampled  
464 points. For each target we use a retain minibatch of size  $|\mathcal{B}_r| = 5$ . Subspace projectors are built  
465 per layer from probe gradients: we draw  $m=100$  training samples to form  $G_{\text{org}}$ ,  $G_u$ , compute thin  
466 SVDs, and keep rank  $k$  preserving 90% gradient energy. We then apply  $\Pi_u^\perp$  and  $\Pi_{\text{org}}$  layerwise to  
467 obtain the filtered target  $\tilde{g}_f$ . The attacker knows the label  $y_f$  and optimizes equation 4 with a TV  
468 regularizer Geiping et al. (2020). The matching loss uses *masked* per-layer gradients: for each layer,  
469 all coordinates are kept and a weighted dot-product alignment is computed Fang et al. (2023).  
470

471 **Effect of teleportation.** Table 2 and Figure 4 compare reconstruction risk under standard NGP  
472 unlearning and its teleported variant using change-of-basis reparameterization. Despite negligible  
473 cost, this symmetry-based randomization disrupts reconstruction: even strong generative-prior  
474 attacks fail to recover meaningful features of forgotten data. Teleportation injects a symmetry  
475 component into  $\Delta\theta$  that is nearly orthogonal to per-sample gradients Armenta et al. (2023), reducing  
476 alignment with the true forget gradient  $g_f$  and driving gradient-matching toward low signal-to-noise  
477 optima. It also undermines our subspace-filtered attack (Eq. 3), since teleportation reshapes gradient  
478 subspaces so  $U_{\text{org}}$  and  $U_u$  overlap little, leaving the residual  $\Pi_{\text{org}}\Pi_u^\perp(-\Delta\theta/\eta)$  small and noisy.  
479 In practice, optimization collapses to the generative prior or class cues, yielding label-consistent  
480 but semantically poor reconstructions (Figure 4). Symmetry moves thus decouple updates from  
481 data-dependent directions, removing the geometric handle exploited by white-box reconstruction.  
482 This motivates examining how teleportation reshapes the information relationship between parameters  
483 and training data (forget-set); a stronger symmetry-aware adaptive reconstruction attack is evaluated  
484 in Appendix N, and Appendix O provides complementary information-theoretic bounds showing how  
485 teleportation expands the symmetry orbit and increases expected reconstruction error.

486 **5 CONCLUSION AND FUTURE WORK**

488 Approximate unlearning provides scalability but introduces privacy risks. We showed that adversaries  
 489 with access to original and unlearned models can mount strong membership inference and reconstruc-  
 490 tion attacks. These risks stem from two properties: parameter proximity and large forget-set gradient  
 491 norms, which amplify leakage.

492 To counter this, we proposed WARP, a symmetry-based defense that interleaves teleportation  
 493 with unlearning. By exploiting network symmetries, WARP reduces forget-set gradient energy  
 494 and displaces parameters in symmetry-preserving directions, weakening both membership and  
 495 reconstruction leakage while preserving retain performance. Across six unlearning algorithms,  
 496 WARP improves privacy, cutting adversarial advantage by up to 64% in black-box and 92% in  
 497 white-box settings. We also stress the need for white-box auditing: methods seemingly robust in  
 498 black-box mode (e.g., SF Huang et al. (2024)) still leak when gradients are exposed. Even simple  
 499 teleportation disrupts reconstruction, reducing quality by  $\sim 45\%$ .

500 Our findings suggest future directions. First, extending Langevin-based privacy analyses to practical  
 501 unlearning with gradient ascent and symmetry moves is promising. Second, recent work shows  
 502 approximate unlearning leaves low-rank weight signals, reversible via re-unlearning Fan et al. (2025)  
 503 or removed by quantization Zhang et al. (2024). Exploring teleportation directly on weights may  
 504 help obscure these signals and mitigate reversals. Finally, as the study of neural network symmetries  
 505 continues to evolve and more efficient estimators and richer invariance families become available,  
 506 WARP can directly inherit these advances by instantiating its symmetry map with stronger or cheaper  
 507 symmetry mechanisms, which further strengthens its resistance to unlearning attacks.

508 **509 REFERENCES**

- 510 Marco Armenta and Pierre-Marc Jodoin. The representation theory of neural networks. *Mathematics*,  
 511 9(24):3216, 2021.
- 513 Marco Armenta, Thierry Judge, Nathan Painchaud, Youssef Skandarani, Carl Lemaire, Gabriel  
 514 Gibeau Sanchez, Philippe Spino, and Pierre-Marc Jodoin. Neural teleportation. *Mathematics*, 11  
 515 (2):480, 2023.
- 516 Sanjeev Arora, Simon Du, Wei Hu, Zhiyuan Li, and Ruosong Wang. Fine-grained analysis of  
 517 optimization and generalization for overparameterized two-layer neural networks. In *International  
 518 conference on machine learning*, pp. 322–332. PMLR, 2019.
- 519 Vijay Badrinarayanan, Bamdev Mishra, and Roberto Cipolla. Symmetry-invariant optimization in  
 520 deep networks. *arXiv preprint arXiv:1511.01754*, 2015.
- 522 Haolei Bai, Siyong Jian, Tuo Liang, Yu Yin, and Huan Wang. Ressvd: Residual compensated svd for  
 523 large language model compression. *arXiv preprint arXiv:2505.20112*, 2025.
- 524 Martin Bertran, Shuai Tang, Michael Kearns, Jamie H Morgenstern, Aaron Roth, and Steven Z Wu.  
 525 Reconstruction attacks on machine unlearning: Simple models are vulnerable. *Advances in Neural  
 526 Information Processing Systems*, 37:104995–105016, 2024.
- 528 Lucas Bourtoule, Varun Chandrasekaran, Christopher A Choquette-Choo, Hengrui Jia, Adelin Travers,  
 529 Baiwu Zhang, David Lie, and Nicolas Papernot. Machine unlearning. In *2021 IEEE symposium  
 530 on security and privacy (SP)*, pp. 141–159. IEEE, 2021.
- 531 Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal,  
 532 Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are  
 533 few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- 534 Xavier F Cadet, Anastasia Borovykh, Mohammad Malekzadeh, Sara Ahmadi-Abhari, and Hamed  
 535 Haddadi. Deep unlearn: Benchmarking machine unlearning. *arXiv preprint arXiv:2410.01276*,  
 536 2024a.
- 538 Xavier F. Cadet, Anastasia Borovykh, Mohammad Malekzadeh, Sara Ahmadi-Abhari, and Hamed  
 539 Haddadi. Deep Unlearn: Benchmarking Machine Unlearning, October 2024b. URL <http://arxiv.org/abs/2410.01276> [cs].

- 540 Yinzh Cao and Junfeng Yang. Towards making systems forget with machine unlearning. In *2015*  
 541 *IEEE symposium on security and privacy*, pp. 463–480. IEEE, 2015.  
 542
- 543 Nicholas Carlini, Steve Chien, Milad Nasr, Shuang Song, Andreas Terzis, and Florian Tramer.  
 544 Membership inference attacks from first principles. In *2022 IEEE symposium on security and*  
 545 *privacy (SP)*, pp. 1897–1914. IEEE, 2022a.
- 546 Nicholas Carlini, Matthew Jagielski, Chiyuan Zhang, Nicolas Papernot, Andreas Terzis, and Florian  
 547 Tramer. The privacy onion effect: Memorization is relative. *Advances in Neural Information*  
 548 *Processing Systems*, 35:13263–13276, 2022b.  
 549
- 550 Sungmin Cha, Sungjun Cho, Dasol Hwang, Honglak Lee, Taesup Moon, and Moontae Lee. Learning  
 551 to unlearn: Instance-wise unlearning for pre-trained classifiers. In *Proceedings of the AAAI*  
 552 *conference on artificial intelligence*, volume 38, pp. 11186–11194, 2024.
- 553 Min Chen, Zhikun Zhang, Tianhao Wang, Michael Backes, Mathias Humbert, and Yang Zhang. When  
 554 machine unlearning jeopardizes privacy. In *Proceedings of the 2021 ACM SIGSAC conference on*  
 555 *computer and communications security*, pp. 896–911, 2021.  
 556
- 557 Eli Chien, Haoyu Wang, Ziang Chen, and Pan Li. Certified machine unlearning via noisy stochastic  
 558 gradient descent. *Advances in Neural Information Processing Systems*, 37:38852–38887, 2024a.  
 559
- 560 Eli Chien, Haoyu Wang, Ziang Chen, and Pan Li. Langevin unlearning: A new perspective of noisy  
 561 gradient descent for machine unlearning. *Advances in neural information processing systems*, 37:  
 562 79666–79703, 2024b.
- 563 Vikram S Chundawat, Ayush K Tarun, Murari Mandal, and Mohan Kankanhalli. Can bad teaching  
 564 induce forgetting? unlearning in deep networks using an incompetent teacher. In *Proceedings of*  
 565 *the AAAI Conference on Artificial Intelligence*, pp. 7210–7217, 2023a.  
 566
- 567 Vikram S Chundawat, Ayush K Tarun, Murari Mandal, and Mohan Kankanhalli. Zero-shot machine  
 568 unlearning. *IEEE Transactions on Information Forensics and Security*, 18:2345–2354, 2023b.
- 569 Ana-Maria Cretu, Daniel Jones, Yves-Alexandre de Montjoye, and Shruti Tople. Investigating  
 570 the effect of misalignment on membership privacy in the white-box setting. *arXiv preprint*  
 571 *arXiv:2306.05093*, 2023.  
 572
- 573 Simon S Du, Wei Hu, and Jason D Lee. Algorithmic regularization in learning deep homogeneous  
 574 models: Layers are automatically balanced. *Advances in neural information processing systems*,  
 575 31, 2018.
- 576 Chongyu Fan, Jiancheng Liu, Yihua Zhang, Eric Wong, Dennis Wei, and Sijia Liu. Salun: Em-  
 577 powering machine unlearning via gradient-based weight saliency in both image classification and  
 578 generation. *arXiv preprint arXiv:2310.12508*, 2023.  
 579
- 580 Chongyu Fan, Jiancheng Liu, Alfred Hero, and Sijia Liu. Challenging forgets: Unveiling the worst-  
 581 case forget sets in machine unlearning. In *European Conference on Computer Vision*, pp. 278–297.  
 582 Springer, 2024.  
 583
- 584 Chongyu Fan, Jinghan Jia, Yihua Zhang, Anil Ramakrishna, Mingyi Hong, and Sijia Liu. Towards  
 585 llm unlearning resilient to relearning attacks: A sharpness-aware minimization perspective and  
 586 beyond. *arXiv preprint arXiv:2502.05374*, 2025.
- 587 Hao Fang, Bin Chen, Xuan Wang, Zhi Wang, and Shu-Tao Xia. Gifd: A generative gradient  
 588 inversion method with feature domain optimization. In *Proceedings of the IEEE/CVF International*  
 589 *Conference on Computer Vision*, pp. 4967–4976, 2023.  
 590
- 591 Stanislav Fort, Gintare Karolina Dziugaite, Maithra Raghu Paul, Surya Ganguli, and David J. Schwab.  
 592 Deep learning versus kernel learning: An empirical study of loss landscape geometry and the time  
 593 evolution of the neural tangent kernel. In *Advances in Neural Information Processing Systems*,  
 volume 33, pp. 5850–5861, 2020.

- 594 Jonas Geiping, Hartmut Bauermeister, Hannah Dröge, and Michael Moeller. Inverting gradients-how  
 595 easy is it to break privacy in federated learning? *Advances in neural information processing*  
 596 *systems*, 33:16937–16947, 2020.
- 597
- 598 Kristian Georgiev, Roy Rinberg, Sung Min Park, Shivam Garg, Andrew Ilyas, Aleksander Madry,  
 599 and Seth Neel. Attribute-to-delete: Machine unlearning via datamodel matching. *arXiv preprint*  
 600 *arXiv:2410.23232*, 2024.
- 601 Behrooz Ghorbani, Shankar Krishnan, and Ying Xiao. An investigation into neural net optimization  
 602 via hessian eigenvalue density. In *Proceedings of the 36th International Conference on Machine*  
 603 *Learning (ICML)*, volume 97 of *Proceedings of Machine Learning Research*, pp. 2232–2241.  
 604 PMLR, 2019.
- 605
- 606 Grzegorz Głuch and Rüdiger Urbanke. Noether: The more things change, the more stay the same.  
 607 *arXiv preprint arXiv:2104.05508*, 2021.
- 608
- 609 Aditya Golatkar, Alessandro Achille, and Stefano Soatto. Forgetting outside the box: Scrubbing deep  
 610 networks of information accessible from input-output observations. In *European Conference on*  
 611 *Computer Vision*, pp. 383–398. Springer, 2020.
- 612
- 613 Gene H. Golub and Charles F. Van Loan. *Matrix Computations*. Johns Hopkins University Press,  
 Baltimore, MD, 4 edition, 2013.
- 614
- 615 Laura Graves, Vineel Nagisetty, and Vijay Ganesh. Amnesiac machine learning. In *Proceedings of*  
 616 *the AAAI Conference on Artificial Intelligence*, volume 35, pp. 11516–11524, 2021.
- 617
- 618 Chiyuan Guo, Tom Goldstein, Awni Hannun, and Laurens van der Maaten. Certified data removal  
 619 from machine learning models. In *International Conference on Machine Learning*, pp. 3832–3842.  
 PMLR, 2020.
- 620
- 621 Guy Gur-Ari, Daniel A. Roberts, and Ethan Dyer. Gradient descent happens in a tiny subspace. *arXiv*  
 622 *preprint arXiv:1812.04754*, 2018.
- 623
- 624 Nathan Halko, Per-Gunnar Martinsson, and Joel A. Tropp. Finding structure with randomness:  
 625 Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM Review*, 53  
 (2):217–288, 2011.
- 626
- 627 Jamie Hayes, Ilia Shumailov, Eleni Triantafillou, Amr Khalifa, and Nicolas Papernot. Inexact  
 628 unlearning needs more careful evaluations to avoid a false sense of privacy. In *2025 IEEE*  
 629 *Conference on Secure and Trustworthy Machine Learning (SaTML)*, pp. 497–519. IEEE, 2025.
- 630
- 631 Tuan Hoang, Santu Rana, Sunil Gupta, and Svetha Venkatesh. Learn to unlearn for deep neural  
 632 networks: Minimizing unlearning interference with gradient projection. In *Proceedings of the*  
 633 *IEEE/CVF Winter Conference on Applications of Computer Vision*, pp. 4819–4828, 2024.
- 634
- 635 Yen-Chang Hsu, Ting Hua, Sungjen Chang, Qian Lou, Yilin Shen, and Hongxia Jin. Language model  
 636 compression with weighted low-rank factorization. *arXiv preprint arXiv:2207.00112*, 2022.
- 637
- 638 Hongsheng Hu, Shuo Wang, Tian Dong, and Minhui Xue. Learn what you want to unlearn: Unlearning  
 639 inversion attacks against machine unlearning. In *2024 IEEE Symposium on Security and Privacy*  
 (SP), pp. 3257–3275. IEEE, 2024.
- 640
- 641 Zhehao Huang, Xinwen Cheng, JingHao Zheng, Haoran Wang, Zhengbao He, Tao Li, and Xiaolin  
 642 Huang. Unified gradient-based machine unlearning with remain geometry enhancement. *Advances*  
 643 *in Neural Information Processing Systems*, 37:26377–26414, 2024.
- 644
- 645 Jinwoo Jeon, Kangwook Lee, Sewoong Oh, Jungseul Ok, et al. Gradient inversion with generative  
 646 image prior. *Advances in neural information processing systems*, 34:29898–29908, 2021.
- 647
- 648 Daniel Kunin, Javier Sagastuy-Brena, Surya Ganguli, Daniel LK Yamins, and Hidenori Tanaka.  
 649 Neural mechanics: Symmetry and broken conservation laws in deep learning dynamics. *arXiv*  
 650 *preprint arXiv:2012.04728*, 2020.

- 648 Meghdad Kurmanji, Peter Triantafillou, Jamie Hayes, and Eleni Triantafillou. Towards unbounded  
 649 machine unlearning. *Advances in neural information processing systems*, 36:1957–1987, 2023.  
 650
- 651 Tobias Leemann, Martin Pawelczyk, and Gjergji Kasneci. Gaussian membership inference privacy.  
 652 *Advances in Neural Information Processing Systems*, 36:73866–73878, 2023.
- 653 Zhuohang Li, Jiaxin Zhang, Luyang Liu, and Jian Liu. Auditing privacy defenses in federated  
 654 learning via generative gradient leakage. In *Proceedings of the IEEE/CVF conference on computer*  
 655 *vision and pattern recognition*, pp. 10132–10142, 2022.
- 656
- 657 Junxu Liu, Mingsheng Xue, Jian Lou, Xiaoyu Zhang, Li Xiong, and Zhan Qin. Muter: Machine  
 658 unlearning on adversarially trained models. In *Proceedings of the IEEE/CVF international*  
 659 *conference on computer vision*, pp. 4892–4902, 2023.
- 660 Ronak Mehta, Sourav Pal, Vikas Singh, and Sathya N Ravi. Deep unlearning via randomized  
 661 conditionally independent hessians. In *Proceedings of the IEEE/CVF Conference on Computer*  
 662 *Vision and Pattern Recognition*, pp. 10422–10431, 2022.
- 663
- 664 Ioannis Mitliagkas, Constantine Caramanis, Prateek Jain, Shiva Prasad Kasiviswanathan, and Sanjiv  
 665 Kumar. Memory limited, streaming PCA. In *Advances in Neural Information Processing Systems*,  
 666 volume 26, 2013.
- 667 Cameron Musco and Christopher Musco. Randomized block krylov methods for stronger and faster  
 668 approximate singular value decomposition. In *Advances in Neural Information Processing Systems*,  
 669 volume 28, 2015.
- 670 Nima Naderloui, Shenao Yan, Binghui Wang, Jie Fu, Wendy Hui Wang, Weiran Liu, and Yuan Hong.  
 671 Rectifying privacy and efficacy measurements in machine unlearning: A new inference attack  
 672 perspective. *arXiv preprint arXiv:2506.13009*, 2025.
- 673
- 674 Milad Nasr, Reza Shokri, and Amir Houmansadr. Comprehensive privacy analysis of deep learning:  
 675 Passive and active white-box inference attacks against centralized and federated learning. In *2019*  
 676 *IEEE symposium on security and privacy (SP)*, pp. 739–753. IEEE, 2019.
- 677 Thanh Tam Nguyen, Thanh Trung Huynh, Zhao Ren, Phi Le Nguyen, Alan Wee-Chung Liew,  
 678 Hongzhi Yin, and Quoc Viet Hung Nguyen. A survey of machine unlearning. *arXiv preprint*  
 679 *arXiv:2209.02299*, 2022.
- 680
- 681 Erkki Oja. A simplified neuron model as a principal component analyzer. *Journal of Mathematical*  
 682 *Biology*, 15(3):267–273, 1982.
- 683
- 684 Xinbao Qiao, Meng Zhang, Ming Tang, and Ermin Wei. Hessian-free online certified unlearning.  
 685 *arXiv preprint arXiv:2404.01712*, 2024.
- 686
- 687 Deepak Ravikumar, Efstathia Souflieri, Abolfazl Hashemi, and Kaushik Roy. Unveiling privacy,  
 688 memorization, and input curvature links. *arXiv preprint arXiv:2402.18726*, 2024a.
- 689
- 690 Deepak Ravikumar, Efstathia Souflieri, and Kaushik Roy. Curvature clues: Decoding deep learning  
 691 privacy with input loss curvature. *Advances in Neural Information Processing Systems*, 37:  
 692 20003–20030, 2024b.
- 693
- 694 Maria Rigaki and Sebastian Garcia. A survey of privacy attacks in machine learning. *ACM Computing*  
 695 *Surveys*, 56(4):1–34, 2023.
- 696
- 697 Charbel Sakr and Brucek Khailany. Espace: Dimensionality reduction of activations for model  
 698 compression. *Advances in Neural Information Processing Systems*, 37:17489–17517, 2024.
- 699
- 700 Ahmed Salem, Apratim Bhattacharya, Michael Backes, Mario Fritz, and Yang Zhang. {Updates-  
 701 Leak}: Data set inference and reconstruction attacks in online learning. In *29th USENIX security*  
 702 *symposium (USENIX Security 20)*, pp. 1291–1308, 2020.
- 703
- 704 Seonguk Seo, Dongwan Kim, and Bohyung Han. Revisiting machine unlearning with dimensional  
 705 alignment. In *2025 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*,  
 706 pp. 3206–3215. IEEE, 2025.

- 702 Nazanin Mohammadi Sepahvand, Anvith Thudi, Berivan Isik, Ashmita Bhattacharyya, Nicolas  
 703 Papernot, Eleni Triantafillou, Daniel M Roy, and Gintare Karolina Dziugaite. Leveraging per-  
 704 instance privacy for machine unlearning. *arXiv preprint arXiv:2505.18786*, 2025.
- 705 Reza Shokri, Marco Stronati, Congzheng Song, and Vitaly Shmatikov. Membership inference attacks  
 706 against machine learning models. In *2017 IEEE symposium on security and privacy (SP)*, pp. 3–18.  
 707 IEEE, 2017.
- 708 Berfin Simsek, François Ged, Arthur Jacot, Francesco Spadaro, Clément Hongler, Wulfram Gerstner,  
 709 and Johanni Brea. Geometry of the loss landscape in overparameterized neural networks: Symme-  
 710 tries and invariances. In *International Conference on Machine Learning*, pp. 9722–9732. PMLR,  
 711 2021.
- 712 Anvith Thudi, Gabriel Deza, Varun Chandrasekaran, and Nicolas Papernot. Unrolling sgd: Under-  
 713 standing factors influencing machine unlearning. In *2022 IEEE 7th European Symposium on  
 714 Security and Privacy (EuroS&P)*, pp. 303–319. IEEE, 2022.
- 715 Marlon Tobaben, Gauri Pradhan, Yuan He, Joonas Jälkö, and Antti Honkela. Understanding practical  
 716 membership privacy of deep learning. In *Privacy Regulation and Protection in Machine Learning*,  
 717 2024.
- 718 Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée  
 719 Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and  
 720 efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023.
- 721 Joel A. Tropp, Alp Yurtsever, Madeleine Udell, and Volkan Cevher. Practical sketching algorithms  
 722 for low-rank matrix approximation. *SIAM Journal on Matrix Analysis and Applications*, 38(4):  
 723 1454–1485, 2017.
- 724 Jana Vatter, Ruben Mayer, and Hans-Arno Jacobsen. The evolution of distributed systems for graph  
 725 neural networks and their origin in graph processing and deep learning: A survey. *ACM Computing  
 726 Surveys*, 56(1):1–37, 2023.
- 727 Qiyuan Wang, Ruiling Xu, Shibo He, Randall Berry, and Meng Zhang. Unlearning incentivizes  
 728 learning under privacy risk. In *Proceedings of the ACM on Web Conference 2025*, pp. 1456–1467,  
 729 2025a.
- 730 Xin Wang, Yu Zheng, Zhongwei Wan, and Mi Zhang. Svd-llm: Truncation-aware singular value  
 731 decomposition for large language model compression. *arXiv preprint arXiv:2403.07378*, 2024.
- 732 Xin Wang, Samiul Alam, Zhongwei Wan, Hui Shen, and Mi Zhang. Svd-llm v2: Optimizing singular  
 733 value truncation for large language model compression. *arXiv preprint arXiv:2503.12340*, 2025b.
- 734 Manfred K. Warmuth and Dima Kuzmin. Randomized online PCA algorithms with regret bounds.  
 735 *Journal of Machine Learning Research*, 9(10):2287–2320, 2008.
- 736 David P. Woodruff. Sketching as a tool for numerical linear algebra. *Foundations and Trends in  
 737 Theoretical Computer Science*, 10(1–2):1–157, 2014.
- 738 Zihao Wu, Juncheng Dong, Ahmed Aloui, and Vahid Tarokh. Teleportation with null space gradient  
 739 projection for optimization acceleration. *arXiv preprint arXiv:2502.11362*, 2025.
- 740 Song Xia, Yi Yu, Wenhan Yang, Meiwen Ding, Zhuo Chen, Ling-Yu Duan, Alex C Kot, and Xudong  
 741 Jiang. Theoretical insights in model inversion robustness and conditional entropy maximization for  
 742 collaborative inference systems. In *Proceedings of the Computer Vision and Pattern Recognition  
 743 Conference*, pp. 8753–8763, 2025.
- 744 Jiayuan Ye, Zhenyu Zhu, Fanghui Liu, Reza Shokri, and Volkan Cevher. Initialization matters:  
 745 Privacy-utility analysis of overparameterized neural networks. *Advances in Neural Information  
 746 Processing Systems*, 36:5419–5446, 2023.
- 747 Hongxu Yin, Arun Mallya, Arash Vahdat, Jose M Alvarez, Jan Kautz, and Pavlo Molchanov. See  
 748 through gradients: Image batch recovery via gradinversion. In *Proceedings of the IEEE/CVF  
 749 conference on computer vision and pattern recognition*, pp. 16337–16346, 2021.

756 Zhiwei Zhang, Fali Wang, Xiaomin Li, Zongyu Wu, Xianfeng Tang, Hui Liu, Qi He, Wenpeng  
 757 Yin, and Suhang Wang. Catastrophic failure of llm unlearning via quantization. *arXiv preprint*  
 758 *arXiv:2410.16454*, 2024.

760 Bo Zhao, Nima Dehmamy, Robin Walters, and Rose Yu. Symmetry teleportation for accelerated  
 761 optimization. *Advances in neural information processing systems*, 35:16679–16690, 2022.

762 Bo Zhao, Robert M Gower, Robin Walters, and Rose Yu. Improving convergence and generalization  
 763 using parameter symmetries. *arXiv preprint arXiv:2305.13404*, 2023.

765 Kairan Zhao, Meghdad Kurmanji, George-Octavian Bărbulescu, Eleni Triantafillou, and Peter Tri-  
 766 antafillou. What makes unlearning hard and what to do about it. *Advances in Neural Information*  
 767 *Processing Systems*, 37:12293–12333, 2024.

## 768 A RELATED WORK

771 **Approximate Unlearning.** The removal of training samples was introduced by Cao & Yang  
 772 (2015) in the context of the “right to be forgotten.” Retraining from scratch guarantees deletion but is  
 773 infeasible for modern networks Vatter et al. (2023). Exact unlearning methods such as SISA Bourtoule  
 774 et al. (2021) and Amnesiac Unlearning Graves et al. (2021) lower cost through partitioning or selective  
 775 retraining but still require storage and scale poorly Nguyen et al. (2022).

777 Approximate unlearning directly updates the trained model to erase the *forget-set* Kurmanji et al.  
 778 (2023); Chundawat et al. (2023a); Golatkar et al. (2020); Thudi et al. (2022). These methods aim to  
 779 match the predictive distribution of retraining while preserving retain accuracy, offering a practical  
 780 forgetting–utility trade-off with large savings in computation and memory. Related methods target  
 781 structured forget sets such as entire classes Chundawat et al. (2023b); Seo et al. (2025), or tackle the  
 782 harder instance-wise setting, where arbitrary samples must be removed Fan et al. (2024); Cha et al.  
 783 (2024); Zhao et al. (2024). Many approaches rely on training-time side information like per-sample  
 784 gradients Qiao et al. (2024); Mehta et al. (2022), or assume specialized regimes with adversarial  
 785 robustness Liu et al. (2023) or differential-privacy noise Chien et al. (2024b;a); Sepahvand et al.  
 786 (2025). While effective, these assumptions add resource overhead, limiting post-hoc use. Our focus,  
 787 therefore, is training-agnostic, instance-wise unlearning that takes only a pretrained classifier and  
 788 a designated *forget-set*, without stored gradients or training modifications Kurmanji et al. (2023);  
 789 Thudi et al. (2022).

790 **Privacy Unlearning.** The effectiveness of approximate unlearning is accessed by two criteria: (I)  
 791 the model should maintain accuracy on non-forgotten data, and (II) its outputs on the *forget-set* should  
 792 be indistinguishable from those of a model with no access to it Naderloui et al. (2025). In practice,  
 793 this is evaluated using MIA Shokri et al. (2017); Carlini et al. (2022a), which test whether a sample  
 794 was part of training. Effective unlearning removes this membership advantage on the *forget-set*.

795 Most prior work evaluates unlearning by comparing outputs of the unlearned model to a retrained  
 796 reference on the *forget-set* Cadet et al. (2024b); Kurmanji et al. (2023); Hayes et al. (2025); Georgiev  
 797 et al. (2024); Naderloui et al. (2025). This black-box view ignores parameters, even though in  
 798 practice—such as MU on edge devices—an adversary may access both original and unlearned models.  
 799 Some studies consider this stronger setting: Chen et al. (2021) showed that output-comparison  
 800 across models can detect unlearning, while others adapted reconstruction to infer forgotten data  
 801 from parameter differences Salem et al. (2020); Hu et al. (2024); Bertran et al. (2024). These works,  
 802 however, are limited to toy models and simplified updates, leaving privacy risk under realistic  
 803 conditions unclear. In particular, they do not capture the robustness of recent multi-step approximate  
 804 methods such as NGP or SCRUB Kurmanji et al. (2023); Chundawat et al. (2023a), where iterative  
 805 updates with retain-set supervision weaken inversion of *forget-set* gradients. We address this gap  
 806 with stronger white-box MIAs (Sec.C) and DRAs (Sec.3.1) tailored to realistic unlearning.

807 **Neural Network Symmetry.** Continuous symmetries in neural networks arise when transformations  
 808 of the weights leave the output unchanged. Such invariances, a byproduct of overparameterization,  
 809 mean that many distinct weight configurations represent the same function Gluch & Urbanke (2021).  
 They appear in homogeneous activations Badrinarayanan et al. (2015); Du et al. (2018) and in

810 components like softmax and batch normalization Kunin et al. (2020), and have been linked to  
 811 both improved optimization and generalization. Neural teleportation leverages these symmetries by  
 812 relocating parameters within the loss-invariant level set, yielding equivalent models that accelerate  
 813 optimization Armenta & Jodoin (2021); Armenta et al. (2023). Building on this idea, Zhao et al.  
 814 (2022) introduced symmetry teleportation, which searches for beneficial relocations while providing  
 815 a framework for analyzing symmetry-induced minima. More recently, teleportation with null-space  
 816 gradient projection Wu et al. (2025) leverages the input null space: moving along projected directions  
 817 leaves the function unchanged, directly aligning with the goal of teleportation.

## 819 B U-LiRA ALGORITHM

820 To evaluate sample-wise privacy leakage, we employ the U-LiRA attack Cadet et al. (2024a); Hayes  
 821 et al. (2025), an adaptation of LiRA Carlini et al. (2022a) to the unlearning setting. The attack relies  
 822 on shadow models to estimate two distributions for a target sample  $(x, y)$ : (i) models trained with  
 823  $(x, y)$  and subsequently unlearned using the same unlearning algorithm, and (ii) models trained from  
 824 scratch without  $(x, y)$ . By fitting simple parametric models (e.g., Gaussians) to the outputs of these  
 825 shadow ensembles, U-LiRA computes the likelihood of the target model’s output under each case  
 826 and classifies membership according to a likelihood ratio test.

827 Crucially, all shadow models are trained with the *same unlearning algorithm and hyperparameters* as  
 828 the audited model. This makes U-LiRA effectively an *adaptive attack*, since it tailors the proxies  
 829 to each specific unlearning method. Such alignment minimizes miscalibration between shadow and  
 830 target models and is known to increase attack success Cretu et al. (2023). Therefore, U-LiRA serves  
 831 as a strong black-box baseline for auditing privacy in unlearning. A complete description of the  
 832 algorithm can be demonstrated in Algorithm 1.

---

### 835 Algorithm 1 U-LiRA (used for auditing unlearning)

---

836 **Require:** Target model  $\theta^*$ , learning algorithm  $A$ , unlearning algorithm  $U$ , number of shadows  $T$ ,  
 837 sample  $(x, y)$

838 **Ensure:** Prediction: is  $(x, y)$  in the *forget-set*?

839 1: Initialize empty lists  $O \leftarrow \{\}$  and  $\hat{O} \leftarrow \{\}$   
 840 2: **for**  $t = 1$  to  $T$  **do**  
 841 3:   Sample dataset  $D$  containing  $(x, y)$   
 842 4:   Train  $\theta^0 \leftarrow A(D)$   
 843 5:   Unlearn  $\theta^f \leftarrow U(\theta^0, \{(x, y)\})$   
 844 6:   Retrain  $\theta^r \leftarrow A(D \setminus \{(x, y)\})$   
 845 7:   Record  $O[t] \leftarrow f(x; \theta^f)_y, \hat{O}[t] \leftarrow f(x; \theta^r)_y$   
 846 8: **end for**  
 847 9: Fit Gaussian  $(\mu, \sigma^2)$  to  $O$ , and  $(\hat{\mu}, \hat{\sigma}^2)$  to  $\hat{O}$   
 848 10: Compute  $o^* \leftarrow f(x; \theta^*)_y$   
 849 11: Compute likelihood ratio:

$$850 \quad p_{\text{member}} = \frac{\mathcal{N}(o^*; \mu, \sigma^2)}{\mathcal{N}(o^*; \mu, \sigma^2) + \mathcal{N}(o^*; \hat{\mu}, \hat{\sigma}^2)}$$

851 12: **if**  $p_{\text{member}} > 0.5$  **then**  
 852 13:   **return** “member of training”  
 853 14: **else**  
 854 15:   **return** “non-member”  
 855 16: **end if**

---

## 859 C WHITE-BOX GAUSSIAN GRADIENT-DIFFERENCE ATTACK ALGORITHM

860 Guided by the GLiR framework of Leemann et al. (2023), we formulate sample-wise MIA in the  
 861 unlearning setting as a binary hypothesis test that uses *both* the pre-unlearning and post-unlearning  
 862 models. Let  $A$  denote the training algorithm,  $U$  the unlearning operator,  $S$  the original training set,

864 and  $F \subseteq S$  the forget subset. For a candidate example  $(x, y)$ , we test  
 865

$$866 \quad H_0 : (x, y) \sim \mathcal{D}_{\text{test}}, \quad (\theta^{\text{org}}, \theta^u) = (A(S), U(A(S), F)) \text{ with } x \notin S, x \notin F, \\ 867 \quad H_1 : (x, y) \in \mathcal{D}_{\text{forg}}, \quad (\theta^{\text{org}}, \theta^u) = (A(S), U(A(S), F)) \text{ with } x \in S \text{ and } x \in F,$$

868 i.e., under  $H_1$  the point participated in the original training and was subsequently targeted by  
 869 unlearning, whereas under  $H_0$  it was never used. With white-box access, we form the gradient-  
 870 difference statistic

$$871 \quad \Delta(x) = \nabla_{\theta} \ell(f(x; \theta^u), y) - \nabla_{\theta} \ell(f(x; \theta^{\text{org}}), y) \in \mathbb{R}^d. \\ 872$$

873 Assuming access to draws from  $\mathcal{D}_{\text{test}}$ , the adversary builds a background set  $B = \{(b_i, \tilde{y}_i)\}_{i=1}^m \sim \mathcal{D}_{\text{test}}^m$   
 874 and estimates the null (non-member) distribution of gradient differences via

$$875 \quad \hat{\mu} = \frac{1}{m} \sum_{i=1}^m \Delta(b_i), \quad \hat{\Sigma} = \frac{1}{m-1} \sum_{i=1}^m (\Delta(b_i) - \hat{\mu})(\Delta(b_i) - \hat{\mu})^{\top}. \\ 876 \\ 877$$

878 Following Leemann et al. (2023), we adopt a Gaussian model for  $\Delta(x)$  under  $H_0$  and compute the  
 879 whitened Mahalanobis statistic

$$880 \quad s(x) = (\Delta(x) - \hat{\mu})^{\top} (\hat{\Sigma} + \lambda I)^{-1} (\Delta(x) - \hat{\mu}), \\ 881$$

882 with a small ridge  $\lambda > 0$  for numerical stability. Under  $H_0$ ,  $s(x)$  is approximately  $\chi_d^2$ -distributed,  
 883 yielding the log- $p$ -value score

$$884 \quad A'(x, y) = -\log(1 - F_{\chi_d^2}(s(x))), \\ 885$$

886 and the final decision rule

$$887 \quad A(x, y) = \mathbb{I}[A'(x, y) > \tau],$$

888 predicting *forgotten* when the score exceeds threshold  $\tau$ . Algorithm 2 provides the full details of the  
 889 proposed attack.

890

891 **Relation to GLiR and unlearning specifics.** GLiR aggregates evidence across training steps by  
 892 comparing per-step sample gradients to a Gaussian background of batch gradients; our adaptation  
 893 replaces the (typically unavailable) per-step trajectory with the two-model contrast  $\Delta(x)$ . The  
 894 geometry is unchanged: Evidence corresponds to the squared norm of the whitened difference,  
 895  $\|(\hat{\Sigma} + \lambda I)^{-1/2} \Delta(x)\|_2^2$ . Unlike standard MIAs that query a single model, the test exploits white-box  
 896 access to  $\theta^{\text{org}}$  and  $\theta^u$  and targets the unlearning-specific alternative  $H_1$  (membership in both  $S$  and  
 897  $F$ ), providing a simple and powerful auditor for residual leakage after unlearning.

---

898 **Algorithm 2** White-box Gaussian Gradient-Difference Attack for Unlearning Audit

---

900 **Require:** Pre-unlearning model  $\theta^{\text{org}}$ , post-unlearning model  $\theta^u$ , candidate sample  $(x, y)$ , loss  $\ell$ ,  
 901 predictor  $f(\cdot; \theta)$ , background sampler  $\mathcal{S}_{\text{test}}(m)$  that returns  $m$  i.i.d. draws from  $\mathcal{D}_{\text{test}}$

902 **Require:** Hyperparameters: background size  $m$ , repetitions  $T$ , ridge  $\lambda > 0$ , decision threshold  $\tau$

1:  $S \leftarrow 0$  ▷ initialize cumulative evidence

2: **for**  $t = 1$  to  $T$  **do**

3:    $B_t = \{(b_i, \tilde{y}_i)\}_{i=1}^m \leftarrow \mathcal{S}_{\text{test}}(m)$  ▷ if labels are unavailable, set  $\tilde{y}_i = \arg \max f(b_i; \theta^{\text{org}})$

4:   **for**  $i = 1$  to  $m$  **do**

5:      $\Delta_i \leftarrow \nabla_{\theta} \ell(f(b_i; \theta^u), \tilde{y}_i) - \nabla_{\theta} \ell(f(b_i; \theta^{\text{org}}), \tilde{y}_i) \in \mathbb{R}^d$

6:   **end for**

7:    $\hat{\mu}_t \leftarrow \frac{1}{m} \sum_{i=1}^m \Delta_i$

8:    $\hat{\Sigma}_t \leftarrow \frac{1}{m-1} \sum_{i=1}^m (\Delta_i - \hat{\mu}_t)(\Delta_i - \hat{\mu}_t)^{\top}$

9:    $\hat{\Sigma}_{t, \lambda} \leftarrow \hat{\Sigma}_t + \lambda I_d$  ▷ ridge for numerical stability

10:    $\Delta_x \leftarrow \nabla_{\theta} \ell(f(x; \theta^u), y) - \nabla_{\theta} \ell(f(x; \theta^{\text{org}}), y)$

11:    $v \leftarrow \Delta_x - \hat{\mu}_t$

12:   Solve  $\hat{\Sigma}_{t, \lambda} w = v$  for  $w$  (e.g., Cholesky);  $s_t \leftarrow v^{\top} w$

13:    $\ell_t \leftarrow -\log(1 - F_{\chi_d^2}(s_t))$  ▷ log tail  $p$ -value under  $H_0$

14:    $S \leftarrow S + \ell_t$

15: **end for**

16: **return** FORGOTTEN if  $S > \tau$ ; else TEST

---

918 D ALTERNATIVE SYMMETRY: CHANGE-OF-BASIS NEURAL TELEPORTATION.  
919920 We also support the “neural teleportation” family of symmetry moves from Armenta et al. (2023). Let  
921  $\tau_a > 0$  be a scale attached to neuron  $a$ . For an edge  $a \rightarrow b$  with weight  $\theta_{ab}$  the teleported weight is  
922

923 
$$\theta'_{ab} = \frac{\tau_b}{\tau_a} \theta_{ab}, \quad (10)$$
  
924

925 and if  $f_d$  is the activation at neuron  $d$  then the teleported activation is  
926

927 
$$g_d(x) = \tau_d f_d \left( \frac{x}{\tau_d} \right), \quad (11)$$
  
928

929 which preserves the function for positively homogeneous activations and extends naturally to batch-  
930 norm scales Armenta et al. (2023). In a subset of experiments, we choose  $\tau$  to further increase  
931 parameter dispersion under loss invariance (outputs unchanged), thereby weakening the differencing  
932 signal and making reconstruction harder; most results rely on the null space instantiation in equation 8.  
933 In the experimental section, it is explicitly indicated when both mechanisms are enabled.  
934936 E BASELINES  
937938 We evaluate our teleportation-based defense as a *plug-and-play* module layered on top of several  
939 state-of-the-art approximate post-hoc unlearning methods. These baselines are representative of the  
940 most widely studied approaches in recent literature, requiring no access to training-time auxiliary  
941 statistics (e.g., per-sample gradients) and operating directly on a pretrained model. Specifically, we  
942 consider:  
943

- 944 1.
- NegGrad+ (NGP)**
- Kurmanji et al. (2023): An improved variant of GA that incorporates
- 
- 945 a regularization term on the retain-set. The method balances ascent on the
- forget-set*
- with
- 
- 946 descent on the retain-set, aiming to preserve model utility while unlearning.
- 
- 947 2.
- SCRUB**
- Kurmanji et al. (2023): A knowledge distillation approach that aligns the unlearned
- 
- 948 model with the original model on the retain-set via a consistency loss, while simultaneously
- 
- 949 removing the
- forget-set*
- ’s influence. SCRUB represents one of the most competitive baselines
- 
- 950 in recent evaluations.
- 
- 951 3.
- SalUn**
- Fan et al. (2023): A saliency-based unlearning method that directs updates to a subset
- 
- 952 of weights deemed
- salient*
- for forgetting, identified via gradient-based weight saliency maps.
- 
- 953 By restricting optimization to these salient weights, SalUn enhances stability and efficiency
- 
- 954 compared to updating the full parameter set, and aims to reduce the gap to exact retraining.
- 
- 955 4.
- Projected Gradient Unlearning (PGU)**
- Hoang et al. (2024): A method that projects the
- 
- 956 gradient ascent update for the
- forget-set*
- onto a subspace orthogonal to retain-set, thereby
- 
- 957 mitigating catastrophic forgetting. PGU is particularly relevant as it addresses gradient-level
- 
- 958 entanglement between forget and retain data.
- 
- 959 5.
- BadTeacher (BT)**
- Chundawat et al. (2023a): A recent distillation-based unlearning method
- 
- 960 where the unlearned model (student) is trained against a deliberately corrupted teacher that
- 
- 961 provides noisy or adversarial labels for the
- forget-set*
- , encouraging the student to erase their
- 
- 962 influence while preserving performance on the retain-set.
- 
- 963 6.
- SRF-ON (SF)**
- Huang et al. (2024): A geometry-aware unlearning method that decomposes
- 
- 964 updates into forget ascent, retain descent, and saliency modulation. By embedding updates
- 
- 965 into the manifold of retain data and approximating Hessian modulation with a fast–slow
- 
- 966 strategy, SRF-ON improves stability–plasticity trade-offs and enables efficient large-scale
- 
- 967 unlearning.
- 
- 968

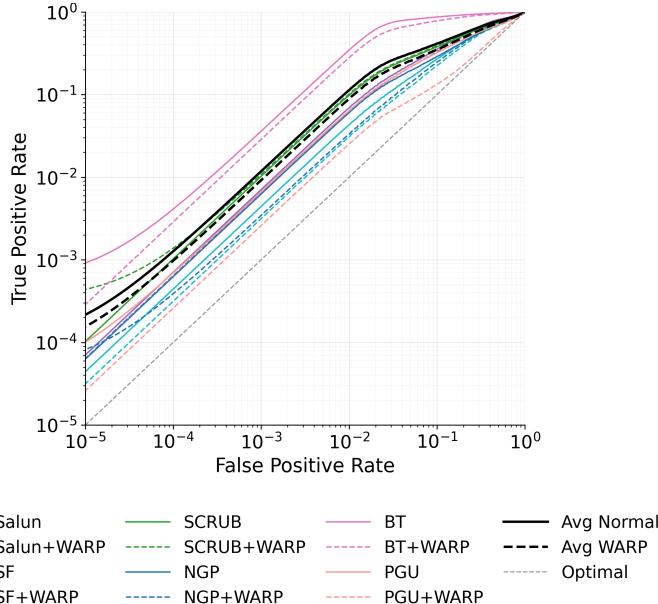
969 These methods span the main paradigms of approximate unlearning—gradient ascent, retain-aware  
970 regularization, distillation, and projection-based updates—making them representative state-of-the-art  
971 baselines.

972 **Table 3: Reconstruction on ImageNet-1K (ResNet-18), NGP (no defense).** Averages over 100  
 973 forgotten samples. Higher is better for PSNR/SSIM; lower is better for LPIPS/MSE.  
 974

Method	PSNR (dB) $\uparrow$	LPIPS (VGG) $\downarrow$	LPIPS (Alex) $\downarrow$	SSIM $\uparrow$	Test MSE $\downarrow$	Feat MSE $\downarrow$
GIFD Fang et al. (2023)	$8.28 \pm 0.28$	$0.630 \pm 0.012$	$0.448 \pm 0.016$	$0.098 \pm 0.007$	$0.174 \pm 0.012$	$6.725 \pm 0.506$
<b>Ours</b> (subspace-filtered + GFID)	<b><math>10.74 \pm 0.31</math></b>	<b><math>0.564 \pm 0.013</math></b>	<b><math>0.345 \pm 0.015</math></b>	<b><math>0.117 \pm 0.008</math></b>	<b><math>0.100 \pm 0.007</math></b>	<b><math>5.388 \pm 0.497</math></b>
<b>Improvement (%)</b>	+29.7	+10.5	+22.9	+19.4	+42.5	+19.9

## 980 F ADDITIONAL WHITE-BOX RESULTS ON CIFAR-10

982 Figure 5 reports the complete ROC curves for the Gaussian gradient-diff test, covering the entire  
 983 FPR range. These correspond to the same 640 unlearned models as in Figure 3, shown here without  
 984 zoom to provide the full view.



1005 **Figure 5: Complete ROC curves for the white-box Gaussian gradient-diff test.** Averaged over  
 1006 640 unlearned models, identical to Figure 3. Lower curves (closer to the random-guess diagonal)  
 1007 indicate stronger privacy.

## 1010 G RECONSTRUCTION ATTACK BASELINES AND COMPARISON.

1012 Table 3 compares three strategies for unlearning: (i) *simple differencing*, directly inverting  $\Delta\theta$  Hu  
 1013 et al. (2024); Bertran et al. (2024); (ii) *generative inversion* (GIFD) Fang et al. (2023) applied to  $\Delta\theta$ ;  
 1014 and (iii) *Ours*, which adds *orthogonal subspace filtering* (Eq. equation 3) to a generative backbone.  
 1015 Results average 100 forgotten samples on ImageNet-1K with ResNet-18 under NGP unlearning.  
 1016

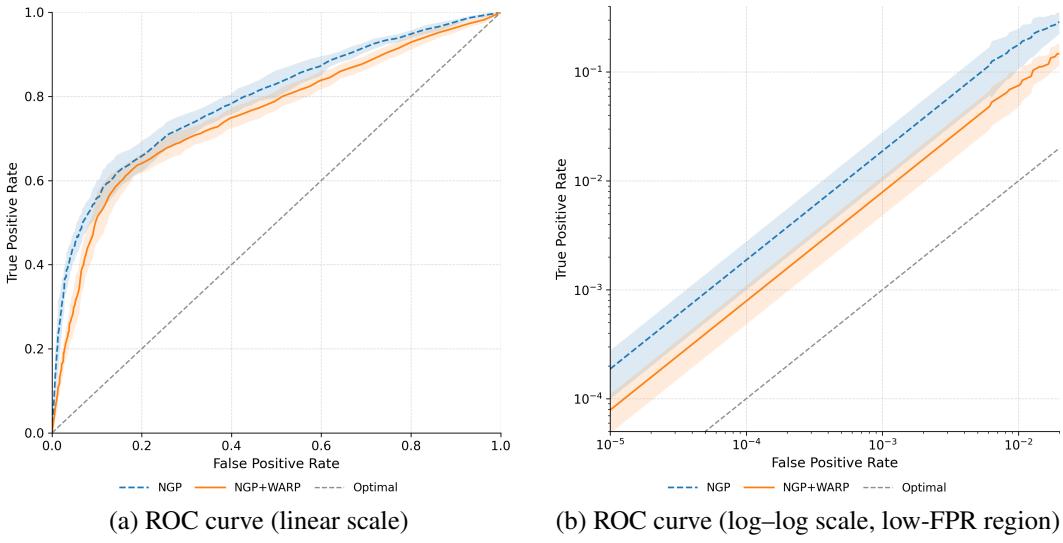
## 1017 H ADDITIONAL RESULTS: ViT ON TINY-IMAGENET

1019 To extend the white-box analysis of Section 4.3, we evaluate Vision Transformer models trained on  
 1020 Tiny-ImageNet. We adopt ViT-B/16 as the base architecture and follow the same setup described  
 1021 in Section 4, with the *forget-set* constructed by randomly sampling 1% of the training data and the  
 1022 *retain-set* consisting of the remainder. All models are trained with SGD and standard augmentations  
 1023 for ViT training. Unlearning is applied with NGP (NGP) and its teleported variant (NGP+WARP).  
 1024

1025 As shown in Table 4 and Figure 6, WARP substantially reduces attack success across all thresholds,  
 with the largest relative gains at low false-positive rates where practical attacks operate. These results

1026  
 1027 **Table 4: White-box membership inference risk with and without teleportation (ViT, Tiny-  
 1028 ImageNet).** Results are reported as mean  $\pm$  standard deviation across five splits. Improvements are  
 1029 computed as advantage reduction over random guessing.

Method	AUC	TPR@0.01%	TPR@0.1%	TPR@1%	TPR@5%
NGP (base)	$0.792 \pm 0.019$	$0.0019 \pm 0.001$	$0.0188 \pm 0.009$	$0.178 \pm 0.072$	$0.444 \pm 0.035$
+ WARP	$0.755 \pm 0.019$	$0.0008 \pm 0.000$	$0.0079 \pm 0.003$	$0.075 \pm 0.027$	$0.302 \pm 0.054$
Improvement (%)	<b>12.7</b>	<b>61.1</b>	<b>61.2</b>	<b>61.2</b>	<b>36.1</b>



1035  
 1036 **Figure 6: White-box ROC for the Gaussian gradient-difference test on ViT-B/16 (Tiny-  
 1037 ImageNet).** Each curve is averaged over five different forget-set splits, with shaded regions showing  
 1038 the standard deviation. Both figures compare NGP and NGP+WARP; (a) presents the full ROC on  
 1039 a linear axis, while (b) zooms into the low-FPR regime on log-log scale, which is the operational  
 1040 region for practical attacks.

1041  
 1042  
 1043  
 1044  
 1045  
 1046  
 1047  
 1048  
 1049  
 1050  
 1051  
 1052  
 1053  
 1054  
 1055  
 1056  
 1057  
 1058  
 1059  
 1060  
 1061  
 1062  
 1063  
 1064  
 1065  
 1066  
 1067  
 1068  
 1069  
 1070  
 1071  
 1072  
 1073  
 1074  
 1075  
 1076  
 1077  
 1078  
 1079  
 1080  
 1081  
 1082  
 1083  
 1084  
 1085  
 1086  
 1087  
 1088  
 1089  
 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  
 1135  
 1136  
 1137  
 1138  
 1139  
 1140  
 1141  
 1142  
 1143  
 1144  
 1145  
 1146  
 1147  
 1148  
 1149  
 1150  
 1151  
 1152  
 1153  
 1154  
 1155  
 1156  
 1157  
 1158  
 1159  
 1160  
 1161  
 1162  
 1163  
 1164  
 1165  
 1166  
 1167  
 1168  
 1169  
 1170  
 1171  
 1172  
 1173  
 1174  
 1175  
 1176  
 1177  
 1178  
 1179  
 1180  
 1181  
 1182  
 1183  
 1184  
 1185  
 1186  
 1187  
 1188  
 1189  
 1190  
 1191  
 1192  
 1193  
 1194  
 1195  
 1196  
 1197  
 1198  
 1199  
 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  
 1242  
 1243  
 1244  
 1245  
 1246  
 1247  
 1248  
 1249  
 1250  
 1251  
 1252  
 1253  
 1254  
 1255  
 1256  
 1257  
 1258  
 1259  
 1260  
 1261  
 1262  
 1263  
 1264  
 1265  
 1266  
 1267  
 1268  
 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  
 1297  
 1298  
 1299  
 1300  
 1301  
 1302  
 1303  
 1304  
 1305  
 1306  
 1307  
 1308  
 1309  
 1310  
 1311  
 1312  
 1313  
 1314  
 1315  
 1316  
 1317  
 1318  
 1319  
 1320  
 1321  
 1322  
 1323  
 1324  
 1325  
 1326  
 1327  
 1328  
 1329  
 1330  
 1331  
 1332  
 1333  
 1334  
 1335  
 1336  
 1337  
 1338  
 1339  
 1340  
 1341  
 1342  
 1343  
 1344  
 1345  
 1346  
 1347  
 1348  
 1349  
 1350  
 1351  
 1352  
 1353  
 1354  
 1355  
 1356  
 1357  
 1358  
 1359  
 1360  
 1361  
 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  
 1405  
 1406  
 1407  
 1408  
 1409  
 1410  
 1411  
 1412  
 1413  
 1414  
 1415  
 1416  
 1417  
 1418  
 1419  
 1420  
 1421  
 1422  
 1423  
 1424  
 1425  
 1426  
 1427  
 1428  
 1429  
 1430  
 1431  
 1432  
 1433  
 1434  
 1435  
 1436  
 1437  
 1438  
 1439  
 1440  
 1441  
 1442  
 1443  
 1444  
 1445  
 1446  
 1447  
 1448  
 1449  
 1450  
 1451  
 1452  
 1453  
 1454  
 1455  
 1456  
 1457  
 1458  
 1459  
 1460  
 1461  
 1462  
 1463  
 1464  
 1465  
 1466  
 1467  
 1468  
 1469  
 1470  
 1471  
 1472  
 1473  
 1474  
 1475  
 1476  
 1477  
 1478  
 1479  
 1480  
 1481  
 1482  
 1483  
 1484  
 1485  
 1486  
 1487  
 1488  
 1489  
 1490  
 1491  
 1492  
 1493  
 1494  
 1495  
 1496  
 1497  
 1498  
 1499  
 1500  
 1501  
 1502  
 1503  
 1504  
 1505  
 1506  
 1507  
 1508  
 1509  
 1510  
 1511  
 1512  
 1513  
 1514  
 1515  
 1516  
 1517  
 1518  
 1519  
 1520  
 1521  
 1522  
 1523  
 1524  
 1525  
 1526  
 1527  
 1528  
 1529  
 1530  
 1531  
 1532  
 1533  
 1534  
 1535  
 1536  
 1537  
 1538  
 1539  
 1540  
 1541  
 1542  
 1543  
 1544  
 1545  
 1546  
 1547  
 1548  
 1549  
 1550  
 1551  
 1552  
 1553  
 1554  
 1555  
 1556  
 1557  
 1558  
 1559  
 1560  
 1561  
 1562  
 1563  
 1564  
 1565  
 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  
 1620  
 1621  
 1622  
 1623  
 1624  
 1625  
 1626  
 1627  
 1628  
 1629  
 1630  
 1631  
 1632  
 1633  
 1634  
 1635  
 1636  
 1637  
 1638  
 1639  
 1640  
 1641  
 1642  
 1643  
 1644  
 1645  
 1646  
 1647  
 1648  
 1649  
 1650  
 1651  
 1652  
 1653  
 1654  
 1655  
 1656  
 1657  
 1658  
 1659  
 1660  
 1661  
 1662  
 1663  
 1664  
 1665  
 1666  
 1667  
 1668  
 1669  
 1670  
 1671  
 1672  
 1673  
 1674  
 1675  
 1676  
 1677  
 1678  
 1679  
 1680  
 1681  
 1682  
 1683  
 1684  
 1685  
 1686  
 1687  
 1688  
 1689  
 1690  
 1691  
 1692  
 1693  
 1694  
 1695  
 1696  
 1697  
 1698  
 1699  
 1700  
 1701  
 1702  
 1703  
 1704  
 1705  
 1706  
 1707  
 1708  
 1709  
 1710  
 1711  
 1712  
 1713  
 1714  
 1715  
 1716  
 1717  
 1718  
 1719  
 1720  
 1721  
 1722  
 1723  
 1724  
 1725  
 1726  
 1727  
 1728  
 1729  
 1730  
 1731  
 1732  
 1733  
 1734  
 1735  
 1736  
 1737  
 1738  
 1739  
 1740  
 1741  
 1742  
 1743  
 1744  
 1745  
 1746  
 1747  
 1748  
 1749  
 1750  
 1751  
 1752  
 1753  
 1754  
 1755  
 1756  
 1757  
 1758  
 1759  
 1760  
 1761  
 1762  
 1763  
 1764  
 1765  
 1766  
 1767  
 1768  
 1769  
 1770  
 1771  
 1772  
 1773  
 1774  
 1775  
 1776  
 1777  
 1778  
 1779  
 1780  
 1781  
 1782  
 1783  
 1784  
 1785  
 1786  
 1787  
 1788  
 1789  
 1790  
 1791  
 1792  
 1793  
 1794  
 1795  
 1796  
 1797  
 1798  
 1799  
 1800  
 1801  
 1802  
 1803  
 1804  
 1805  
 1806  
 1807  
 1808  
 1809  
 1810  
 1811  
 1812  
 1813  
 1814  
 1815  
 1816  
 1817  
 1818  
 1819  
 1820  
 1821  
 1822  
 1823  
 1824  
 1825  
 1826  
 1827  
 1828  
 1829  
 1830  
 1831  
 1832  
 1833  
 1834  
 1835  
 1836  
 1837  
 1838  
 1839  
 1840  
 1841  
 1842  
 1843  
 1844  
 1845  
 1846  
 1847  
 1848  
 1849  
 1850  
 1851  
 1852  
 1853  
 1854  
 1855  
 1856  
 1857  
 1858  
 1859  
 1860  
 1861  
 1862  
 1863  
 1864  
 1865  
 1866  
 1867  
 1868  
 1869  
 1870  
 1871  
 1872  
 1873  
 1874  
 1875  
 1876  
 1877  
 1878  
 1879  
 1880  
 1881  
 1882  
 1883  
 1884  
 1885  
 1886  
 1887  
 1888  
 1889  
 1890  
 1891  
 1892  
 1893  
 1894  
 1895  
 1896  
 1897  
 1898  
 1899  
 1900  
 1901  
 1902  
 1903  
 1904  
 1905  
 1906  
 1907  
 1908  
 1909  
 1910  
 1911  
 1912  
 1913  
 1914  
 1915  
 1916  
 1917  
 1918  
 1919  
 1920  
 1921  
 1922  
 1923  
 1924  
 1925  
 1926  
 1927  
 1928  
 1929  
 1930  
 1931  
 1932  
 1933  
 1934  
 1935  
 1936  
 1937  
 1938  
 1939  
 1940  
 1941  
 1942  
 1943  
 1944  
 1945  
 1946  
 1947  
 1948  
 1949  
 1950  
 1951  
 1952  
 1953  
 1954  
 1955  
 1956  
 1957  
 1958  
 1959  
 1960  
 1961  
 1962  
 1963  
 1964  
 1965  
 1966  
 1967  
 1968  
 1969  
 1970  
 1971  
 1972  
 1973  
 1974  
 1975  
 1976  
 1977  
 1978  
 1979  
 1980  
 1981  
 1982  
 1983  
 1984  
 1985  
 1986  
 1987  
 1988  
 1989  
 1990  
 1991  
 1992  
 1993  
 1994  
 1995  
 1996  
 1997  
 1998  
 1999  
 2000  
 2001  
 2002  
 2003  
 2004  
 2005  
 2006  
 2007  
 2008  
 2009  
 2010  
 2011  
 2012  
 2013  
 2014  
 2015  
 2016  
 2017  
 2018  
 2019  
 2020  
 2021  
 2022  
 2023  
 2024  
 2025  
 2026  
 2027  
 2028  
 2029  
 2030  
 2031  
 2032  
 2033  
 2034  
 2035  
 2036  
 2037  
 2038  
 2039  
 2040  
 2041  
 2042  
 2043  
 2044  
 2045  
 2046  
 2047  
 2048  
 2049  
 2050  
 2051  
 2052  
 2053  
 2054  
 2055  
 2056  
 2057  
 2058  
 2059  
 2060  
 2061  
 2062  
 2063  
 2064  
 2065  
 2066  
 2067  
 2068  
 2069  
 2070  
 2071  
 2072  
 2073  
 2074  
 2075  
 2076  
 2077  
 2078  
 2079  
 2080  
 2081  
 2082  
 2083  
 2084  
 2085  
 2086  
 2087  
 2088  
 2089  
 2090  
 2091  
 2092  
 2093  
 2094  
 2095  
 2096  
 2097  
 2098  
 2099  
 2100  
 2101  
 2102  
 2103  
 2104  
 2105  
 2106  
 2107  
 2108  
 2109  
 2110  
 2111  
 2112  
 2113  
 2114  
 2115  
 2116  
 2117  
 2118  
 2119  
 2120  
 2121  
 2122  
 2123  
 2124  
 2125  
 2126  
 2127  
 2128  
 2129  
 2130  
 2131  
 2132  
 2133  
 2134  
 2135  
 2136  
 2137  
 2138  
 2139  
 2140  
 2141  
 2142  
 2143  
 2144  
 2145  
 2146  
 2147  
 2148  
 2149  
 2150  
 2151  
 2152  
 2153  
 2154  
 2155  
 2156  
 2157  
 2158  
 2159  
 2160  
 2161  
 2162  
 2163  
 2164  
 2165  
 2166  
 2167  
 2168  
 2169  
 2170  
 2171  
 2172  
 2173  
 2174  
 2175  
 2176  
 2177  
 2178  
 2179  
 2180  
 2181  
 2182  
 2183  
 2184  
 2185  
 2186  
 2187  
 2188  
 2189  
 2190  
 2191  
 2192  
 2193  
 2194  
 2195  
 2196  
 2197  
 2198  
 2199  
 2200  
 2201  
 2202  
 2203  
 2204  
 2205  
 2206  
 2207  
 2208  
 2209  
 2210  
 2211  
 2212  
 2213  
 2214  
 2215  
 2216  
 2217  
 2218  
 2219  
 2220  
 2221  
 2222  
 2223  
 2224  
 2225  
 2226  
 2227  
 2228  
 2229  
 2230  
 2231  
 2232  
 2233  
 2234  
 2235  
 2236  
 2237  
 2238  
 2239  
 2240  
 2241  
 2242  
 2243  
 2244  
 2245  
 2246  
 2247  
 2248  
 2249  
 2250  
 2251  
 2252  
 2253  
 2254  
 2255  
 2256  
 2257  
 2258  
 2259  
 2260  
 2261  
 2262  
 2263  
 2264  
 2265  
 2266  
 2267  
 2268  
 2269  
 2270  
 2271  
 2272  
 2273  
 2274  
 2275  
 2276  
 2277  
 2278  
 2279  
 2280  
 2281  
 2282  
 2283  
 2284  
 2285  
 2286  
 2287  
 2288  
 2289  
 2290  
 2291  
 2292  
 2293  
 2294  
 2295  
 2296  
 2297  
 2298  
 2299  
 2300  
 2301  
 2302  
 2303  
 2304  
 2305  
 2306  
 2307  
 2308  
 2309  
 2310  
 2311  
 2312  
 2313  
 2314  
 2315  
 2316  
 2317  
 2318<br

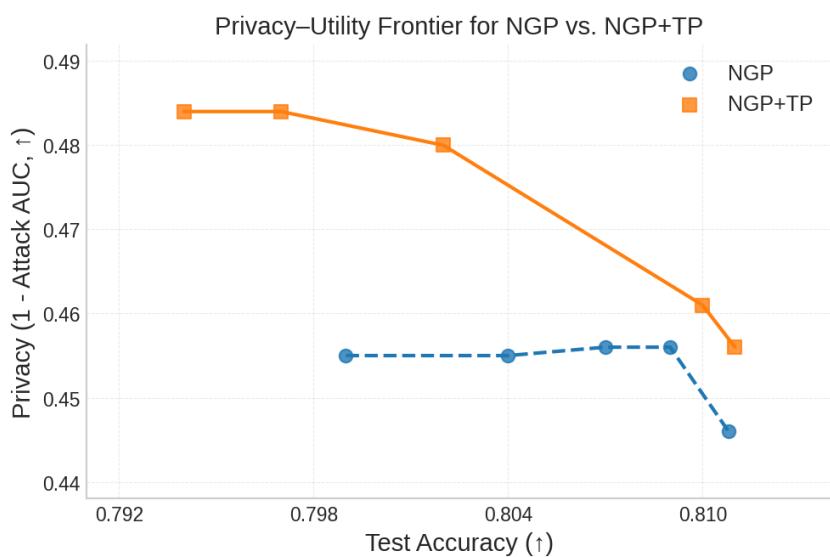


Figure 7: **Privacy–utility trade-off for NGP with and without WARP.** Each point is a hyperparameter trial, with privacy (1-AUC) averaged over 640 shadow models (64 shadows  $\times$  10 forget sets) under the U-LiRA protocol. Points further to the right (higher accuracy) and upward (higher privacy) indicate better trade-offs.

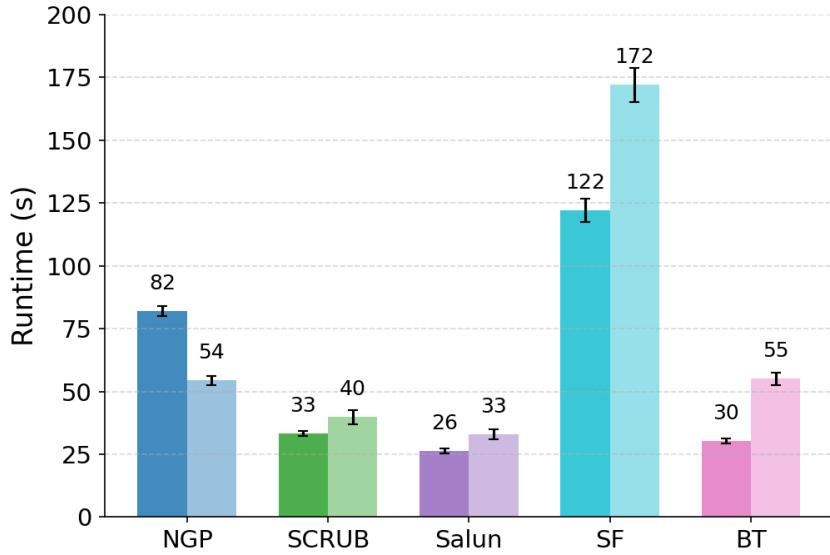


Figure 8: **Runtime overhead of teleportation.** Average runtimes (seconds) of unlearning algorithms with and without the WARP plugin, evaluated on CIFAR-10 with ResNet-18. Each bar reports the mean over five runs, with error bars showing standard deviations.

gains. Across the frontier, improvements remain stable, confirming that teleportation meaningfully reshapes the privacy–utility boundary in favor of the defender.

## J RUNTIME ANALYSIS

In this appendix we focus on the retain–null-space instantiation of  $T_\phi$ , which is the only variant that requires explicit SVDs; the change-of-basis teleportation in Appendix D is SVD-free and without its

1134 computational overhead as a result. Moreover, Section L introduces FastWARP, which replaces full  
 1135 SVD with randomized low-rank approximations and further reduces this overhead.  
 1136

1137 We benchmark the runtime of our teleportation defense across unlearning algorithms on a machine  
 1138 equipped with an NVIDIA GeForce RTX 4090 GPU (24 GB memory) and an Intel 13th Gen Core  
 1139 i9-13900KF CPU (24 cores, 32 threads, base 3.0 GHz, boost up to 5.8 GHz). Each experiment  
 1140 was repeated five times, and Figure 8 reports averages with standard deviations in the caption. All  
 1141 algorithms were run with the hyperparameters used in Table 1 and Figure 3, ensuring runtime reflects  
 1142 the same conditions as our privacy–utility evaluations.  
 1143

1144 For this particular SVD-based instantiation, teleportation increases runtime by approximately +27%  
 1145 relative to the baseline on average, reflecting the overhead of constructing the retain subspace. The  
 1146 main exception is NGP, where teleportation reduces runtime by about -32%, due to more stable  
 1147 updates that in turn lower the required number of unlearning epochs. Since subspace computation can  
 1148 be pre-computed offline and does not need to be repeated after every teleportation step, this overhead  
 1149 can be amortized in practice. While updating the retain subspace less frequently can reduce cost, the  
 1150 primary computational overhead from full SVD is addressed directly by an approximate low-rank  
 1151 implementation (Appendix L), which removes the per-step bottleneck entirely.  
 1152

## K TELEPORTATION-BASED UNLEARNING ALGORITHM

1153 In Algorithm 3,  $T_\phi$  denotes an abstract symmetry operator; in our experiments we instantiate it  
 1154 either with retain–null-space teleportation or with change-of-basis teleportation, but any other loss-  
 1155 preserving symmetry could be used in its place.  
 1156

---

1162 **Algorithm 3** WARP (retain–null-space instantiation): teleportation-augmented gradient-based un-  
 1163 learning.

---

1164 **Require:**  $\theta^{\text{org}}$ ,  $\mathcal{D}_f$ ,  $\mathcal{D}_r$ ,  $\ell_f$ ,  $\ell_r$ ,  $\lambda$ ,  $\beta$ ,  $\{\eta_t\}$ ,  $\eta_{\text{tel}}$ ,  $k$ ,  $S$  or  $\tau_{\text{grad}}$ ,  $\sigma^2$ ,  $\varepsilon$ ,  $T$   
 1165 1:  $\theta_0 \leftarrow \theta^{\text{org}}$   
 1166 2: **for**  $t = 0, \dots, T - 1$  **do**  
 1167 3:   sample  $\mathcal{B}_f \subset \mathcal{D}_f$ ,  $\mathcal{B}_r \subset \mathcal{D}_r$   
 1168 4:    $\theta_{t+\frac{1}{2}} \leftarrow \theta_t - \eta_t (\nabla_\theta \ell_f(\theta_t \mid \mathcal{B}_f) + \lambda \nabla_\theta \ell_r(\theta_t \mid \mathcal{B}_r))$   
 1169 5:   **if**  $(t \bmod S = 0) \vee \|\nabla_\theta \ell_f(\theta_{t+\frac{1}{2}} \mid \mathcal{B}_f)\|_2 > \tau_{\text{grad}}$  **then**  
 1170 6:     **for** layer  $\ell$  **do**  
 1171       build  $R_\ell(\mathcal{B}_r)$ ;  $R_\ell = U_\ell \Sigma_\ell V_\ell^\top$  (SVD)  
 1172        $B_\ell \leftarrow U_{\ell,1:k}$ ;  $\Pi_\ell^\perp \leftarrow I - B_\ell B_\ell^\top$   
 1173     **end for**  
 1174      $\mathcal{L}_{\text{tel}}(\theta) = \frac{1}{2} \sum_{(x,y) \in \mathcal{B}_f} \|\nabla_\theta \ell(f(x; \theta), y)\|_2^2 - \frac{\beta}{2} \|\theta - \theta^{\text{org}}\|_2^2$   
 1175     **for** layer  $\ell$  **do**  
 1176        $W_\ell^{t+1} \leftarrow W_\ell^{t+\frac{1}{2}} - \eta_{\text{tel}} \Pi_\ell^\perp (\nabla_{W_\ell} \mathcal{L}_{\text{tel}}(\theta_{t+\frac{1}{2}})) + \sqrt{2 \eta_{\text{tel}} \sigma^2} \varepsilon_{\ell,t}$   
 1177        $\varepsilon_{\ell,t} \sim \mathcal{N}(0, I)$   
 1178     **end for**  
 1179      $\theta_{t+1} \leftarrow \{W_\ell^{t+1}\}_\ell$   
 1180     **if**  $\ell_r(\theta_{t+1} \mid \mathcal{B}_r) > \ell_r(\theta_t \mid \mathcal{B}_r) + \varepsilon$  **then**  
 1181        $\theta_{t+1} \leftarrow \theta_{t+\frac{1}{2}}$   
 1182     **end if** ▷ backtrack/safeguard  
 1183     **else**  
 1184        $\theta_{t+1} \leftarrow \theta_{t+\frac{1}{2}}$   
 1185     **end if**  
 1186     **end for**  
 1187 23: **return**  $\theta^u \leftarrow \theta_T$

---

---

1188 **L APPROXIMATE NULL-SPACE TELEPORTATION**  
1189

1190 **Low-rank structure of retain representations.** For a retain minibatch  $\mathcal{B}_r$  and layer  $\ell$ , let  
1191  $R_\ell(\mathcal{D}_r) \in \mathbb{R}^{|\mathcal{B}_r| \times d_\ell}$  denote the matrix whose rows collect the layer- $\ell$  inputs  $\{\phi_\ell(x)\}_{x \in \mathcal{B}_r}$ . Em-  
1192 pirically,  $R_\ell(\mathcal{D}_r)$  exhibits strong spectral decay: its spectrum is dominated by a small number of  
1193 singular values, and most of the energy lies in a low-dimensional subspace. Such low-rank structure  
1194 of activations, gradients and Hessians has been observed repeatedly in modern deep networks (Arora  
1195 et al., 2019; Ghorbani et al., 2019; Fort et al., 2020; Gur-Ari et al., 2018), and is often attributed to  
1196 overparameterisation and the implicit regularisation of SGD. In WARP, the retain subspace at layer  $\ell$   
1197 is defined by the top- $k$  left singular vectors of  $R_\ell(\mathcal{D}_r)$ :  
1198

$$R_\ell(\mathcal{D}_r) = U_\ell \Sigma_\ell V_\ell^\top, \quad B_\ell = U_{\ell,1:k}, \quad \Pi_\ell^\perp = I - B_\ell B_\ell^\top.$$

1199 Since only the span of these dominant directions matters for teleportation, *exact* SVD is not required:  
1200 any procedure that recovers a good approximation to the top- $k$  principal subspace suffices.  
1201

1202 **Covariance-based PCA and subspace iteration.** Instead of computing a full thin SVD of  $R_\ell(\mathcal{D}_r)$ ,  
1203 FASTWARP estimates  $B_\ell$  via a covariance eigen-decomposition and a small number of subspace-  
1204 iteration updates, following classical PCA and online PCA methods (Golub & Van Loan, 2013; Oja,  
1205 1982; Warmuth & Kuzmin, 2008; Mitliagkas et al., 2013). We first form the covariance  
1206

$$C_\ell = X_\ell X_\ell^\top \in \mathbb{R}^{d_\ell \times d_\ell},$$

1207 where  $X_\ell \in \mathbb{R}^{d_\ell \times N}$  is a layer-wise input matrix constructed from  $\mathcal{B}_r$  (for convolutional layers we  
1208 use unfolded patches; for batch-norm we aggregate per-channel features). We then compute the  
1209 eigen-decomposition  $C_\ell = Q_\ell \Lambda_\ell Q_\ell^\top$  and retain the smallest  $k$  such that the cumulative explained  
1210 variance exceeds a threshold  $\tau$ :  
1211

$$k = \min \left\{ j : \frac{\sum_{i=1}^j \max(\lambda_{\ell,i}, 0)}{\sum_{i=1}^{d_\ell} \max(\lambda_{\ell,i}, 0)} \geq \tau \right\}, \quad B_\ell = Q_{\ell,1:k},$$

1213 optionally capped by a user-specified  $k_{\max}$ . For subsequent teleportation steps, we update  $B_\ell$  using a  
1214 few iterations of subspace iteration (Golub & Van Loan, 2013; Halko et al., 2011; Musco & Musco,  
1215 2015; Tropp et al., 2017; Woodruff, 2014):  
1216

$$Y \leftarrow C_\ell B_\ell, \quad [B_\ell, \_] \leftarrow \text{qr}(Y),$$

1217 which amounts to an Oja-style streaming PCA update (Oja, 1982) with QR re-orthogonalisation.  
1218 This reduces the cost of updating  $B_\ell$  for a new minibatch from the  $\mathcal{O}(|\mathcal{B}_r| d_\ell^2)$  cost of a fresh thin  
1219 SVD to  $\mathcal{O}(|\mathcal{B}_r| d_\ell k)$  for the covariance application plus  $\mathcal{O}(d_\ell k^2)$  for QR, with  $k \ll d_\ell$ . The resulting  
1220 projector  $\Pi_\ell^\perp = I - B_\ell B_\ell^\top$  is then used exactly as in the original WARP update.  
1221

---

1222 **Algorithm 4** FASTWARP basis update at layer  $\ell$   
1223

1224 **Require:**  $d_\ell$ , retain minibatch  $\mathcal{B}_r$ ,  $B_\ell^{\text{prev}}$  (or **NONE**),  $\tau \in (0, 1]$ ,  $k_{\max}$ ,  $T_{\text{track}}$   
1225 1: build  $X_\ell \in \mathbb{R}^{d_\ell \times N}$  from  $\mathcal{B}_r$   
1226 2:  $C_\ell \leftarrow X_\ell X_\ell^\top$ ;  $C_\ell \leftarrow \frac{1}{2}(C_\ell + C_\ell^\top)$   
1227 3: **if**  $B_\ell^{\text{prev}} = \text{NONE}$  **then**  
1228 4:    $C_\ell = Q_\ell \Lambda_\ell Q_\ell^\top$   
1229 5:   sort  $\Lambda_\ell$  in descending order, permute  $Q_\ell$  accordingly  
1230 6:    $k \leftarrow \min \left\{ k_{\max}, \min \left\{ k : \frac{\sum_{i=1}^k \Lambda_{\ell,ii}}{\sum_i \Lambda_{\ell,ii}} \geq \tau \right\} \right\}$   
1231 7:    $B_\ell \leftarrow Q_\ell[:, 1:k]$   
1232 8: **else**  
1233 9:    $B_\ell \leftarrow B_\ell^{\text{prev}}$   
1234 10:   **for**  $t = 1, \dots, T_{\text{track}}$  **do**  
1235 11:      $Y \leftarrow C_\ell B_\ell$   
1236 12:      $[B_\ell, \_] \leftarrow \text{qr}(Y)$   
1237 13:      $B_\ell \leftarrow B_\ell[:, 1:k]$   
1238 14:   **end for**  
1239 15: **end if**  
1240 16:  $\Pi_\ell^\perp \leftarrow I_{d_\ell} - B_\ell B_\ell^\top$   
1241 17: **return**  $B_\ell, \Pi_\ell^\perp$   


---

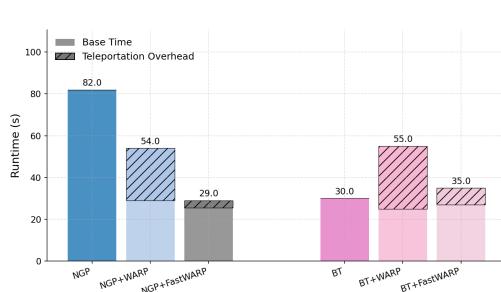


Figure 9: Runtime of the WARP plug-in on CIFAR-10 with ResNet-18. Each bar reports the mean over five runs. The top hatched segments correspond to the additional teleportation time; the solid base is the runtime of the underlying MU algorithm.

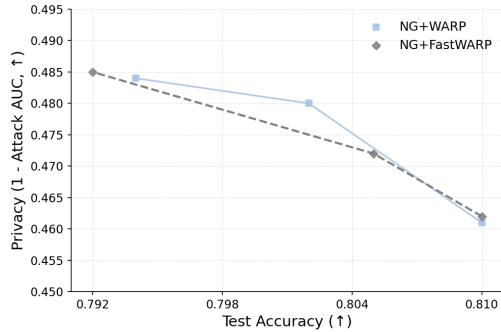


Figure 10: Privacy–utility comparison of NG+WARP and NG+FASTWARP. The approximate teleportation method (FASTWARP) matches the privacy–utility frontier of the exact variant, achieving nearly identical privacy and test accuracy.

**Runtime and privacy–utility impact.** Figure 9 reports the runtime for NG and BT with and without teleportation on CIFAR-10/ResNet-18. The hatched segments correspond to the teleportation component. Using full SVD yields a moderate yet visible overhead (e.g., BT+WARP increases runtime from 30 s to 55 s). Replacing full SVD with the covariance-based PCA and subspace iteration of Algorithm 4 (FASTWARP) shrinks this overhead substantially: total runtime drops to 29 s and 35 s for NG+FastWARP and BT+FastWARP, corresponding to a 2×–3× reduction in the teleportation time. The teleportation component becomes only a small fraction of the overall MU cost.

To measure the effect of this approximation on privacy and accuracy, Figure 10 compares NG+WARP and NG+FASTWARP along the privacy–utility frontier. The two curves are nearly indistinguishable: privacy ( $1 - \text{AUC}$ ) differs by at most 0.3–0.6% across operating points, and test accuracy changes by at most 0.2–0.3 percentage points. We also track retain-set loss during teleportation and observe that the relative drift under FASTWARP remains below 2%, indicating that the approximate projector continues to enforce practical loss invariance. In some configurations, the additional numerical noise introduced by the approximation yields slightly *higher* privacy for the same utility. Overall, these results show that the privacy gains of WARP are robust to approximate PCA, and that FASTWARP preserves the empirical privacy–utility trade-off while significantly reducing computational overhead.

**Scalability to LLMs and calibration of the retain subspace.** A natural concern is whether null-space teleportation remains practical and stable at LLM scale, where layer widths reach  $d_\ell \sim 10^3\text{--}10^4$  and a single minibatch may not span the retain subspace. Empirically, recent compression work shows that truncated SVD and related low-rank factorizations are already applied efficiently to full LLM weight matrices with comparable or larger dimensions: SVD-LLM Wang et al. (2024; 2025b) optimizes singular-value truncation for LLaMA Touvron et al. (2023)- and GPT Brown et al. (2020)-class models while preserving perplexity and throughput, demonstrating that rank- $k$  SVD with  $k \ll d_\ell$  is tractable in practice on modern hardware. Complementary methods such as ResSVD Bai et al. (2025) leverage the residual matrix left by truncation to correct the approximation, further reducing the effective loss of expressivity at fixed rank. Orthogonal lines of work, e.g., weighted low-rank factorization for LMs, explicitly introduce data-dependent weights in the covariance (or Gram) operator to bias the recovered subspace toward high-importance tokens or examples, and report competitive compression ratios on transformer-based LMs Hsu et al. (2022); Sakr & Khailany (2024). In our setting, we can adopt the same design principles: instead of forming  $R_\ell(\mathcal{B}_r)$  from an arbitrary minibatch, we maintain a small buffer of retain batches with large gradient norm Sakr & Khailany (2024) or Fisher information, and construct the activation matrix  $X_\ell$  from this “high-influence” pool. This yields a weighted or importance-sampled covariance  $C_\ell = X_\ell X_\ell^\top$  whose top- $k$  eigenspace more faithfully captures the retain subspace seen over the full retain stream, while keeping the per-teleportation cost at  $\mathcal{O}(|\mathcal{B}_r|d_\ell k)$ . Combined with low-rank SVD implementations that are already optimized for LLM compression, these heuristics make the FastWARP projector construction compatible with large transformer architectures without breaking the retain loss invariance enforced

1296 by WARP. We leave the adaptation to large language models for future research. Our contributions  
 1297 target symmetry-based defenses for generic neural networks and established MU baselines, and do  
 1298 not address LLM-specific challenges in unlearning, which constitute a distinct line of investigation.  
 1299

## 1300 M COMPARISON WITH DP–LANGEVIN NOISE DEFENCES

1302 While our goal is to make neural networks more resilient to privacy attacks *post hoc*, a natural  
 1303 question is how WARP compares with defences based on differential privacy (DP). DP is the strongest  
 1304 known framework for providing indistinguishability guarantees between neighbouring datasets, and  
 1305 a small number of recent unlearning methods have attempted to translate these guarantees into  
 1306 *certified* machine unlearning. Among these, noisy-gradient (Langevin) approaches provide the closest  
 1307 analogue to our setting; we therefore include them as a comparison point.

1308 Certified unlearning methods such as Guo et al. (2020); Chien et al. (2024b) formalise unlearning  
 1309 as an indistinguishability requirement between (i) a model obtained by training on the full dataset,  
 1310 and (ii) a counterfactual model that has never seen the forget set. These works build on the principle  
 1311 that if the training algorithm is itself DP, then suitable post-processing can yield certified removal of  
 1312 training points. Such guarantees make DP–Langevin the strongest known *general-purpose* defence  
 1313 with explicit indistinguishability guarantees, hence a meaningful baseline to evaluate privacy–utility  
 1314 trade-offs.

1316 **What the DP guarantees actually require.** The formal guarantees in Guo et al. (2020); Chien  
 1317 et al. (2024b) rely on assumptions that do *not* hold in the deep, non-convex MU regime we consider:

- 1319 **Convexity and strong dissipativity.** Both works require (strongly) convex,  $\ell_2$ –regularised  
 1320 objectives to bound the stationary distribution of the noisy dynamics. Deep convolutional  
 1321 networks trained with cross-entropy fundamentally violate these assumptions.
- 1322 **DP-trained initial model required.** The certified-unlearning guarantee requires that the *original*  
 1323 model be obtained using *the same* noisy-gradient mechanism (noisy SGD or Langevin) applied  
 1324 throughout training on the full dataset. This is explicitly stated as a necessary condition in Chien  
 1325 et al. (2024b). In contrast, our setting begins from a standard ERM-trained model, which is  
 1326 non-DP and therefore outside the scope of their certification theorem.

1327 As a result, the “ $\varepsilon$ ” obtained from the RDP accountant in our experiments should be interpreted  
 1328 purely as a calibrated *noise level*, not as a valid DP guarantee. Our use of Langevin noise is therefore  
 1329 a *strong noise-based defence*, not a certified mechanism.

1331 **Adapting projected Langevin unlearning to MU.** Following Chien et al. (2024b), we implement  
 1332 projected Langevin dynamics on top of the same MU objective used throughout the paper. For a  
 1333 per-sample clipped gradient with radius  $C$  and loss

$$1334 \mathcal{L}_{\text{MU}}(\theta) = \alpha (\ell_r(\theta) + \lambda \|\theta - \theta_p\|_2^2) - (1 - \alpha) \ell_f(\theta),$$

1336 the DP–Langevin update is

$$1337 g_t = \text{clip}(\nabla_\theta \mathcal{L}_{\text{MU}}(\theta_t), C), \quad (12)$$

$$1339 \theta_{t+1} = \theta_t - \eta_t g_t + \sqrt{2 \eta_t \lambda} \xi_t, \quad \xi_t \sim \mathcal{N}(0, I), \quad (13)$$

1340 where  $\lambda$  is the regularisation parameter entering the RDP privacy analysis. Given a target privacy  
 1341 level  $\varepsilon$ , we follow the exact Rényi-DP accounting of Chien et al. (2024b) to compute the Gaussian  
 1342 noise standard deviation  $\sigma$  required by their Langevin update. In our implementation, three quantities  
 1343 act as tunable hyperparameters: the learning rate  $\eta$ , the per-sample gradient-clipping radius  $C$ , and  
 1344 the regularisation coefficient  $\lambda$  that appears in the RDP analysis. For any chosen  $(\eta, C, \lambda)$  and target  
 1345  $\varepsilon$ , the formulas of Chien et al. (2024b) uniquely determine the corresponding noise scale  $\sigma$ . To ensure  
 1346 fairness across baselines, we run the same number of hyperparameter-search trials as for the MU  
 1347 baselines, jointly sweeping  $(\eta, C, \lambda)$  to obtain the set of reported results in Table 5.

1348 **Interpretation under non-convexity.** Although the privacy accountant yields a numerical  $\varepsilon$ , none  
 1349 of the formal conditions needed for DP-certified unlearning hold for our deep ResNet models.

1350 Table 5: **NGP+WARP vs. Langevin noise (U-LiRA, black-box)**. Reported are risks on *all forget samples* and on the *most-memorized* subset (top 5%), plus test accuracy. U-LiRA AUC and  
 1351 TPR@0.1% (FPR) are shown for each setting.  
 1352

Method	All samples (BB)		Most-memorized (top 5%)		Acc. Test
	AUC	TPR@0.1	AUC	TPR@0.1	
Langevin ( $\varepsilon = 1$ )	0.523	0.004	0.671	0.029	0.682
Langevin ( $\varepsilon = 4$ )	0.571	0.006	0.766	0.048	0.718
Langevin ( $\varepsilon = 8$ )	0.627	0.020	0.912	0.166	0.771
Langevin ( $\varepsilon = 16$ )	0.650	0.027	0.935	0.224	<b>0.798</b>
NGP + WARP	<b>0.516</b>	<b>0.003</b>	<b>0.598</b>	<b>0.015</b>	0.797

1363  
 1364 Consequently, we reiterate that the resulting values should not be interpreted as DP guarantees but  
 1365 rather as a systematic way of calibrating the magnitude of injected noise. The comparison therefore  
 1366 isolates the *empirical* effect of noise injection on forgetting, retention, and attack success.  
 1367

1368  
 1369 **Empirical privacy–utility trade-off.** Table 5 reveals a clear tension between nominal DP guar-  
 1370 antees and empirical membership privacy. As the target privacy budget for Langevin is relaxed  
 1371 from  $\varepsilon = 1$  to  $\varepsilon = 16$ , test accuracy gradually recovers (from 0.682 up to 0.798), but U-LiRA risk  
 1372 monotonically *increases*: the all-sample AUC rises from 0.523 to 0.650, and the AUC on the top-5%  
 1373 most memorised points grows from 0.671 to 0.935, with TPR@0.1% FPR increasing from 0.029  
 1374 to 0.224. In contrast, NGP+WARP simultaneously achieves competitive utility and strictly lower  
 1375 attack success: on all forget samples it attains the best AUC and TPR@0.1% (0.516 and 0.003), and  
 1376 on the most–memorised subset it reduces AUC to 0.598 and TPR@0.1% to 0.015, outperforming  
 1377 every Langevin configuration by a wide margin. Notably, relative to the lowest-noise setting ( $\varepsilon = 16$ ),  
 1378 NGP+WARP matches accuracy (0.797 vs. 0.798) while cutting the memorised AUC from 0.935 to  
 1379 0.598 and TPR@0.1% from 0.224 to 0.015. For stronger nominal privacy ( $\varepsilon = 1$  or 4), Langevin  
 1380 noise severely degrades accuracy (down to 0.682) yet still leaves substantially higher attack AUC  
 1381 and TPR than WARP. Overall, these results suggest that isotropic DP noise is poorly aligned with  
 1382 the specific memorization patterns exploited by U-LiRA: it injects substantial randomness into all  
 1383 updates, harming utility without reliably protecting the most vulnerable examples, whereas WARP  
 1384 reshapes the parameter space in a targeted way that yields a markedly better empirical privacy–utility  
 1385 frontier.

1386 Taken together, these observations clarify the roles of the two approaches. Langevin noise offers a  
 1387 principled mechanism for *certified* unlearning in the restricted setting of convex, DP-trained models,  
 1388 but its guarantees do not extend to the non-convex MU regime nor to pretrained models obtained  
 1389 without DP noise. Consequently, applying Langevin updates post hoc to deep networks provides  
 1390 no formal protection and yields an unfavourable privacy–utility trade-off in practice. By contrast,  
 1391 WARP operates directly on arbitrary pretrained models, targets the directions most responsible for  
 1392 memorization, and empirically achieves substantially stronger resistance to membership inference at  
 1393 comparable accuracy. A compelling direction for future work is to investigate whether the geometric  
 1394 structure exploited by WARP can be combined with, or serve as a foundation for, certified unlearning  
 1395 mechanisms that simultaneously handle non-convex objectives and non-DP initialisation—a capability  
 1396 not supported by current DP-Langevin frameworks.

## N ADAPTIVE RECONSTRUCTION WITH SYMMETRY–AWARE ATTACKER

1397 Teleportation acts by composing the unlearning update with a symmetry transform that preserves  
 1398 predictions but redistributes parameter mass along loss–invariant directions (Section 3.2). This raises  
 1399 a natural question: can a stronger white-box adversary, aware of the teleportation family, *invert* or  
 1400 compensate for these symmetry moves and recover the residual forget gradient? More concretely, if  
 1401 the attacker can parameterise and optimize over the change-of-basis (COB) scales  $\tau$  used in neural  
 1402 teleportation (Armenta et al., 2023), does this restore reconstruction quality and defeat WARP?  
 1403

1404 It is worth noting that our privacy evaluation already includes two adaptive-attack families: U-LiRA  
 1405 and GLiR, both of which instantiate adaptive membership-inference attacks by optimising proxy  
 1406 models or surrogate loss landscapes. However, the reconstruction attack considered in Section 3.1—which  
 1407 directly targets instance-level recovery of the forgotten data—was *not* adaptive: the attacker optimised  
 1408 only over the dummy image while keeping the teleportation parameters fixed. To fully test the  
 1409 robustness of symmetry-based teleportation, we now consider a strictly stronger attacker that *jointly*  
 1410 optimizes both the dummy image and the teleportation parameters themselves.

1411 Concretely, we study whether an attacker who can parameterise and optimize over the change-of-basis  
 1412 (COB) symmetry scales  $\tau$  used in neural teleportation (Armenta et al., 2023) can undo the defender’s  
 1413 symmetry moves, thereby restoring the clean gradient geometry required for successful reconstruction.  
 1414 This experiment directly probes whether teleportation is merely hiding the forget gradient behind  
 1415 a reversible reparameterisation, or whether it fundamentally reshapes the inverse problem faced by  
 1416 reconstruction attacks.

1417 **Attack formulation.** In the adaptive setting, we give the attacker full knowledge of the teleportation  
 1418 family and let them *shadow* the defender’s operations. Specifically, starting from the original  
 1419 pretrained weights  $\theta_{\text{org}}$ , the attacker first applies a change-of-basis symmetry parametrised by COB  
 1420 scales  $\tau = \{\tau_a > 0\}$ , obtaining

$$\theta_{\text{org}}^{(\tau)} = T_{\tau}(\theta_{\text{org}}), \quad (14)$$

1421 where  $T_{\tau}$  is the COB teleportation map (Appendix D). They then perform a single gradient step in  
 1422 parameter space using a dummy image–label pair  $(x, y)$ :

$$\theta^{(\tau)}(x, y) = \theta_{\text{org}}^{(\tau)} + \eta_{\text{att}} \nabla_{\theta} \ell(f(x; \theta_{\text{org}}^{(\tau)}), y), \quad (15)$$

1423 with attack step size  $\eta_{\text{att}} > 0$ . The attacker’s goal is to choose  $(x, \tau)$  so that the shadowed update  
 1424 in equation 15 closely matches the actual unlearned parameters  $\theta_u$  produced by WARP. Formally, we  
 1425 solve

$$\hat{x}_f, \hat{\tau} \in \arg \min_{x, \tau} \left[ D(\theta^{(\tau)}(x, y), \theta_u) + \lambda_{\text{TV}} \text{TV}(x) + \lambda_{\tau} \Omega(\tau) \right], \quad (16)$$

1426 where  $D(\cdot, \cdot)$  is a parameter-space discrepancy (we use  $\ell_2$  distance over all weights),  $\text{TV}(x)$  is the  
 1427 total-variation regulariser on the image, and  $\Omega(\tau)$  implements a Gaussian prior  $\tau_a \sim \mathcal{N}(1, \sigma_{\text{cob}}^2)$  on  
 1428 each COB scale. We optimize equation 16 by alternating gradient steps on  $x$  and  $\tau$ , with  $\tau$  clipped to  
 1429 a bounded interval around 1 to avoid degenerate scalings.

1430 **Experimental setup.** For a fair comparison, we reuse exactly the reconstruction protocol of  
 1431 Section 4.4 (same model, dataset, forgotten examples, optimizer, and image priors), and only extend  
 1432 the attack to optimize over the COB parameters  $\tau$  via equation 16. We vary the COB prior variance  
 1433  $\sigma_{\text{cob}}$  that defines  $\Omega(\tau)$ , treating each  $\tau_a$  as a scalar random variable centred at 1 with variance  $\sigma_{\text{cob}}$ .  
 1434 We sweep  $\sigma_{\text{cob}} \in \{0, 0.1, 0.2, 0.4, 0.8\}$ , where  $\sigma_{\text{cob}} = 0$  recovers the non-adaptive attack with fixed  
 1435  $\tau \equiv 1$ , and larger values correspond to stronger dispersion along the symmetry orbit induced by  
 1436 WARP. Following the evaluation protocol of Table 2, we quantify reconstruction quality using PSNR,  
 1437 SSIM, LPIPS, and feature MSE, reporting averages over 30 randomly drawn forget examples.

1438 **Results and connection to theory.** Figure 11 shows how reconstruction quality changes as we  
 1439 increase the COB prior std  $\sigma_{\text{cob}}$  that controls the spread of admissible symmetry scales. When  
 1440  $\sigma_{\text{cob}} = 0$  the symmetry prior collapses around  $\tau_a \approx 1$ , so the attacker effectively searches over a  
 1441 narrow neighbourhood of the defender’s true teleportation and can partially recover the forgotten  
 1442 signal: the adaptive attack achieves substantially lower feature MSE and LPIPS than the non-adaptive  
 1443 WARP attack (dashed line). However, the attacker never observes the ground-truth COB scales  
 1444 used by the defender; as  $\sigma_{\text{cob}}$  grows and the symmetry orbit broadens, the optimisation over  $(x, \tau)$   
 1445 quickly becomes unstable. Both metrics deteriorate almost monotonically with  $\sigma_{\text{cob}}$ : already at  
 1446 moderate variance the gains over the non-adaptive attack largely disappear, and for the largest tested  
 1447  $\sigma_{\text{cob}}$  the adaptive reconstructions are statistically indistinguishable from (or slightly worse than) the  
 1448 non-adaptive baseline. Importantly, the COB standard deviation is a defender-controlled knob: in  
 1449 this symmetry family we can raise  $\sigma_{\text{cob}}$  up to 1.0 without changing the realised network function,  
 1450 and in our main reconstruction experiments in Table 2 we set  $\sigma_{\text{cob}} = 0.8$ , already placing the attacker  
 1451 in a high-variance regime where adaptive reconstruction is strongly impaired.

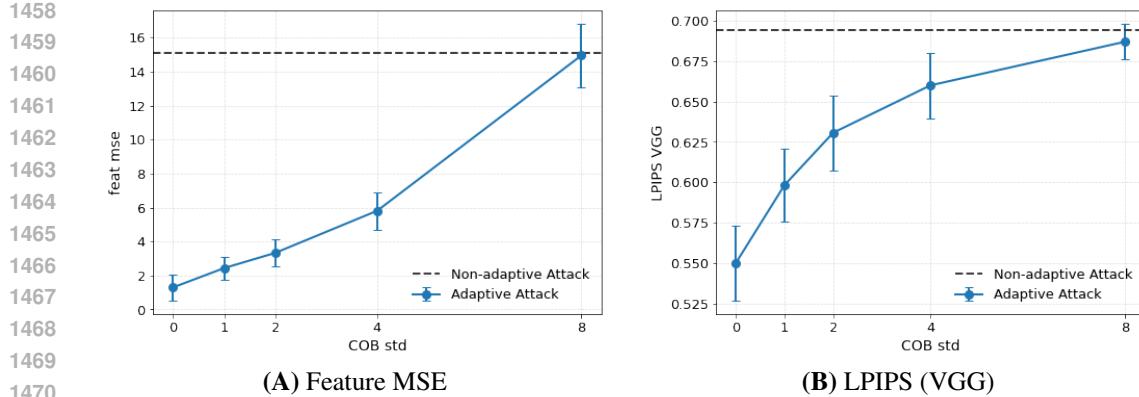


Figure 11: **Adaptive reconstruction under change-of-basis teleportation (NGP, ImageNet-1K).** (A) Feature MSE and (B) LPIPS (VGG) as a function of the COB standard deviation  $\sigma_{\text{cob}}$ . Increasing the symmetry variance consistently worsens reconstruction quality across both metrics.

This trend is consistent with our theoretical analysis in Appendix O, which shows that the expected reconstruction error increases with the variance of the COB scales. Larger  $\sigma_{\text{cob}}$  expands the symmetry orbit of  $\theta_{\text{org}}$  and  $\theta_u$ , so the update  $\Delta\theta$  admits many symmetry-equivalent decompositions whose gradients are nearly orthogonal to the true forget gradient  $g_f$ . The optimisation problem in equation 16 thus becomes a highly ill-posed inverse problem over the joint space  $(x, \tau)$ , where many different configurations of  $(x, \tau)$  produce similar matches in parameter space. Empirically, the adaptive optimiser drifts toward such low-signal-to-noise solutions that satisfy the symmetry constraints but no longer encode the specific forgotten example, explaining the systematic degradation in reconstruction quality as symmetry variance (or std) increases.

**Takeaway.** Even under a strong white-box threat model—where the attacker knows the teleportation family and jointly adapts both the dummy input and the symmetry parameter—teleportation continues to disrupt reconstruction effectively. The injected symmetry components become entangled with the forget-induced update  $\Delta\theta$ , enlarging the attacker’s search space and destroying the geometric alignment between parameter differences and the underlying forgotten example. Thus, teleportation does not merely reparameterise the model in a way that can be inverted; instead, by injecting symmetry variance into the update, it structurally increases reconstruction error and removes the clean gradient-based signal that standard reconstruction attacks depend on. This provides empirical and theoretical evidence that symmetry-based teleportation fundamentally hardens the inverse problem faced by adaptive adversaries.

## O TELEPORTATION-AWARE INFORMATION-THEORETIC BOUNDS ON GRADIENT-BASED RECONSTRUCTION

### O.1 OVERVIEW OF THE THEORETICAL ANALYSIS

This appendix develops an information-theoretic lower bound on the minimal reconstruction mean-squared error (MSE) achievable by a gradient-based inversion adversary within a shared probabilistic model for gradients. We first adapt standard entropy-MSE relationships to the case where the attacker observes gradients rather than intermediate features, closely following the spirit of the analysis in Xia et al. (2025). We then introduce a Gaussian-mixture model (GMM) for gradient features and derive a parametric lower bound on the conditional entropy  $H(x | g)$ , analogous to the intermediate-feature analysis in Xia et al. (2025) but specialized to gradients. Finally, we incorporate teleportation (change-of-basis) noise as private randomness in the training dynamics and analyze its impact on the *same* lower-bound pipeline, under an explicit diagonal approximation and an energy-preserving design assumption on the change-of-basis distribution. Throughout, we keep the modelling assumptions identical between the teleported and non-teleported channels, so any improvement we prove directly reflects a genuine tightening of the analytic lower bound on

1512 reconstruction error—and hence a provable gain in information-theoretic privacy *within this common*  
 1513 *generative framework*. We emphasize that  $H(x)$  is fixed by the dataset distribution, so only *relative*  
 1514 differences between the channels are meaningful.  
 1515

1516 **O.2 SETUP AND THREAT MODEL**

1518 **Data and model.** Let  $x \in \mathbb{R}^d$  denote the  $d$ -dimensional input random variable, distributed according  
 1519 to some unknown data distribution on a measurable subset  $\mathcal{X} \subseteq \mathbb{R}^d$ . We assume throughout that  $x$   
 1520 admits a density w.r.t. Lebesgue measure and has finite second moment. (If one wishes to model  
 1521 discrete or manifold-supported data, the analysis can be recovered by adding an arbitrarily small  
 1522 Gaussian perturbation to  $x$  as is standard in differential-entropy arguments; we implicitly assume  
 1523 such smoothing has been applied so that conditional covariances below are positive definite.)

1524 Consider a deep network with parameters  $W \in \mathbb{R}^{m \times d}$  and first-layer pre-activations

$$z = Wx \in \mathbb{R}^m,$$

1525 and a subsequent decoder  $F_d$ . Let  $\ell(\cdot, y)$  be a loss for a label  $y$ , and define the gradient with respect  
 1526 to  $z$ :

$$g_z = \nabla_z \ell(F_d(z), y) \in \mathbb{R}^m.$$

1527 In the analysis below, the attacker’s observation will be a gradient-based signal  $g$  (not necessarily  
 1528 equal to  $g_z$  directly) that is deterministically related to  $(x, y, W)$  plus noise. In a white-box setting,  
 1529 for instance, the adversary can observe weight differences across steps, which are affine functions  
 1530 of the underlying gradient features; since mutual information and our entropy-based bounds are  
 1531 invariant under fixed invertible affine reparametrizations, it is without loss of generality to work with  
 1532 a canonical gradient feature  $g$ .  
 1533

1534 **Adversarial objective.** An inversion adversary aims to reconstruct  $x$  from the observable  $g$ . Given  
 1535 an estimator  $\hat{x}(g)$ , we measure reconstruction quality by the mean-squared error (MSE)

$$\xi_g(\hat{x}) := \frac{1}{d} \mathbb{E}[\|x - \hat{x}(g)\|_2^2]. \quad (17)$$

1536 The *minimal* MSE  $\xi_g$  is the infimum of equation 17 over all measurable estimators  $\hat{x}(\cdot)$ . We interpret  
 1537 “information-theoretic robustness” as the regime where the attacker is Bayes-optimal under the  
 1538 assumed generative model, i.e. has access to the true posterior  $P(x | g)$  induced by that model and  
 1539 implements the Minimum Mean Square Error (MMSE) estimator.

1540 *Assumption 1* (Basic regularity). We assume:

- 1541 (i)  $x$  has a density on  $\mathbb{R}^d$  and finite second moment;
- 1542 (ii) for the observation  $g$ , the conditional distribution  $P(x | g)$  admits a density with finite  
 1543 second moment, and its covariance matrix  $\text{Cov}(x | g)$  is positive definite almost surely;
- 1544 (iii) all entropies, mutual informations and expectations used below are finite.

1545 These conditions are standard in information-theoretic MMSE analysis (see, e.g., Xia et al. (2025))  
 1546 and ensure that all quantities are well-defined and that the maximum-entropy characterization for  
 1547 Gaussians can be applied without degeneracy.  
 1548

1549 **O.3 MINIMAL MSE FROM GRADIENTS AND AN ENTROPY-BASED LOWER BOUND**

1550 **Bayes-optimal reconstruction from gradients** We first recall the standard MMSE characterization.

1551 **Proposition 1** (Minimal reconstruction MSE from gradients). *Let  $x \in \mathbb{R}^d$  and an observation  $g$   
 1552 satisfy Assumption 1. Consider estimators  $\hat{x}(g)$  of  $x$  based on  $g$  and define  $\xi_g(\hat{x})$  as in equation 17.  
 1553 Then:*

- 1554 (i) *The estimator that minimizes  $\xi_g(\hat{x})$  is the conditional mean  $\hat{x}^*(g) = \mathbb{E}[x | g]$ .*
- 1555 (ii) *The corresponding minimal MSE is*

$$\xi_g := \inf_{\hat{x}} \xi_g(\hat{x}) = \frac{1}{d} \mathbb{E}_g \left[ \text{Tr}(\text{Cov}(x | g)) \right], \quad (18)$$

1556 *where  $\text{Cov}(x | g)$  denotes the conditional covariance of  $x$  given  $g$  and  $\mathbb{E}_g$  is expectation  
 1557 w.r.t.  $g$ .*

1566 *Proof.* For any fixed  $g$ , the conditional risk  $\mathbb{E}[\|x - \hat{x}(g)\|_2^2 \mid g]$  is uniquely minimized by  $\hat{x}^*(g) =$   
 1567  $\mathbb{E}[x \mid g]$  (standard MMSE theory, cf. Xia et al. (2025)). The minimal conditional risk at  $g$  is  
 1568

$$\mathbb{E}[\|x - \mathbb{E}[x \mid g]\|_2^2 \mid g] = \text{Tr}(\text{Cov}(x \mid g)),$$

1570 since for any random vector  $X$  with mean  $\mu$  and covariance  $\Sigma$  one has  $\mathbb{E}\|X - \mu\|_2^2 = \text{Tr}(\Sigma)$ . Taking  
 1571 expectation over  $g$  and dividing by  $d$  yields equation 18.  $\square$

1572 Thus, when we refer to the “minimal MSE achievable by an attacker” for a given observation model,  
 1573 we mean  $\xi_g$  as given in equation 18, corresponding to a Bayes-optimal adversary within that model.  
 1574

1575 **An entropy-based lower bound on the minimal MSE** We now relate the minimal MSE  $\xi_g$  to  
 1576 the conditional entropy  $H(x \mid g)$ , generalizing standard entropy–MMSE inequalities (cf. Xia et al.  
 1577 (2025)).

1578 **Theorem 1** (Entropy-based lower bound on gradient reconstruction). *Under Assumption 1, let*  
 1579  *$H(x \mid g)$  be the conditional differential entropy of  $x$  given the observation  $g$ . Then the minimal*  
 1580 *reconstruction MSE  $\xi_g$  in equation 18 satisfies*

$$\xi_g \geq \frac{1}{2\pi e} \exp\left(\frac{2}{d} H(x \mid g)\right). \quad (19)$$

1584 *Proof.* Fix  $g$  and define  $\Sigma(g) := \text{Cov}(x \mid g)$ . Under Assumption 1,  $\Sigma(g)$  is symmetric and positive  
 1585 definite almost surely. For each such  $g$ , the conditional distribution of  $x$  given  $g$  has entropy bounded  
 1586 above by that of a Gaussian with the same covariance:

$$H(x \mid g = g) \leq \frac{1}{2} \log((2\pi e)^d \det(\Sigma(g))),$$

1587 with equality iff  $x \mid g$  is Gaussian. This is the usual maximum entropy property of Gaussians. Taking  
 1588 expectation over  $g$  gives

$$H(x \mid g) = \mathbb{E}_g[H(x \mid g = g)] \leq \mathbb{E}_g\left[\frac{1}{2} \log((2\pi e)^d \det(\Sigma(g)))\right]. \quad (20)$$

1590 Let  $\lambda_1(g), \dots, \lambda_d(g)$  be the eigenvalues of  $\Sigma(g)$  (all positive). Then

$$\det(\Sigma(g)) = \prod_{j=1}^d \lambda_j(g), \quad \text{Tr}(\Sigma(g)) = \sum_{j=1}^d \lambda_j(g).$$

1591 By the Arithmetic Mean–Geometric Mean (AM–GM) inequality,

$$\prod_{j=1}^d \lambda_j(g) \leq \left(\frac{1}{d} \sum_{j=1}^d \lambda_j(g)\right)^d = \left(\frac{\text{Tr}(\Sigma(g))}{d}\right)^d,$$

1592 so

$$\log \det(\Sigma(g)) \leq d \log\left(\frac{\text{Tr}(\Sigma(g))}{d}\right).$$

1593 Substituting into equation 20,

$$H(x \mid g) \leq \mathbb{E}_g\left[\frac{1}{2} \log((2\pi e)^d \det(\Sigma(g)))\right] \leq \mathbb{E}_g\left[\frac{d}{2} \log\left(2\pi e \frac{\text{Tr}(\Sigma(g))}{d}\right)\right].$$

1594 Since  $\log(\cdot)$  is concave, Jensen’s inequality yields

$$\mathbb{E}_g\left[\log\left(2\pi e \frac{\text{Tr}(\Sigma(g))}{d}\right)\right] \leq \log\left(2\pi e \frac{\mathbb{E}_g[\text{Tr}(\Sigma(g))]}{d}\right).$$

1595 Therefore

$$H(x \mid g) \leq \frac{d}{2} \log\left(2\pi e \frac{\mathbb{E}_g[\text{Tr}(\Sigma(g))]}{d}\right). \quad (21)$$

1596 By Proposition 1,  $\mathbb{E}_g[\text{Tr}(\Sigma(g))] = d \xi_g$ , so equation 21 becomes

$$H(x \mid g) \leq \frac{d}{2} \log(2\pi e \xi_g).$$

1597 Rearranging,

$$\log(2\pi e \xi_g) \geq \frac{2}{d} H(x \mid g), \quad 2\pi e \xi_g \geq \exp\left(\frac{2}{d} H(x \mid g)\right),$$

1598 which yields equation 19.  $\square$

1620 Note that  $H(x)$ —and hence the absolute scale of these lower bounds—is fully determined by the  
 1621 underlying dataset distribution and does not depend on teleportation. In our comparisons between  
 1622 teleported and non-teleported channels,  $H(x)$  cancels and only *differences* or ratios matter.  
 1623

#### 1624 0.4 A PARAMETRIC LOWER BOUND ON $H(x|g)$ VIA GAUSSIAN MIXTURES

1626 We now introduce a specific probabilistic model for the gradient signal and derive a tractable  
 1627 parametric lower bound on  $H(x|g)$ . The modelling choices mirror those used for intermediate  
 1628 features in Xia et al. (2025), but here are applied to gradients.  
 1629

##### 1630 0.4.1 GRADIENT FEATURE AND OBSERVATION MODEL

1632 **Clean gradient feature.** Let  $G : \mathbb{R}^d \rightarrow \mathbb{R}^m$  be a deterministic mapping producing a *clean* gradient  
 1633 feature from input  $x$ . Specifically, let  $u = G(x) \in \mathbb{R}^m$  denote a feature derived deterministically  
 1634 from  $(x, y, W)$  (e.g., the gradient with respect to first-layer pre-activations, or a flattened stack of  
 1635 first-layer weight gradients). Thus  $u$  is a deterministic function of  $x$  once the model and label are  
 1636 fixed.

1637 *Assumption 2* (Gaussian Mixture Model (GMM) for  $u$ ). We assume that the marginal distribution of  
 1638  $u$  can be well approximated by a Gaussian mixture

$$1639 \quad u \sim \sum_{i=1}^K \pi_i \mathcal{N}(\mu_i, \Sigma_i), \quad \sum_{i=1}^K \pi_i = 1, \quad \pi_i > 0, \quad \Sigma_i \succ 0. \quad (22)$$

1643 This GMM assumption is standard in information-theoretic analyses of representations Xia et al.  
 1644 (2025) and serves as our common surrogate model for gradient features.  
 1645

1646 **Noisy gradient observation.** We model the attacker’s baseline observation as a noisy version of  $u$ :

$$1648 \quad g_0 = u + \varepsilon, \quad \varepsilon \sim \mathcal{N}(0, \Sigma_g), \quad \varepsilon \perp (x, u), \quad (23)$$

1650 where  $\Sigma_g \succ 0$  is a fixed positive-definite covariance matrix. This captures gradient perturbations due  
 1651 to stochastic training, subsampling, or other noise sources;  $\Sigma_g$  is assumed known to the attacker, as  
 1652 in Xia et al. (2025). We use this Gaussian channel as the standard abstraction of gradient perturbations  
 1653 for the subsequent information-theoretic analysis.

##### 1655 0.4.2 A MUTUAL-INFORMATION IDENTITY FOR DETERMINISTIC FEATURES

1656 We will repeatedly use the following simple lemma for deterministic features.

1658 **Lemma 1** (Mutual information for deterministic feature maps). *Let  $u = G(x)$  be a deterministic  
 1659 function of  $x$ , and let  $g$  be a random variable such that  $p(g|x, u) = p(g|u)$  (i.e.,  $g$  depends on  
 1660  $(x, u)$  only through  $u$ ). Then*

$$1661 \quad I(x; g) = I(u; g).$$

1663 Where  $I(x; g)$  denotes the mutual information between  $x$  and  $g$ .

1665 *Proof.* Since  $u$  is a deterministic function of  $x$ , we have  $H(u|x) = 0$  and  $H(x, u) = H(x)$ .  
 1666 Moreover,  $p(g|x) = p(g|u)$  by the conditional-independence assumption, so  
 1667

$$1668 \quad H(g|x) = \mathbb{E}_x H(g|x=x) = \mathbb{E}_x H(g|u=G(x)) = H(g|u).$$

1670 Therefore

$$1671 \quad I(x; g) = H(g) - H(g|x) = H(g) - H(g|u) = I(u; g). \quad \square$$

1673 We will apply this lemma to both the baseline channel  $g_0$  and the teleported channel  $g$  below.

1674 O.4.3 PARAMETRIC GMM-BASED LOWER BOUND ON  $H(x | g_0)$   
1675

1676 We now adapt the mixture-entropy bound used in Xia et al. (2025) to gradients.

1677 **Theorem 2** (Parametric lower bound on  $H(x | g_0)$ ). *Under Assumption 1 and Assumption 2 and the  
1678 channel equation 23, the conditional entropy  $H(x | g_0)$  satisfies*

1679  
1680 
$$H(x | g_0) \geq H(x) - \sum_{i=1}^K \pi_i \left( -\log \pi_i + \frac{1}{2} \log \frac{|\Sigma_i + \Sigma_g|}{|\Sigma_g|} \right). \quad (24)$$
  
1681  
1682

1683 *Proof.* Because  $u = G(x)$  is deterministic given  $x$ , and  $g_0$  depends on  $(x, u)$  only through  $u$   
1684 via equation 23, we have  $g_0 \perp x | u$  and the conditions of Lemma 1 hold. Thus

1685 
$$I(x; g_0) = I(u; g_0),$$
  
1686 and

1687 
$$H(x | g_0) = H(x) - I(x; g_0) = H(x) - I(u; g_0).$$

1688 We bound  $I(u; g_0)$  from above using the GMM model. We have

1689 
$$I(u; g_0) = H(g_0) - H(g_0 | u).$$
  
1690

1691 From equation 23,  $g_0 | u \sim \mathcal{N}(u, \Sigma_g)$ , so

1692 
$$H(g_0 | u) = \frac{1}{2} \log((2\pi e)^m |\Sigma_g|).$$

1693 Marginally,  $g_0$  is the convolution of the GMM  $u$  with the Gaussian  $\varepsilon$ , hence

1694  
1695 
$$g_0 \sim \sum_{i=1}^K \pi_i \mathcal{N}(\mu_i, \Sigma_i + \Sigma_g).$$
  
1696

1697 For any mixture density  $p(z) = \sum_i \pi_i p_i(z)$  with components  $p_i$ , the differential entropy satisfies the  
1698 standard upper bound

1699  
1700 
$$H(p) \leq H(\pi) + \sum_i \pi_i H(p_i),$$

1701 where  $H(\pi) = -\sum_i \pi_i \log \pi_i$  is the discrete entropy of the mixture weights (this follows by  
1702 considering the joint entropy of the component index and the sample). Applying this with Gaussian  
1703 components  $p_i = \mathcal{N}(\mu_i, \Sigma_i + \Sigma_g)$  yields

1704  
1705 
$$H(g_0) \leq \sum_{i=1}^K \pi_i \left( -\log \pi_i + \frac{1}{2} \log((2\pi e)^m |\Sigma_i + \Sigma_g|) \right),$$
  
1706

1707 as in Xia et al. (2025). Therefore

1708  
1709  
1710 
$$I(u; g_0) \leq \sum_{i=1}^K \pi_i \left( -\log \pi_i + \frac{1}{2} \log((2\pi e)^m |\Sigma_i + \Sigma_g|) \right) - \frac{1}{2} \log((2\pi e)^m |\Sigma_g|)$$
  
1711  
1712  
1713 
$$= \sum_{i=1}^K \pi_i \left( -\log \pi_i + \frac{1}{2} \log((2\pi e)^m |\Sigma_i + \Sigma_g|) \right) + \sum_{i=1}^K \pi_i \left( -\frac{1}{2} \log((2\pi e)^m |\Sigma_g|) \right)$$
  
1714  
1715  
1716 
$$= \sum_{i=1}^K \pi_i \left( -\log \pi_i + \frac{1}{2} \log((2\pi e)^m |\Sigma_i + \Sigma_g|) - \frac{1}{2} \log((2\pi e)^m |\Sigma_g|) \right)$$
  
1717  
1718  
1719 
$$= \sum_{i=1}^K \pi_i \left( -\log \pi_i + \frac{1}{2} \log \frac{|\Sigma_i + \Sigma_g|}{|\Sigma_g|} \right).$$
  
1720  
1721

1722 where the  $(2\pi e)^m$  terms cancel. Substituting into  $H(x | g_0) = H(x) - I(u; g_0)$  yields equation 24.  
1723  $\square$ 1724  
1725 Theorem 2 yields a parametric lower bound on  $H(x | g_0)$ —parametric in the GMM and noise  
1726 covariances. Via Theorem 1, this in turn induces a lower bound on the minimal reconstruction MSE  
1727 for an attacker observing  $g_0$ . Our teleportation analysis will reuse exactly the same ingredients (GMM  
approximation and mixture-entropy bound) so comparisons are on equal footing.

1728 O.5 TELEPORTATION / CHANGE-OF-BASIS NOISE ON GRADIENTS  
17291730 We now incorporate teleportation (change-of-basis; CoB) symmetry as a source of private randomness  
1731 in the gradient dynamics and analyze its impact on the *same* lower-bound pipeline used for  $g_0$ .  
17321733 O.5.1 TELEPORTATION AS PRIVATE MULTIPLICATIVE NOISE  
17341735 **Teleportation structure.** For each layer  $\ell$ , let  $\tau^{[\ell]}$  denote the corresponding CoB vector (with all  
1736 entries nonzero). The teleported gradient at layer  $\ell$  is obtained by column-scaling with  $\tau^{[\ell-1]}$  and  
1737 row-scaling with  $1/\tau^{[\ell]}$ , i.e.  
1738

1739 
$$dV^{[\ell]} = \tau^{[\ell-1]} \bullet dW^{[\ell]} \bullet (1/\tau^{[\ell]}),$$

1740 
$$dV_{ij}^{[\ell]} = \tau_j^{[\ell-1]} dW_{ij}^{[\ell]} (1/\tau_i^{[\ell]}),$$

1741 where the left operation multiplies each column of  $dW^{[\ell]}$  by the corresponding coordinate of  $\tau^{[\ell-1]}$ ,  
1742 and the right operation multiplies each row by the corresponding coordinate of  $1/\tau^{[\ell]}$ . Consequently,  
1743 each gradient entry acquires a multiplicative factor equal to a ratio of CoB coordinates. As such, each  
1744 gradient entry picks up a multiplicative factor equal to a ratio of CoB entries. Flattening all gradient  
1745 parameters into a single vector, we write the clean gradient feature as  $u$  and its teleported version as  
1746

1747 
$$\tilde{u} = R(\tau) u, \quad (25)$$

1748 where  $R(\tau)$  is a diagonal matrix with entries  $r_j(\tau) = \tau_{b(j)}/\tau_{a(j)}$  corresponding to the appropriate  
1749 input/output channels  $(a(j), b(j))$  of coordinate  $j$ . In practice, these ratios are constrained by the  
1750 underlying channel-wise  $\tau^{[\ell]}$  structure; our analysis below treats  $\{r_j(\tau)\}$  as effective per-coordinate  
1751 scalings induced by that structure.  
17521753 **Threat model for teleportation.** We adopt the following threat model.  
17541755 *Assumption 3* (Teleportation threat model).

- 1756 (i) The CoB parameters
- $\tau$
- are sampled from a distribution
- $P_\tau$
- that is independent of
- $(x, u)$
- .
- 
- 1757 (ii) Teleportation is applied internally in the training update rule, so that the observable gradient
- 
- 1758 feature (e.g., weight differences across a step) is a function of
- $\tilde{u}$
- rather than
- $u$
- . Algebraically,
- 
- 1759 this yields an observation of the form equation 26 below.
- 
- 1760 (iii) The adversary has white-box access to the model architecture and weights but
- does not*
- 
- 1761 observe
- $\tau$
- directly. They know the distribution
- $P_\tau$
- .
- 
- 1762

1763 Under Assumption 3, the teleported observation channel is  
1764

1765 
$$g = \tilde{u} + \varepsilon = R(\tau) u + \varepsilon, \quad \varepsilon \sim \mathcal{N}(0, \Sigma_g), \quad \varepsilon \perp (x, u, \tau). \quad (26)$$

1766 This is the same additive-noise form as in equation 23, applied to a multiplicatively perturbed feature  
1767  $R(\tau)u$ .  
17681769 O.5.2 TELEPORTATION-AWARE ENTROPY LOWER BOUND  
17701771 We now derive the teleportation-aware counterpart of Theorem 2, using the same GMM approximation  
1772 for  $u$ . Here the relevant mutual-information identity is again supplied by Lemma 1.  
17731774 **Theorem 3** (Teleportation-aware lower bound on  $H(x | g)$ ). *Under Assumption 1, Assumption 2,  
1775 Assumption 3 and the teleported channel equation 26, the conditional entropy  $H(x | g)$  satisfies*  
1776

1777 
$$H(x | g) \geq H(x) - \sum_{i=1}^K \pi_i \left( -\log \pi_i + \frac{1}{2} \mathbb{E}_\tau \log \frac{|R(\tau)\Sigma_i R(\tau)^\top + \Sigma_g|}{|\Sigma_g|} \right), \quad (27)$$

1778 where the expectation is taken w.r.t.  $\tau \sim P_\tau$ .  
17791780 *Proof.* As before,  $u = G(x)$  is deterministic given  $x$ , and  $g$  depends on  $(x, u)$  only through  $(u, \tau)$   
1781 via equation 26. In particular, we have the Markov chain

1782 
$$x \rightarrow u \rightarrow (g, \tau) \rightarrow g,$$

1782 and  $g \perp x \mid (u, \tau)$ . Integrating over the independent  $\tau$  yields  $p(g \mid x, u) = p(g \mid u)$ , and hence the  
 1783 conditions of Lemma 1 hold, giving  
 1784

$$1785 \quad I(x; g) = I(u; g), \quad H(x \mid g) = H(x) - I(x; g) = H(x) - I(u; g).$$

1787 We bound  $I(u; g)$  from above. By the chain rule and independence of  $u$  and  $\tau$ ,

$$1789 \quad I(u; g) = I(u; g, \tau) - I(u; \tau \mid g) = I(u; g \mid \tau) - I(u; \tau \mid g) \leq I(u; g \mid \tau),$$

1790 since  $I(u; \tau \mid g) \geq 0$ . Here  $I(u; g \mid \tau)$  is conditional mutual information and can be written as  
 1791

$$1792 \quad I(u; g \mid \tau) = \mathbb{E}_\tau [I(u; g \mid \tau = t)].$$

1794 For a fixed realization  $\tau = t$ , the channel is linear with Gaussian noise:

$$1795 \quad g \mid \tau = t = R(t)u + \varepsilon.$$

1797 Conditionally on mixture component  $i$ ,  $u \mid i \sim \mathcal{N}(\mu_i, \Sigma_i)$ , so  
 1798

$$1799 \quad g \mid (i, \tau = t) \sim \mathcal{N}(R(t)\mu_i, R(t)\Sigma_i R(t)^\top + \Sigma_g),$$

1800 and  $g \mid \tau = t$  is a GMM with components indexed by  $i$ . For this fixed  $t$ ,

$$1802 \quad I(u; g \mid \tau = t) = H(g \mid \tau = t) - H(g \mid u, \tau = t).$$

1804 Since  $g \mid (u, \tau = t) \sim \mathcal{N}(R(t)u, \Sigma_g)$ , we obtain

$$1806 \quad H(g \mid u, \tau = t) = \frac{1}{2} \log((2\pi e)^m |\Sigma_g|).$$

1807 Using the same mixture-entropy bound as before, applied to the GMM  $g \mid \tau = t$ , we have  
 1808

$$1809 \quad H(g \mid \tau = t) \leq \sum_{i=1}^K \pi_i \left( -\log \pi_i + \frac{1}{2} \log((2\pi e)^m |R(t)\Sigma_i R(t)^\top + \Sigma_g|) \right).$$

1813 Therefore

$$1814 \quad I(u; g \mid \tau = t) \leq \sum_{i=1}^K \pi_i \left( -\log \pi_i + \frac{1}{2} \log((2\pi e)^m |R(t)\Sigma_i R(t)^\top + \Sigma_g|) \right) - \frac{1}{2} \log((2\pi e)^m |\Sigma_g|)$$

$$1817 \quad = \sum_{i=1}^K \pi_i \left( -\log \pi_i + \frac{1}{2} \log \frac{|R(t)\Sigma_i R(t)^\top + \Sigma_g|}{|\Sigma_g|} \right).$$

1820 Taking expectation over  $\tau$  yields  
 1821

$$1822 \quad I(u; g \mid \tau) = \mathbb{E}_\tau I(u; g \mid \tau = t) \leq \sum_{i=1}^K \pi_i \left( -\log \pi_i + \frac{1}{2} \mathbb{E}_\tau \log \frac{|R(\tau)\Sigma_i R(\tau)^\top + \Sigma_g|}{|\Sigma_g|} \right).$$

1825 Combining  $I(u; g) \leq I(u; g \mid \tau)$  with  $H(x \mid g) = H(x) - I(u; g)$  gives equation 27.  $\square$   
 1826

1828 Theorem 3 is the teleportation analogue of Theorem 2, obtained via the same steps, with  $\Sigma_i$  replaced  
 1829 by  $R(\tau)\Sigma_i R(\tau)^\top$  and an additional expectation over  $\tau$ .  
 1830

## 1831 O.6 DIAGONAL APPROXIMATION AND THE ROLE OF THE COB DISTRIBUTION

1833 To make the teleportation effect more interpretable at a per-coordinate level, we now adopt a diagonal  
 1834 approximation. This is a modelling simplification, similar in spirit to Xia et al. (2025), and all  
 1835 comparisons between teleported and baseline channels will be made *within* this shared surrogate  
 approximation.

1836 O.6.1 DIAGONAL APPROXIMATION  
18371838 *Assumption 4* (Diagonal covariance approximation). We work in the canonical channel basis in which  
1839 teleportation is defined and posit that, in this basis,

1840 
$$\Sigma_i = \text{diag}(\sigma_{i,1}^2, \dots, \sigma_{i,m}^2), \quad \Sigma_g = \text{diag}(\gamma_1^2, \dots, \gamma_m^2),$$
  
1841

1842 and the teleportation matrix has the form

1843 
$$R(\tau) = \text{diag}(r_1(\tau), \dots, r_m(\tau)).$$
  
1844

1845 That is, we adopt a surrogate model in which gradient covariance, observation noise and CoB  
1846 factors act coordinatewise in the natural channel basis, rather than attempting to diagonalize arbitrary  
1847 covariances and then reinterpret teleportation in that rotated frame. This is not claimed to be an exact  
1848 description of real networks, but a structured approximation for per-coordinate interpretation.  
1849

1850 Under Assumption 4,

1851 
$$R(\tau)\Sigma_i R(\tau)^\top + \Sigma_g = \text{diag}(\gamma_1^2 + r_1(\tau)^2\sigma_{i,1}^2, \dots, \gamma_m^2 + r_m(\tau)^2\sigma_{i,m}^2),$$
  
1852

1853 and hence

1854 
$$\frac{|R(\tau)\Sigma_i R(\tau)^\top + \Sigma_g|}{|\Sigma_g|} = \prod_{j=1}^m \left(1 + \alpha_{i,j} r_j(\tau)^2\right), \quad \alpha_{i,j} := \frac{\sigma_{i,j}^2}{\gamma_j^2}. \quad (28)$$
  
1855  
1856

1857 Taking logs and expectation in equation 27, we obtain

1858 
$$\mathbb{E}_\tau \log \frac{|R(\tau)\Sigma_i R(\tau)^\top + \Sigma_g|}{|\Sigma_g|} = \sum_{j=1}^m \psi_{i,j},$$
  
1859  
1860

1861 where we define the per-coordinate quantities  
1862

1863 
$$\psi_{i,j} := \mathbb{E}_\tau [\log(1 + \alpha_{i,j} r_j(\tau)^2)]. \quad (29)$$
  
1864

1865 Thus Theorem 3 becomes, under Assumption 4,

1866 
$$H(x | g) \geq H(x) - \sum_{i=1}^K \pi_i \left( -\log \pi_i + \frac{1}{2} \sum_{j=1}^m \psi_{i,j} \right). \quad (30)$$
  
1867  
1868  
1869

1870 O.6.2 BASELINE (NON-TELEPORTED) DIAGONAL BOUND  
18711872 For comparison, if no teleportation is applied, we have  $R(\tau) \equiv I$  and  $r_j(\tau)^2 \equiv 1$ . Under the same  
1873 diagonal surrogate,

1874 
$$\frac{|\Sigma_i + \Sigma_g|}{|\Sigma_g|} = \prod_{j=1}^m (1 + \alpha_{i,j}),$$
  
1875  
1876

1877 and the GMM-based entropy lower bound equation 24 reduces to

1878 
$$H(x | g_0) \geq H_0^{\text{lb}} := H(x) - \sum_{i=1}^K \pi_i \left( -\log \pi_i + \frac{1}{2} \sum_{j=1}^m \log(1 + \alpha_{i,j}) \right). \quad (31)$$
  
1879  
1880  
1881

1882 We explicitly introduce  $H_0^{\text{lb}}$  to denote the analytic lower bound on  $H(x | g_0)$  obtained under the  
1883 GMM and diagonal surrogate.1884 Similarly, in the teleported diagonal setting equation 30 we define  
1885

1886 
$$H(x | g) \geq H_{\text{tele}}^{\text{lb}} := H(x) - \sum_{i=1}^K \pi_i \left( -\log \pi_i + \frac{1}{2} \sum_{j=1}^m \psi_{i,j} \right). \quad (32)$$
  
1887  
1888  
1889

Both  $H_0^{\text{lb}}$  and  $H_{\text{tele}}^{\text{lb}}$  are computed from exactly the same modelling ingredients and diagonal surrogate.

1890 O.6.3 ENERGY-PRESERVING CoB AND IMPROVEMENT OF THE BOUND  
1891

1892 To isolate teleportation as a pure source of *randomization* (rather than a trivial global rescaling of  
1893 gradient energy), we consider energy-preserving CoB distributions at the level of the per-coordinate  
1894 effective ratios.

1895 *Assumption 5* (Energy-preserving CoB marginals). For each coordinate  $j$ , the marginal distribution  
1896 of  $r_j(\tau)^2$  satisfies  $\mathbb{E}_\tau[r_j(\tau)^2] = 1$ .

1897 This condition enforces that, on average, teleportation does not inflate or shrink per-coordinate  
1898 gradient energy; it only redistributes it stochastically. In practice, the defender controls the sampling  
1899 of  $\tau$  and hence the induced distribution of ratios  $\{r_j(\tau)\}$ , subject to architectural constraints (shared  
1900 channels, etc.). We do not model those constraints explicitly here; we treat  $\{r_j(\tau)\}$  as effective  
1901 per-coordinate scalings whose marginals can be chosen to satisfy Assumption 5.

1902 We do not assume independence of  $r_j(\tau)$  across  $j$ , only these marginals.

1903 Define, for each  $(i, j)$ ,

$$1905 \Delta\psi_{i,j} := \log(1 + \alpha_{i,j}) - \psi_{i,j} = \log(1 + \alpha_{i,j}) - \mathbb{E}_\tau[\log(1 + \alpha_{i,j}r_j(\tau)^2)]. \quad (33)$$

1906 Subtracting equation 31 from equation 32 yields an exact relation between the two analytic entropy  
1907 lower bounds under the diagonal surrogate.

1908 **Corollary 1** (Exact relation between diagonal entropy lower bounds). *Under Assumption 4, the  
1909 diagonal entropy lower bounds equation 31–equation 32 satisfy*

$$1911 H_{\text{tele}}^{\text{lb}} = H_0^{\text{lb}} + \frac{1}{2} \sum_{i=1}^K \pi_i \sum_{j=1}^m \Delta\psi_{i,j}, \quad (34)$$

1914 with  $\Delta\psi_{i,j}$  defined in equation 33. If, in addition, Assumption 5 holds, then each  $\Delta\psi_{i,j}$  is non-  
1915 negative, and hence

$$1916 H_{\text{tele}}^{\text{lb}} \geq H_0^{\text{lb}}. \quad (35)$$

1918 *Proof.* Equation equation 34 is obtained by direct subtraction of equation 31 from equation 32 and  
1919 using equation 33. For the sign of  $\Delta\psi_{i,j}$ , fix  $\alpha > 0$  and define  $\phi_\alpha(t) := \log(1 + \alpha t)$ , which is  
1920 concave on  $t > 0$ . Under Assumption 5,

$$1921 \psi_{i,j} = \mathbb{E}_\tau[\phi_{\alpha_{i,j}}(r_j(\tau)^2)] \leq \phi_{\alpha_{i,j}}(\mathbb{E}_\tau[r_j(\tau)^2]) = \phi_{\alpha_{i,j}}(1) = \log(1 + \alpha_{i,j}),$$

1922 so  $\Delta\psi_{i,j} \geq 0$  for all  $(i, j)$ , implying equation 35.  $\square$

1924 *Remark 1* (Scope and strength of the entropy result). Within the shared modelling assumptions  
1925 (GMM, diagonal surrogate, energy-preserving CoB), Corollary 1 shows that teleportation *never*  
1926 *decreases* the analytic entropy lower bound:

$$1927 H(x | g_0) \geq H_0^{\text{lb}}, \quad H(x | g) \geq H_{\text{tele}}^{\text{lb}} \geq H_0^{\text{lb}}.$$

1929 We stress that we do *not* claim  $H(x | g) \geq H(x | g_0)$  for the true channels. Rather, we compare the  
1930 surrogate quantities  $H_0^{\text{lb}}$  and  $H_{\text{tele}}^{\text{lb}}$  arising under the same generative model; under this common lens,  
1931 teleportation strictly improves the analytic lower bound on uncertainty about  $x$ .

1932 O.7 TELEPORTATION-AWARE RECONSTRUCTION LOWER BOUND  
1933

1934 We now translate the entropy bounds into reconstruction MSE lower bounds using Theorem 1.

1936 O.7.1 BASELINE AND TELEPORTED MSE LOWER BOUNDS  
1937

1938 From equation 31–equation 32 and Theorem 1, we obtain analytic lower bounds on the minimal  
1939 reconstruction MSE for the baseline and teleported channels:

$$1940 \xi_0 := \frac{1}{2\pi e} \exp\left(\frac{2}{d} H_0^{\text{lb}}\right), \quad (36)$$

$$1942 \xi_{\text{tele}} := \frac{1}{2\pi e} \exp\left(\frac{2}{d} H_{\text{tele}}^{\text{lb}}\right). \quad (37)$$

1944 By construction and monotonicity of the exponential,  
 1945  
 1946  $\xi_{g0} \geq \underline{\xi}_0, \quad \xi_g \geq \underline{\xi}_{\text{tele}},$  (38)

1947 where  $\xi_{g0}$  and  $\xi_g$  are the true minimal MSEs for the baseline and teleported channels, respectively.  
 1948 Again,  $H(x)$  is common to both channels and cancels in all *relative* statements about  $\underline{\xi}_{\text{tele}}/\underline{\xi}_0$ .  
 1949

### 1950 O.7.2 IMPROVEMENT FACTOR ON THE ANALYTIC MSE BOUND

1951 Combining the definitions, the ratio between the teleported and baseline MSE *lower bounds* satisfies

1953  
 1954 
$$\frac{\underline{\xi}_{\text{tele}}}{\underline{\xi}_0} = \exp\left(\frac{2}{d}(H_{\text{tele}}^{\text{lb}} - H_0^{\text{lb}})\right) = \exp\left(\frac{1}{d} \sum_{i=1}^K \pi_i \sum_{j=1}^m \Delta\psi_{i,j}\right),$$
 (39)

1956 with  $\Delta\psi_{i,j}$  as in equation 33.

1958 Under the energy-preserving assumption (Assumption 5),  $\Delta\psi_{i,j} \geq 0$ , hence the exponential factor  
 1959 in equation 39 is at least 1, and the analytic teleportation-aware MSE lower bound is never smaller  
 1960 than the baseline one. In other words, teleportation provably raises the information-theoretic floor on  
 1961 reconstruction accuracy *as captured by this shared surrogate model*. We do not assert any ordering  
 1962 between the true minimal MSEs  $\xi_{g0}$  and  $\xi_g$ .

1963 *Remark 2* (Interpretation for privacy). Equation equation 39 provides a quantitative, distribution-  
 1964 aware guarantee: under the shared assumptions (GMM, diagonal surrogate, energy-preserving CoB),  
 1965 teleportation inflates the analytic lower bound on the attacker's reconstruction MSE by a factor given  
 1966 by the RHS of equation 39. This factor depends on the CoB distribution only through  $\Delta\psi_{i,j}$ , which in  
 1967 turn are functions of the per-coordinate signal-to-noise ratios  $\alpha_{i,j}$  and the marginals of  $r_j(\tau)^2$ . Thus  
 1968 teleportation is not merely a heuristic perturbation: for any attacker whose behaviour is dominated by  
 1969 this generative model (in essentially the same sense as in Xia et al. (2025)), there is a formal lower  
 1970 bound on how accurately they can reconstruct  $x$ .

## 1971 O.8 LOG-NORMAL COB FAMILY (EFFECTIVE MODEL)

1973 We now specialize the general diagonal analysis to an effective log-normal model for the CoB-induced  
 1974 per-coordinate scalings  $r_j(\tau)^2$ , to make the dependence on CoB variance explicit in the analytic MSE  
 1975 lower bounds.

1976 **Log-normal marginal model.** We model each per-coordinate scaling as

1977 
$$r_j(\tau)^2 = \exp(Y_j),$$

1979 where

1980 
$$Y_j \sim \mathcal{N}\left(-\frac{1}{2}s_j^2, s_j^2\right),$$

1982 so that

1983 
$$\mathbb{E}[r_j(\tau)^2] = \mathbb{E}[e^{Y_j}] = \exp\left(-\frac{1}{2}s_j^2 + \frac{1}{2}s_j^2\right) = 1.$$

1984 This ensures the energy-preserving condition  $\mathbb{E}[r_j(\tau)^2] = 1$  (Assumption 5), while the parameter  
 1985  $s_j^2 \geq 0$  controls the strength of teleportation-induced variability on coordinate  $j$ . Practically, the  
 1986 defender can aim to implement such marginals by sampling  $\tau$  so that the induced ratios  $r_j(\tau)^2$   
 1987 are approximately log-normal; we do not model the exact mapping from channel-wise  $\tau^{[\ell]}$  to ratio  
 1988 marginals. We emphasize that this is an *effective parametric family* for  $r_j^2$ , chosen for analytical  
 1989 clarity; our rigorous inequalities rely only on Assumption 5, while log-normality is used to express  
 1990 the dependence on a small number of variance parameters.

1991 Under this model, the per-coordinate quantities  $\psi_{i,j}$  and  $\Delta\psi_{i,j}$  admit explicit expressions.

1992 **Corollary 2** (Log-normal teleportation and analytic MSE bound improvement). *Under the log-normal*  
 1993 *CoB marginal model above, for each mixture component  $i$  and coordinate  $j$ ,*

1995 
$$\psi_{i,j}(s_j^2) = \mathbb{E}_{Y_j \sim \mathcal{N}\left(-\frac{1}{2}s_j^2, s_j^2\right)} [\log(1 + \alpha_{i,j} e^{Y_j})],$$
 (40)

1996 and

1997 
$$\Delta\psi_{i,j}(s_j^2) = \log(1 + \alpha_{i,j}) - \mathbb{E}_{Y_j \sim \mathcal{N}\left(-\frac{1}{2}s_j^2, s_j^2\right)} [\log(1 + \alpha_{i,j} e^{Y_j})].$$
 (41)

1998 Let  $\underline{\xi}_0$  and  $\underline{\xi}_{\text{tele}}$  denote the analytic lower bounds on the minimal reconstruction MSE for the  
 1999 non-teleported and teleported channels, respectively, as defined in equation 36–equation 37. Then  
 2000

$$\frac{\underline{\xi}_{\text{tele}}(s^2)}{\underline{\xi}_0} = \exp\left(\frac{1}{d} \sum_{i=1}^K \pi_i \sum_{j=1}^m [\log(1 + \alpha_{i,j}) - \psi_{i,j}(s_j^2)]\right), \quad (42)$$

2004 where  $s^2 = (s_1^2, \dots, s_m^2)$  collects the log-variance parameters across coordinates.  
 2005

2007 *Proof.* The identities equation 40–equation 41 are obtained by substituting  $r_j^2 = e^{Y_j}$  with  
 2008  $Y_j \sim \mathcal{N}(-\frac{1}{2}s_j^2, s_j^2)$  into the definition equation 29 of  $\psi_{i,j}$  and the definition equation 33 of  
 2009  $\Delta\psi_{i,j}$ . The ratio equation 42 then follows immediately by plugging  $\Delta\psi_{i,j}(s_j^2)$  into equation 39,  
 2010 which relates the analytic MSE lower bounds  $\underline{\xi}_{\text{tele}}$  and  $\underline{\xi}_0$  to the  $\Delta\psi_{i,j}$ .  $\square$   
 2011

2012 *Remark 3 (Local small-variance expansion (heuristic)).* To gain intuition about the dependence on  
 2013 teleportation strength, it is useful to consider the regime  $s_j^2 \ll 1$ , where the log-normal marginals are  
 2014 close to the degenerate case  $r_j^2 \equiv 1$ . This section provides a local Taylor expansion for intuition; it is  
 2015 not used in our rigorous inequalities, which already follow from Assumption 5.  
 2016

2017 For  $t_j := r_j^2 = e^{Y_j}$  with  $Y_j \sim \mathcal{N}(-\frac{1}{2}s_j^2, s_j^2)$  we have

$$\mathbb{E}[t_j] = 1, \quad \text{Var}(t_j) = \mathbb{E}[t_j^2] - 1 = \exp(s_j^2) - 1.$$

2018 Thus  $\text{Var}(t_j) = s_j^2 + O(s_j^4)$  as  $s_j^2 \rightarrow 0$ . Writing  $t_j = 1 + \delta_j$ , we have  $\mathbb{E}[\delta_j] = 0$  and  
 2019  $\text{Var}(\delta_j) = \text{Var}(t_j)$ .  
 2020

2022 Since log-normal marginals have finite moments of all orders, a second-order Taylor expansion of  
 2023  $\phi_\alpha(t) := \log(1 + \alpha t)$  around  $t = 1$  yields  
 2024

$$\phi_\alpha(1 + \delta) = \log(1 + \alpha) + \frac{\alpha}{1 + \alpha} \delta - \frac{\alpha^2}{2(1 + \alpha)^2} \delta^2 + R_\alpha(\delta),$$

2025 with  $|R_\alpha(\delta)| \leq C_\alpha |\delta|^3$  for some constant  $C_\alpha$  depending on  $\alpha$ . Taking expectations with  $\mathbb{E}[\delta] = 0$   
 2026 and  $\mathbb{E}[\delta^2] = \text{Var}(t_j)$  gives  
 2027

$$\mathbb{E}[\phi_\alpha(1 + \delta)] = \log(1 + \alpha) - \frac{\alpha^2}{2(1 + \alpha)^2} \text{Var}(t_j) + O(\mathbb{E}[|\delta|^3]).$$

2033 Applying this with  $\alpha = \alpha_{i,j}$  and  $t_j = r_j^2$ , and recalling that  $\psi_{i,j} = \mathbb{E}[\log(1 + \alpha_{i,j} r_j^2)]$ , we obtain the  
 2034 local approximation  
 2035

$$\Delta\psi_{i,j}(s_j^2) = \log(1 + \alpha_{i,j}) - \psi_{i,j}(s_j^2) \approx \frac{\alpha_{i,j}^2}{2(1 + \alpha_{i,j})^2} \text{Var}(t_j),$$

2038 with an error term controlled by  $\mathbb{E}[|\delta_j|^3]$ . For the log-normal model,  $\text{Var}(t_j) = \exp(s_j^2) - 1$ , so we  
 2039 arrive at the heuristic expression  
 2040

$$\Delta\psi_{i,j}(s_j^2) \approx \frac{\alpha_{i,j}^2}{2(1 + \alpha_{i,j})^2} (\exp(s_j^2) - 1), \quad s_j^2 \ll 1.$$

2043 Substituting this into equation 42 yields the corresponding small-variance approximation for the  
 2044 logarithm of the analytic MSE bound ratio:  
 2045

$$\log \frac{\underline{\xi}_{\text{tele}}(s^2)}{\underline{\xi}_0} \approx \frac{1}{2d} \sum_{i=1}^K \pi_i \sum_{j=1}^m \frac{\alpha_{i,j}^2}{(1 + \alpha_{i,j})^2} (\exp(s_j^2) - 1), \quad s_j^2 \ll 1.$$

2050 This approach highlight that, in the small-variance regime, the teleportation-induced improvement in  
 2051 the analytic reconstruction MSE lower bound grows approximately linearly in  $s_j^2$  (via  $\exp(s_j^2) - 1$ ),  
 with a slope governed by the per-coordinate signal-to-noise ratios  $\alpha_{i,j}$  and the mixture weights  $\pi_i$

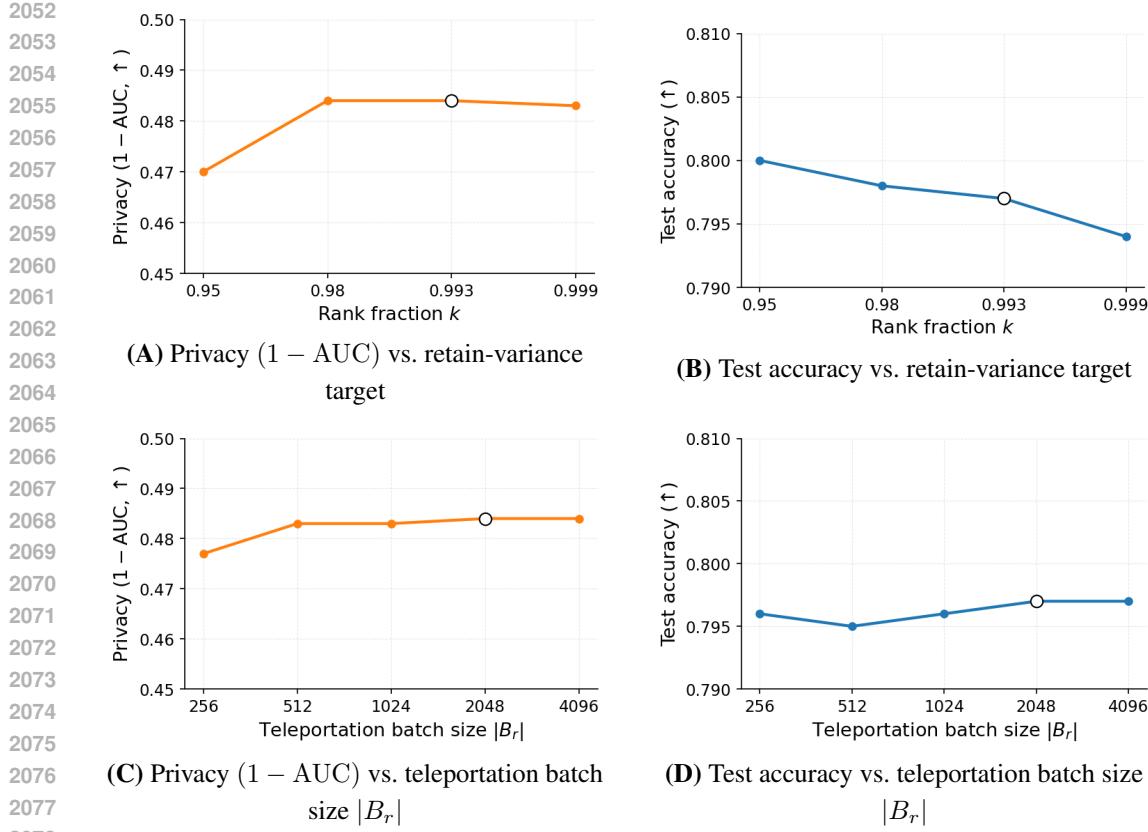


Figure 12: **Sensitivity of teleportation hyperparameters.** Plots (A,B) vary the target retain-variance level used to set the per-layer rank  $k_\ell$ ; plots (C,D) vary the retain minibatch size  $|B_r|$  used to estimate the retain subspace. Privacy is measured as  $1 - \text{AUC}$  of U-LiRA (higher is better). Markers highlight the configuration used in our main experiments (95.3% retain variance and  $|B_r| = 2048$ ).

## P ABLATION: SENSITIVITY OF TELEPORTATION HYPERPARAMETERS

Teleportation introduces a small set of additional hyperparameters that control how strongly we move along symmetry directions. In this section we study the sensitivity of WARP to two core choices: (i) the target retain-variance fraction used to choose the per-layer rank  $k$  in the SVD projector (Section 3.2), and (ii) the size of the retain minibatch  $B_r$  used to estimate the retain subspace. Both directly govern the geometry of the retain null space and the amount of stochasticity in the teleportation step, and were explicitly highlighted as potential sources of instability.

**Setup.** We perform a controlled sweep on CIFAR-10 with ResNet-18 and NGP+WARP under the U-LiRA black-box auditor (Section 4.2). For the SVD projector, we vary the target retain-variance level from 95% to 99.9%, which induces different per-layer ranks  $k_\ell$  such that the top singular vectors of  $R_\ell(D_r)$  capture the chosen fraction of retain energy. For the retain minibatch, we vary the teleportation batch size  $|B_r| \in \{256, 512, 1024, 2048, 4096\}$  while keeping the forget minibatch and unlearning hyperparameters fixed. For each configuration we run the full unlearning pipeline and record test accuracy as well as privacy measured by  $(1 - \text{AUC})$  of U-LiRA (higher is better).

**Results and discussion.** Figure 12 shows that teleportation is *remarkably insensitive* to both hyperparameters in the regime we consider.

**Retain-variance target.** Increasing the target retain-variance from 95% to 99.9% changes privacy  $(1 - \text{AUC})$  by less than 0.015 in absolute terms, while test accuracy varies in a narrow band of  $\approx 0.79\text{--}0.80$ . Privacy slightly improves as we move from 95% to around 99.3%, after which the curve flattens: very high targets effectively make the retain projector full-rank, leaving less room for

2106 teleportation to move in symmetry directions and yielding diminishing returns. The configuration  
 2107 used in the main experiments (target retain-variance  $\approx 99.3\%$ ) lies near this plateau, indicating that  
 2108 our chosen rank provides a good privacy–utility compromise.

2109 *Retain minibatch size  $|B_r|$ .* Varying  $|B_r|$  over an order of magnitude has only a minor effect: privacy  
 2110 ( $1 - \text{AUC}$ ) shifts by at most  $\sim 0.01$ , and test accuracy remains within  $\pm 0.2\%$  points of 0.796. Even  
 2111 relatively small batches ( $|B_r| = 256$ ) already provide a sufficiently representative retain subspace for  
 2112 teleportation, and larger batches only yield a slight, saturating gain in privacy. This suggests that the  
 2113 random retain minibatch need not tightly approximate the full retain set to obtain a stable projector  
 2114 and effective defense; in practice, a modest  $|B_r|$  balances computational cost with stable subspace  
 2115 estimation.

2116 Overall, these ablations show that WARP’s performance does not hinge on fragile hyperparameter  
 2117 choices: both privacy and utility are stable across wide ranges of the SVD rank and retain minibatch  
 2118 size. Moreover, the small spread in test accuracy ( $< 0.6\%$  across all settings) empirically confirms  
 2119 that teleportation remains approximately loss-preserving on the retain set, providing an implicit  
 2120 bound on worst-case retain-loss drift in our experiments.

2121

2122

2123

2124

2125

2126

2127

2128

2129

2130

2131

2132

2133

2134

2135

2136

2137

2138

2139

2140

2141

2142

2143

2144

2145

2146

2147

2148

2149

2150

2151

2152

2153

2154

2155

2156

2157

2158

2159