

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 MEMBERSHIP PRIVACY RISKS OF SHARPNESS AWARE MINIMIZATION

005 **Anonymous authors**

006 Paper under double-blind review

ABSTRACT

011 Optimization algorithms that seek flatter minima, such as Sharpness-Aware Mini-
012 mization (SAM), are credited with improved generalization and robustness to noise.
013 We ask whether such gains impact membership privacy. Surprisingly, we find
014 that SAM is more prone to Membership Inference Attacks (MIA) than classical
015 SGD across multiple datasets and attack methods, despite achieving lower test
016 error. This suggests that the very geometric mechanism of SAM that improves
017 generalization simultaneously exacerbates membership leakage. We investigate
018 this phenomenon through extensive analysis of memorization and influence scores.
019 Our results reveal that SAM is more capable of capturing atypical subpatterns,
020 leading to higher memorization scores of samples. Conversely, SGD depends more
021 heavily on majority features, exhibiting worse generalization on atypical subgroups
022 and lower memorization. Crucially, this characteristic of SAM can be linked to
023 lower variance in the output confidence for unseen samples, thereby amplifying
024 membership signals. Finally, we model SAM under a perfectly interpolating linear
025 regime and theoretically prove that geometric mechanism of SAM inherently re-
026 duces variance, guaranteeing a higher MIA advantage for confidence and likelihood
027 ratio attacks.

1 INTRODUCTION

031 Sharpness-Aware Minimization (SAM) has emerged as a prominent optimization technique for
032 improving generalization in deep learning by encouraging flatter minima – i.e., similar loss values for
033 weight perturbations of certain degree around the optima – in the loss landscape (Norton & Royset,
034 2021; Foret et al., 2020; Wu et al., 2020; Kim et al., 2022; Du et al., 2022; Kwon et al., 2021). Flatter
035 optima have been linked to robustness to noise and improved test performance (Chen et al., 2023;
036 Foret et al., 2020; Baek et al., 2024), while a tension exists between whether SAM’s implicit bias
037 is geared more towards diversity (Springer et al., 2024) or simplicity (Andriushchenko et al., 2023;
Chang & Khanna, 2025) of features.

038 Models that generalize well are thought to rely less on memorizing specific training examples, which
039 should also improve privacy. Consider membership inference attacks (MIAs) in which an attacker
040 exploits the model behavior gap between training and unseen data to infer if a data point was part of
041 the training data or not (Shokri et al., 2017). Intuitively, when a model strongly overfits (training error
042 \ll test error), MIA would become easier. Yeom et al. (2018) formally showed that, under certain
043 assumptions, the advantage of a threshold-based MIA is upper bounded by the model’s generalization
044 error. In light of this, one would naturally expect that a technique like SAM – which demonstrably
045 improves generalization – should also decrease a model’s susceptibility to MIAs.

046 Contrary to this expectation, we find that models trained with SAM are actually *more* vulnerable
047 to MIAs than SGD consistently across diverse datasets and attack methods, even as they achieve
048 better generalization (see Tables 1 and 2). Furthermore, this finding challenges the notion that *flatter*
049 *minima=good* from a privacy standpoint, calling for a deeper investigation into the relationship
050 between generalization, memorization, and privacy to unearth this phenomenon both empirically
051 and theoretically. Our work is the first to systematically demonstrate higher membership privacy
052 leakage for a sharpness-based algorithm known to generalize better, connecting loss sharpness to MIA
053 vulnerability. We note that there have been previous works that exhibit utility–privacy tradeoffs (Long
et al., 2018; Carlini et al., 2021; Chen et al., 2022a; Liu et al., 2024) or demonstrate that higher

054
 055 Table 1: Attack accuracy of direct threshold MIA on SGD and SAM showing tradeoffs in test
 056 accuracy and MIA attacks. In **blue** we highlight the best performing model on the test set and in **red**
 057 the model against which MIA is more successful. SAM is more prone to direct threshold attacks.

Dataset	Algo	NN	Confidence	Entropy	M-entropy	Test Acc
CIFAR-100	SGD	76.62%	77.19%	76.61%	77.30%	80.30%
	SAM	77.99%	79.10%	78.66%	79.25%	81.60%
CIFAR-10	SGD	50%	59.37%	59.09%	59.51%	96.00%
	SAM	50.08%	61.64%	61.64%	61.70%	96.72%
Purchase-100	SGD	66.00%	66.76%	64.78%	67.13%	85.50%
	SAM	66.62%	67.30%	65.35%	67.54%	85.54%
Texas-100	SGD	59.81%	65.20%	55.74%	65.13%	50.83%
	SAM	59.56%	66.59%	57.14%	65.42%	51.34%
EyePacs	SGD	73.62%	73.40%	68.50%	73.40%	73.67%
	SAM	77.73%	77.07%	73.37%	77.36%	75.41%

070
 071
 072
 073
 074 generalization does not necessarily decrease privacy leakage (Kaya & Dumitras, 2021; Del Grosso
 075 et al., 2023).

076 The difference in model behavior on data points that were part of the training set versus those that were
 077 not can be more precisely quantified using *memorization scores* (Feldman, 2020; Feldman & Zhang,
 078 2020), defined via Leave-One-Out (LOO) error. Memorization scores measure the change in model
 079 performance when a specific training sample is removed, and thus serve as a proxy for how much the
 080 model has memorized that sample. Motivated by this connection, we analyze the memorization scores
 081 of samples trained with SAM and find that *SAM exhibits more memorization than SGD*, indicating
 082 a stronger reliance on individual training samples. While this increased memorization provides a
 083 plausible explanation for SAM’s heightened vulnerability to MIA, it raises a key question: “How can
 084 a model that memorizes more generalize better?”

085 We hypothesize that the answer lies in *what* is being memorized. Under overparameterization
 086 (Allen-Zhu et al., 2019), models can learn not only noise in the data, but also atypical patterns in
 087 under-represented subpopulations through memorization—i.e., few white tiger images with numerous
 088 yellow tiger images. This distinction is important as real world datasets are known to have a long tail
 089 of such rare subclasses (Feldman, 2020). We conjecture that SAM is capable of doing more structured
 090 memorization, selectively focusing on atypical subclass patterns, which contributes positively to
 091 generalization. Corroborating this hypothesis, we observe that SAM’s memorization score distribution
 092 is concentrated in the mid range, rather than the high end—which is typically associated with noise
 093 memorization (refer to Section 4.1). This suggests that SAM emphasizes samples that are neither
 094 trivially learned nor purely noisy, but instead represent rare, but generalizable sub-patterns.

095 To further validate this finding, we analyze influence scores—which measure the impact of individual
 096 training samples on test predictions (see Section 4.2). Our results show that, for SAM, samples
 097 corresponding to moderate memorization exert higher influence on test predictions compared to SGD,
 098 confirming that SAM’s generalization gains derive from its ability to better capture rare sub-patterns.
 099 Conversely, for SGD, the lower influence of such points implies that the majority pattern is learned
 100 dominantly: since this feature is redundant across many samples, the marginal influence of any single
 101 point is diluted. These results seem to suggest more that SAM’s implicit bias is towards diversity as
 102 opposed to simplicity.

103 We support our intuition further by introducing a novel metric that quantifies the degree of memo-
 104 rization involved in predicting a test sample in Section 4.3. Using this metric, we dissect SAM’s
 105 performance gains and identify that SAM’s improvements mostly stem from its performance on
 106 atypical test samples that depend heavily on a handful of memorized training points. Meanwhile,
 107 SGD performs slightly better on typical samples that rely more on broadly learned features. Stronger
 108 influence of minority samples can lead to greater membership privacy risk due to the increased
 109 retention of information an attacker can exploit.

108
 109 Table 2: Comparison of online shadow-model MIA on SGD and SAM. In **blue** we highlight the best
 110 performing model on the test set, and in **red** the model with higher privacy leakage (higher AUC,
 111 Attack Accuracy, and TPR@0.1%FPR). SAM is more prone to shadow-model attacks.

Dataset	Attack	SGD				SAM			
		Test Acc	AUC	Attack Acc	TPR@.1	Test Acc	AUC	Attack Acc	TPR@.1
CIFAR-100	RMIA	67.7%	90.4%	80.8%	21.0%	69.1%	91.6%	82.2%	23.4%
	LiRA	92.6%	92.6%	82.9%	27.0%		93.7%	84.1%	31.0%
CIFAR-10	RMIA	92.3%	71.4%	63.5%	4.8%	93.1%	74.9%	65.9%	6.7%
	LiRA	73.0%	73.0%	64.2%	8.8%		76.4%	66.7%	12.5%
Purchase100	RMIA	76.5%	68.8%	62.7%	1.5%	77.4%	70.2%	63.7%	1.7%
	LiRA	68.9%	68.9%	62.6%	1.4%		70.2%	63.4%	1.6%
Texas100	RMIA	46.9%	79.8%	70.8%	2.9%	49.2%	80.6%	71.5%	2.8%
	LiRA	80.8%	80.8%	71.3%	6.9%		81.6%	72.0%	8.2%

123
 124
 125 However, what intrinsic mechanism drives this structured memorization? Results for confidence
 126 threshold attack indicates that SGD produces more predictions with extreme confidence compared to
 127 SAM. These samples reside beyond the threshold and are source of attacker’s error. This suggests a
 128 distinct geometric effect: *SAM induces a shrinkage in the variance of the model’s output predictions*.
 129 This property translates to structured memorization and suppression of majority sub-feature. Relying
 130 heavily on a single/few majority subclass feature to classify diverse inputs requires the model to
 131 assign large weights to that feature. Geometrically, this creates a steep decision boundary and,
 132 consequently, high output variance under perturbation. By penalizing this sharpness, SAM prohibits
 133 the amplification of the majority feature, effectively forcing the model to distribute its reliance across
 134 diverse, subclass-specific features.

135 Completing this conjecture, we provide a theoretical foundation for this mechanism in Section 5. We
 136 prove that the interpolating solution favored by sharpness-aware geometry inherently reduces the
 137 variance of the output logits. We then demonstrate how this variance reduction amplifies the attacker’s
 138 advantage for both confidence-based and likelihood ratio attacks. Our proofs highlight a strong result:
 139 *SAM is more vulnerable at any fixed threshold*. We empirically corroborate this with ROC curves in
 140 Figure 7. Lastly, we theoretically analyze a dataset with majority and minority subclasses and show
 141 how capturing minority feature better leads to enhanced generalization in Section C.

142
 143 **Contributions** In summary, our contributions are the following: *(i)* we are the first to empirically
 144 demonstrate that SAM-trained models exhibit higher membership privacy risk than SGD-trained
 145 models, serving as a cautionary tale against *flatter minima=good* notion from a privacy standpoint;
 146 *(ii)* we offer a detailed and conceptually grounded analysis of the root causes of SAM’s generaliza-
 147 tion-memorization relationship, suggesting SAM’s implicit bias towards diversity; *(iii)* we introduce
 148 a novel methodology to dissect generalization gains, proving that SAM’s generalization gains stem
 149 from its performance on unseen atypical samples; *(iv)* we theoretically show variance shrinkage effect
 150 of interpolating sharpness-aware solutions and how it increases MIA risk for both confidence and
 151 likelihood ratio attacks; and *(v)* we theoretically formulate a data distribution composed of subclasses
 152 where stronger alignment with minority subclass features enhances generalization.

153 2 BACKGROUND & PRELIMINARIES

154
 155
 156 **Memorization & Influence scores** For a training algorithm \mathcal{A} that is used to train the model
 157 $f(\cdot)$ using dataset $\mathcal{D} = ((\mathbf{x}_1, y_1), \dots, (\mathbf{x}_n, y_n))$, the amount of label memorization by \mathcal{A} on a sample
 158 $(\mathbf{x}_i, y_i) \in \mathcal{D}$ is defined by Equation (1). The probability is taken over randomness of the algorithm
 159 such as weight initialization.

$$160 \quad mem(\mathcal{A}, \mathcal{D}, i) := \Pr_{f \leftarrow \mathcal{A}(\mathcal{D})} [f(\mathbf{x}_i) = y_i] - \Pr_{f \leftarrow \mathcal{A}(\mathcal{D} \setminus (\mathbf{x}_i, y_i))} [f(\mathbf{x}_i) = y_i] \quad (1)$$

162

163 Influence score of a training example (\mathbf{x}_i, y_i) on test example (\mathbf{x}'_j, y'_j) is defined by:

164
165
$$infl(\mathcal{A}, \mathcal{D}, i, j) = \Pr_{f \leftarrow \mathcal{A}(\mathcal{D})} [f(\mathbf{x}'_j) = y'_j] - \Pr_{f \leftarrow \mathcal{A}(\mathcal{D} \setminus (\mathbf{x}_i, y_i))} [f(\mathbf{x}'_j) = y'_j] \quad (2)$$

166

167 **Sharpness Aware Minimization (SAM)** Consider a model $f : X \rightarrow Y$ parameterized by a
168 weight vector \mathbf{w} and a per-sample loss function $l: W \times X \times Y \rightarrow R_+$. Given a dataset $S = \{(\mathbf{x}_1,$
169 $y_1), \dots, (\mathbf{x}_n, y_n)\}$ sampled i.i.d. from a data distribution, the training loss is defined as $L_S(\mathbf{w}) =$
170 $\sum_{i=1}^n l(y_i, f(\mathbf{x}_i, \mathbf{w}))/n$. Sharpness Aware Minimization combines traditional loss with sharpness
171 term to minimize the difference between maximum loss in the vicinity (a Ball of radius ρ : $B(\rho)$) of
172 the current minima. Formally, it is defined as the following:

173
$$L_{SAM}(\mathbf{w}) = \min_{\mathbf{w}} L_S(\mathbf{w}) + [\max_{\epsilon \in B(\rho)} L_S(\mathbf{w} + \epsilon) - L_S(\mathbf{w})] = \min_{\mathbf{w}} \max_{\epsilon \in B(\rho)} L_S(\mathbf{w} + \epsilon) \quad (3)$$

174

175

176

2.1 MEMBERSHIP INFERENCE ATTACKS

177 Consider a victim model f_v trained on dataset $D \sim \mathcal{D}$ and attack model f_a . In a black-box setting,
178 an attacker infers whether a sample (\mathbf{x}, y) belongs to D (IN) or not (OUT). In this paper, we consider
179 two types of attacks: *Direct threshold attacks*, which directly learns a threshold from obtained
180 member/non-member data, such as confidence and entropy attacks; and *Shadow model attacks*, which
181 train proxy models to calibrate membership scores, such as Likelihood Ratio Attack (LiRA) and
182 Robust Membership Inference Attack (RMIA) (Carlini et al., 2022; Zarifzadeh et al., 2024).183 We quantify privacy risk using the empirical attack accuracy, defined as the average of the true
184 positive (TPR) and true negative rates (1-FPR):
185

186
$$Acc_{MIA} = \frac{1}{2n_m} \sum_{i=1}^{n_m} \mathbb{1}[f_a(\mathbf{x}_i, y_i) = 1] + \frac{1}{2n_{nm}} \sum_{j=1}^{n_{nm}} \mathbb{1}[f_a(\mathbf{x}_j, y_j) = 0], \quad (4)$$

187
188

189 where n_m, n_{nm} are the counts of IN and OUT samples. Additionally, metrics such as Area Under
190 ROC Curve (AUC) and TPR at low FPR are employed to characterize vulnerability. Further details
191 are provided in Appendix E.2.192

3 PRIVACY RISKS OF SAM

193 Inspired by the link between SAM and generalization and how MIAs should exploit poor generalization-
194 we here scrutinize the membership privacy risk of SAM by comparing the membership attack
195 accuracy (see Equation (4) and Table 2) of different MIAs against SAM- and SGD-trained models
196 across five different benchmark datasets for direct threshold attacks and four different benchmark
197 datasets for shadow model attacks.198 We utilize datasets and target models that are widely employed in studies on MIAs and defenses
199 (Yeom et al., 2018; Fang & Kim, 2024; Chen et al., 2022b; Jia et al., 2019). Furthermore, we assume
200 that the attacker has access to some portion of the training data and non-training data that it uses to
201 train the attack models—a common assumption in the MIA literature.

202

203 **Datasets** We use CIFAR-10, CIFAR-100 and Purchase-100 along with two medical datasets
204 Texas-100 and EyePacs. For direct threshold attacks, we follow Tang et al. (2022) to determine the
205 partition between training and test data and to determine the subset that constitutes the attacker’s prior
206 knowledge¹. For shadow model attacks, we use a different dataset split to account for shadow model
207 training². More details about the datasets and the experimental setup can be found in Appendix H.

208

209 **Methods** For direct threshold attacks, we train a set of models and choose the one achieving highest
210 validation accuracy. We then employ different MIA methods – namely NN-based, confidence-based,
211 entropy-based and modified entropy-based attacks (see Appendix E.2 for a detailed formulation of
212 each MIA) – to evaluate the attack accuracy on the target model. For shadow model attacks, we213
214 ¹We adopt and extend the code in <https://github.com/inspire-group/MIAdefenseSELENA>215 ²We adopt and extend the code in <https://github.com/orientino/lira-pytorch>

216 generate 256 random half-splits of the training dataset into member set and non-member set and
 217 train a model for each split. One model is chosen as the target model and all the other models are
 218 used as shadow models for reference. More details about the experimental settings can be found in
 219 Appendix I.
 220

221 **Results** *Direct Threshold Attacks* We report the attack accuracy and test accuracy for each model
 222 in Table 1. We observe that while SAM achieves higher generalization performance, it also incurs
 223 highest attack accuracy for almost all settings. To further investigate the connection between flatness
 224 of minima and membership privacy, we report the results for other sharpness-aware optimizers and
 225 custom designed optimizer that explicitly aims to find sharper minima in Section F. The results support
 226 a relationship between loss landscape geometry and membership privacy, with other optimizers
 227 exhibiting similar behavior. For confidence threshold attack, we observe that SGD model incur
 228 higher number of extremely confident predictions compared to SAM for both correctly classified and
 229 wrongly classified non-members. Because the threshold is typically set near a high value (i.e. 0.92),
 230 non-member samples with high confidence are missed by the attacker. From this observation, we
 231 conjecture that variance of the model output is an important factor in MIA risk. To confirm that these
 232 results are not model-dependent, we report an ablation study in Appendix J and verify that similar
 233 findings can be observed for different model architectures over the same datasets.
 234

235 *Shadow Model Attacks* Table 2 summarizes the results for online Robust MIA (RMIA) and Likelihood
 236 Ratio Attack (LiRA), reporting Attack Accuracy, AUC, and TPR at 0.1% FPR averaged over 10
 237 attack splits (Zarifzadeh et al., 2024; Carlini et al., 2022). Consistent with the previous results, SAM
 238 achieves higher generalization while incurring non-trivially higher privacy leakage across nearly all
 239 settings. We note that the reported test accuracies differ from Table 1 due to difference in data splits
 240 for shadow model training.
 241

242 The results for offline attacks are presented in Table 4. Analogous to the online setting, SAM exhibits
 243 superior generalization paired with heightened MIA vulnerability. Since shadow model attacks
 244 represent the state-of-the-art in membership inference, these results suggest that SAM’s vulnerability
 245 is not merely an artifact of global threshold shifts. Instead, it points to an intrinsic geometric property
 246 of SAM that persists even under rigorous, sample-specific SOTA attacks.
 247

248 4 SAM LEARNS ATYPICAL SUBCLASS FEATURES MORE

249 To investigate the source of SAM’s increased membership privacy risk, we analyze its optimization
 250 behavior through the lens of sample memorization and influence. We follow the procedure of Feldman
 251 & Zhang (2020) to compute the memorization and influence scores for SAM-trained models on
 252 CIFAR-100. We then compare these scores against the publicly available scores for SGD-trained
 253 models on the same dataset³, enabling a direct comparison of sample-level behavior between the two
 254 optimizers.
 255

256 4.1 SAM MEMORIZES ATYPICAL SUB-PATTERNS MORE

257 We first focus on comparing the memorization behavior of SAM and SGD. Figure 1(a) shows kernel
 258 density estimates of memorization scores for both models. Although the overall shapes of the
 259 distributions are similar – reflecting the long-tailed nature of the dataset –, SAM exhibits a lower
 260 density at the lowest end of the spectrum, with the mass redistributed more evenly across the rest of
 261 the range. This indicates that SAM assigns higher memorization scores more broadly, suggesting a
 262 structured memorization of more diverse patterns compared to SGD.
 263

264 To further investigate this behavior, we plot the memorization scores of individual CIFAR-100
 265 samples under both SAM and SGD in Figure 1(b). Each sample is represented as a blue dot, with its
 266 x- and y-coordinates corresponding to its memorization score under SGD and SAM, respectively.
 267 The red diagonal line denotes equal memorization across both optimizers. Samples above this line
 268 and to the left (top-left quadrant) are more memorized by SAM, while those below and to the right
 269 (bottom-right quadrant) are more memorized by SGD.
 270

³<https://pluskid.github.io/influence-memorization/>

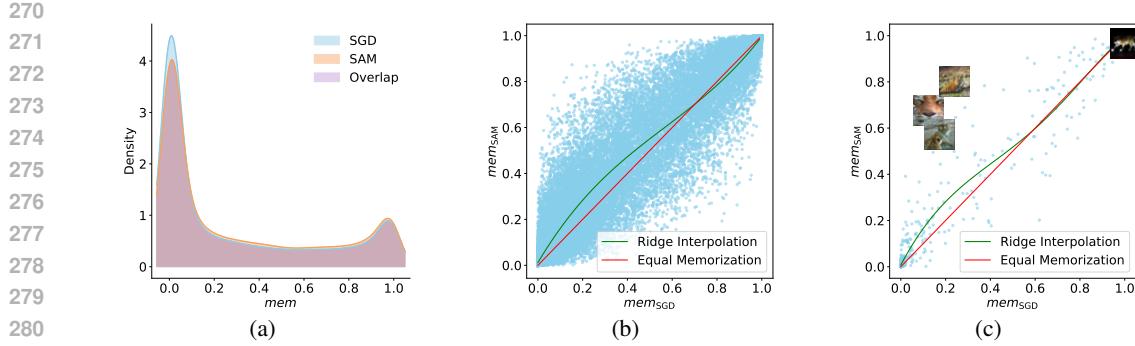


Figure 1: (a): Memorization score density plot for SAM vs SGD. SAM has less density in the lowest range, but more density spread evenly across the remaining range. (b): Memorization scores of CIFAR-100 training samples under SAM and SGD. The regression curve (in green) shows a consistent deviation from the identity line (in red), indicating that SAM memorizes a larger subset of samples in the lower score range which are likely to be atypical subclass samples. (c): Visualization of samples more memorized by SAM for the tiger class following the same setting of (b).

A regression analysis over all samples – shown via the green curve – reveals a consistent deviation from the identity – red – line, skewed towards the top-left quadrant. This indicates a systematic increase in memorization for a large subset of samples under SAM. Crucially, this deviation is not concentrated at the high end of the memorization spectrum. This finding – together with the kernel density plot – supports our hypothesis that SAM does not simply memorize pure noise samples, but rather focuses on non-dominant, atypical subclass samples that are underrepresented in the training distribution. Indeed, if SAM were picking up sample-specific noise, we would expect a sharp concentration of kernel density at the highest end of the spectrum and a deviation of the regression curve in the top-right quadrant.

Figure 1(c) illustrates this phenomenon within the *tiger* class. Samples with higher SAM memorization relative to SGD (top-left region) tend to depict clean samples containing atypical sub-patterns—e.g., close-ups of tiger heads, tigers in water, or multiple tigers in a single image. These are visually distinct yet semantically consistent with the class label. In contrast, samples with high memorization under both SAM and SGD (top-right region) often contain sample-specific noise, – e.g., a tiger with shiny paws on a pitch-black background –, which are less likely to generalize.

4.2 SAM INCREASES INFLUENCE OF HIGH MEMORIZATION SAMPLES

We here analyze how memorization affects the influence of training samples on test predictions, using the influence metric defined in Equation (2). Following the setup of Feldman & Zhang (2020), we first filter training–test sample pairs with influence scores above 0.2 to exclude non-influential cases. We then group training samples by memorization intervals – defined as $l < \text{mem}(\mathcal{A}, \mathcal{D}, i) < u$, with l and u ranging from 0 to 1 in steps of 0.05 — and, for each interval, we select the 20 training samples achieving the highest influence score on test data. This yields a distribution of influence scores of the most influential training samples conditioned on their memorization levels.

Figures 2(a) and 2(b) show the resulting distributions for SGD and SAM, respectively. As in Figure 1(b), we fit a regression curve (green line) to highlight the trends. For SGD, influence scores incur in a steep transition from lowly influent samples to highly influent points at the upper end of the memorization spectrum. This indicates SGD’s reliance on a very narrow set of highly memorized samples. In contrast, SAM exhibits a smoother transition curve, with a larger set of high (and mid-to-high) memorization samples contributing more consistently to test predictions. This supports our earlier finding that SAM emphasizes a set of atypical, non-dominant subclass patterns.

To further validate this, we examine the difference in influence scores between SAM and SGD as a function of their memorization score differences (Figure 2(c)). Training samples which are more memorized by SGD tend to have lower influence under SAM, suggesting that SAM down-weights the influence of its low-memorized samples. Conversely, samples with similar memorization under both

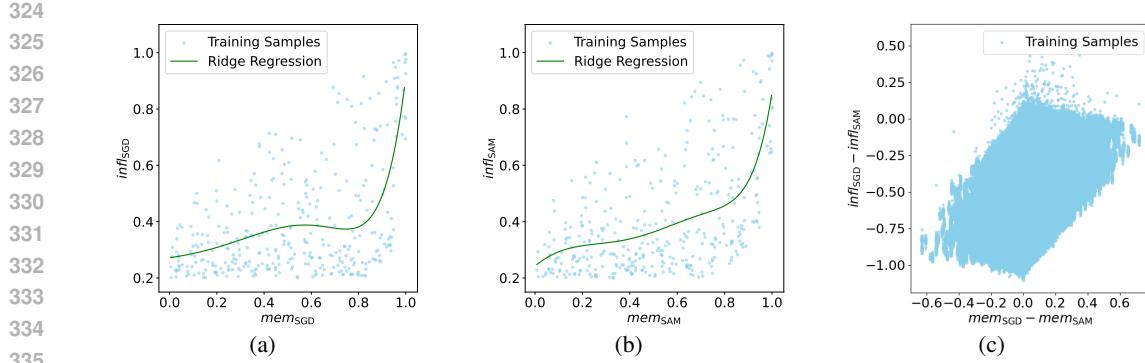
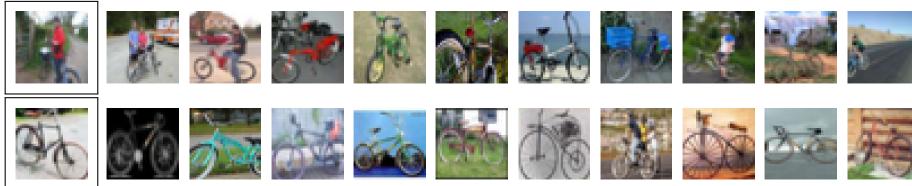


Figure 2: (a) and (b): Distribution of the influence scores of the 20 most influential training samples over each memorization interval for SGD (a) and SAM (b). The regression analysis (green lines) shows that SAM maintains a smoother influence distribution, relying more on mid-to-high (0.6 - 0.85) memorization samples (subclass features) than SGD. (c) Difference in influence scores between SAM and SGD as a function of memorization score differences. SAM downweights influence for low-memorized samples and selectively amplifies the influence of mid-to-high memorization samples.



(a) Bicycle



(b) Tiger

Figure 3: Test images (boxed) from buckets 1 and 5 and their respective top-10 influential training images. For each object the top row is an image from bucket 1 and the bottom row is an image from bucket 5. For bucket 1 images (higher memorization, top row), notice that the images are atypical for their classes, and there is a near duplicate in the training data that was important for generalizing on this test image. For bucket 5 images, on the other hand, the top influential images are reminiscent of the test image at a conceptual level.

optimizers but higher influence under SAM tend to lie in the mid-to-high memorization range. These are precisely the samples containing atypical subpatterns that SAM selectively amplifies, confirming our intuitions.

4.3 SAM’S GENERALIZATION GAIN COMES FROM HIGHER MEMORIZATION OF SUB-PATTERNS

In this section, we dissect the generalization gains of SAM at a finer granularity by constructing a metric that divides the test data points into groups based on the amount of memorization used for predicting them. We then compare the performance on each group between SGD and SAM.

We measure the typicality of a test data point as the entropy of its corresponding training samples’ influence scores. We rely on this measure since, in practice, the prediction of a typical unseen sample

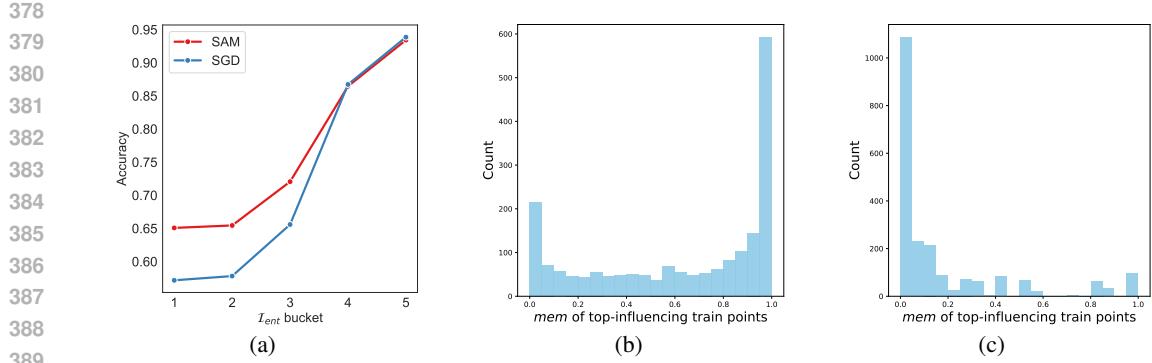


Figure 4: (a): Test accuracy on \mathcal{I}_{ent} groups as evaluated by equation 5. SAM’s performance gains comes from it correctly predicting more atypical data points that need memorization of atypical sub-patterns to be classified correctly. (b) and (c): Distribution of top-1 most influential training point’s memorization scores for \mathcal{I}_{ent} buckets 1 and 5. Testing samples falling in the lower (higher) numbered buckets are influenced by training points with higher (lower) memorization.

would be evenly influenced by numerous training data points within the same class, while atypical counterparts would be heavily influenced by a handful of training samples that are themselves also atypical. To measure the even spread, we normalize the influence scores and leverage entropy. Even influence spread would follow a more uniform distribution resulting in high entropy, while uneven spread incurs low entropy. Formally, for each test data point i , let \mathcal{S}_i be the set of influence scores of all the training points in the same class. Let $S_{i,j}$ be the influence score of j^{th} training point and m be the number of training points in the same class. Then, our entropy metric \mathcal{I}_{ent} is defined as:

$$\mathcal{I}_{ent}[i] = \sum_{j=1}^m -p_{i,j} * \log p_{i,j}, \text{ where, } p_{i,j} = \frac{S_{i,j}}{\sum_j S_{i,j}} \quad (5)$$

We group test data points into 5 buckets in the order of lowest \mathcal{I}_{ent} to highest \mathcal{I}_{ent} . We present some test images and their top-10 influential training images in Figure 3 from bucket 1 and bucket 5. The figure illustrates that images from bucket 1 tend to be atypical images – e.g., bicycle alongside people, white tiger, etc. – for their respective labels while images from bucket 5 tend to be more typical images—e.g., typical bicycle and yellow tiger. For quantitative verification, we plot the distribution of memorization scores of the highest influencing training points from each bucket. We observe that lower numbered buckets are associated with high memorization and vice versa (see Figures 4(b) and 4(c)). The results for other buckets interpolate between those of bucket 1 and 5, and are skipped for brevity.

We compare the generalization gains of SAM against SGD on each of these buckets and show the results in Figure 4(a)⁴. For test data points in bucket 5, SGD achieves a negligible performance gain, while for bucket 1 SAM achieves a significant gain w.r.t. SGD. Thus, the performance gains of SAM can be attributed to it correctly predicting more atypical data points which need more memorization of atypical sub-patterns to be classified correctly. For further validation, we generate a synthetic dataset and illustrate SAM’s capability of learning atypical subclasses better than SGD in Appendix B. We theoretically analyze how higher minority subclass alignment can lead to better generalization in Section C.

Summary of experimental findings These results – together with those in Sections 4.1 and 4.2 – provide strong empirical evidence that *SAM’s increased performance derives from better capturing atypical but informative sub-patterns, allowing predictions to be less dominated by the majority feature*. In other words, SAM is capturing more diverse features. Connecting these findings with the observation in Section 3, we posit that SGD’s reliance on the majority feature necessitates amplifying that feature’s weight to classify diverse inputs. This results in a model with high output sensitivity: it

⁴These results do not consider image transformations (e.g. random crop, rotations), however we have also replicated the experiment with transformations, obtaining a similar trend.

432 tends to yield high-confidence predictions that fluctuate significantly based on alignment with the
 433 majority feature, thereby inflating the output variance. Paradoxically, this high variance on unseen
 434 data acts as a cloak for membership privacy, as it mimics the high confidence usually reserved for
 435 members. In the next section, we theoretically analyze how SAM inherently suppresses variance on
 436 unseen data.

437

438 5 THEORETICAL ANALYSIS

439

440 In this section, we provide a theoretical foundation for the variance–shrinkage effect of sharpness-
 441 aware minimization and its implication for membership inference risk. To obtain a clean and
 442 analyzable characterization of SAM’s geometric bias, we study overparameterized linear models
 443 trained to interpolation. Closest to our analysis in spirit is that of Tan et al. (2022), which explains
 444 how increasing the number of parameters or ridge regression affects membership risk via a variance
 445 gap in Gaussian linear regression. In contrast, we introduce a curvature-aligned geometry modeling
 446 SAM and show that it provably increases membership risk.

447

448 Although our proofs are written in the regression form $X\theta = y$, the connection to classification is
 449 direct. For separable binary classification with losses such as logistic or exponential, gradient descent
 450 is known to implicitly maximize the ℓ_2 margin, and the resulting predictor is geometrically equivalent
 451 to the minimum- ℓ_2 -norm interpolator of a corresponding regression problem (Soudry et al., 2024;
 452 Gunasekar et al., 2017; Muthukumar et al., 2021). Motivated by this equivalence, we can consider
 453 classification as regression to high-magnitude targets $y_i \in \{-M, +M\}$ for a large constant $M \gg 1$
 454 under MSE loss. In this regime, training points are interpolated to $\pm M$ and therefore have fixed high
 455 confidence, whereas test points yield outputs that fluctuate around the decision boundary (0).

456

457 **Model and notation.** We work in finite dimension $d \gg n$. Let the population feature covariance
 458 be $\Sigma \in \mathbb{R}^{d \times d}$ with $\Sigma \succ 0$. Draw i.i.d. samples

$$459 \quad x_i \sim \mathcal{N}(0, \Sigma), \quad y_i = \theta^{*\top} x_i + \xi_i, \quad \xi_i \sim \mathcal{N}(0, \sigma_y^2), \text{ independent of } x_i, \quad i = 1, \dots, n.$$

460 Let $X \in \mathbb{R}^{n \times d}$ have rows x_i^\top , $y = (y_1, \dots, y_n)^\top$, and assume $\text{rank}(X) = n$. Define the orthogonal
 461 projector onto the data span and the covariance matrix

$$462 \quad P := X^\top (X X^\top)^{-1} X, \quad P^2 = P = P^\top, \quad \widehat{\Sigma} := \frac{1}{n} X^\top X, \quad \Sigma := \mathbb{E}[x x^\top].$$

463 The model is defined as $f_G(x) := \widehat{\theta}_G^\top x$. In the squared-loss linear model, the population Hessian H
 464 equals Σ and the empirical Hessian \widehat{H} equals $\widehat{\Sigma}$.

465 Now we introduce the geometries we employ to model SGD and SAM.

466

467 **Min- G interpolation.** For any symmetric positive definite matrix $G \succ 0$, consider the minimum-
 468 G -norm interpolant:

$$469 \quad \widehat{\theta}_G := \arg \min_{\theta \in \mathbb{R}^d} \frac{1}{2} \theta^\top G \theta \quad \text{s.t.} \quad X\theta = y. \quad (6)$$

470 We compare the standard Euclidean case,

$$471 \quad G_0 := I_d$$

472 with the *sharpness-aware* geometry,

$$473 \quad G_\eta := I + \eta H = I + \eta \Sigma, \quad \eta > 0,$$

474

475 Classical implicit-bias results identify standard SGD with the Euclidean case under step size and
 476 weight initialization assumptions (Zhang et al., 2017). Sharpness-aware geometry, on the other hand,
 477 reflects SAM’s penalty on local sharpness, which can be shown by expanding the minimization
 478 objective. We prove that, relative to the Euclidean interpolation, the geometry G_η strictly reduces
 479 the variance of non-member outputs, thereby enlarging the separation between members and non-
 480 members and increasing the advantage of membership inference attacks.

481

482 We first establish that SAM strictly reduces the variance of the model’s output on unseen data. By
 483 Lemma 3, for a non-member sample X_{out} , and the output follows $f_G(X_{\text{out}}) \sim \mathcal{N}(0, \sigma_G^2)$.

486
487 **Theorem 1** (Variance strictly decreases for SAM geometry). *Under Assumption 1, there exists $\eta_0 > 0$
488 such that for all $\eta \in (0, \eta_0]$,*

$$489 \quad \sigma_{G_\eta}^2 < \sigma_{G_0}^2 \quad \text{with probability } 1 - o(1).$$

490
491 *Remark 1.* Theorem 1 formalizes a simple geometric picture. The minimum- G interpolant balances
492 fitting the training span against penalizing components in directions where G is large. Moving from
493 G_0 to G_η increases the penalty precisely along high-curvature directions of the Hessian. Lemma 2
494 shows that, under the overlap condition in Assumption 1, this reweighting strictly suppresses the
495 $(I - P)\Sigma\hat{\theta}_{G_0}$ component of the interpolant. Since non-member predictions depend on $\hat{\theta}_G$ only
496 through the quadratic form $\hat{\theta}_G^\top \Sigma \hat{\theta}_G$, this suppression translates directly into a strict decrease of the
497 output variance on unseen data.

498 Next, we consider confidence-threshold attack and likelihood ratio attack.

500 **Confidence-threshold attack.** The attacker uses the confidence score $\text{Conf}_G(x) := |f_G(x)|$ and
501 predicts “member” iff $\text{Conf}_G(x) \geq \tau$. Let $(X_{\text{in}}, Y_{\text{in}})$ be a random training pair and $X_{\text{out}} \sim \mathcal{N}(0, \Sigma)$
502 an independent non-member. Define attack advantage

$$503 \quad \text{Adv}_G^{\text{conf}} := \sup_{\tau \geq 0} (\text{TPR}_G(\tau) - \text{FPR}_G(\tau)).$$

504 **Theorem 2** (SAM strictly increases confidence-based MIA advantage). *Under Assumption 1, for all
505 sufficiently small $\eta > 0$,*

$$506 \quad \text{Adv}_{G_\eta}^{\text{conf}} > \text{Adv}_{G_0}^{\text{conf}} \quad \text{with probability } 1 - o(1).$$

507 *Remark 2.* For the confidence-threshold attack, interpolation implies that member confidences are
508 geometry-invariant: the training logits are fixed (up to label noise) for any choice of G . Thus,
509 changing the geometry from G_0 to G_η leaves $\text{TPR}_G(\tau)$ unchanged for every threshold τ , while
510 Theorem 1 strictly reduces the non-member variance and hence lowers $\text{FPR}_G(\tau)$ at every $\tau > 0$.
511 In other words, SAM pushes non-member scores closer to the decision boundary, sharpening the
512 separation between the two confidence distributions.

513 **Likelihood Ratio (LR) attack.** The oracle LRA predicts “member” iff $\Lambda_G(f_G(x)) \geq t$. Define

$$514 \quad \Lambda_G(s) := \log \frac{p_{\text{in}}(s)}{p_{\text{out}}(s; G)}, \quad \text{Adv}_G^{\text{LR}} := \sup_{t \in \mathbb{R}} (\Pr(\Lambda_G(f_G(X_{\text{in}})) \geq t) - \Pr(\Lambda_G(f_G(X_{\text{out}})) \geq t)).$$

515 **Theorem 3** (SAM strictly increases LR-attack advantage). *Under Assumption 1, for all sufficiently
516 small $\eta > 0$,*

$$517 \quad \text{Adv}_{G_\eta}^{\text{LR}} > \text{Adv}_{G_0}^{\text{LR}} \quad \text{with probability } 1 - o(1),$$

518 where $G_\eta = I + \eta\Sigma$.

519 *Remark 3.* The argument is similar to Theorem 2. Theorem 3 analyzes an oracle LR attacker that
520 uses a single global IN/OUT distribution. Corollary 5 strengthens this to sample-adaptive IN/OUT
521 distributions, matching the per-query calibration used by LiRA/RMIA. Theorem 2 and Theorem 3
522 prove an improvement in the $\sup_{\tau} (\text{TPR} - \text{FPR})$ advantage, but the proof in fact shows a stronger
523 statement: for all sufficiently small $\eta > 0$, *the entire ROC curve of the confidence-based attack under
524 G_η strictly dominates that under G_0 .* This is empirically verified in Figure 7, where SAM’s ROC
525 curve is above that of SGD’s for nearly the *entire* range across most settings.

526 6 CONCLUSION AND FUTURE WORK

527 This work investigates the mechanism behind SAM’s dual effect: superior generalization coupled
528 with heightened privacy leakage. As algorithms seeking flatter minima are widely employed, our
529 work serves as an important cautionary tale users should be aware of. Furthermore, our findings
530 suggest that exploring sharper minima could serve as a novel defense against Membership Inference
531 Attacks. Finally, future research is encouraged to further reconcile the apparent dichotomy between
532 whether SAM’s implicit bias is geared more towards diversity or simplicity of features.

540 REFERENCES
541

542 Zeyuan Allen-Zhu, Yuanzhi Li, and Yingyu Liang. Learning and generalization in overparameterized
543 neural networks, going beyond two layers. *Advances in neural information processing systems*, 32,
544 2019.

545 Maksym Andriushchenko, Dara Bahri, Hossein Mobahi, and Nicolas Flammarion. Sharpness-
546 aware minimization leads to low-rank features. *arXiv preprint arXiv:2305.16292*, 2023. URL
547 <https://arxiv.org/abs/2305.16292>. NeurIPS 2023 poster version.

548 Christina Baek, Zico Kolter, and Aditi Raghunathan. Why is sam robust to label noise?, 2024. URL
549 <https://arxiv.org/abs/2405.03676>.

550 Nicholas Carlini, Florian Tramer, Eric Wallace, Matthew Jagielski, Ariel Herbert-Voss, Katherine
551 Lee, Adam Roberts, Tom B Brown, Dawn Song, Ulfar Erlingsson, et al. Extracting training data
552 from large language models. In *USENIX Security Symposium*, volume 6, 2021.

553 Nicholas Carlini, Steve Chien, Milad Nasr, Shuang Song, Andreas Terzis, and Florian Tramer.
554 Membership inference attacks from first principles, 2022. URL <https://arxiv.org/abs/2112.03570>.

555 Junbum Cha, Sanghyuk Chun, Kyungjae Lee, Han-Cheol Cho, Seunghyun Park, Yunsung Lee,
556 and Sungrae Park. Swad: Domain generalization by seeking flat minima. *Advances in Neural
557 Information Processing Systems*, 34:22405–22418, 2021.

558 Wei-Kai Chang and Rajiv Khanna. A unified stability analysis of sam vs sgd: Role of data coherence
559 and emergence of simplicity bias, 2025. URL <https://arxiv.org/abs/2511.17378>.

560 Dingfan Chen, Ning Yu, and Mario Fritz. Relaxloss: Defending membership inference attacks
561 without losing utility, 2022a. URL <https://arxiv.org/abs/2207.05801>.

562 Dingfan Chen, Ning Yu, and Mario Fritz. Relaxloss: Defending membership inference attacks
563 without losing utility, 2022b. URL <https://arxiv.org/abs/2207.05801>.

564 Zixiang Chen, Junkai Zhang, Yiwen Kou, Xiangning Chen, Cho-Jui Hsieh, and Quanquan Gu. Why
565 does sharpness-aware minimization generalize better than sgd?, 2023. URL <https://arxiv.org/abs/2310.07269>.

566 Ganesh Del Grosso, Georg Pichler, Catuscia Palamidessi, and Pablo Piantanida. Bounding information
567 leakage in machine learning. *Neurocomputing*, 534:1–17, May 2023. ISSN 0925-2312. doi:
568 10.1016/j.neucom.2023.02.058. URL <http://dx.doi.org/10.1016/j.neucom.2023.02.058>.

569 Jiawei Du, Zhou Daquan, Jiashi Feng, Vincent Tan, and Joey Tianyi Zhou. Sharpness-aware training
570 for free. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho (eds.), *Advances
571 in Neural Information Processing Systems*, 2022.

572 Xingli Fang and Jung-Eun Kim. Center-based relaxed learning against membership inference attacks,
573 2024. URL <https://arxiv.org/abs/2404.17674>.

574 Vitaly Feldman. Does learning require memorization? a short tale about a long tail. In *Proceedings
575 of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*, pp. 954–959, 2020.

576 Vitaly Feldman and Chiyuan Zhang. What neural networks memorize and why: Discovering the long
577 tail via influence estimation. *Advances in Neural Information Processing Systems*, 33:2881–2891,
578 2020.

579 Pierre Foret, Ariel Kleiner, Hossein Mobahi, and Behnam Neyshabur. Sharpness-aware minimization
580 for efficiently improving generalization. *arXiv preprint arXiv:2010.01412*, 2020.

581 Suriya Gunasekar, Blake Woodworth, Srinadh Bhojanapalli, Behnam Neyshabur, and Nathan Srebro.
582 Implicit regularization in matrix factorization, 2017. URL <https://arxiv.org/abs/1705.09280>.

594 Pavel Izmailov, Dmitrii Podoprikhin, Timur Garipov, Dmitry Vetrov, and Andrew Gordon Wilson. Averaging weights leads to wider optima and better generalization. *arXiv preprint arXiv:1803.05407*,
 595 2018.

596

597 Jinyuan Jia, Ahmed Salem, Michael Backes, Yang Zhang, and Neil Zhenqiang Gong. Memguard: De-
 598 fending against black-box membership inference attacks via adversarial examples. In *Proceedings*
 599 *of the 2019 ACM SIGSAC conference on computer and communications security*, pp. 259–274,
 600 2019.

601

602 Yigitcan Kaya and Tudor Dumitras. When does data augmentation help with membership inference
 603 attacks? In Marina Meila and Tong Zhang (eds.), *Proceedings of the 38th International Conference*
 604 *on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 5345–5355.
 605 PMLR, 18–24 Jul 2021. URL <https://proceedings.mlr.press/v139/kaya21a.html>.

606

607 Minyoung Kim, Da Li, Shell X Hu, and Timothy Hospedales. Fisher SAM: Information geometry
 608 and sharpness aware minimisation. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba
 609 Szepesvari, Gang Niu, and Sivan Sabato (eds.), *Proceedings of the 39th International Conference*
 610 *on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pp. 11148–
 611 11161. PMLR, 17–23 Jul 2022.

612

613 Yiwen Kou, Zixiang Chen, Yuanzhou Chen, and Quanquan Gu. Benign overfitting for two-layer relu
 614 convolutional neural networks, 2023. URL <https://arxiv.org/abs/2303.04145>.

615

616 Jungmin Kwon, Jeongseop Kim, Hyunseo Park, and In Kwon Choi. Asam: Adaptive sharpness-aware
 617 minimization for scale-invariant learning of deep neural networks. In Marina Meila and Tong
 618 Zhang (eds.), *Proceedings of the 38th International Conference on Machine Learning*, volume 139
 619 of *Proceedings of Machine Learning Research*, pp. 5905–5914. PMLR, 18–24 Jul 2021.

620

621 Zhenlong Liu, Lei Feng, Huiping Zhuang, Xiaofeng Cao, and Hongxin Wei. Mitigating privacy risk
 622 in membership inference by convex-concave loss, 2024. URL <https://arxiv.org/abs/2402.05453>.

623

624 Yunhui Long, Vincent Bindschaedler, Lei Wang, Diyue Bu, Xiaofeng Wang, Haixu Tang, Carl A
 625 Gunter, and Kai Chen. Understanding membership inferences on well-generalized learning models.
 626 *arXiv preprint arXiv:1802.04889*, 2018.

627

628 Vidya Muthukumar, Adhyayan Narang, Vignesh Subramanian, Mikhail Belkin, Daniel Hsu, and Anant
 629 Sahai. Classification vs regression in overparameterized regimes: Does the loss function matter?,
 2021. URL <https://arxiv.org/abs/2005.08054>.

630

631 Milad Nasr, Reza Shokri, and Amir Houmansadr. Machine learning with membership privacy using
 632 adversarial regularization. In *Proceedings of the 2018 ACM SIGSAC conference on computer and*
 633 *communications security*, pp. 634–646, 2018.

634

635 Milad Nasr, Reza Shokri, and Amir Houmansadr. Comprehensive privacy analysis of deep learning:
 636 Passive and active white-box inference attacks against centralized and federated learning. In *2019*
IEEE symposium on security and privacy (SP), pp. 739–753. IEEE, 2019.

637

638 Matthew D Norton and Johannes O Royset. Diametrical risk minimization: Theory and computations.
Machine Learning, pp. 1–19, 2021.

639

640 Ahmed Salem, Yang Zhang, Mathias Humbert, Pascal Berrang, Mario Fritz, and Michael Backes.
 641 MI-leaks: Model and data independent membership inference attacks and defenses on machine
 642 learning models. *arXiv preprint arXiv:1806.01246*, 2018.

643

644 Reza Shokri, Marco Stronati, Congzheng Song, and Vitaly Shmatikov. Membership inference attacks
 645 against machine learning models. In *2017 IEEE symposium on security and privacy (SP)*, pp. 3–18.
 646 IEEE, 2017.

647

Liwei Song and Prateek Mittal. Systematic evaluation of privacy risks of machine learning models.
 In *USENIX Security Symposium*, volume 1, pp. 4, 2021.

648 Daniel Soudry, Elad Hoffer, Mor Shpigel Nacson, Suriya Gunasekar, and Nathan Srebro. The implicit
 649 bias of gradient descent on separable data, 2024. URL <https://arxiv.org/abs/1710.10345>.

650

651 Jacob Mitchell Springer, Vaishnavh Nagarajan, and Aditi Raghunathan. Sharpness-aware minimiza-
 652 tion enhances feature quality via balanced learning, 2024. URL <https://arxiv.org/abs/2405.20439>.

653

654

655 Jasper Tan, Blake Mason, Hamid Javadi, and Richard G. Baraniuk. Parameters or privacy: A
 656 provable tradeoff between overparameterization and membership inference, 2022. URL <https://arxiv.org/abs/2202.01243>.

657

658

659 Xinyu Tang, Saeed Mahloujifar, Liwei Song, Virat Shejwalkar, Milad Nasr, Amir Houmansadr, and
 660 Prateek Mittal. Mitigating membership inference attacks by {Self-Distillation} through a novel
 661 ensemble architecture. In *31st USENIX Security Symposium (USENIX Security 22)*, pp. 1433–1450,
 662 2022.

663

664 Dongxian Wu, Shu-Tao Xia, and Yisen Wang. Adversarial weight perturbation helps robust general-
 665 ization. *Advances in Neural Information Processing Systems*, 33:2958–2969, 2020.

666

667 Samuel Yeom, Irene Giacomelli, Matt Fredrikson, and Somesh Jha. Privacy Risk in Machine
 668 Learning: Analyzing the Connection to Overfitting . In *2018 IEEE 31st Computer Security*
 669 *Foundations Symposium (CSF)*, pp. 268–282, Los Alamitos, CA, USA, July 2018. IEEE Computer
 670 Society. doi: 10.1109/CSF.2018.00027. URL <https://doi.ieee.org/10.1109/CSF.2018.00027>.

671

672 Samuel Yeom, Irene Giacomelli, Alan Menaged, Matt Fredrikson, and Somesh Jha. Overfitting,
 673 robustness, and malicious algorithms: A study of potential causes of privacy risk in machine
 674 learning. *Journal of Computer Security*, 28(1):35–70, 2020.

675

676 Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. *arXiv preprint arXiv:1605.07146*,
 677 2016.

678

679 Sajjad Zarifzadeh, Philippe Liu, and Reza Shokri. Low-cost high-power membership inference
 680 attacks, 2024. URL <https://arxiv.org/abs/2312.03262>.

681

682 Chiyuan Zhang, Qianli Liao, Alexander Rakhlin, Brando Miranda, Noah Golowich, and Tomaso
 683 Poggio. Musings on deep learning: Properties of sgd. Technical Report CBMM Memo No. 067,
 684 Center for Brains, Minds and Machines (CBMM), MIT, 2017. URL <https://cbmm.mit.edu/sites/default/files/publications/CBMM-Memo-067-v4.pdf>.

685

686 Juntang Zhuang, Boqing Gong, Liangzhe Yuan, Yin Cui, Hartwig Adam, Nicha Dvornek, Sekhar
 687 Tatikonda, James Duncan, and Ting Liu. Surrogate gap minimization improves sharpness-aware
 688 training, 2022. URL <https://arxiv.org/abs/2203.08065>.

689

690

691

692

693

694

695

696

697

698

699

700

701

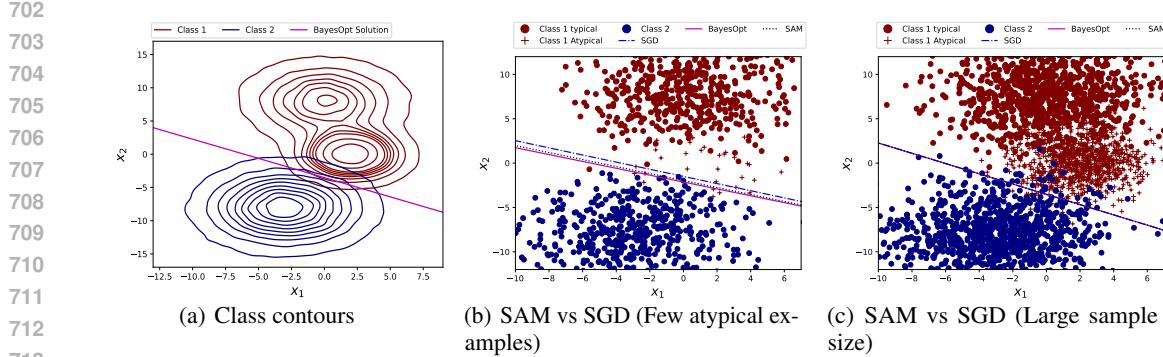


Figure 5: A synthetic construction illustrating the generalization ability of SAM over SGD for atypical examples. Fig (a) shows class density contours of a two-class, 2-dimensional classification problem, along with the Bayes Optimal solution. The red class has two ‘clusters’, one representing typical examples and one representing atypical examples. Fig (b) shows an instance of data sampled from densities shown in (a); the larger cluster of red dots represent typical examples in the red class, and the red ‘+’ points represent a lot fewer atypical examples. SAM generalizes better than SGD in this case. Fig (c) shows that if there are enough samples generated from both typical and atypical clusters, SAM and SGD coincide with the Bayes Optimal classifier.

A LLM USAGE

LLMs were used to assist with grammar correction and sentence-level proofreading throughout the manuscript.

B SYNTHETIC DATASET

In this section, we provide a simple synthetic construction that illustrates how SAM can achieve better generalization performance vs vanilla SGD. The example is illustrated in Figure 5. The data is generated from two-dimensional densities illustrated in Figure 5(a). The densities are supported in two dimensions labelled as x_1 and x_2 . There are two classes - the red class and the blue class. Figure 5(a) also shows the Bayes Optimal classifier. The red class has two ‘clusters’, one representing the typical examples (e.g. yellow tigers), and the other representing the atypical examples (e.g. white tigers). The data is sampled in such a way that we have several samples from the typical cluster, while there are only a few samples from the atypical cluster in the red class. This is shown in Figure 5(b). Figure 5(b) further shows that seeking flatter minima using the SAM optimizer learns a classifier that is closer to the Bayes Optimal classifier than the classifier learnt using vanilla SGD, and thus the former generalizes better. This difference in performance vanishes in Figure 5(c) when we have a large sample size for the atypical examples as well.

This synthetic construction shows that one possible reason that SAM can perform better is if it tends to memorize atypical samples more than vanilla SGD. In other words, the gain in generalization could potentially come from those atypical data subgroups. In the next subsection, we empirically verify this conjecture for the CIFAR-100 dataset.

C BETTER MINORITY SUBCLASS ALIGNMENT LEADS TO HIGHER GENERALIZATION

In this section, we theoretically analyze a setting involving majority subclass samples, minority subclass samples, and pure noise samples to illustrate how an overfitting model generalizes better. Driven by the motivations in the previous sections, we discern a model by the amount of minority subclass feature it captures. Theorem 4 shows that this property leads to higher generalization. The proofs are in Section D.3.

756 **Definition 1** (Data Model). *Training dataset $D = \{(\mathbf{x}_i, y_i)\}_{i=1}^n$ is sampled i.i.d. from the data*
 757 *distribution \mathcal{D} . We define $\mathbf{x} = [\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3]$, $y \in \{\pm 1\}$, $\mathbf{x}_1 \in \mathbb{R}^{d_1}$, $\mathbf{x}_2 \in \mathbb{R}^{d_2}$, $\mathbf{x}_3 \in \mathbb{R}^{d_3}$. We*
 758 *consider a overparameterized regime where $d_1, d_2, d_3 \gg n$. We consider a linear predictor for*
 759 *classification, $f_{\mathbf{w}}(\mathbf{x}) = \text{sign}\{\langle \mathbf{w}, \mathbf{x} \rangle\}$ with $\mathbf{w} = [\mathbf{w}_1, \mathbf{w}_2, \mathbf{w}_3]$. We assume perfect interpolation*
 760 *satisfying finite margin $\forall i 0 \leq m_0 \leq m_i \leq M < \infty$, where margin $m_i = y_i \langle \mathbf{w}, \mathbf{x}_i \rangle$. Let $\mathcal{M}, \mathcal{S}, \mathcal{N}$*
 761 *denote the majority, minority, and noise subsets, respectively. $D = \mathcal{M} \cup \mathcal{S} \cup \mathcal{N}$. We define*
 762 *$\sigma(z) = 1/(1 + \exp^{-z})$.*

763 \mathcal{D} is a mixture distribution

$$\mathcal{D} = p_{\mathcal{M}} \mathcal{D}_{\mathcal{M}} + p_{\mathcal{S}} \mathcal{D}_{\mathcal{S}} + p_{\mathcal{N}} \mathcal{D}_{\mathcal{N}},$$

764 with mixture weights $p_{\mathcal{M}} + p_{\mathcal{S}} + p_{\mathcal{N}} = 1$, and $p_{\mathcal{M}} \gg p_{\mathcal{S}}, p_{\mathcal{N}}$. Here, $\mathcal{D}_{\mathcal{M}}$ generates majority
 765 samples with $\mathbf{x}_1 = y\boldsymbol{\mu}_1$, $\mathbf{x}_2 \sim \mathcal{N}(0, I_{d_2})$, $\mathbf{x}_3 \sim \mathcal{N}(0, I_{d_3})$ where $\boldsymbol{\mu}_1$ is a fixed vector; $\mathcal{D}_{\mathcal{S}}$ generates
 766 minority samples with $\mathbf{x}_1 = \boldsymbol{\nu}$, $\mathbf{x}_2 = y\boldsymbol{\mu}_2$, $\mathbf{x}_3 \sim \mathcal{N}(0, I_{d_3})$ where $\boldsymbol{\mu}_2$ is a fixed vector and $\boldsymbol{\nu}$ is a
 767 random vector anti-aligned with \mathbf{w}_1 ($\langle \mathbf{w}_1, \boldsymbol{\nu} \rangle < 0$); and $\mathcal{D}_{\mathcal{N}}$ generates pure noise samples with
 768 $\mathbf{x} \sim \mathcal{N}(0, I_{d_1+d_2+d_3})$. Let $(n_{\mathcal{M}}, n_{\mathcal{S}}, n_{\mathcal{N}})$ denote the counts of majority, minority, and noise samples
 769 in S .

770 Our model easily translates to a logistic regression model if we set $f_{\mathbf{w}}(\mathbf{x}) = \sigma(\langle \mathbf{w}, \mathbf{x} \rangle)$ and change
 771 the label to $y \in \{0, 1\}$. The setup in which a part of the input contains the true signal has been
 772 commonly used in previous works (Chen et al., 2023; Kou et al., 2023) but we generalize their setup
 773 by including and analyzing subclasses in the design. Concretely, one can think of a minority sample
 774 as belonging to a long-tail subgroup of the class that requires a different feature to be recognized.
 775 We formalize an anti-alignment condition: the minority features are arranged such that any model
 776 that heavily prioritizes the majority feature u will perform poorly on the minority. Intuitively, fitting
 777 the minority subgroup requires the model to memorize an alternative pattern that is independent
 778 of (or even interfering with) the main decision boundary for the majority. We assume the minority
 779 subgroup is very small, so that by default a standard empirical risk minimizer might deem it negligible.
 780 This resonates with the long-tail phenomena observed in practice – a handful of unusual examples
 781 exist that a model could easily ignore without significant impact on overall training loss. However,
 782 those examples are crucial for tail generalization: they represent rare yet valid concepts that an ideal
 783 model should learn. Our assumptions reflect prior findings that real datasets contain such long-tailed
 784 subpopulations which must be memorized to achieve the best possible accuracy Feldman (2020).
 785 Formally,

786 **Condition 1.** *For each minority point $i \in \mathcal{S}$, define*

$$B_i := -y_i \langle \mathbf{w}_1, \boldsymbol{\nu}_i \rangle > 0, \quad A := \langle \mathbf{w}_2, \boldsymbol{\mu}_2 \rangle.$$

787 Let B be a random variable with CDF $F_B(A) = \Pr(B < A)$ such that $B \stackrel{d}{=} B_i$ (i.e., F_B is the
 788 law/distribution of the B_i 's when i is drawn uniformly from \mathcal{S}). We assume

$$A < B_{\max}, \quad B_{\max} := \sup\{b \in \mathbb{R} \mid F_B(b) < 1\},$$

789 This condition assures that the majority feature still dominates globally. Furthermore, we formulate
 790 $\boldsymbol{\nu}$ as a random variable to effectively capture multiple atypical subclasses. For example, there can be
 791 a wide range of tigers that are purely white or yellowish-white. By modeling $\boldsymbol{\nu}$ as a random variable,
 792 it provides variation on the strength of anti-alignment with the majority feature.

793 Driven by the empirical motivations in Section 4, we define an ordering of the models as how much
 794 minority subclass alignment (MSA) they achieved. Formally,

795 **Definition 2** (Minority Subclass Alignment Order). *Given two interpolating solutions $\mathbf{w}^{(A)}, \mathbf{w}^{(B)}$
 796 trained on the same S , define*

$$A^{(A)} := \langle \mathbf{w}_2^{(A)}, y\boldsymbol{\mu}_2 \rangle, \quad A^{(B)} := \langle \mathbf{w}_2^{(B)}, y\boldsymbol{\mu}_2 \rangle.$$

797 We write

$$\mathbf{w}^{(A)} \succsim^{\text{MSA}} \mathbf{w}^{(B)}$$

800 and say that $\mathbf{w}^{(A)}$ has higher minority subclass alignment than $\mathbf{w}^{(B)}$ if

$$A^{(A)} \geq A^{(B)} \quad \text{and} \quad A^{(A)}, A^{(B)} < B_{\max}.$$

Now, we formally define attack setup and attacker's advantage for MIA. We focus on the confidence threshold based attack as it is empirically one of the most effective attacks and naturally connected to entropy and modified entropy.

Definition 3 (Confidence-threshold attacker). *For $t > 0$ (probabilistic confidence $\tau = \sigma(t)$), define $\mathcal{A}_t(x, y) = \mathbb{1}\{y\langle \mathbf{w}, \mathbf{x} \rangle \geq t\}$ and the attacker's advantage*

$$\text{Adv}^M(w, t) := \Pr(\mathcal{A}_t=1 \mid (\mathbf{x}, y) \in D) - \Pr(\mathcal{A}_t=1 \mid (\mathbf{x}, y) \sim \mathcal{D}).$$

where the probability is taken over S .

Note that the attacker's advantage is the difference in the True Positive Rate (TPR) and the False Positive Rate (FPR). Now we state our main results.

Theorem 4 (Higher MSA \implies Better Generalization). *Let $\mathbf{w}^{(A)}, \mathbf{w}^{(B)}$ be two interpolating solutions trained on the same D . Under Definition 1 and regulatory conditions, if $\mathbf{w}^{(A)} \succ_{\text{MSA}} \mathbf{w}^{(B)}$, then*

$$\Pr(y\langle \mathbf{w}^{(A)}, \mathbf{x} \rangle > 0 \mid (\mathbf{x}, y) \sim \mathcal{D}) \geq \Pr(y\langle \mathbf{w}^{(B)}, \mathbf{x} \rangle > 0 \mid (\mathbf{x}, y) \sim \mathcal{D}),$$

with strict inequality if $F_B((A^{(B)}, A^{(A)}]) > 0$.

Theorem 4 shows that the interpolating model that aligns with the minority subclass feature more generalizes better.

D PROOFS

Assumption 1. We make the following assumptions:

(i) (Kernel overlap) Let $P := X^\top (X X^\top)^{-1} X$ be the orthogonal projector onto $\text{span}(X^\top)$. The population-sharpness direction has nontrivial mass outside the data span:

$$\|(I - P) \Sigma \hat{\theta}_{G_0}\| \geq c > 0$$

with probability $1 - o(1)$ for some constant c .

(ii) (Bounded interpolator) The Euclidean interpolator satisfies $\|\hat{\theta}_{G_0}\| = O_{\mathbb{P}}(1)$.

Remark 4 (Justification of assumptions). Assumption 1(i) is generic in the overparameterized regime ($d \gg n$): the vector $\Sigma \hat{\theta}_{G_0}$ is a population signal direction, while $\text{span}(X^\top)$ is a random n -dimensional subspace. Under mild spectral regularity of Σ , their alignment is not perfect with high probability, yielding a non-negligible $(I - P)$ component.

Assumption 1(ii) is standard for benign overfitting and ensures finite test variance (e.g., Bartlett et al., 2020). If $\|\hat{\theta}_{G_0}\|$ diverged, the Euclidean interpolator would have exploding prediction variance and thus be unstable on unseen data.

Lemma 1. *The unique solution of equation 6 is*

$$\hat{\theta}_G = G^{-1} X^\top \left(X G^{-1} X^\top \right)^{-1} y, \quad \text{and} \quad X \hat{\theta}_G = y.$$

Proof. Form the Lagrangian $L(\theta, \lambda) = \frac{1}{2} \theta^\top G \theta + \lambda^\top (X \theta - y)$. The KKT conditions are $G\theta + X^\top \lambda = 0$ and $X\theta = y$. Eliminate $\theta = -G^{-1} X^\top \lambda$ to get $-XG^{-1} X^\top \lambda = y$, hence $\lambda = -(XG^{-1} X^\top)^{-1} y$. Substituting back yields the stated $\hat{\theta}_G$. Since $G \succ 0$, the objective is strictly convex and the solution is unique. \square

Lemma 2. *Let $\hat{\theta}(\eta) := \hat{\theta}_{G_\eta}$ with $G_\eta = I + \eta \Sigma$. Then*

$$\hat{\theta}'(0) := \frac{d}{d\eta} \hat{\theta}(\eta) \Big|_{\eta=0} = -(I - P) \Sigma \hat{\theta}_{G_0}.$$

864 *Proof.* By Lemma 1,

$$\hat{\theta}(\eta) = G_{\eta}^{-1} X^{\top} (X G_{\eta}^{-1} X^{\top})^{-1} y.$$

865 Differentiate at $\eta = 0$.

866 (i) $\frac{d}{d\eta} G_{\eta}^{-1} = -G_{\eta}^{-1} \Sigma G_{\eta}^{-1}$, so $\frac{d}{d\eta} G_{\eta}^{-1} \big|_{\eta=0} = -\Sigma$.

867 (ii) Let $A(\eta) := X G_{\eta}^{-1} X^{\top}$. Then $A'(0) = -X \Sigma X^{\top}$ and

$$\frac{d}{d\eta} A(\eta)^{-1} \big|_{\eta=0} = A(0)^{-1} (X \Sigma X^{\top}) A(0)^{-1}.$$

873 Combining (i)–(ii),

$$\begin{aligned} \hat{\theta}'(0) &= -\Sigma X^{\top} (X X^{\top})^{-1} y + X^{\top} (X X^{\top})^{-1} (X \Sigma X^{\top}) (X X^{\top})^{-1} y \\ &= -(I - P) \Sigma \hat{\theta}_{G_0}. \end{aligned}$$

878 \square

879 **Lemma 3** (Distribution of model output on non-member data). *Let $X_{\text{out}} \sim \mathcal{N}(0, \Sigma)$ independently*
 880 *of the training set. Conditioned on $\hat{\theta}_G$, the prediction $f_G(X_{\text{out}}) = \hat{\theta}_G^{\top} X_{\text{out}}$ satisfies*

$$f_G(X_{\text{out}}) \sim \mathcal{N}(0, \sigma_G^2), \quad \sigma_G^2 := \hat{\theta}_G^{\top} \Sigma \hat{\theta}_G.$$

884 *Proof.* Condition on $\hat{\theta}_G$. Since X_{out} is Gaussian and f_G is linear in X_{out} , the claim follows with
 885 variance $\hat{\theta}_G^{\top} \Sigma \hat{\theta}_G$. \square

887 **Lemma 4** (First derivative of non-member variance at $\eta = 0$). *With $G_{\eta} = I + \eta \Sigma$, let $\sigma^2(\eta) :=$*
 888 *$\hat{\theta}(\eta)^{\top} \Sigma \hat{\theta}(\eta)$. Then*

$$\sigma^2(0) = 2 \hat{\theta}_{G_0}^{\top} \Sigma \hat{\theta}'(0) = -2 \|(I - P) \Sigma \hat{\theta}_{G_0}\|^2.$$

891 *Proof.* Differentiate: $\sigma^2(0) = 2 \hat{\theta}_{G_0}^{\top} \Sigma \hat{\theta}'(0)$ since Σ is symmetric. Insert Lemma 2:

$$\sigma^2(0) = -2 \hat{\theta}_{G_0}^{\top} \Sigma (I - P) \Sigma \hat{\theta}_{G_0} = -2 \|(I - P) \Sigma \hat{\theta}_{G_0}\|^2,$$

894 using symmetry and idempotence of $(I - P)$. \square

896 Proof of Theorem 1

898 *Proof.* By Lemma 4,

$$\sigma^2(0) = -2 \|(I - P) \Sigma \hat{\theta}_{G_0}\|^2.$$

900 Assumption 1(i) yields $\|(I - P) \Sigma \hat{\theta}_{G_0}\| \geq c$, hence $\sigma^2(0) \leq -2c^2 < 0$ with probability $1 - o(1)$.

902 Since $\hat{\theta}(\eta)$ is smooth in η for $G_{\eta} \succ 0$, $\sigma^2(\eta)$ is differentiable at 0. Therefore, on the same high-
 903 probability event, there exists $\eta_0 > 0$ such that for all $\eta \in (0, \eta_0]$,

$$\sigma^2(\eta) = \sigma^2(0) + \eta \sigma^2(0) + o(\eta) < \sigma^2(0).$$

906 This implies $\sigma_{G_{\eta}}^2 < \sigma_{G_0}^2$ for all sufficiently small $\eta > 0$ with probability $1 - o(1)$. \square

908 D.1 CONFIDENCE-THRESHOLD ATTACK

910 Let \mathcal{I}_{in} be the index set of member training points, and let $\{(x_i, y_i)\}_{i \in \mathcal{I}_{\text{in}}}$ denote the training samples.
 911 For a geometry $G \succ 0$, let $\hat{\theta}_G \in \mathbb{R}^d$ be the (interpolating) solution of equation 6. Define the signed
 912 score and confidence by

$$f_G(x) := \hat{\theta}_G^{\top} x, \quad \text{Conf}_G(x) := |f_G(x)| = |\hat{\theta}_G^{\top} x|.$$

913 We model a black-box confidence-threshold attacker as follows. Let I be a random index drawn from
 914 \mathcal{I}_{in} (e.g., uniformly), and let the random member pair be

$$(X_{\text{in}}, Y_{\text{in}}) := (x_I, y_I).$$

918 Let $X_{\text{out}} \sim \mathcal{N}(0, \Sigma)$ be an independent non-member (test) input, independent of the training set and
 919 algorithmic randomness. For a threshold $\tau \geq 0$, define

$$\text{TPR}_G(\tau) := \Pr(\text{Conf}_G(X_{\text{in}}) \geq \tau), \quad (7)$$

$$\text{FPR}_G(\tau) := \Pr(\text{Conf}_G(X_{\text{out}}) \geq \tau), \quad (8)$$

$$\text{Adv}_G^{\text{conf}} := \sup_{\tau \geq 0} (\text{TPR}_G(\tau) - \text{FPR}_G(\tau)). \quad (9)$$

925 **Lemma 5.** *For any geometry $G \succ 0$,*

$$\text{Conf}_G(X_{\text{in}}) = |Y_{\text{in}}| \quad \text{almost surely.}$$

928 *Proof.* Because $\widehat{\theta}_G$ interpolates the training data, we have $\widehat{\theta}_G^\top x_i = y_i$ for every $i \in \mathcal{I}_{\text{in}}$. Therefore
 929 for the random member index I ,

$$\text{Conf}_G(X_{\text{in}}) = |\widehat{\theta}_G^\top x_I| = |y_I| = |Y_{\text{in}}| \quad \text{a.s.}$$

□

933 **Lemma 6.** *Conditioned on $\widehat{\theta}_G$,*

$$\text{Conf}_G(X_{\text{out}}) \stackrel{d}{=} |Z|, \quad Z \sim \mathcal{N}(0, \sigma_G^2),$$

936 *Proof.* By Lemma 3 and absolute value. □

938 **Lemma 7.** *Let $0 < \sigma_1 < \sigma_2$ and let $Z_k \sim \mathcal{N}(0, \sigma_k^2)$ for $k \in \{1, 2\}$. Then for every $\tau > 0$,*

$$\Pr(|Z_1| \geq \tau) < \Pr(|Z_2| \geq \tau).$$

941 *Proof.* Let $U \sim \mathcal{N}(0, 1)$. Then $Z_k \stackrel{d}{=} \sigma_k U$, so

$$\Pr(|Z_k| \geq \tau) = \Pr\left(|U| \geq \frac{\tau}{\sigma_k}\right).$$

945 The function $t \mapsto \Pr(|U| \geq t)$ is strictly decreasing on $(0, \infty)$. Since $\tau/\sigma_1 > \tau/\sigma_2$, the claim
 946 follows. □

947 **Lemma 8.** *Assume $\Pr(|Y_{\text{in}}| > 0) > 0$ (true whenever labels have any continuous noise). Then*

$$\text{Adv}_G^{\text{conf}} = \sup_{\tau > 0} (\text{TPR}_G(\tau) - \text{FPR}_G(\tau)).$$

951 *Proof.* At $\tau = 0$, $\text{TPR}_G(0) = \text{FPR}_G(0) = 1$, so the gap equals 0. Because $\Pr(|Y_{\text{in}}| > 0) > 0$, we
 952 have $\text{TPR}_G(\tau) > 0$ for some $\tau > 0$. Also, by Lemma 6, $\text{FPR}_G(\tau) \rightarrow 0$ as $\tau \rightarrow \infty$. Hence there
 953 exists a $\tau > 0$ such that $\text{TPR}_G(\tau) - \text{FPR}_G(\tau) > 0$. Therefore the supremum cannot be attained at
 954 $\tau = 0$, and we may restrict to $\tau > 0$. □

955 Proof of Theorem 2

958 *Proof.* By Lemma 5, member confidences are geometry-invariant, so

$$\text{TPR}_{G_\eta}(\tau) = \text{TPR}_{G_0}(\tau) \quad \text{for all } \tau \geq 0.$$

960 By Theorem 1, for all sufficiently small $\eta > 0$,

$$\sigma_{G_\eta}^2 < \sigma_{G_0}^2 \quad \text{with probability } 1 - o(1).$$

963 Condition on this high-probability event. Lemma 6 implies $\text{Conf}_{G_\eta}(X_{\text{out}}) \stackrel{d}{=} |Z_\eta|$ with $Z_\eta \sim$
 964 $\mathcal{N}(0, \sigma_{G_\eta}^2)$, and similarly for G_0 . Then Lemma 7 yields that for every $\tau > 0$,

$$\text{FPR}_{G_\eta}(\tau) = \Pr(|Z_\eta| \geq \tau) < \Pr(|Z_0| \geq \tau) = \text{FPR}_{G_0}(\tau).$$

966 Therefore, for every $\tau > 0$,

$$\text{TPR}_{G_\eta}(\tau) - \text{FPR}_{G_\eta}(\tau) > \text{TPR}_{G_0}(\tau) - \text{FPR}_{G_0}(\tau).$$

969 Taking the supremum over $\tau > 0$ and using Lemma 8,

$$\text{Adv}_{G_\eta}^{\text{conf}} = \sup_{\tau > 0} (\text{TPR}_{G_\eta}(\tau) - \text{FPR}_{G_\eta}(\tau)) > \sup_{\tau > 0} (\text{TPR}_{G_0}(\tau) - \text{FPR}_{G_0}(\tau)) = \text{Adv}_{G_0}^{\text{conf}}.$$

971 This holds with probability $1 - o(1)$. □

972 D.2 LIKELIHOOD-RATIO ATTACK
973

974 We consider an oracle likelihood-ratio attack that uses the model output $f_G(x) = \hat{\theta}_G^\top x$. The attacker
975 knows the true member and non-member score distributions and performs the Neyman–Pearson
976 likelihood-ratio test.

977 **Score distributions.** Recall the data model $x \sim \mathcal{N}(0, \Sigma)$ and $y = \theta^{*\top} x + \xi$ with $\xi \sim \mathcal{N}(0, \sigma_y^2)$
978 independent of x . Let $(X_{\text{in}}, Y_{\text{in}})$ be a random member training pair obtained by sampling a random
979 index from the training set. Unconditionally over the training data draw, we have the same marginal
980 law as a fresh sample:

$$982 \quad X_{\text{in}} \sim \mathcal{N}(0, \Sigma), \quad Y_{\text{in}} = \theta^{*\top} X_{\text{in}} + \xi.$$

983 **Lemma 9** (Member score is geometry-invariant Gaussian). *For any geometry G , the member score
984 satisfies*

$$985 \quad f_G(X_{\text{in}}) = Y_{\text{in}} \quad \text{a.s.} \quad \text{and hence} \quad f_G(X_{\text{in}}) \sim \mathcal{N}(0, v_{\text{in}}),$$

986 where $v_{\text{in}} := \theta^{*\top} \Sigma \theta^* + \sigma_y^2$.

987 *Proof.* Interpolation gives $X \hat{\theta}_G = y$, so for every training point $\hat{\theta}_G^\top x_i = y_i$. Thus for a random
988 member index I ,

$$989 \quad f_G(X_{\text{in}}) = \hat{\theta}_G^\top x_I = y_I = Y_{\text{in}} \quad \text{a.s.}$$

990 Unconditionally, $Y_{\text{in}} = \theta^{*\top} X_{\text{in}} + \xi$ with $X_{\text{in}} \sim \mathcal{N}(0, \Sigma)$ and $\xi \sim \mathcal{N}(0, \sigma_y^2)$ independent, so it is
991 mean-zero Gaussian with variance v_{in} . \square

992 **Lemma 10** (Non-member score is geometry-dependent Gaussian). *Let $X_{\text{out}} \sim \mathcal{N}(0, \Sigma)$ be independent
993 of the training set. Conditioned on $\hat{\theta}_G$,*

$$994 \quad f_G(X_{\text{out}}) = \hat{\theta}_G^\top X_{\text{out}} \sim \mathcal{N}(0, v_{\text{out}}(G)), \quad v_{\text{out}}(G) := \hat{\theta}_G^\top \Sigma \hat{\theta}_G.$$

1001 *Proof.* Same argument as Lemma 3. \square

1002 **Oracle likelihood-ratio test.** Let p_{in} and p_{out} be the densities of $\mathcal{N}(0, v_{\text{in}})$ and $\mathcal{N}(0, v_{\text{out}}(G))$,
1003 respectively. The oracle LR score is

$$1004 \quad \Lambda_G(s) := \log \frac{p_{\text{in}}(s)}{p_{\text{out}}(s)}.$$

1005 **Lemma 11.** *For zero-mean Gaussians with variances $v_{\text{in}} > 0$ and $v_{\text{out}}(G) > 0$, the LR test
1006 $\Lambda_G(s) \geq t$ is equivalent to $|s| \geq \tau$ for some $\tau \geq 0$. Moreover, the optimal test for any fixed FPR is of
1007 this form.*

1008 *Proof.* For $s \in \mathbb{R}$,

$$1009 \quad \Lambda_G(s) = -\frac{1}{2} \log v_{\text{in}} + \frac{1}{2} \log v_{\text{out}}(G) - \frac{s^2}{2v_{\text{in}}} + \frac{s^2}{2v_{\text{out}}(G)} = C_G + \frac{s^2}{2} \left(\frac{1}{v_{\text{out}}(G)} - \frac{1}{v_{\text{in}}} \right),$$

1010 where C_G does not depend on s . Thus $\Lambda_G(s) \geq t$ is equivalent to $s^2 \geq \tau^2$ for some $\tau \geq 0$, i.e.
1011 $|s| \geq \tau$. Neyman–Pearson gives optimality of the LR test. Note that as we are in a setting where
1012 members are highly confident, $v_{\text{in}} > v_{\text{out}}$. \square

1013 Define the LR-attack advantage as the best achievable TPR–FPR gap over all two-sided thresholds:

$$1014 \quad \text{Adv}_G^{\text{LR}} := \sup_{\tau \geq 0} \left(\Pr(|f_G(X_{\text{in}})| \geq \tau) - \Pr(|f_G(X_{\text{out}})| \geq \tau) \right).$$

1026 **Proof of Theorem 3**

1028 *Proof.* By Lemma 9,

1030 $f_{G_\eta}(X_{\text{in}}) \sim \mathcal{N}(0, v_{\text{in}})$ for all $\eta \geq 0$,

1031 so the member tail $\Pr(|f_{G_\eta}(X_{\text{in}})| \geq \tau)$ is geometry-invariant for every $\tau \geq 0$.

1032 By Lemma 10,

1034 $f_{G_\eta}(X_{\text{out}}) \sim \mathcal{N}(0, v_{\text{out}}(G_\eta)), \quad v_{\text{out}}(G_\eta) = \sigma_{G_\eta}^2$.

1036 Theorem 1 gives that, for all sufficiently small $\eta > 0$,

1037 $v_{\text{out}}(G_\eta) < v_{\text{out}}(G_0)$ with probability $1 - o(1)$.

1039 Condition on this high-probability event. Then Lemma 7 (applied to $|f_G(X_{\text{out}})|$) yields that for every
1040 $\tau > 0$,

1041 $\Pr(|f_{G_\eta}(X_{\text{out}})| \geq \tau) < \Pr(|f_{G_0}(X_{\text{out}})| \geq \tau)$.

1042 Therefore, for every $\tau > 0$,

1044 $\Pr(|f_{G_\eta}(X_{\text{in}})| \geq \tau) - \Pr(|f_{G_\eta}(X_{\text{out}})| \geq \tau) > \Pr(|f_{G_0}(X_{\text{in}})| \geq \tau) - \Pr(|f_{G_0}(X_{\text{out}})| \geq \tau)$.

1046 Taking the supremum over $\tau > 0$ on both sides gives $\text{Adv}_{G_\eta}^{\text{LR}} > \text{Adv}_{G_0}^{\text{LR}}$. (As in Lemma 8, $\tau = 0$
1047 yields zero gap, so the supremum is attained for some $\tau > 0$.) This holds with probability $1 - o(1)$. \square

1049 **Corollary 5** (Sample-adaptive LR monotonicity). *Fix a query point z :*

1051 $f_{\text{in},z} \sim \mathcal{N}(0, v_{\text{in},z}), \quad f_{\text{out},z}(G) \sim \mathcal{N}(0, v_{\text{out},z}(G))$.

1052 Let G_a, G_b be two geometries such that $v_{\text{out},z}(G_a) < v_{\text{out},z}(G_b)$. Then for every false positive rate
1053 $\alpha \in (0, 1)$, the optimal likelihood-ratio test at level α attains strictly larger true positive rate under
1054 G_a than under G_b . Equivalently, the per-sample ROC curve under G_a strictly dominates that under
1055 G_b , so any MIA advantage is strictly larger under G_a .

1057 *Proof.* Fix z and $k \in \{a, b\}$. The LR between $\mathcal{N}(0, v_{\text{in},z})$ and $\mathcal{N}(0, v_{\text{out},z}(G_k))$ is

1059
$$\Lambda_z(s; G_k) = \frac{1}{2} \log \frac{v_{\text{out},z}(G_k)}{v_{\text{in},z}} + \frac{s^2}{2} \left(\frac{1}{v_{\text{out},z}(G_k)} - \frac{1}{v_{\text{in},z}} \right).$$

1062 If $v_{\text{in},z} \neq v_{\text{out},z}(G_k)$ then $\Lambda_z(s; G_k)$ is a strictly monotone function of $|s|$, so by the Neyman-
1063 Pearson lemma the optimal level- α test is equivalent to a two-sided magnitude test $|s| \geq \tau_k(\alpha)$ for
1064 some unique threshold $\tau_k(\alpha) > 0$.

1066 Write $f_{\text{out},z}(G_k) \stackrel{d}{=} \sqrt{v_{\text{out},z}(G_k)} U$ with $U \sim \mathcal{N}(0, 1)$. The constraint $\Pr(|f_{\text{out},z}(G_k)| \geq \tau_k(\alpha)) =$
1067 α is then

1068
$$\alpha = \Pr\left(|U| \geq \frac{\tau_k(\alpha)}{\sqrt{v_{\text{out},z}(G_k)}}\right).$$

1070 Since $u \mapsto \Pr(|U| \geq u)$ is strictly decreasing on $(0, \infty)$ and $v_{\text{out},z}(G_a) < v_{\text{out},z}(G_b)$, this forces

1072
$$\frac{\tau_a(\alpha)}{\sqrt{v_{\text{out},z}(G_a)}} = \frac{\tau_b(\alpha)}{\sqrt{v_{\text{out},z}(G_b)}} \Rightarrow \tau_a(\alpha) < \tau_b(\alpha).$$

1075 The member distribution $f_{\text{in},z}$ is the same under G_a and G_b , so

1077
$$\Pr(|f_{\text{in},z}| \geq \tau_a(\alpha)) > \Pr(|f_{\text{in},z}| \geq \tau_b(\alpha)).$$

1079 Thus at every FPR level α the optimal LR/LiRA test has strictly larger TPR under G_a , which implies
strict ROC dominance and the claimed advantage comparison. \square

1080 D.3 HIGHER SUBCLASS ALIGNMENT LEADS TO HIGHER GENERALIZATION
10811082 **Lemma 12** (High-dimensional near-orthogonality). *Let $\mathbf{x}_1, \dots, \mathbf{x}_N \in \mathbb{R}^d$ have i.i.d. $\mathcal{N}(0, 1)$ entries.
1083 Then, as $d \rightarrow \infty$,*

1084
$$\|\mathbf{x}_i\|^2 = d(1 + o(1)) \quad \text{and} \quad \frac{\langle \mathbf{x}_i, \mathbf{x}_j \rangle}{\|\mathbf{x}_i\| \|\mathbf{x}_j\|} = o(1)$$

1085

1086 *for each fixed $i \neq j$, with probability tending to 1. Moreover, the two conclusions hold uniformly
1087 over all $i \neq j$ with probability tending to 1 provided $\log N = o(d)$.*1088 *Proof.* For norms, $\|\mathbf{x}_i\|^2 \sim \chi^2(d)$. Laurent–Massart’s inequality implies that for all $t > 0$,

1089
$$\Pr(\|\mathbf{x}_i\|^2 - d \geq 2\sqrt{dt} + 2t) \leq e^{-t}, \quad \Pr(d - \|\mathbf{x}_i\|^2 \geq 2\sqrt{dt}) \leq e^{-t}.$$

1090

1092 Taking $t = \varepsilon^2 d$ gives $\|\mathbf{x}_i\|^2 = d(1 \pm O(\varepsilon))$ with probability at least $1 - 2e^{-\varepsilon^2 d}$; hence $\|\mathbf{x}_i\|^2 =$
1093 $d(1 + o(1))$ w.h.p. and $\|\mathbf{x}_i\| = \sqrt{d} + \mathcal{O}_p(1)$.
10941095 For inner products, write $\langle \mathbf{x}_i, \mathbf{x}_j \rangle = \sum_{k=1}^d Z_k$ with $Z_k := \mathbf{x}_{i,k} \mathbf{x}_{j,k}$, which are i.i.d., mean 0, and
1096 sub-exponential. Bernstein’s inequality yields

1097
$$\Pr(|\langle \mathbf{x}_i, \mathbf{x}_j \rangle| \geq t) \leq 2 \exp\left(-c \min\{t^2/d, t\}\right)$$

1098

1099 for a universal $c > 0$. Taking $t = C\sqrt{d}$ shows $|\langle \mathbf{x}_i, \mathbf{x}_j \rangle| = \mathcal{O}_p(\sqrt{d})$. Combining with $\|\mathbf{x}_i\| \|\mathbf{x}_j\| =$
1100 $d(1 + o_p(1))$,

1102
$$\frac{|\langle \mathbf{x}_i, \mathbf{x}_j \rangle|}{\|\mathbf{x}_i\| \|\mathbf{x}_j\|} = \mathcal{O}_p(d^{-1/2}) = o_p(1).$$

1103

1104 A union bound over the $\binom{N}{2}$ pairs then gives the uniform statement whenever $\log N = o(d)$, since
1105 both tails are $\exp(-\Theta(d))$. \square 1106 **Noise weights** By the representer theorem, write

1107
$$\mathbf{w}_3 = \sum_{j=1}^n \beta_j y_j \mathbf{x}_{3,j}.$$

1108

1109 Then, for any training point i ,

1110
$$y_i \langle \mathbf{w}_3, \mathbf{x}_{3,i} \rangle = \beta_i \|\mathbf{x}_{3,i}\|^2 + \zeta_i, \quad \zeta_i := \sum_{j \neq i} \beta_j y_i y_j \langle \mathbf{x}_{3,j}, \mathbf{x}_{3,i} \rangle.$$

1111

1112 Under Lemma 12, $\|\mathbf{x}_{3,i}\|^2 = (1 \pm o(1)) d_3$ and the cross inner products are $o(\|\mathbf{x}_{3,i}\| \|\mathbf{x}_{3,j}\|) = o(d_3)$ 1113 **Condition 1.** For each minority point $i \in \mathcal{S}$, define

1114
$$B_i := -y_i \langle \mathbf{w}_1, \boldsymbol{\nu}_i \rangle > 0, \quad A := \langle \mathbf{w}_2, \boldsymbol{\mu}_2 \rangle.$$

1115

1116 Let B be a random variable with CDF $F_B(A) = \Pr(B < A)$ such that $B \stackrel{d}{=} B_i$ (i.e., F_B is the
1117 law/distribution of the B_i ’s when i is drawn uniformly from \mathcal{S}). We assume

1118
$$A < B_{\max}, \quad B_{\max} \triangleq \sup\{b \in \mathbb{R} \mid F_B(b) < 1\},$$

1119

1120 **Remark 5** (Distribution F_B and why $F_B(A)$ appears). For each minority sample i , the anti-alignment
1121 magnitude $B_i \triangleq -y_i \langle \mathbf{w}_1, \boldsymbol{\nu}_i \rangle > 0$ summarizes how strongly the majority anchor opposes the
1122 minority anchor for that sample. We assume $\{B_i\}_{i \in \mathcal{S}}$ are i.i.d. draws from a common distribution F_B
1123 supported on $(0, B_{\max}]$. Unseen minority test point has margin $m' = -B + A + \zeta'$, with $B \sim F_B$,
1124 so $\Pr_{\mathcal{S}}(m' > 0) = \Pr(B < A) = F_B(A)$ (up to the $o(1)$ fluctuation ζ' from Lemma 12). Intuitively,
1125 $F_B(A)$ is the fraction of minority subclasses whose majority anti-alignment is not too strong relative
1126 to the learned shared minority signal A .1127 **Definition 4** (Generalization gap). *The generalization gap for a training point i from distribution
1128 $\mathcal{D}_{\mathcal{K}} \in \mathcal{D}_{\mathcal{S}}, \mathcal{D}_{\mathcal{M}}, \mathcal{D}_{\mathcal{N}}$ is defined as*

1129
$$R_i^{\mathcal{K}}(w) := \Pr\{y_i \langle \mathbf{w}, \mathbf{x}_i \rangle > 0\} - \Pr\{y' \langle \mathbf{w}, \mathbf{x}' \rangle > 0\}, \quad (\mathbf{x}', y') \sim \mathcal{D}_{\mathcal{K}}. \quad (10)$$

1130

1134 *Assumption 2* (Majority alignment). We consider models whose first block weights align with the
 1135 majority subclass feature:

$$1136 \quad \mathbf{w}_1 = \alpha \boldsymbol{\mu} \quad (\alpha > 0),$$

1137 This captures the implicit bias of common ERM procedures (e.g., logistic regression trained by
 1138 gradient descent, or minimum- ℓ_2 -norm interpolation) to align with the dominant signal in the data.

1139 *Assumption 3* (Majority dominance). We assume the majority signal dominates the stochastic parts:

$$1141 \quad \frac{\langle \mathbf{w}_1, \boldsymbol{\mu}_1 \rangle^2}{\|\mathbf{w}_2\|^2 + \|\mathbf{w}_3\|^2} \rightarrow \infty.$$

1143 *Assumption 4* (Perfect interpolation and finite-margin). A trained model \mathbf{w} has finite margins. That
 1144 is, there exist $0 < m_0 \leq M < \infty$ such that

$$1146 \quad m_0 \leq y_i \langle \mathbf{w}, \mathbf{x}_i \rangle \leq M \quad \forall i.$$

1147 **Lemma 13** (Generalization gap of majority subclass samples). *Let $i \in \mathcal{M}$. Then with probability*
 1148 $1 - o(1)$,

$$1149 \quad R_i^{\mathcal{M}}(\mathbf{w}) = o(1)$$

1151 *Proof.* By Assumption 4, $y_i \langle \mathbf{w}, \mathbf{x}_i \rangle > 0$, hence $\Pr\{y_i \langle \mathbf{w}, \mathbf{x}_i \rangle > 0\} = 1$.

1153 Draw $(\mathbf{x}', y') \sim \mathcal{D}_{\mathcal{M}}$, so $\mathbf{x}'_1 = y' \boldsymbol{\mu}_1$, $\mathbf{x}'_2 \sim \mathcal{N}(0, I_{d_2})$, $\mathbf{x}'_3 \sim \mathcal{N}(0, I_{d_3})$, independent of y' . Then

$$1155 \quad y' \langle \mathbf{w}, \mathbf{x}' \rangle = \langle \mathbf{w}_1, \boldsymbol{\mu}_1 \rangle + y' \langle \mathbf{w}_2, \mathbf{x}'_2 \rangle + y' \langle \mathbf{w}_3, \mathbf{x}'_3 \rangle =: \langle \mathbf{w}_1, \boldsymbol{\mu}_1 \rangle + Z,$$

1156 where Z is a mean-zero sub-Gaussian random variable with variance proxy $\text{VarProxy}(Z) = \|\mathbf{w}_2\|^2 +$
 1157 $\|\mathbf{w}_3\|^2$, since $y' \langle \mathbf{w}_2, \mathbf{x}'_2 \rangle \sim \mathcal{N}(0, \|\mathbf{w}_2\|^2)$ and $y' \langle \mathbf{w}_3, \mathbf{x}'_3 \rangle \sim \mathcal{N}(0, \|\mathbf{w}_3\|^2)$ are independent and
 1158 centered.

1159 For any $a > 0$, a standard sub-Gaussian tail bound yields

$$1161 \quad \Pr\{Z \leq -a\} \leq \exp\left(-\frac{a^2}{2(\|\mathbf{w}_2\|^2 + \|\mathbf{w}_3\|^2)}\right).$$

1163 Taking $a = \langle \mathbf{w}_1, \boldsymbol{\mu}_1 \rangle$,

$$1165 \quad \Pr\{y' \langle \mathbf{w}, \mathbf{x}' \rangle \leq 0\} = \Pr\{Z \leq -\langle \mathbf{w}_1, \boldsymbol{\mu}_1 \rangle\} \leq \exp\left(-\frac{\langle \mathbf{w}_1, \boldsymbol{\mu}_1 \rangle^2}{2(\|\mathbf{w}_2\|^2 + \|\mathbf{w}_3\|^2)}\right).$$

1167 Hence

$$1168 \quad \Pr\{y' \langle \mathbf{w}, \mathbf{x}' \rangle > 0\} \geq 1 - \exp\left(-\frac{\langle \mathbf{w}_1, \boldsymbol{\mu}_1 \rangle^2}{2(\|\mathbf{w}_2\|^2 + \|\mathbf{w}_3\|^2)}\right).$$

1170 By Assumption 3,

$$1172 \quad R_i^{\mathcal{M}}(\mathbf{w}) = 1 - \Pr\{y' \langle \mathbf{w}, \mathbf{x}' \rangle > 0\} \leq \exp\left(-\frac{\langle \mathbf{w}_1, \boldsymbol{\mu}_1 \rangle^2}{2(\|\mathbf{w}_2\|^2 + \|\mathbf{w}_3\|^2)}\right) = o(1)$$

1175 \square

1177 **Lemma 14** (Generalization gap of minority subclass samples). *Let $i \in \mathcal{S}$.*

$$1178 \quad R_i^{\mathcal{S}}(\mathbf{x}) = 1 - F_B(A) \quad (\text{up to } o(1) \text{ terms}).$$

1180 *Proof.* By Assumption 4, $y_i \langle \mathbf{w}, \mathbf{x}_i \rangle > 0$, hence $\Pr\{y_i \langle \mathbf{w}, \mathbf{x}_i \rangle > 0\} = 1$.

1182 Draw $(\mathbf{x}', y') \sim \mathcal{D}_{\mathcal{S}}$. \mathbf{x}'_3 is independent of $\{\mathbf{x}_{3,j}\}$ and mean-zero; hence $y \langle \mathbf{w}_3, \mathbf{x}'_3 \rangle =$
 1183 $\sum_j \beta_j y y_j \langle \mathbf{x}_{3,j}, \mathbf{x}'_3 \rangle$ is a mean-zero fluctuation with variance vanishing relative to $\|\mathbf{x}'_3\|^2$; set this
 1184 fluctuation to $\zeta' = o_{\mathbb{P}}(1)$ using Lemma 12. Then, $y' \langle \mathbf{w}, \mathbf{x}' \rangle = -B + A + \zeta'$.

$$1186 \quad \Pr\{y' \langle \mathbf{w}, \mathbf{x}' \rangle > 0\} = \Pr(B < A - \zeta') \in (F_B(A - \varepsilon), F_B(A + \varepsilon)).$$

1187 Letting $\varepsilon \downarrow 0$ gives the stated identities up to $o(1)$. \square

1188 **Lemma 15** (Generalization gap for pure noise samples). *Let $i \in \mathcal{N}$ (pure noise).*

1189
1190
$$R_i^{\mathcal{N}}(\mathbf{w}) = \frac{1}{2}$$

1191

1192 *Proof.* By Assumption 4, $y_i \langle \mathbf{w}, \mathbf{x}_i \rangle > 0$, hence $\Pr\{y_i \langle \mathbf{w}, \mathbf{x}_i \rangle > 0\} = 1$.

1193 For an unseen noise $(\mathbf{x}', y') \sim \mathcal{D}_{\mathcal{N}}$, each term $y' \langle \mathbf{w}_k, \mathbf{w}'_k \rangle$ ($k = 1, 2, 3$) is a centered continuous
1194 symmetric random variable (linear form of a mean-zero isotropic vector, independent of y'). The
1195 sum remains centered and symmetric; hence $\Pr(y' \langle \mathbf{w}, \mathbf{x}' \rangle > 0) = 1/2$. \square

1196
1197 **Proof of Theorem 4**

1198 By Lemma 14, minority test accuracy equals $F_B(A)$. This is monotone increasing in A . The
1199 majority subclass and noise samples yield the same result for both models by Lemmas 13 and 15. \square

1200
1201 **E ADDITIONAL RELATED WORKS**

1202
1203 **E.1 CONNECTION OF FLATTER MINIMA WITH GENERALIZATION GAP**

1204 There have been numerous studies (Foret et al., 2020; Izmailov et al., 2018; Cha et al., 2021; Norton &
1205 Royset, 2021; Wu et al., 2020) which account for the worst-case empirical risks within neighborhoods
1206 in parameter space. Diametrical Risk Minimization (DRM) was first proposed by (Norton & Royset,
1207 2021) and they asserted that the practical and theoretical performance of Empirical Risk Minimization
1208 (ERM) tends to suffer when dealing with loss functions that exhibit poor behavior characterized
1209 by large Lipschitz moduli and spurious sharp minimizers. They tackled this concern by employing
1210 DRM, which offers generalization bounds that are unaffected by Lipschitz moduli, applicable to
1211 both convex and non-convex problems. Another algorithm that improves generalization is Sharpness
1212 Aware Minimization (SAM) (Foret et al., 2020) which performs gradient descent while regularizing
1213 for the highest loss in the neighborhood of radius ρ of the parameter space. (Izmailov et al., 2018)
1214 proposed Stochastic Weight Averaging (SWA) that performs averaging of weights with a cyclical
1215 or constant learning rate which leads to better generalization than conventional training. They also
1216 prove that the optima chosen by the single model is in fact a flatter minima than the SGD solution.
1217 Further, (Cha et al., 2021) argues that simply performing the Empirical Risk Minimization (ERM)
1218 is not enough to achieve a good generalization, in particular, domain generalization. Hence, they
1219 introduce SWAD which seeks for flatter optima and hence, will generalize well across domain shifts.
1220

1221
1222 **E.2 DIFFERENT MEMBERSHIP INFERENCE ATTACKS**

1223 There are many variants of Direct Single-query attacks (DSQ) based on the approach of the attack
1224 and below we describe the ones used in our experiments:

1225 **NN-based attack (Shokri et al., 2017; Tang et al., 2022; Nasr et al., 2018)** This is the first MIA
1226 proposed by Shokri et al. (2017) where they use a binary classifier to distinguish between the training
1227 members and the non-members using the victim model's behavior on these data points. The adversary
1228 can utilize the prediction vectors from the target model and incorporate them along with the one-hot
1229 encoded ground truth labels as inputs. Then, they can construct a neural network (I_{NN}) called attack
1230 model.

1231 **Confidence-based attack (Yeom et al., 2020; Salem et al., 2018; Song & Mittal, 2021)** If the
1232 highest prediction confidence of an input record exceeds a predetermined threshold, the adversary
1233 considers it a member; otherwise, it is inferred as a non-member. This approach is based on the
1234 understanding that the target model is trained to minimize prediction loss using its training data,
1235 implying that the maximum confidence score of a prediction vector for a training member should be
1236 near 1. The attack I_{conf} is defined as follows:

1237
1238
$$I_{conf} \hat{p}(y|\mathbf{x}) = \mathbb{1}(\max \hat{p}(y|\mathbf{x}) \geq \tau) \quad (11)$$

1239

1240 Here, $\mathbb{1}(\cdot)$ is an indicator function which returns 1 if the predicate inside it holds True else the
1241 function evaluates to 0.

Entropy-based attack (Nasr et al., 2019; Song & Mittal, 2021; Tang et al., 2022) When the prediction entropy of an input record falls below a predetermined threshold, the adversary considers it a member. Conversely, if the prediction entropy exceeds the threshold, the adversary infers that the record is a non-member. This inference is based on the observation that there are notable disparities in the prediction entropy distributions between training and test data. Typically, the target model exhibits higher prediction entropy on its test data compared to its training data. The entropy of a prediction vector $p(\hat{y}|x)$ is defined as follows:

$$H(p(\hat{y}|x)) = - \sum_i (p_i \log(p_i)) \quad (12)$$

where p_i is the confidence score in $p(\hat{y}|x)$. Then, the attack I_{entr} is given as:

$$I_{entr}(\hat{p}(y|x), y) = \mathbb{1}(H(p(\hat{y}|x)) \leq \tau) \quad (13)$$

Modified entropy-based attack (Song & Mittal, 2021) Song et al.[15] introduced an enhanced prediction entropy metric that integrates both the entropy metric and the ground truth labels. The modified entropy metric tends to yield lower values for training samples compared to testing samples. To infer membership, either a class-dependent threshold τ_y or a class-independent threshold τ_{attack} is applied.

$$I_{Mentr}(\hat{p}(y|x), y) = \mathbb{1}(Mentr(p(\hat{y}|x)) \leq \tau_y) \quad (14)$$

where $Mentr(p(\hat{y}|x))$ for (x,y) data sample is given by combination of entropy information and ground truth label as:

$$Mentr(p(\hat{y}|x)) = -((1 - p(\hat{y}|x)_y) \log(p(\hat{y}|x)_y) - \sum_{i \neq y} (p(\hat{y}|x)_i \log(1 - p(\hat{y}|x)_i))) \quad (15)$$

Likelihood Ratio Attack (LiRA) (Carlini et al., 2022) LiRA is a shadow-model based single-query attack that explicitly models the distributions of a scalar score for members and non-members and then performs a likelihood ratio test. For a sample (x, y) , the attacker first defines a one-dimensional score $s(x, y)$ from the target model, typically the negative cross-entropy loss or the (log-)confidence on the true label y . Using multiple shadow models trained with and without (x, y) , the attacker estimates two score distributions: one for members (IN) and one for non-members (OUT). In practice, LiRA fits parametric Gaussians

$$s(x, y) \mid \text{IN} \sim \mathcal{N}(\mu_{\text{in}}, \sigma_{\text{in}}^2), \quad s(x, y) \mid \text{OUT} \sim \mathcal{N}(\mu_{\text{out}}, \sigma_{\text{out}}^2),$$

and computes the log-likelihood ratio

$$\Lambda_{\text{LiRA}}(x, y) = \log \frac{\phi(s(x, y); \mu_{\text{in}}, \sigma_{\text{in}}^2)}{\phi(s(x, y); \mu_{\text{out}}, \sigma_{\text{out}}^2)},$$

where $\phi(\cdot; \mu, \sigma^2)$ denotes the Gaussian density. The LiRA decision rule is then

$$I_{\text{LiRA}}(\hat{p}(y|x), y) = \mathbb{1}(\Lambda_{\text{LiRA}}(x, y) \geq \tau_{\text{LiRA}}), \quad (16)$$

for some threshold τ_{LiRA} chosen to trade off between TPR and FPR. In the “online” variant, the attacker fits both IN and OUT distributions from shadow models; in the “offline” variant, only the OUT distribution is estimated and low likelihood under the OUT model is treated as evidence of membership.

Robust MIA (RMIA) (Zarifzadeh et al., 2024) RMIA reframes membership inference as a calibrated hypothesis test based on a *pairwise likelihood ratio* between a query (x, y) and many population samples (z, y_z) . For each pair (x, z) , RMIA compares how the (approximate) probability of x and z change when conditioning on the event that x was used to train the target model. Concretely, RMIA defines

$$\text{LR}(x, z) = \frac{\text{Pr}(x \mid \theta)}{\text{Pr}(z \mid \theta)} / \frac{\text{Pr}(x)}{\text{Pr}(z)},$$

where $\text{Pr}(\cdot \mid \theta)$ denotes the target model’s likelihood and $\text{Pr}(\cdot)$ is a population prior. Intuitively, if including x in the target training set fits x disproportionately better than many other population points

1296
 1297 Table 3: Attack accuracy of direct threshold MIA on SGD, Sharp, ASAM, and GSAM. Analogous to
 1298 SAM, optimization methods that improve generalization (ASAM, GSAM) through finding flatter
 1299 minima tend to be more prone to direct threshold attacks, while optimization that looks for sharp
 1300 minima instead is more robust to MIA attack while being worse off in generalization.

Dataset	Algo	NN	Confidence	Entropy	M-entropy	Test Acc
CIFAR-100	SGD	76.62%	77.19%	76.61%	77.30%	80.30%
	Sharp	57.62%	59.69%	57.88%	59.69%	76.14%
	ASAM	78.92%	79.22%	78.86%	79.31%	81.80%
	GSAM	78.63%	79.23%	79.00%	79.23%	82.16%
CIFAR-10	SGD	50.00%	59.37%	59.09%	59.51%	96.00%
	Sharp	50.22%	52.86%	52.47%	52.78%	92.86%
	ASAM	50.48%	61.39%	61.20%	61.32%	96.66%
	GSAM	50.00%	61.46%	61.38%	61.54%	96.64%
Purchase100	SGD	66.00%	66.76%	64.78%	67.13%	85.50%
	Sharp	59.58%	60.96%	58.04%	61.16%	84.31%
	ASAM	66.85%	66.84%	65.39%	67.03%	85.54%
	GSAM	67.45%	67.72%	66.51%	67.87%	85.82%
Texas100	SGD	59.81%	65.20%	55.74%	65.13%	50.83%
	Sharp	51.11%	59.89%	53.46%	59.36%	49.97%
	ASAM	60.92%	67.50%	58.80%	67.10%	53.17%
	GSAM	54.89%	67.07%	57.93%	67.13%	52.04%

1319 \mathbf{z} , the ratio $\text{LR}(\mathbf{x}, \mathbf{z})$ becomes large. RMIA samples many \mathbf{z} from the population and defines a robust
 1320 membership score

$$R(\mathbf{x}) = \frac{1}{|n_z|} \sum_{\mathbf{z}} \mathbb{1}(\text{LR}(\mathbf{x}, \mathbf{z}) > \gamma),$$

1325 where γ is a fixed pairwise LR threshold and n_z is number of population (non-member) samples. The
 1326 attack then declares membership if $R(\mathbf{x})$ exceeds a global threshold τ ; by sweeping τ one obtains a
 1327 calibrated ROC curve, and for a chosen FPR one can directly pick the corresponding τ . In the *offline*
 1328 mode, all reference models are OUT models trained once on population data; in the *online* mode, the
 1329 attacker additionally trains IN reference models that explicitly include \mathbf{x} in their training set, which
 1330 yields a more accurate approximation of the conditional likelihoods but is more computationally
 1331 expensive.

F OTHER SHARPNESS-AWARE OPTIMIZERS

1335 In this section, we discuss other variants of SAM, namely Adaptive SAM (ASAM) (Kwon et al.,
 1336 2021), Guided SAM (GSAM) (Zhuang et al., 2022), and custom designed optimizer, namely Sharp.
 1337 Sharp objective is designed to explicitly find a sharper minima. The objective function of Sharp is,
 1338

$$\mathcal{L}_{\text{Sharp}}(\mathbf{w}) = L(\mathbf{w}) - \beta \max_{\epsilon \in B(\rho)} L(\mathbf{w} + \epsilon). \quad (17)$$

1343 This objective can be seen as minimizing the loss at current \mathbf{w} while maximizing the loss in the
 1344 vicinity. We empirically verify that this objective does lead to a sharper minima measuring its hessian
 1345 trace. Results and discussion about Sharp are available in Appendix I.4.

1347 The results on CIFAR10, CIFAR100, Purchase100, and Texas100 are reported in Table 3. Other
 1348 sharpness-aware optimizers are shown to achieve similar generalization gain, albeit at the cost of
 1349 higher membership attack accuracy. On the other hand, optimizer that explicitly looks for a sharp
 minima does worse in terms of generalization, but has better membership privacy.

1350
 1351 Table 4: Comparison of **offline** shadow model MIA on SGD and SAM. In **blue** we highlight the best
 1352 performing model on the test set, and in **red** the model with higher privacy leakage (higher AUC,
 1353 Attack Accuracy, and TPR@0.1%FPR).

Dataset	Attack	SGD				SAM			
		Test Acc	AUC	Attack Acc	TPR@.1	Test Acc	AUC	Attack Acc	TPR@.1
CIFAR-100	RMIA	67.7%	86.8%	77.3%	17.3%	69.1%	87.7%	78.1%	18.9%
	LiRA	76.2%	76.2%	71.9%	16.9%		77.8%	73.2%	19.4%
CIFAR-10	RMIA	92.3%	69.4%	62.3%	4.3%	93.1%	72.7%	64.6%	5.7%
	LiRA	54.1%	54.1%	55.9%	4.1%		58.7%	58.3%	7.0%
Purchase100	RMIA	76.5%	68.7%	62.6%	1.7%	77.4%	70.2%	63.6%	1.9%
	LiRA	52.9%	52.9%	53.4%	0.1%		53.7%	54.2%	0.1%
Texas100	RMIA	46.9%	74.8%	67.4%	3.6%	49.2%	78.8%	69.8%	5.6%
	LiRA	56.9%	56.9%	58.3%	0.8%		61.7%	61.5%	2.6%

G OFFLINE SHADOW MODEL ATTACKS

1369 We report the results for offline shadow model attacks in Table Table 4. The results are in line and
 1370 support the finding that SAM tends to incur higher membership privacy leakage. Excluding cases
 1371 for tabular datasets where TPR at 0.1% FPR is near zero for both models, SAM has higher values
 1372 for all other attack metrics. Consistent with the literature, RMIA is more effective for offline setting
 1373 compared to LiRA (Zarifzadeh et al., 2024). For online setting, we use a different experimental setup
 1374 (WideResNet targets and shadows, our own training pipeline, and a different choice of auxiliary
 1375 z-points). In this setting, LiRA is slightly stronger than RMIA across most metrics (Table 3), but
 1376 the gaps are modest and much smaller than those reported for the offline comparison in Zarifzadeh
 1377 et al. (2024). We therefore view our results as broadly consistent with prior work: RMIA and
 1378 LiRA are competitive state-of-the-art shadow-model attacks, and their exact ranking can depend
 1379 on architectural and training choices. Our main conclusions—in particular, that SAM consistently
 1380 increases vulnerability to both LiRA and RMIA compared to SGD—are unaffected by these small
 1381 differences.

H DATASETS

1384 Here we introduce the four benchmark datasets used in the experiments and they have been widely
 1385 used in prior works on MI attacks:

1387 **CIFAR-10** ⁵ This is a benchmark dataset for image classification task. The dataset consists of
 1388 60,000 color images of 32x32 size. There are 6,000 images from 10 classes where 5,000 images per
 1389 class belong to the training dataset and 1,000 images per class belong to the test dataset.

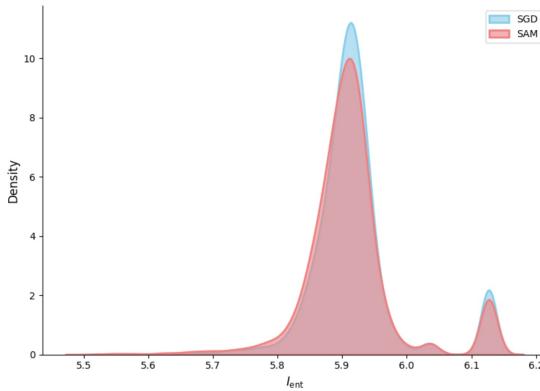
1390 **CIFAR-100** ⁶ The dataset is designed to be more challenging than CIFAR-10 as it contains a greater
 1391 number of classes and more fine-grained distinctions between objects. There are a total of 60,000
 1392 images from 100 classes. Each subclass consists of 600 images, and within each subclass, there are
 1393 500 training images and 100 testing images. This distribution ensures a balanced representation of
 1394 each class in both the training and testing sets.

1396 **Purchase-100** ⁷ This a 100 class classification task with 197,324 data samples and consists of 600
 1397 binary feature; each dimension corresponds to a product and its value states if corresponding customer
 1398 purchased the product; the corresponding label represents the shopping habit of the customer. We
 1399 use the pre-processed and simplified version provided by Shokri et al. (2017) and used by Tang et al.
 1400 (2022).

⁵<https://www.cs.toronto.edu/~kriz/cifar.html>

⁶<https://www.cs.toronto.edu/~kriz/cifar.html>

⁷<https://www.kaggle.com/c/acquire-valued-shoppers-challenge>

Figure 6: \mathcal{I}_{ent} distribution excluding bucket 5 for SGD and SAM

Texas-100⁸ This dataset is based on the Hospital Discharge Data public files with information about inpatients stays in several health facilities released by the Texas Department of State Health Services from 2006 to 2009. We used a prepossessed and simplified version of this dataset provided by (Shokri et al., 2017) and used by (Tang et al., 2022) which is composed of 67,330 data samples with 6,170 binary features. Each feature represents a patient’s medical attribute like the external causes of injury, the diagnosis and other generic information. The classification task is to classify patients into 100 output classes which represent the main procedure that was performed on the patient.

EyePacs⁹ The pre-processed version of this dataset is obtained from Kaggle and it was originally used for a Diabetic Retinopathy Detection challenge. The dataset consists of 88,702 colour fundus images, including 35,126 samples for training and 53,576 samples for testing. The images were captured under various conditions by various devices at multiple primary care sites throughout California and elsewhere. For each subject, two images of the left and right eyes were collected, with the same resolution. A clinician was asked to rate each image for the presence of DR with a scale of 0–4 according to the Early Treatment Diabetic Retinopathy Study (ETDRS) scale. Note that for this dataset only training set (35k images) is used since the labels for testing set is not publicly available. The images in the dataset vary in their image resolution and we resized all the images to 128x128 pixels for our experiments.

I EXPERIMENTAL SETUP

I.1 MODELS

For CIFAR-100 and CIFAR-10, we use WideResNet (WRN) (Zagoruyko & Komodakis, 2016) with 16 layer depth and 8 as width factor. For Purchase-100 and Texas-100, we follow the setting in Tang et al. (2022) and use a 4-layer fully connected neural network with layer sizes [1024, 512, 256, 100]. For EyePacs, we use ResNet-18.

I.2 \mathcal{I}_{ent} EXPERIMENT

Here we discuss how test data points were grouped into 5 buckets according to different \mathcal{I}_{ent} levels. Bucket 5 contains highest \mathcal{I}_{ent} level, and is composed of test points where all 500 training points have 0 influence score. This means that the prediction output for that test point does not change had the model been trained without any one particular training data point. Because influence scores for all training points are equal, these test points have highest \mathcal{I}_{ent} ¹⁰. Figure 6 displays distribution of \mathcal{I}_{ent} for remaining test data points. We group those above 6.1 into bucket 4. For the rest of the points,

⁸<https://www.dshs.texas.gov/THCIC/Hospitals/Download.shtm>.

⁹<https://www.kaggle.com/datasets/mariaherrerot/eyepacspreprocess>

¹⁰When actually calculating \mathcal{I}_{ent} with Equation (5), this evaluates to 0 due to probability normalization, but represents highest value.

1458 we calculate the mean and standard deviation and use them for grouping. We group points below
 1459 -0.4σ from the mean into bucket 1, points between -0.4σ and 0.4σ into bucket 2, and points above
 1460 0.4σ into bucket 3. Final number of test points in each buckets are [Bucket 1: 1924, Bucket 2: 2996,
 1461 Bucket 3: 2392, Bucket 4: 535, Bucket 5: 2153]. For SAM’s buckets, final number of test points
 1462 are [Bucket 1: 1913, Bucket 2: 3181, Bucket 3: 2548, Bucket 4: 502, Bucket 5: 1856]. Number of
 1463 overlapping indices were [Bucket 1: 1116, Bucket 2: 1625, Bucket 3: 1199, Bucket 4: 133, Bucket 5:
 1464 1678].

1465

1466

I.3 ATTACK SETUP & SIZE OF DATA SPLITS

1467
 1468 We adopt the attack setting from (Tang et al., 2022; Nasr et al., 2018) to determine the partition
 1469 between training data and test data and to determine the subset of the training and test data that
 1470 constitutes attacker’s prior knowledge for CIFAR-100, Purchase-100 and Texas-100 datasets. We
 1471 use similar strategy to determine the data split for CIFAR-10. Specifically, the attacker’s knowledge
 1472 corresponds to half of the training and test data, and the MIA success is evaluated over the remaining
 1473 half. For shadow model attacks, the total sample pool used is 50000 for CIFAR10 and CIFAR100,
 1474 40000 for Purchase100, and 20000 for Texas100. For RMIA attack, we used $\gamma = 1$ and selected all
 1475 of the z samples within the training pool that were not part of the target model’s training set. On
 1476 CIFAR10, for example, number of z samples was 25000.

1477

1478

I.4 HYPERPARAMETER TUNING AND EMPIRICAL VALIDATION OF FLATNESS FOR SHARP

1479
 1480 For the Sharp objective, we fine-tuned β and ρ that resulted in a model that exhibited sufficient
 1481 difference in test accuracy and sharpness of the minima compared to SAM and SGD. The final
 1482 hyperparameters of the model reported were $\rho = 0.01, \beta = 0.6818$ for CIFAR-100 and CIFAR-10,
 1483 $\rho = 0.01, \beta = 0.83$ for Purchase-100, $\rho = 0.001, \beta = 0.513$ for Texas-100, and $\rho = 0.001, \beta =$
 1484 0.18 for EyePacs.

1485 To verify that Sharp actually finds a sharper minima, we computed the trace of the hessian matrix
 1486 using Hutchinson’s method for SGD, SAM, and Sharp models on CIFAR-100. The results are in
 1487 Table 5. Higher trace indicates a sharper minima and vice versa. The trace is the largest for Sharp
 1488 and smallest for SAM.

1489

Method	Trace of the Hessian
Sharp	1556.54
SGD	307.87
SAM	84.18

1490

1491 Table 5: Comparison of Hessian trace values across methods.

1492

1493

1494

1495

1496

1497

1498

1499

1500

I.4.1 BALL OF RADIUS ρ

1501

1502

1503

1504

1505

1506

1507

1508

1509

1510

1511

1512

I.4.2 HYPERPARAMETER TUNING FOR CIFAR-10 & CIFAR-100

1513 For SAM loss, sharp minima loss, and our proposed loss, we approximate the maximum loss in the
 1514 ball of radius ρ around the minima. Norton & Royset (2021) have found that the type of norm that is
 1515 used for defining the ball has large impact along with actual ρ value. For all our experiments, we use
 1516 L2 norm for our ball of radius ρ .

1512

1513 Table 6: Privacy vs Generalization tradeoff for SAM and SGD using InceptionV4 and Resnet18

1514	Dataset	Model	Optimizer	Test Acc	Best Attack Acc
1515	CIFAR-100	Resnet18	SGD	78.42%	74.31%
1516			SAM	78.74%	77.45%
1517		InceptionV4	SGD	77.44%	77.22%
1518			SAM	79.60%	80.82%
1519	CIFAR-10	Resnet18	SGD	95.18%	57.90%
1520			SAM	96.16%	60.05%
1521		InceptionV4	SGD	94.26%	61.60%
1522			SAM	95.76%	64.41%
1523					
1524					
1525	I.4.3 HYPERPARAMETER TUNING FOR TEXAS-100 & PURCHASE-100				
1526					
1527	We chose the best model as discussed before for CIFAR-10/100. We trained models with a learning				
1528	rate of 0.1 with weight decay 0.0005 and Nesterov momentum of 0.9. We trained the models on				
1529	Purchase-100 for a total of 100 epochs and on Texas-100 for a total on 75 epochs. During training, we				
1530	employed a batch size of 512 for the Purchase-100 dataset and a batch size of 128 for the Texas-100				
1531	dataset.				
1532					
1533	I.4.4 HYPERPARAMETER TUNING FOR EYEPACS				
1534					
1535	We trained ResNet-18 with SGD, SAM and our proposed loss using EyePacs dataset for 100 epochs.				
1536	Since, the dataset is highly imbalanced with about 25k data points out of 35k training data points				
1537	belonging to one of the five classes, we used the balanced batch sampling strategy and a lower				
1538	learning rate of 0.01 with learning rate decay of 0.2 at 60th epoch. As before, we also used weight				
1539	decay 0.0005 and Nesterov momentum of 0.9. For our experiments, we utilized a batch size of 100,				
1540	consisting of 12 samples from each of the 5 classes.				
1541	J ABLATION STUDY: COMPARISON OF DIFFERENT ARCHITECTURES				
1542					
1543	To validate consistency across different model architectures, we report direct threshold attack results				
1544	in Table 6 using InceptionV4 ¹¹ and resnet18 ¹² for CIFAR-100 and CIFAR-10. We kept our ρ the				
1545	same across all model architectures with value 0.1. The results are consistent with our findings that				
1546	SAM tends to have higher test accuracy while having higher membership attack accuracy at the same				
1547	time. Overall best attack accuracy is higher for SAM for all the cases although we find mixed findings				
1548	for multi-query attack accuracy specifically.				
1549					
1550					
1551					
1552					
1553					
1554					
1555					
1556					
1557					
1558					
1559					
1560					
1561					
1562					
1563					
1564					
1565	¹¹ https://github.com/weiaicunzai/pytorch-cifar100/blob/master/models/				

¹²<https://github.com/inspire-group/MIAdefenseSELENA/tree/main>

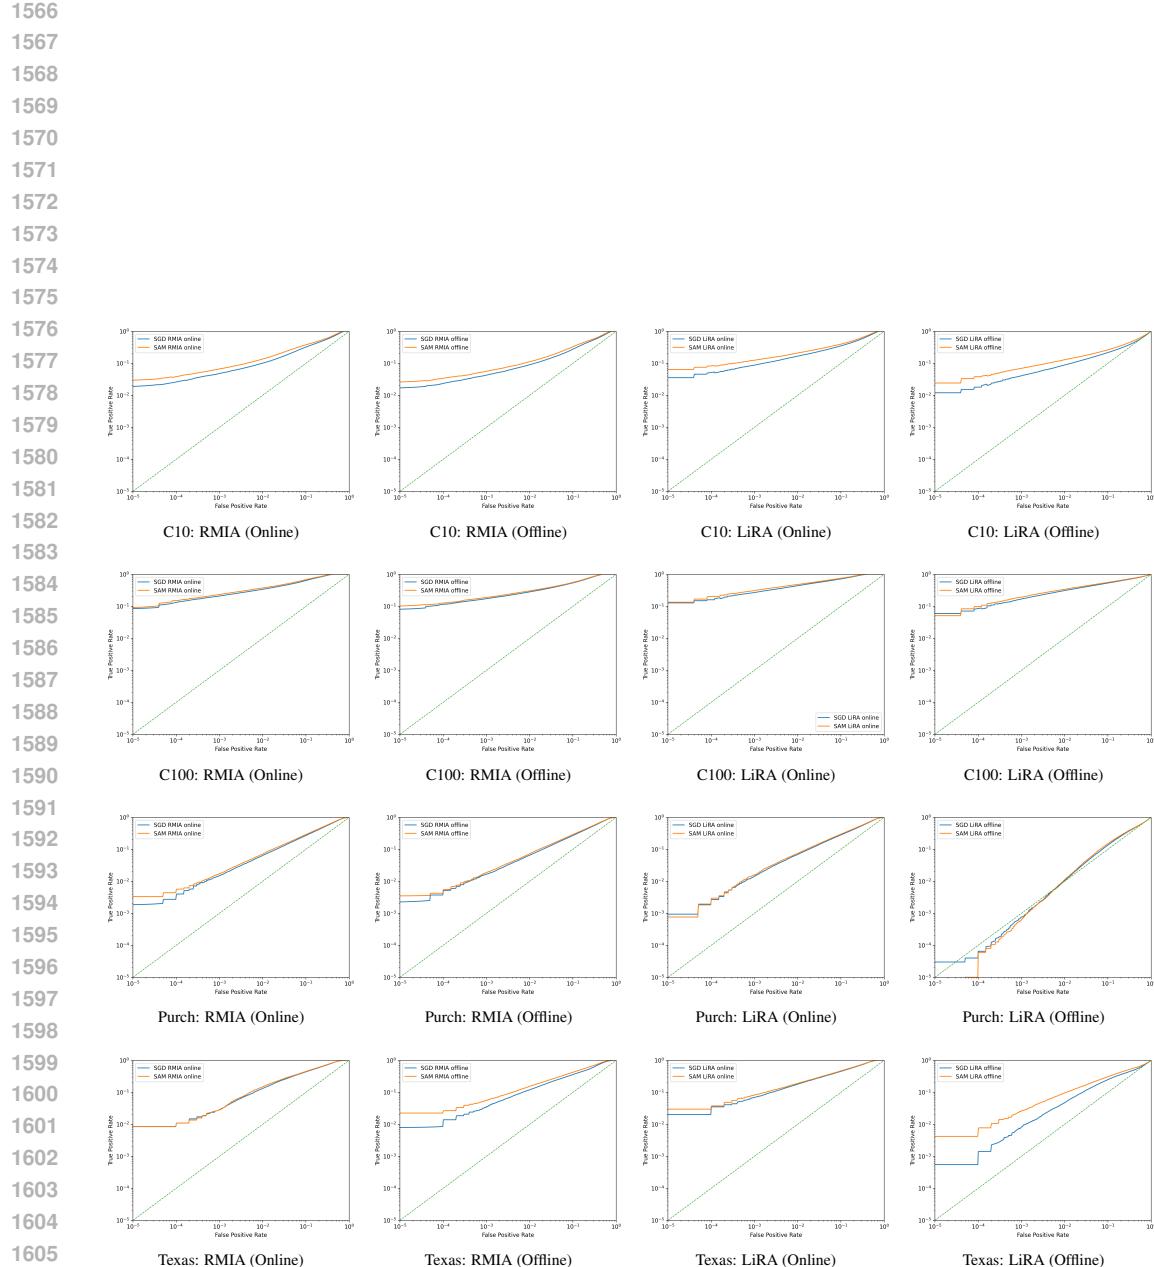


Figure 7: ROC curves comparing SAM (Orange) vs. SGD (Blue) across all datasets and attack modes on log-log scale. Rows represent datasets (CIFAR-10, CIFAR-100, Purchase100, Texas100). Columns represent the attack configuration. The ROC curve for SAM (orange) is above the ROC curve for SGD (blue) for nearly the entire range for most settings.