

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  

# GENERALIZATION AWARE MINIMIZATION

**Anonymous authors**

Paper under double-blind review

## ABSTRACT

Sharpness-Aware Minimization (SAM) optimizers have improved neural network generalization relative to stochastic gradient descent (SGD). The goal of SAM is to steer model parameters away from sharp regions of the training loss landscape, which are believed to generalize poorly. However, the underlying mechanisms of SAM – including whether its bias toward flatter regions is why it improves generalization – are not fully understood. In this work, we introduce Generalization-Aware Minimization (GAM), derived by directly applying the goal of guiding model parameters toward regions of the landscape that generalize better. We do so by showing mathematically through a Bayesian derivation that the landscape of expected true (test) loss is a rescaled version of the observed training loss landscape, and that a sequence of perturbative updates in place of SAM’s single perturbative update can optimize the expected test loss. We present a practical online algorithm to implement GAM’s perturbative steps during training. Finally, we empirically demonstrate that GAM has superior performance over SAM, improving generalization performance on a range of benchmarks. We believe that GAM provides valuable insights into how sharpness-based algorithms improve generalization, is a superior optimizer for generalization, and may inspire the development of still-better optimizers.

## 1 INTRODUCTION

Generalization is a fundamental challenge in training deep neural networks, where the goal is to perform well on unseen data rather than just fitting the training set. One promising approach to enhance generalization is *Sharpness-Aware Minimization* (SAM) (Foret et al., 2021), which has empirically demonstrated success by guiding model parameters away from sharp minima in the training loss landscape. The underlying intuition is that flatter minima correspond to solutions that are less sensitive to perturbations and thus generalize better to new data.

Despite its empirical effectiveness, the theoretical understanding of why SAM improves generalization remains limited. Recent studies have questioned whether SAM’s bias toward flatter regions is the primary reason for its success (Wen et al., 2023). This ambiguity highlights the need for a deeper exploration of the mechanisms through which SAM and similar algorithms enhance generalization. Understanding the mechanisms behind SAM is crucial for developing more effective optimization algorithms that consistently improve generalization across various architectures and datasets.

In this work, we adopt a Bayesian perspective to investigate what the observed training loss landscape reveals about the expected test loss. Analytically, we derive a relationship between the training and test loss landscapes under the assumption of general quadratic loss functions, a reasonable assumption in many conditions. Our analysis reveals that the expected test loss landscape is a rescaled version of the training loss landscape. This rescaling result implies that it may be possible to directly optimize the expected test loss, instead of using the indirect hypothesis that promoting flatness is better for generalization.

Building on this insight, we introduce *Generalization-Aware Minimization* (GAM), a generalization of SAM that employs multiple perturbation steps designed to transform the observed training loss landscape to the rescaled expected test loss landscape. GAM moves beyond SAM’s heuristic of flatness by directly targeting the expected test loss, thereby enabling better generalization. Moreover, we develop a practical online algorithm that adapts the perturbation sizes during training by using

054 the training loss on auxiliary minibatches as a proxy for the test loss and demonstrate superior  
 055 performance.  
 056

057 Our contributions are as follows:

058

- 059 • **Theoretical Insight:** We demonstrate that for quadratic loss functions, the expected test  
 060 loss landscape is a rescaled version of the observed training loss landscape. This provides  
 061 a direct link between training dynamics and generalization performance.
- 062 • **Gradient Transformation:** We show that the gradient of the expected test loss can be  
 063 obtained by evaluating the gradient of the training loss after applying a specific sequence  
 064 of parameter perturbations. This finding bridges the gap between optimizing for training  
 065 loss and directly targeting test loss.
- 066 • **Algorithm Design:** Based on our theoretical results, we propose GAM, an algorithm that  
 067 extends SAM by using multiple perturbation steps with higher-order derivatives and by  
 068 tuning perturbation sizes online during training. This makes GAM practical for use in large-  
 069 scale neural network training. We recover SAM as the one-step perturbation specialization  
 070 of GAM.
- 071 • **Empirical Validation:** We empirically validate GAM on benchmark datasets including  
 072 MNIST, CIFAR-10, SVHN and ImageNet. Our results show that GAM consistently leads  
 073 to better generalization than baselines.

074

## 2 RELATED WORK

### 2.1 SHARPNESS-AWARE MINIMIZATION (SAM) ALGORITHMS

075 Sharpness-Aware Minimization (SAM) algorithms were introduced to improve the generalization of  
 076 neural networks by favoring solutions in flatter regions of the training loss landscape, which have  
 077 empirically been linked to better generalization performance (Foret et al., 2021). SAM perturbs  
 078 model parameters in the direction of the loss gradient and then optimizes using a second gradient  
 079 step, effectively minimizing the sharpness of the loss function. Numerous extensions and variants of  
 080 SAM have since been proposed, focusing on improving computational efficiency and generalization  
 081 (Mi et al., 2022; Liu et al., 2022a;b; Du et al., 2022a;b; Li et al., 2024; Wu et al., 2024).  
 082

083 Despite its empirical success, the theoretical understanding of SAM remains limited and an active  
 084 area of research (Andriushchenko et al., 2023; Zhuang et al., 2022; Chen et al., 2024; Si & Yun,  
 085 2024; Dai et al., 2024). Recent studies have raised questions about whether SAM’s generalization  
 086 improvements stem directly from its bias toward flatter regions of the loss landscape. For instance,  
 087 some works argue that SAM’s effectiveness may not always be directly attributable to sharpness,  
 088 but instead to other implicit regularization effects introduced by the perturbation procedure (Wen  
 089 et al., 2023; Andriushchenko & Flammarion, 2022). Our work builds on this debate by introducing  
 090 a generalized framework that moves beyond the sharpness heuristic and directly targets the expected  
 091 test loss.  
 092

### 2.2 BAYESIAN OPTIMIZATION

093 Bayesian optimization is a well-established framework for optimizing functions that are expensive  
 094 to evaluate, and it has been successfully applied in hyperparameter tuning and low-dimensional op-  
 095 timization problems (Snoek et al., 2012; Frazier, 2018). The fundamental principle of Bayesian  
 096 optimization is to maintain a probabilistic model of the objective function and update it using new  
 097 observations, guiding the search toward areas of the input space that are likely to yield better out-  
 098 comes.  
 099

100 While Bayesian optimization has shown promise in various applications, its applicability to high-  
 101 dimensional settings, such as neural network training, has been limited. Methods that rely on Gaus-  
 102 sian processes or other surrogate models struggle to scale due to the curse of dimensionality and high  
 103 computational costs. Although some efforts have extended Bayesian optimization to use gradient-  
 104 based information for more scalable updates (Wu & Frazier, 2016; Wu et al., 2017; Shekhar &  
 105 Javidi, 2021), these approaches have yet to achieve widespread practical adoption in deep learning  
 106 beyond hyperparameter optimization.  
 107

---

108 

### 3 GENERALIZATION AWARE MINIMIZATION

109

110 In this section, we present our theoretical framework and introduce *Generalization-Aware Minimization* (GAM),
111 a novel optimization algorithm designed to directly improve generalization by aligning
112 the training loss landscape with the expected test loss landscape.
113

114 

#### 3.1 PROBLEM SETUP AND NOTATION

115

116 Consider a parametric model with parameters  $\theta \in \mathbb{R}^d$ . Let  $L(\theta)$  denote the true (test) loss function,
117 which measures the expected loss over the data distribution  $\mathcal{D}$ . In practice, we have access only
118 to the empirical training loss  $\tilde{L}(\theta)$  computed over a finite training dataset sampled from  $\mathcal{D}$ . Our
119 objective is to find the parameter vector  $\theta$  that minimizes  $L(\theta)$ , even though we can only observe
120 and optimize  $\tilde{L}(\theta)$ .
121

122 To formalize our analysis, we consider a quadratic loss function. Note that any general smooth loss
123 landscape can be approximated locally as a quadratic, so we expect our analysis to hold for general
124 losses within a small enough local neighborhood. In what follows, we will derive an optimization
125 rule assuming that the quadratic approximations are local, and therefore may vary over the course of
126 training. Specifically, we consider the true loss function  $L(\theta)$  and the observed training loss function
127  $\tilde{L}(\theta)$  given by:
128

$$L(\theta) = \frac{1}{2}(\theta - \theta^*)^T M(\theta - \theta^*) + c, \quad (1)$$

$$\tilde{L}(\theta) = \frac{1}{2}(\theta - \tilde{\theta}^*)^T \tilde{M}(\theta - \tilde{\theta}^*) + \tilde{c}, \quad (2)$$

131 where  $\theta^*, \tilde{\theta}^* \in \mathbb{R}^d$  represent the minima of the loss functions,  $M, \tilde{M} \in \mathbb{R}^{d \times d}$  are symmetric matrices
132 characterizing the curvature of the loss landscapes, and  $c, \tilde{c} \in \mathbb{R}$  are constants. The parameters
133  $\theta^*, M, c$  of the test loss are unknown, while  $\tilde{\theta}^*, \tilde{M}, \tilde{c}$  of the training loss can be estimated from data.
134 Our goal is to understand how the expected test loss landscape relates to the observed training loss
135 landscape and to devise an optimization strategy that minimizes  $L(\theta)$  by appropriately manipulating
136  $\tilde{L}(\theta)$ .
137

141 

#### 3.2 EXPECTED TEST LOSS LANDSCAPE RESCALES THE TRAINING LOSS LANDSCAPE

142

143 We begin by examining the relationship between the expected test loss landscape and the observed
144 training loss landscape under the assumption of quadratic losses. We show that, under certain conditions,
145 the expected test loss can be expressed as a rescaled version of the training loss.
146

147 The intuition behind this result is that, while the training loss provides an estimate of the true loss,
148 it is subject to sampling variability and noise. By modeling the loss functions as random quadratics,
149 we can analyze how the expected test loss relates to the observed training loss. Specifically, we aim
150 to determine how the curvature (represented by the Hessian matrices) and the minima of the two
151 loss functions are related in expectation.
152

153 **Theorem 1.** *Consider an unknown quadratic loss function:*

$$L(\theta) = \frac{1}{2}(\theta - \theta^*)^T M(\theta - \theta^*) + c \quad (3)$$

154 where  $\theta^* \in \mathbb{R}^d$ ,  $M \in \mathbb{R}^{d \times d}$  and  $c \in \mathbb{R}$  are drawn from a known distribution. Without loss of
155 generality, we assume  $M$  is symmetric. Suppose we observe another random quadratic loss  $\tilde{L}(\theta)$ :
156

$$\tilde{L}(\theta) = \frac{1}{2}(\theta - \tilde{\theta}^*)^T \tilde{M}(\theta - \tilde{\theta}^*) + \tilde{c} \quad (4)$$

157 where  $\tilde{\theta}^*, \tilde{M}$  and  $\tilde{c}$  are random variables dependent on  $\theta^*, M$  and  $c$ . Again, suppose  $\tilde{M}$  is symmetric.
158 Assume,  $\theta^*, \tilde{\theta}^* \perp M, \tilde{M}, c, \tilde{c}$  and  $M \perp \tilde{c} | \tilde{M}$ , where  $\perp$  indicates independence. Furthermore,
159

162 assume  $p_{\theta^*, \tilde{\theta}^*} = p_{\tilde{\theta}^*, \theta^*}$ , where  $p$  denotes probability density. We also assume the following rotation  
163 invariance conditions:

$$164 \quad p_{M|\tilde{M}}(UMU^T|U\tilde{M}U^T) = p_{M|\tilde{M}}(M|\tilde{M}) \quad (5)$$

166 for all orthogonal matrices  $U$  and  $\mathbb{E}[M|\tilde{M}]$  is diagonal when  $\tilde{M}$  is diagonal. Suppose for all  $\theta$ ,

$$168 \quad \mathbb{E}[\tilde{L}(\theta)|\theta^*, M, c] = L(\theta) \quad (6)$$

169 Denote  $\tilde{Q}\tilde{\Lambda}\tilde{Q}^T$  the diagonalization of  $\tilde{M}$  for some diagonal matrix  $\tilde{\Lambda}$  and orthogonal matrix  $\tilde{Q}$ .  
170 Then,

$$172 \quad \mathbb{E}[L(\theta)|\tilde{\theta}^*, \tilde{M}, \tilde{c}] = \frac{1}{2}(\theta - \tilde{\theta}^*)^T \tilde{Q}D(\tilde{\Lambda})\tilde{Q}^T(\theta - \tilde{\theta}^*) + C(\tilde{\theta}^*, \tilde{M}, \tilde{c}) \quad (7)$$

173 for some function  $D$  that outputs a diagonal matrix and function  $C$  outputting a scalar.  
174

175 Please see Appendix A for a proof and Appendix C for justifications of our theoretical assumptions  
176 (including the rotational invariance of conditional distributions and independence assumptions). The  
177 proof involves leveraging the rotational invariance and the independence assumptions to show that  
178 the expected test loss maintains the same eigenvectors as the training loss but with rescaled eigen-  
179 values. This implies that the curvature (Hessian) of the expected test loss is a rescaled version of  
180 that of the training loss, aligned along the same principal directions.

181 Theorem 1 intuitively suggests that the curvature directions of the expected test loss are aligned with  
182 those of the training loss but rescaled in each direction. This rescaling affects the sharpness of the  
183 loss landscape in different directions, suggesting that optimizing for flatter regions in the training  
184 loss (as in SAM algorithms) may not necessarily correspond to optimizing the expected test loss  
185 without additional assumptions.

### 187 3.3 A SERIES OF PERTURBATIONS TRANSFORMS THE TRAINING LOSS TO THE TEST LOSS

189 Having established the relationship between the expected test loss and the training loss, we now  
190 explore how to better approximate the gradient of the test loss using the training loss. We show  
191 that a sequence of perturbations applied to the parameters allows us to transform the training loss  
192 gradient into an approximation of the test loss gradient.

193 The key idea is that higher-order derivatives of the training loss can capture information about the  
194 curvature of the true loss landscape. By recursively computing these derivatives through pertur-  
195 bations, we can construct a series that approximates the effect of rescaling the eigenvalues in the  
196 Hessian of the training loss, effectively transforming it into the Hessian of the test loss.

197 **Theorem 2.** Consider two quadratic loss functions:

$$198 \quad \tilde{L}(\theta) = \frac{1}{2}(\theta - \tilde{\theta}^*)^T \tilde{M}(\theta - \tilde{\theta}^*) + \tilde{c} \quad (8)$$

$$201 \quad \bar{L}(\theta) = \frac{1}{2}(\theta - \tilde{\theta}^*)^T \tilde{Q}f(\tilde{\Lambda})\tilde{Q}^T(\theta - \tilde{\theta}^*) + \bar{c} \quad (9)$$

203 where  $\tilde{M} \in \mathbb{R}^{d \times d}$  is symmetric with eigendecomposition  $\tilde{M} = \tilde{Q}\tilde{\Lambda}\tilde{Q}^T$ , and  $f$  is an elementwise  
204 continuous function satisfying  $f(0) = 0, f'(0) = 1$ . Suppose elements of  $\tilde{\Lambda}$  are bounded. Define  
205  $D^t(\theta) \in \mathbb{R}^d$  recursively as:

$$206 \quad D^1(\theta) = \nabla \tilde{L}(\theta) \quad (10)$$

$$208 \quad D^t(\theta) = \frac{\partial}{\partial \zeta} D^1(\theta + \zeta D^{t-1}(\theta))|_{\zeta=0} \quad (11)$$

210 for  $t > 1$ . Then, for all  $\epsilon > 0$ , there exists a sequence  $\gamma_1, \gamma_2, \dots, \gamma_T \in \mathbb{R}$  such that:

$$212 \quad \|\nabla \bar{L}(\theta) - \nabla \tilde{L}(\hat{\theta})\| \leq \epsilon \|\theta - \tilde{\theta}^*\| \quad (12)$$

213 where

$$215 \quad \hat{\theta} = \theta + \sum_{t=1}^T \gamma_t D^t(\theta) \quad (13)$$

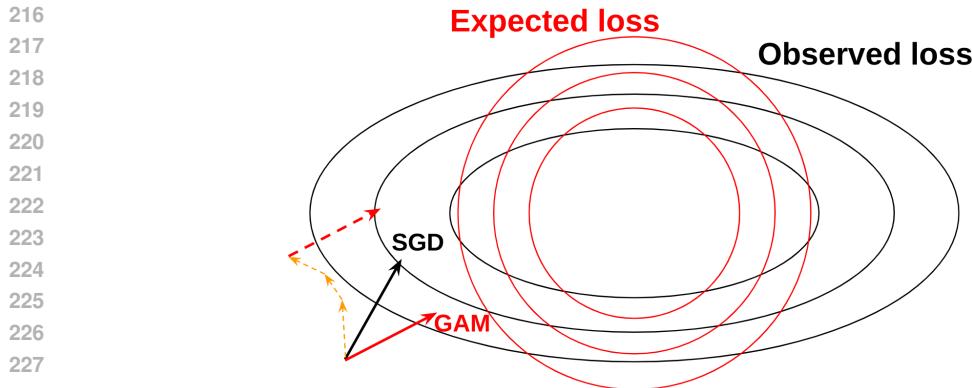


Figure 1: Schematic diagram of generalization aware minimization (GAM) vs stochastic gradient descent (SGD). Black and red contour lines indicate the observed and expected loss landscapes. Note that the expected loss landscape corresponds to the observed loss landscape with rescaled contour lines. SGD takes a gradient step directly against the gradient of the observed loss (black arrow). GAM first makes a series of perturbations in the parameters (orange dashed arrows), computes the update step from the observed loss at the perturbed value (red dashed arrow), and applies the step at the original parameter value (red arrow). This direction corresponds to gradient descent on the expected loss.

Please see Appendix B for a proof. The proof constructs the perturbation coefficients  $\gamma_t$  to approximate the effect of the function  $f$  on the eigenvalues of  $\tilde{M}$ . By iteratively applying directional derivatives along the sequence  $D^t(\theta)$ , we adjust the gradient of the training loss to closely match that of the transformed loss  $\bar{L}(\theta)$ . See Appendix C for justifications of our theoretical assumptions (including assumptions on the function  $f$ ). Notably, we assume the eigenvalue transformation  $f$  is elementwise, differing from the more general eigenvalue transformation derived in Theorem 1; in essence, this assumes orthogonal directions in the parameter space can be treated independently.

Theorem 2 provides a method to approximate the gradient of the expected test loss by applying a series of specific perturbations to the parameters and computing higher-order derivatives of the training loss. This result suggests that we can design an optimization algorithm that adjusts the parameter updates based on these perturbations to directly minimize the expected test loss.

### 3.4 A PRACTICAL ONLINE ALGORITHM TO LEARN PERTURBATIONS

We now extend the theoretical approach to general loss functions by considering local quadratic approximations. We propose a practical algorithm that continually adapts the perturbation sizes over the course of training to improve generalization.

**Algorithm Design** Based on Theorem 2, we design the *Generalization-Aware Minimization* (GAM) algorithm. GAM uses multiple perturbation steps with higher-order derivatives and tunes the perturbation coefficients  $\gamma_t$  online during training. This allows GAM to approximate gradients of the expected test loss instead of using training loss gradients as illustrated in Figure 1. The key components of GAM are:

- **Higher-Order Perturbations:** We compute a sequence of directional derivatives  $D^t(\theta)$  to capture higher-order information about the loss landscape.
- **Adaptive Perturbation Sizes:** We update the perturbation coefficients  $\gamma_t$  by minimizing a discrepancy function  $\Delta$  between the gradient on perturbed parameters and an estimate of the test loss gradient; in practice, we use negative dot product as our discrepancy measure.
- **Auxiliary Minibatches:** We use the training loss on auxiliary minibatches as a proxy for the test loss to guide the adaptation of  $\gamma_t$ .

Notably, if the number of perturbation steps is fixed at one and  $\gamma_1$  is fixed at a constant value, we recover SAM exactly: thus, SAM is a special case of GAM.

270 **Algorithm Details** Algorithm 1 outlines the steps of GAM.  
 271

---

272 **Algorithm 1** Generalization Aware Minimization

273 **Require:** Initial parameters  $\theta^0$ , training set  $\mathcal{D}$ , GAM steps  $T$ , training iterations  $N$ , gradient dis-  
 274 crepancy function  $\Delta$ , small constant  $\epsilon > 0$   
 275 1: Initialize  $\gamma_1, \gamma_2, \dots, \gamma_T = 0$   
 276 2: Initialize  $\theta = \theta^0$   
 277 3: Sample minibatch  $\bar{X}, \bar{Y} \sim \mathcal{D}$   
 278 4: **for** iteration = 1, ...,  $N$  **do**  
 279 5:   Sample minibatch  $X, Y \sim \mathcal{D}$   
 280 6:    $d^1 = \nabla L(\theta, (X, Y))$   
 281 7:   **for**  $t = 2, \dots, T$  **do**  
 282 8:      $d^t = \frac{1}{\epsilon}(\nabla L(\theta + \epsilon d^{t-1}, (X, Y)) - d^1)$   
 283 9:   **end for**  
 284 10:    $\hat{\theta} = \theta + \sum_{t=1}^T \gamma_t d^t$   
 285 11:    $g_\theta = \nabla L(\hat{\theta}, (X, Y))$   
 286 12:    $\bar{g}_\theta = \nabla L(\theta, (\bar{X}, \bar{Y}))$   
 287 13:    $g_\gamma = \frac{\partial}{\partial \gamma_1, \gamma_2, \dots, \gamma_T} \Delta(g_\theta, \bar{g}_\theta)$   
 288 14:   Update  $\gamma$  following  $-g_\gamma$   
 289 15:   Update  $\theta$  following  $-g_\theta$   
 290 16: **end for**  
 291 17: Return  $\theta$

---

292  
 293 **Explanation of Steps**

294

- 295 • **Lines 6–9 (Higher-Order Derivatives):** We recursively compute the directional derivatives  $d^t$  using finite differences. The small constant  $\epsilon$  ensures numerical stability.
- 296 • **Line 10 (Perturbed Parameters):** We obtain the perturbed parameter vector  $\hat{\theta}$  by combin-  
 297 ing the directional derivatives weighted by the coefficients  $\gamma_t$ .
- 298 • **Lines 11–12 (Gradient Computation):** We compute the gradient at the perturbed param-  
 299 eters  $g_\theta$  and the gradient on the auxiliary minibatch  $\bar{g}_\theta$ , which serves as a proxy for the test  
 300 loss gradient.
- 301 • **Line 13 (Perturbation Coefficient Update):** We compute the gradient of the discrepancy  
 302 between  $g_\theta$  and  $\bar{g}_\theta$  with respect to the perturbation coefficients and update  $\gamma_t$  accordingly.
- 303 • **Line 14 (Parameter Update):** We perform a standard gradient descent update on the pa-  
 304 rameters using  $g_\theta$ .

305  
 306 **Practical Considerations**

307

- 308 • **Computational Overhead:** Computing higher-order derivatives increases computational  
 309 cost. However, since we use finite differences and a small number of perturbation steps  $T$ ,  
 310 the overhead remains manageable.
- 311 • **Stability and Convergence:** The choice of  $\epsilon$  and the learning rates  $\eta_\gamma, \eta_\theta$  can affect the  
 312 stability of the algorithm. In practice, these hyperparameters are tuned based on validation  
 313 performance.
- 314 • **Extension to Non-Quadratic Losses:** While theoretically motivated for quadratic losses,  
 315 GAM can be applied to general loss functions by assuming local quadraticity. This allows  
 316 GAM to be used with complex neural networks and loss functions encountered in deep  
 317 learning.

318  
 319 **Empirical Validation** To empirically validate that GAM effectively learns the perturbation coef-  
 320 ficients  $\gamma_t$ , we conduct experiments on a synthetic quadratic optimization problem where the exact  
 321 relationship between the observed (training) loss and the expected (test) loss is known. In Ap-  
 322 pendix D, we find that GAM can indeed approximate the correct transformation mapping from the  
 323 training loss landscape to the expected test loss landscape.

324  
 325 Table 1: Test set accuracies of various network architectures trained on the MNIST, CIFAR-10  
 326 and SVHN datasets with different methods: stochastic gradient descent (SGD), sharpness aware  
 327 minimization (SAM) with different parameter values  $\gamma_1$ , and generalization aware minimization  
 328 (GAM). Mean results and standard errors are reported over 5 trials where applicable. Best results  
 329 are bolded.

330 Method	331 MNIST			332 CIFAR-10			333 SVHN			334 Imagenet
	335 3-layer MLP	336 3-layer CNN	337 3-layer MLP	338 5-layer CNN	339 14-layer CNN	340 3-layer MLP	341 5-layer CNN	342 14-layer CNN	343 ResNet-50	
SGD	0.97368 ± 0.00072	0.96216 ± 0.00117	0.54500 ± 0.00216	0.66334 ± 0.00104	0.83430 ± 0.00237	0.79729 ± 0.00223	0.87824 ± 0.00055	0.93852 ± 0.00084	0.64928	
SAM 0.001	0.97356 ± 0.00082	0.96224 ± 0.00111	0.54552 ± 0.00241	0.66716 ± 0.00123	0.83138 ± 0.00231	0.79909 ± 0.00246	0.87919 ± 0.00033	0.93937 ± 0.00069	-	
SAM 0.01	0.97352 ± 0.00032	0.96244 ± 0.00117	0.54404 ± 0.00216	0.66974 ± 0.00232	0.83328 ± 0.00213	0.80234 ± 0.00169	0.87853 ± 0.00033	0.94156 ± 0.00048	-	
SAM 0.1	0.97466 ± 0.00046	0.96298 ± 0.00088	0.55394 ± 0.0017	0.68100 ± 0.0014	0.84288 ± 0.00345	0.80627 ± 0.0022	0.88334 ± 0.00095	0.94252 ± 0.00143	0.65232	
CRSAM	0.97038 ± 0.00149	0.93852 ± 0.01727	0.55402 ± 0.00148	<b>0.69444</b> ± 0.00117	<b>0.85892</b> ± 0.00059	0.80181 ± 0.00221	<b>0.89792</b> ± 0.00119	<b>0.95383</b> ± 0.00113	-	
GAM	<b>0.97518</b> ± 0.00063	<b>0.96392</b> ± 0.00036	<b>0.56356</b> ± 0.00162	0.69396 ± 0.00154	0.85074 ± 0.00122	<b>0.81175</b> ± 0.00217	0.88357 ± 0.00057	0.94299 ± 0.00025	<b>0.65924</b>	

## 344 4 RESULTS

### 345 4.1 EXPERIMENTAL SETUP

346 In this section, we evaluate the performance of Generalization-Aware Minimization (GAM) in  
 347 comparison to Sharpness-Aware Minimization (SAM), Curvature-regularized SAM (CRSAM) Wu et al.  
 348 (2024), and standard stochastic gradient descent (SGD) on standard image classification bench-  
 349 marks: MNIST (Deng, 2012), CIFAR-10 (Krizhevsky et al., 2009), SVHN (Netzer et al., 2011) and  
 350 ImageNet (Deng et al., 2009).

351 For SAM, which corresponds to a special case of GAM with  $T = 1$  and fixed perturbation size  $\gamma_1$ ,  
 352 we experiment with perturbation magnitudes  $\gamma_1 \in \{0.001, 0.01, 0.1\}$  to evaluate its sensitivity to  
 353 this hyperparameter. GAM adaptively learns the perturbation coefficients  $\gamma_t$  during training using  
 354 multiple perturbation steps ( $T > 1$ ) and higher-order gradient information.

355 All methods are trained using the same optimization settings, including learning rates, batch sizes,  
 356 and training epochs, to ensure a fair comparison. Detailed architectures, hyperparameter settings  
 357 and training procedures are provided in Appendix E. Appendix F Figure 6 shows the performance of  
 358 GAM under different hyperparameter settings; in summary, we find that GAM can become unstable  
 359 under large  $T$ , performs best at small batch sizes, and is relatively insensitive to  $\epsilon$ .

### 360 4.2 GAM OUTPERFORMS SAM ON BENCHMARKS

361 Table 1 presents the test accuracies achieved by each method across different datasets and network  
 362 architectures. The results demonstrate that GAM consistently and statistically significantly outper-  
 363 forms both SAM and standard SGD on all benchmarks. For instance, for ResNet-50 trained on  
 364 ImageNet, GAM achieves a test accuracy of **65.92%**, surpassing SAM’s performance of 65.23%.  
 365 Nevertheless, on certain architectures, CRSAM outperforms GAM by a notable margin indicating  
 366 GAM may not be universally better than all SAM variants.

367 Appendix F Figure 5 shows the test error over the course of training for each method. Notably, GAM  
 368 may underperform relative to other methods during the early stages of training as it learns the  
 369 optimal perturbation coefficients  $\gamma_t$ . However, as training progresses, GAM adjusts these coefficients  
 370 effectively, leading to superior generalization performance by the end of training. This adaptive  
 371 learning of perturbations allows GAM to fine-tune its optimization strategy based on the evolving  
 372 loss landscape. These results suggest that GAM’s ability to adaptively learn perturbation sizes and  
 373 use higher-order gradient information contributes to its enhanced generalization performance across  
 374 different models and datasets.

### 375 4.3 ANALYZING THE TRANSFORMATION FROM TRAINING TO TEST LOSS

376 To gain further insight into how GAM improves generalization, we analyze the transformation ap-  
 377 plied by GAM to the loss landscape. Specifically, we examine how GAM modifies the eigenvalues

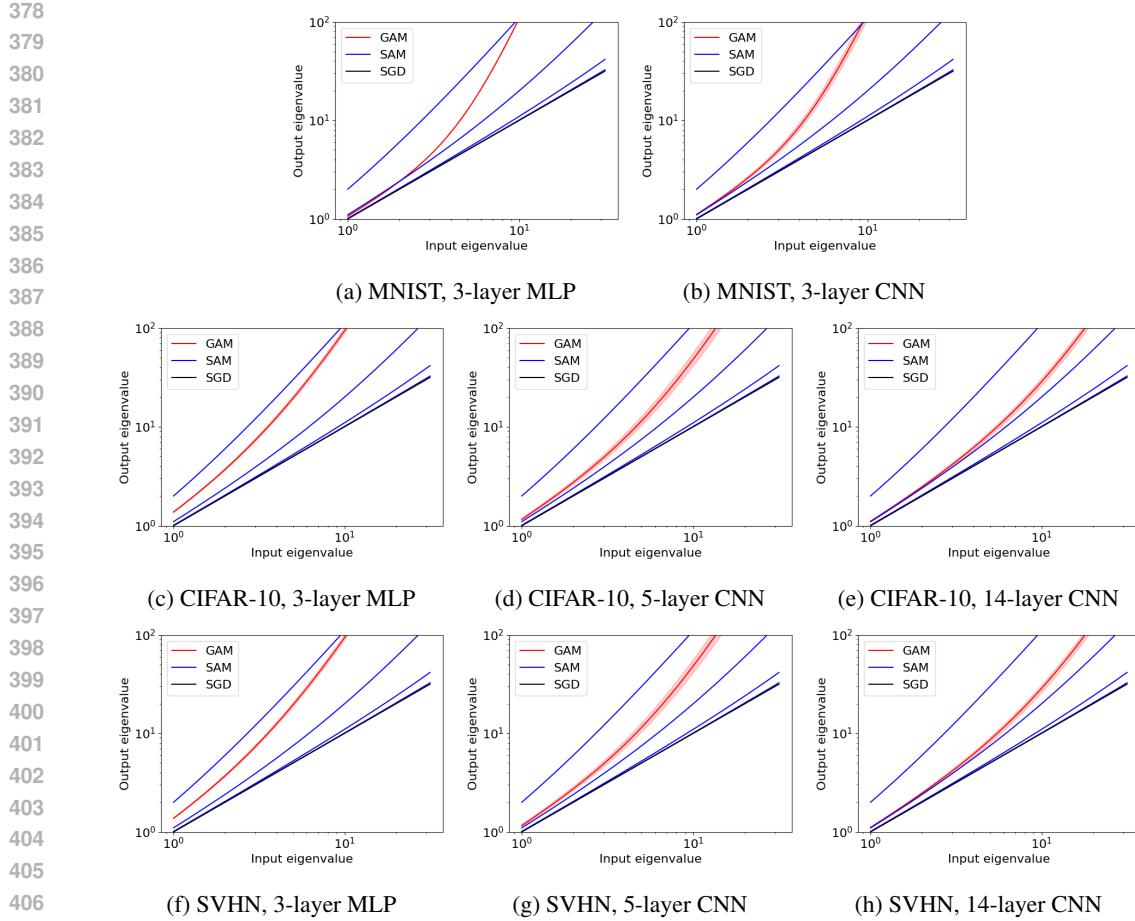


Figure 2: Visualization for different dataset-architecture combinations of Hessian eigenvalue transformation of different training methods: stochastic gradient descent (SGD), sharpness aware minimization (SAM) with different parameter values  $\gamma_1$ , and generalization aware minimization (GAM). Input eigenvalue corresponds to the observed loss Hessian while the output eigenvalue corresponds to the transformed loss Hessian. Margins for GAM indicate standard errors over 5 trials.

of the Hessian of the training loss, which correspond to the curvature along different parameter directions. Figure 2 visualizes the relationship between the original Hessian eigenvalues (input eigenvalues) and the transformed eigenvalues (output eigenvalues) for different methods. For GAM, we compute the effective transformation induced by the learned perturbation coefficients  $\gamma_t$ . We observe that GAM tends to sharpen already sharp directions (i.e., directions with large eigenvalues), qualitatively similar to SAM. As GAM is tuned specifically to optimize for generalization, the results suggest that SAM generalizes well because of its similarity to GAM. Interestingly, GAM’s transformation exhibits a higher contrast between small and large curvature directions than SAM, selectively sharpening sharp directions while maintaining others. This more complex behavior arises from the higher-order gradient information used by GAM, ultimately yielding better generalization.

#### 4.4 MITIGATING GAM’S COMPUTATIONAL COST

One of GAM’s disadvantages is its computational cost relative to SGD and SAM: it requires computing the derivative of parameter updates with respect to perturbation coefficients  $\gamma_t$  which can be quite costly. We propose mitigating this cost by updating the  $\gamma_t$  periodically instead of at each training step as done in Algorithm 1. As shown in Figure 3, when updating the  $\gamma_t$  every time step, GAM’s computational cost is roughly  $4\times$  that of SGD (relative to roughly  $1.3\times$  for SAM). However, this cost can be reduced to roughly  $3\times$  when updating the perturbation coefficients periodically. Although this reduces test accuracy, the accuracy of GAM still exceeds that of SAM or SGD.

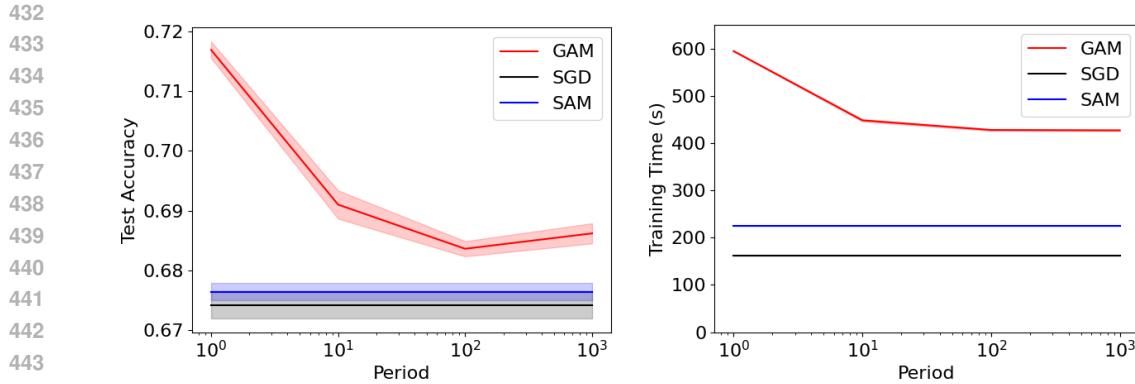


Figure 3: Test set accuracies and training time of a 5-layer CNN model trained on the CIFAR-10 dataset with stochastic gradient descent (SGD), sharpness aware minimization (SAM) with  $\gamma_1 = 0.01$ , and generalization aware minimization (GAM). The x-axis indicates the period at which GAM updates its perturbation coefficients. Mean results and standard errors are reported over 5 trials. Training times are on an NVIDIA RTX 4090 GPU.

## 5 DISCUSSION

In this work, we introduced Generalization-Aware Minimization (GAM), a novel optimization algorithm that directly targets the expected test loss by using higher-order gradient information and adaptive perturbations. Unlike Sharpness-Aware Minimization (SAM) algorithms, which rely on the heuristic that flatter regions of the loss landscape generalize better, GAM is grounded in a theoretical framework that aligns the optimization process with the expected test loss. By demonstrating that the expected test loss landscape is a rescaled version of the observed training loss landscape for quadratic losses, we provided a principled approach to improve generalization.

The surprising similarity between the update mechanisms of SAM and GAM offers a new perspective on why SAM improves generalization. Our analysis suggests that SAM may implicitly approximate the expected test loss through its single-step perturbations, which could explain its empirical success. However, GAM’s use of multiple perturbation steps and higher-order derivatives allows it to more accurately capture the transformation between the training and test loss landscapes. Our empirical results on benchmark datasets confirm that GAM consistently outperforms SAM, highlighting the benefits of directly optimizing for generalization.

GAM shows promising results in real settings, however it relies on using higher-order derivatives, which may be computationally challenging for large or non-differentiable networks. Future work could explore approximations or scalable implementations of higher-order derivatives. We also note the possibility of integrating GAM with newer variants of SAM such as CRSAM which in some cases can outperform GAM.

We believe that GAM opens new avenues for developing optimization algorithms that can further enhance generalization in deep learning models, potentially leading to more robust and reliable AI systems. By incorporating higher-order gradient information and adaptive strategies, future optimizers can more effectively navigate the loss landscape to find solutions that generalize well. We hope that our work inspires further research into optimization techniques that bridge theoretical insights and practical performance, ultimately contributing to the advancement of generalization in machine learning.

## REFERENCES

MakSYM Andriushchenko and Nicolas Flammarion. Towards understanding sharpness-aware minimization. In *ICML*, pp. 639–668. PMLR, 2022.

MakSYM Andriushchenko, Dara Bahri, Hossein Mobahi, and Nicolas Flammarion. Sharpness-aware minimization leads to low-rank features. In *NeurIPS*, volume 36, pp. 47032–47051, 2023.

486 Zixiang Chen, Junkai Zhang, Yiwen Kou, Xiangning Chen, Cho-Jui Hsieh, and Quanquan Gu. Why  
 487 does sharpness-aware minimization generalize better than sgd? In *NeurIPS*, volume 36, 2024.  
 488

489 Yan Dai, Kwangjun Ahn, and Suvrit Sra. The crucial role of normalization in sharpness-aware  
 490 minimization. In *NeurIPS*, volume 36, 2024.

491 Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hi-  
 492 erarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*,  
 493 pp. 248–255. Ieee, 2009.

494

495 Li Deng. The mnist database of handwritten digit images for machine learning research. *IEEE  
 496 Signal Processing Magazine*, 29(6):141–142, 2012.

497

498 Jiawei Du, Hanshu Yan, Jiashi Feng, Joey Tianyi Zhou, Liangli Zhen, Rick Siow Mong Goh, and  
 499 Vincent YF Tan. Efficient sharpness-aware minimization for improved training of neural net-  
 500 works. In *ICLR*, 2022a.

501 Jiawei Du, Daquan Zhou, Jiashi Feng, Vincent Tan, and Joey Tianyi Zhou. Sharpness-aware training  
 502 for free. In *NeurIPS*, volume 35, pp. 23439–23451, 2022b.

503

504 Pierre Foret, Ariel Kleiner, Hossein Mobahi, and Behnam Neyshabur. Sharpness-aware minimiza-  
 505 tion for efficiently improving generalization. In *ICLR*, 2021.

506

507 Peter I Frazier. A tutorial on bayesian optimization. *arXiv preprint*, 2018.

508

509 Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recog-  
 510 nition, 2016.

511

512 Tao Li, Pan Zhou, Zhengbao He, Xinwen Cheng, and Xiaolin Huang. Friendly sharpness-aware  
 513 minimization. In *CVPR*, pp. 5631–5640, 2024.

514

515 Yong Liu, Siqi Mai, Xiangning Chen, Cho-Jui Hsieh, and Yang You. Towards efficient and scalable  
 516 sharpness-aware minimization. In *CVPR*, pp. 12360–12370, June 2022a.

517

518 Yong Liu, Siqi Mai, Minhao Cheng, Xiangning Chen, Cho-Jui Hsieh, and Yang You. Random  
 519 sharpness-aware minimization. In *NeurIPS*, volume 35, pp. 24543–24556. Curran Associates,  
 520 Inc., 2022b. URL [https://proceedings.neurips.cc/paper\\_files/paper/2022/file/9b79416c0dc4b09fea169ed5cdd63d4-Paper-Conference.pdf](https://proceedings.neurips.cc/paper_files/paper/2022/file/9b79416c0dc4b09fea169ed5cdd63d4-Paper-Conference.pdf).

521

522 Peng Mi, Li Shen, Tianhe Ren, Yiyi Zhou, Xiaoshuai Sun, Rongrong Ji, and Dacheng  
 523 Tao. Make sharpness-aware minimization stronger: A sparsified perturbation ap-  
 524 proach. In *NeurIPS*, volume 35, pp. 30950–30962. Curran Associates, Inc., 2022. URL  
 525 [https://proceedings.neurips.cc/paper\\_files/paper/2022/file/c859b99b5d717c9035e79d43dfd69435-Paper-Conference.pdf](https://proceedings.neurips.cc/paper_files/paper/2022/file/c859b99b5d717c9035e79d43dfd69435-Paper-Conference.pdf).

526

527 Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bissacco, Baolin Wu, Andrew Y Ng, et al.  
 528 Reading digits in natural images with unsupervised feature learning. In *NIPS workshop on deep  
 529 learning and unsupervised feature learning*, pp. 4. Granada, 2011.

530

531 Shubhanshu Shekhar and Tara Javidi. Significance of gradient information in bayesian optimization.  
 532 In *ICAIS*, pp. 2836–2844. PMLR, 2021.

533

534 Dongkuk Si and Chulhee Yun. Practical sharpness-aware minimization cannot converge all the way  
 535 to optima. In *NeurIPS*, volume 36, 2024.

536

537 Jasper Snoek, Hugo Larochelle, and Ryan P Adams. Practical bayesian optimization of machine  
 538 learning algorithms. *NeurIPS*, 25, 2012.

540 Karl Weierstrass. Über die analytische darstellbarkeit sogenannter willkürlicher functionen einer  
 541 reellen veränderlichen. *Sitzungsberichte der Königlich Preußischen Akademie der Wissenschaften*  
 542 zu Berlin, 2(633-639):364, 1885.

544 Kaiyue Wen, Tengyu Ma, and Zhiyuan Li. How does sharpness-aware minimization minimize  
 545 sharpness? In *ICLR*, 2023.

546 Jian Wu and Peter Frazier. The parallel knowledge gradient method for batch bayesian optimization.  
 547 *NeurIPS*, 29, 2016.

549 Jian Wu, Matthias Poloczek, Andrew G Wilson, and Peter Frazier. Bayesian optimization with  
 550 gradients. *NeurIPS*, 30, 2017.

551 Tao Wu, Tie Luo, and Donald C Wunsch II. Cr-sam: Curvature regularized sharpness-aware min-  
 552 imization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pp.  
 553 6144–6152, 2024.

555 Juntang Zhuang, Boqing Gong, Liangzhe Yuan, Yin Cui, Hartwig Adam, Nicha Dvornek, Sekhar  
 556 Tatikonda, James Duncan, and Ting Liu. Surrogate gap minimization improves sharpness-aware  
 557 training. In *ICLR*, 2022.

## 560 A PROOF OF THEOREM 1

562 *Proof.* First, observe that

$$564 \mathbb{E}[L(\theta)|\tilde{\theta}^*, \tilde{M}, \tilde{c}] = \mathbb{E}\left[\frac{1}{2}(\theta - \theta^*)^T M(\theta - \theta^*) + c|\tilde{\theta}^*, \tilde{M}, \tilde{c}\right] \\ 565 = \frac{1}{2}\theta^T \mathbb{E}[M|\tilde{\theta}^*, \tilde{M}, \tilde{c}]\theta - \theta^T \mathbb{E}[M\theta^*|\tilde{\theta}^*, \tilde{M}, \tilde{c}] + \mathbb{E}[c|\tilde{\theta}^*, \tilde{M}, \tilde{c}] \quad (14)$$

568 Since  $\theta^*$  and  $\tilde{\theta}^*$  are independent of  $\tilde{M}$  and  $\tilde{c}$ , and  $M$  is independent of  $\tilde{\theta}^*$ , we have:

$$570 = \frac{1}{2}\theta^T \mathbb{E}[M|\tilde{M}, \tilde{c}]\theta - \theta^T \mathbb{E}[M|\tilde{M}, \tilde{c}]\mathbb{E}[\theta^*|\tilde{\theta}^*] + \mathbb{E}[c|\tilde{\theta}^*, \tilde{M}, \tilde{c}] \quad (15)$$

572 Factoring:

$$574 = \frac{1}{2}(\theta - \mathbb{E}[\theta^*|\tilde{\theta}^*])^T \mathbb{E}[M|\tilde{M}, \tilde{c}](\theta - \mathbb{E}[\theta^*|\tilde{\theta}^*]) + \frac{1}{2}\mathbb{E}[\theta^*|\tilde{\theta}^*]^T \mathbb{E}[M|\tilde{M}, \tilde{c}]\mathbb{E}[\theta^*|\tilde{\theta}^*] + \mathbb{E}[c|\tilde{\theta}^*, \tilde{M}, \tilde{c}] \quad (16)$$

576 Letting  $C(\tilde{\theta}^*, \tilde{M}, \tilde{c}) = \frac{1}{2}\mathbb{E}[\theta^*|\tilde{\theta}^*]^T \mathbb{E}[M|\tilde{M}, \tilde{c}]\mathbb{E}[\theta^*|\tilde{\theta}^*] + \mathbb{E}[c|\tilde{\theta}^*, \tilde{M}, \tilde{c}]$ :

$$578 \mathbb{E}[L(\theta)|\tilde{\theta}^*, \tilde{M}, \tilde{c}] = \frac{1}{2}(\theta - \mathbb{E}[\theta^*|\tilde{\theta}^*])^T \mathbb{E}[M|\tilde{M}, \tilde{c}](\theta - \mathbb{E}[\theta^*|\tilde{\theta}^*]) + C(\tilde{\theta}^*, \tilde{M}, \tilde{c}) \quad (17)$$

580 Next, we use the fact that  $\mathbb{E}[\tilde{L}(\theta)|\theta^*, M, c] = L(\theta)$ . Expanding using the definition of  $\tilde{L}(\theta)$  and  
 581  $L(\theta)$ :

$$583 \frac{1}{2}(\theta - \theta^*)^T M(\theta - \theta^*) + c = \frac{1}{2}\theta^T \mathbb{E}[\tilde{M}|\theta^*, M, c]\theta - \theta^T \mathbb{E}[\tilde{M}\theta^*|\theta^*, M, c] + \mathbb{E}[\tilde{c}|\theta^*, M, c] \quad (18)$$

585 Once again using the independence between  $\theta^*$  and  $\tilde{\theta}^*$  from  $M$  and  $c$ , and the independence of  $\tilde{M}$   
 586 is  $\theta^*$ :

$$588 \frac{1}{2}(\theta - \theta^*)^T M(\theta - \theta^*) + c = \frac{1}{2}\theta^T \mathbb{E}[\tilde{M}|M, c]\theta - \theta^T \mathbb{E}[\tilde{M}|M, c]\mathbb{E}[\tilde{\theta}^*|\theta^*] + \mathbb{E}[\tilde{c}|\theta^*, M, c] \quad (19)$$

590 Since this holds for all  $\theta$ , we may equate coefficients:

$$592 M = \mathbb{E}[\tilde{M}|M, c] \quad (20)$$

$$593 \theta^* = \mathbb{E}[\tilde{\theta}^*|\theta^*] \quad (21)$$

594 Next, note that  $p_{\theta^*, \tilde{\theta}^*} = p_{\tilde{\theta}^*, \theta^*}$  implies that  $\theta^*$  and  $\tilde{\theta}^*$  have the same marginal distributions, and  
 595 same conditional distributions conditioned on each other. Since  $\theta^* = \mathbb{E}[\tilde{\theta}^* | \theta^*]$ , by symmetry, we  
 596 must have:  
 597

$$\tilde{\theta}^* = \mathbb{E}[\theta^* | \tilde{\theta}^*] \quad (22)$$

598 Thus, we may write the expectation of  $L(\theta)$  as:  
 599

$$\mathbb{E}[L(\theta) | \tilde{\theta}^*, \tilde{M}, \tilde{c}] = \frac{1}{2}(\theta - \tilde{\theta}^*)^T \mathbb{E}[M | \tilde{M}, \tilde{c}] (\theta - \tilde{\theta}^*) + C(\tilde{\theta}^*, \tilde{M}, \tilde{c}) \quad (23)$$

600 Next, we consider  $\mathbb{E}[M | \tilde{M}, \tilde{c}]$ . Since  $M \perp \tilde{c} | \tilde{M}$ , we have  $\mathbb{E}[M | \tilde{M}, \tilde{c}] = \mathbb{E}[M | \tilde{M}]$ . Expanding:  
 601

$$\mathbb{E}[M | \tilde{M}, \tilde{c}] = \sum_M M p_{M | \tilde{M}}(M | \tilde{M}) \quad (24)$$

602 We denote the eigendecomposition of  $\tilde{M} = \tilde{Q} \tilde{\Lambda} \tilde{Q}^T$ . Note that since  $\tilde{M}$  is symmetric,  $\tilde{Q}$  is orthogonal.  
 603 Substituting:  
 604

$$\mathbb{E}[M | \tilde{M}, \tilde{c}] = \sum_M M p_{M | \tilde{M}}(M | \tilde{Q} \tilde{\Lambda} \tilde{Q}^T) \quad (25)$$

605 By the rotation invariance of  $p_{M | \tilde{M}}$ , we have:  
 606

$$\mathbb{E}[M | \tilde{M}, \tilde{c}] = \sum_M M p_{M | \tilde{M}}(\tilde{Q}^T M \tilde{Q} | \tilde{\Lambda}) \quad (26)$$

607 Making a change of variables in the summation,  $M' = \tilde{Q}^T M \tilde{Q}$ :  
 608

$$\mathbb{E}[M | \tilde{M}, \tilde{c}] = \tilde{Q} \left[ \sum_{M'} M' p_{M | \tilde{M}}(M' | \tilde{\Lambda}) \right] \tilde{Q}^T \quad (27)$$

609 Note that the term in the brackets is simply  $\mathbb{E}[M | \tilde{\Lambda}]$  which is diagonal by assumption. Thus, for  
 610 some diagonal matrix  $D(\tilde{\Lambda})$ , we may write  
 611

$$\mathbb{E}[M | \tilde{M}, \tilde{c}] = \tilde{Q} D(\tilde{\Lambda}) \tilde{Q}^T \quad (28)$$

612 Finally, the expectation of  $L(\theta)$  becomes:  
 613

$$\mathbb{E}[L(\theta) | \tilde{\theta}^*, \tilde{M}, \tilde{c}] = \frac{1}{2}(\theta - \tilde{\theta}^*)^T \tilde{Q} D(\tilde{\Lambda}) \tilde{Q}^T (\theta - \tilde{\theta}^*) + C(\tilde{\theta}^*, \tilde{M}, \tilde{c}) \quad (29)$$

614  $\square$

## 630 B PROOF OF THEOREM 2

631 *Proof.* First, observe that  
 632

$$D^t(\theta) = \tilde{M}^t(\theta - \tilde{\theta}^*) \quad (30)$$

633 We may see this by induction.  $D^1(\theta) = \nabla \tilde{L}(\theta) = \tilde{M}(\theta - \tilde{\theta}^*)$ . If  $D^t(\theta) = \tilde{M}^t(\theta - \tilde{\theta}^*)$ , then  
 634

$$\begin{aligned} 635 D^{t+1}(\theta) &= \frac{\partial}{\partial \zeta} D^1(\theta + \zeta D^t(\theta))|_{\zeta=0} = \frac{\partial}{\partial \zeta} \tilde{M}[\theta + \zeta \tilde{M}^t(\theta - \tilde{\theta}^*) - \tilde{\theta}^*]|_{\zeta=0} \\ 636 &= \frac{\partial}{\partial \zeta} \tilde{M}(\theta - \tilde{\theta}^*) + \zeta \tilde{M}^{t+1}(\theta - \tilde{\theta}^*)|_{\zeta=0} = M^{t+1}(\theta - \tilde{\theta}^*) \end{aligned} \quad (31)$$

637 Now consider  $\nabla \tilde{L}(\theta)$  and  $\nabla \tilde{L}(\hat{\theta})$ .  
 638

$$\nabla \tilde{L}(\theta) = \tilde{Q} f(\tilde{\Lambda}) \tilde{Q}^T (\theta - \tilde{\theta}^*) \quad (32)$$

639 and  
 640

$$\nabla \tilde{L}(\hat{\theta}) = \tilde{M}(\hat{\theta} - \tilde{\theta}^*) = \tilde{M}(\theta - \tilde{\theta}^* + \sum_{t=1}^T \gamma_t \tilde{M}^t(\theta - \tilde{\theta}^*)) = (\tilde{M} + \sum_{t=1}^T \gamma_t \tilde{M}^{t+1})(\theta - \tilde{\theta}^*) \quad (33)$$

648 Using the eigendecomposition of  $\tilde{M}$ , we have:  
 649

$$650 \quad \nabla \tilde{L}(\hat{\theta}) = \tilde{Q}(\tilde{\Lambda} + \sum_{t=1}^T \gamma_t \tilde{\Lambda}^{t+1}) \tilde{Q}^T (\theta - \tilde{\theta}^*) \quad (34)$$

$$651$$

$$652$$

653 Now, we compare  $f(\tilde{\Lambda})$  to  $\tilde{\Lambda} + \sum_{t=1}^T \gamma_t \tilde{\Lambda}^{t+1}$ . Observe that the function  $P(\tilde{\lambda}) = \tilde{\lambda} + \sum_{t=1}^T \gamma_t \tilde{\lambda}^{t+1}$   
 654 can represent any polynomial with intercept  $P(0) = 0$  and slope  $P'(0) = 1$ . By the Weierstrass  
 655 approximation theorem (Weierstrass, 1885), because the elements of  $\tilde{\Lambda}$  are bounded and  $f$  is contin-  
 656 uous, we may construct the following uniform bound:  
 657

$$658 \quad \|\tilde{\Lambda} + \sum_{t=1}^T \gamma_t \tilde{\Lambda}^{t+1} - f(\tilde{\Lambda})\|_F \leq \epsilon \quad (35)$$

$$659$$

$$660$$

661 for all  $\epsilon > 0$  and diagonal  $\Lambda$ , for some choice of sequence  $\gamma_1, \gamma_2, \dots, \gamma_T$ . By the rotation invariance  
 662 of the Frobenius norm, we have:

$$663 \quad \|\tilde{Q}(\tilde{\Lambda} + \sum_{t=1}^T \gamma_t \tilde{\Lambda}^{t+1}) \tilde{Q}^T - \tilde{Q}f(\tilde{\Lambda}) \tilde{Q}^T\|_F \leq \epsilon \quad (36)$$

$$664$$

$$665$$

$$666$$

667 Finally, since the Frobenius norm is an upper bound on the maximum eigenvalue of a matrix, we  
 668 have:  
 669

$$670 \quad \|\tilde{Q}(\tilde{\Lambda} + \sum_{t=1}^T \gamma_t \tilde{\Lambda}^{t+1}) \tilde{Q}^T (\theta - \tilde{\theta}^*) - \tilde{Q}f(\tilde{\Lambda}) \tilde{Q}^T (\theta - \tilde{\theta}^*)\| = \|\nabla \tilde{L}(\hat{\theta}) - \nabla \bar{L}(\theta)\| \leq \epsilon \|\theta - \tilde{\theta}^*\| \quad (37)$$

$$671$$

$$672 \quad \square$$

$$673$$

$$674$$

$$675$$

## C JUSTIFICATION OF THEORETICAL ASSUMPTIONS

676 In this section, we provide practical justifications for the theoretical assumptions made in Theorems  
 677 1 and 2. These assumptions are critical for the validity of our theoretical results and are grounded in  
 678 common practices and observations in machine learning.  
 679

680 **Assumption 1:  $M$  and  $\tilde{M}$  are symmetric matrices.** **Explanation:** In the context of quadratic  
 681 loss functions,  $M$  and  $\tilde{M}$  represent the Hessian matrices (second derivatives) of the true loss  $L(\theta)$   
 682 and the training loss  $\tilde{L}(\theta)$ , respectively. By definition, Hessian matrices of scalar-valued functions  
 683 are symmetric because mixed partial derivatives commute (i.e.,  $\frac{\partial^2 L}{\partial \theta_i \partial \theta_j} = \frac{\partial^2 L}{\partial \theta_j \partial \theta_i}$ ) when the function  
 684 is twice continuously differentiable. In practice, loss functions used in machine learning, such as  
 685 mean squared error and cross-entropy loss, satisfy these smoothness conditions. Therefore, assum-  
 686 ing that  $M$  and  $\tilde{M}$  are symmetric is both standard and justifiable.  
 687

688 **Assumption 2:  $\theta^*$  and  $\tilde{\theta}^*$  are independent of  $M$ ,  $\tilde{M}$ ,  $c$ , and  $\tilde{c}$ .** **Explanation:** This assumption  
 689 simplifies the analysis by decoupling the location of the minima from the curvature and offset of  
 690 the loss functions. In practical terms, it means that the position of the minimum (i.e., the parameter  
 691 values that minimize the loss) does not influence the curvature of the loss landscape or the constant  
 692 term. This is a reasonable approximation when considering local behavior around  $\theta$ , especially in  
 693 high-dimensional parameter spaces where the curvature is determined by the structure of the model  
 694 and the data distribution rather than the specific parameter values.  
 695

696 **Assumption 3:  $M$  is independent of  $\tilde{c}$  given  $\tilde{M}$ , i.e.,  $M \perp \tilde{c} | \tilde{M}$ .** **Explanation:** This assump-  
 697 tion asserts that, conditioned on the training loss curvature  $\tilde{M}$ , the curvature of the true loss  $M$  is  
 698 independent of the constant offset  $\tilde{c}$  of the training loss. In practical scenarios, the constant term  
 699  $\tilde{c}$  does not affect the gradient or Hessian of the loss function and, therefore, does not influence the  
 700 optimization process. Since  $\tilde{c}$  merely shifts the loss landscape vertically without changing its shape  
 701 or curvature, it is reasonable to consider  $M$  independent of  $\tilde{c}$  given  $\tilde{M}$ .

**Assumption 4: The joint distribution of  $\theta^*$  and  $\tilde{\theta}^*$  is symmetric, i.e.,  $p_{\theta^*, \tilde{\theta}^*} = p_{\tilde{\theta}^*, \theta^*}$ .** **Explanation:** This symmetry assumption implies that the statistical relationship between the true minimum  $\theta^*$  and the observed training minimum  $\tilde{\theta}^*$  is bidirectional and unbiased. In practical terms, it suggests that there is no preferential direction in the estimation errors between  $\theta^*$  and  $\tilde{\theta}^*$ . This is a reasonable assumption when the training data is a representative sample of the underlying data distribution, and there are no systematic biases affecting the estimation of the minima. It facilitates the theoretical analysis by ensuring consistent behavior regardless of the direction of estimation.

**Assumption 5: Rotation invariance condition:**  $p_{M|\tilde{M}}(UMU^T|U\tilde{M}U^T) = p_{M|\tilde{M}}(M|\tilde{M})$  **for all orthogonal matrices  $U$ .** **Explanation:** The rotation invariance assumption states that the conditional distribution of the true loss curvature  $M$  given the training loss curvature  $\tilde{M}$  is invariant under orthogonal transformations (rotations) of the parameter space. Practically, this means that the orientation of the coordinate system does not affect the statistical relationship between  $M$  and  $\tilde{M}$ . This assumption is justified in many machine learning models where the parameter space does not have a natural orientation, especially in isotropic settings where all directions are treated equally. It allows us to generalize results without loss of generality and simplifies the analysis by enabling diagonalization of matrices through rotations.

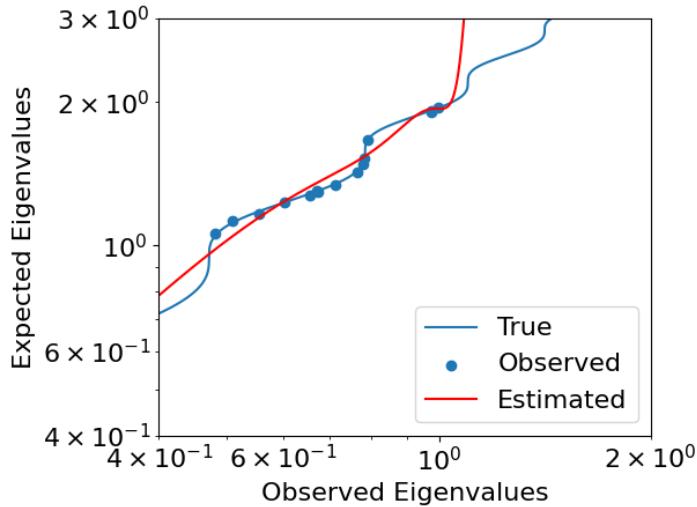
**Assumption 6: The conditional expectation  $\mathbb{E}[M|\tilde{M}]$  is diagonal when  $\tilde{M}$  is diagonal.** **Explanation:** This assumption suggests that if the training loss curvature matrix  $\tilde{M}$  is diagonal (indicating no interaction between different parameters), then the expected test loss curvature matrix  $M$  conditioned on  $\tilde{M}$  is also diagonal. In practical terms, when the training loss landscape exhibits axis-aligned curvature, it is reasonable to expect that the true loss landscape will have similar properties in expectation. Without this assumption, we must break the symmetry provided by the training loss landscape by assuming the expected test loss has a different and arbitrary set of curvature axes.

**Assumption 7: For all  $\theta$ , the expected training loss equals the true loss, i.e.,  $\mathbb{E}[\tilde{L}(\theta)|\theta^*, M, c] = L(\theta)$ .** **Explanation:** This assumption embodies the idea that, conditioned on the true loss parameters, the training loss is an unbiased estimator of the true loss at any point  $\theta$ . Practically, this means that the training data provides an accurate reflection of the true loss landscape on average. This assumption is justified when the training data is an independent and identically distributed (i.i.d.) sample from the same distribution as the test data, and there are no systematic errors or biases in the data collection process. It underpins the validity of using the training loss to make inferences about the true loss.

**Assumption 8: The function  $f$  is element-wise.** **Explanation:** In Theorem 1, recall that the eigenvalues of the observed Hessian and the expected test loss Hessian are related by an arbitrary function  $D(\tilde{\Lambda})$ . Here, we make the assumption that the function is applied *independently* to each eigenvalue. We believe this is reasonable because in many machine learning models, especially those with large numbers of parameters, the interactions between different parameters can often be approximated as negligible. This means that curvature transformation between train and test landscapes for one principle parameter direction is independent of the transformation for another principle parameter direction.

**Assumption 9: The function  $f$  is element-wise continuous with  $f(0) = 0$  and  $f'(0) = 1$ .** **Explanation:** In Theorem 2,  $f$  is used to modify the eigenvalues of the training loss curvature matrix  $\tilde{M}$  to approximate the curvature of the true loss. The conditions  $f(0) = 0$  and  $f'(0) = 1$  ensure that  $f$  behaves smoothly near zero and that small eigenvalues are not disproportionately affected, which is important for stability. In practice, we believe it is reasonable to expect that very flat directions of the expected test loss correspond to similarly flat directions of the training loss and vice versa, which is what these conditions on  $f$  imply.

**Assumption 10: The elements of  $\tilde{\Lambda}$  are bounded.** **Explanation:** The boundedness of the eigenvalues in  $\tilde{\Lambda}$  prevents extreme curvature in the training loss landscape. In practice, deep learning models are initialized with weights near 0 and activation functions with bounded derivatives; thus, it is reasonable to expect curvature to be practically boundable.

756 D EMPIRICAL VALIDATION OF GAM  
757  
758  
759770  
771  
772  
773  
774  
775  
776  
777  
778 Figure 4: True versus inferred eigenvalue transformation: Red line: Estimated transformation of  
779 eigenvalues from training to true loss landscape. Blue line: Actual Hessian eigenvalue transforma-  
780 tion. Blue dots: eigenvalues in the observed loss.781  
782 To empirically validate that GAM effectively learns the perturbation coefficients  $\gamma_t$ , we conduct ex-  
783 periments on a synthetic quadratic optimization problem where we control the relationship between  
784 the observed (training) loss and the expected (test) loss. This setup allows us to directly assess  
785 whether GAM can learn the true transformation from the training loss landscape to the expected test  
786 loss landscape.787 We consider observed and true quadratic loss landscapes with the same Hessian eigenvectors, but  
788 different Hessian eigenvalues. To simulate the observed training Hessian  $\tilde{M}$ , we define a nonlinear  
789 transformation  $f$  that relates the observed eigenvalues  $\tilde{\lambda}$  to the true eigenvalues  $\lambda$ :

790  
791  
792 
$$\tilde{\lambda} = f(\lambda) = \frac{1}{2}\lambda + \frac{1}{20}\sin(10\lambda). \quad (38)$$
  
793  
794

795 This transformation introduces both scaling and oscillatory behavior, mimicking complex discrep-  
796 ancies between the training and test loss landscapes. Our objective is to learn the perturbation  
797 coefficients  $\gamma_t$  such that the gradient of the perturbed training loss  $\nabla \tilde{L}(\hat{\theta})$  closely approximates the  
798 gradient of the true loss  $\nabla L(\theta)$ ; we use squared error as our discrepancy metric  $\Delta$ . See Appendix E  
799 for further details.800 As shown in Figure 4, the estimated transformation approximates the true transformation across  
801 the range of Hessian eigenvalues in the observed loss. This indicates that the learned perturbation  
802 coefficients effectively capture the nonlinear mapping between the training and test loss landscapes.  
803 We highlight, however, that the transformation may be inaccurate outside of the range of observed  
804 Hessian eigenvalues.805  
806 E EXPERIMENTAL DETAILS  
807  
808

809 All experiments were run on a 24 GB GPU.

810  
811 E.1 SYNTHETIC QUADRATIC PROBLEM

812 We set the parameter dimension to  $d = 15$  and use  $T = 12$  perturbation steps. We use a finite  
 813 difference constant of  $\epsilon = 0.1$  to approximate higher order derivatives. The perturbation coefficients  
 814  $\gamma_t$  are optimized using the Adam (Kingma, 2015) optimizer with a learning rate of  $10^{-3}$  over 100000  
 815 training iterations.

816 We generate the true Hessian  $M$  by sampling:  
 817

- 818 • A random orthogonal matrix  $Q \in \mathbb{R}^{d \times d}$  via QR decomposition of a random Gaussian  
 819 matrix.
- 820 • True eigenvalues  $\lambda \in \mathbb{R}^d$  sampled uniformly from  $[1, 2]$ , ensuring positive definiteness.  
 821

822 The true Hessian is then constructed as  $M = Q \text{diag}(\lambda)Q^T$ .  
 823

824 To simulate noise in the observed Hessian (as would occur due to sampling variability in real  
 825 datasets), we add Gaussian noise to the true eigenvalues:  
 826

$$827 \quad \lambda_{\text{noisy}} = \lambda + \sigma \cdot \eta, \quad (39)$$

828 where  $\eta \sim \mathcal{N}(0, I)$  and  $\sigma = 0.01$ . The observed eigenvalues are then computed as  $\tilde{\lambda} = f(\lambda_{\text{noisy}})$   
 829 using the transformation in Equation 38.  
 830

831  
832 E.2 MNIST

833 We consider two networks, 1) a softplus-activated MLP network with 3 fully-connected layers of  
 834 hidden layer size 256, 2) a softplus-activated CNN with 2 stride-2, kernel-3 convolutional layers  
 835 with channel sizes 32 and 64, followed by global average pooling and a final linear layer. Each  
 836 learnable weight layer is preceeded by batch normalization.  
 837

838 We train all methods for 10 epochs with batch size 100 using Adam optimizer (Kingma, 2015) at  
 839 learning rate  $10^{-3}$ . For GAM, we use  $T = 3$  perturbation steps and tune  $\gamma_s$  using Adam at learning  
 840 rate  $10^{-3}$ . All experiments are conducted over 5 random seeds. For GAM, we use the following  
 841 discrepancy function:  $\Delta(g_\theta, \bar{g}_\theta) = -g_\theta^T \bar{g}_\theta$  and set  $\epsilon = 10^{-3}$ . For CRSAM, we set step size to 0.1,  
 842  $\alpha = 0.1$ ,  $\beta = 0.01$ .  
 843

MNIST is made available to us via a Creative Commons license.  
 844

845  
846 E.3 CIFAR-10 AND SVHN

847 We consider three networks, 1) a softplus-activated MLP network with 3 fully-connected layers of  
 848 hidden layer size 1024, 2) a softplus-activated CNN with 4 stride-2, kernel-3 convolutional layers  
 849 with channel sizes 32, 64, 128 and 256, followed by global average pooling and a final linear layer,  
 850 3) a softplus-activated convolutional neural network with 13 convolutional layers followed by global  
 851 average pooling and a final linear layer. In the 14-layer CNN, all convolutions have stride 1 except  
 852 for the sixth and tenth, and have channel sizes 16, 32, 32, 32, 32, 64, 64, 64, 64, 128, 128, 128, 128.  
 853 Each learnable weight layer is preceeded by batch normalization.  
 854

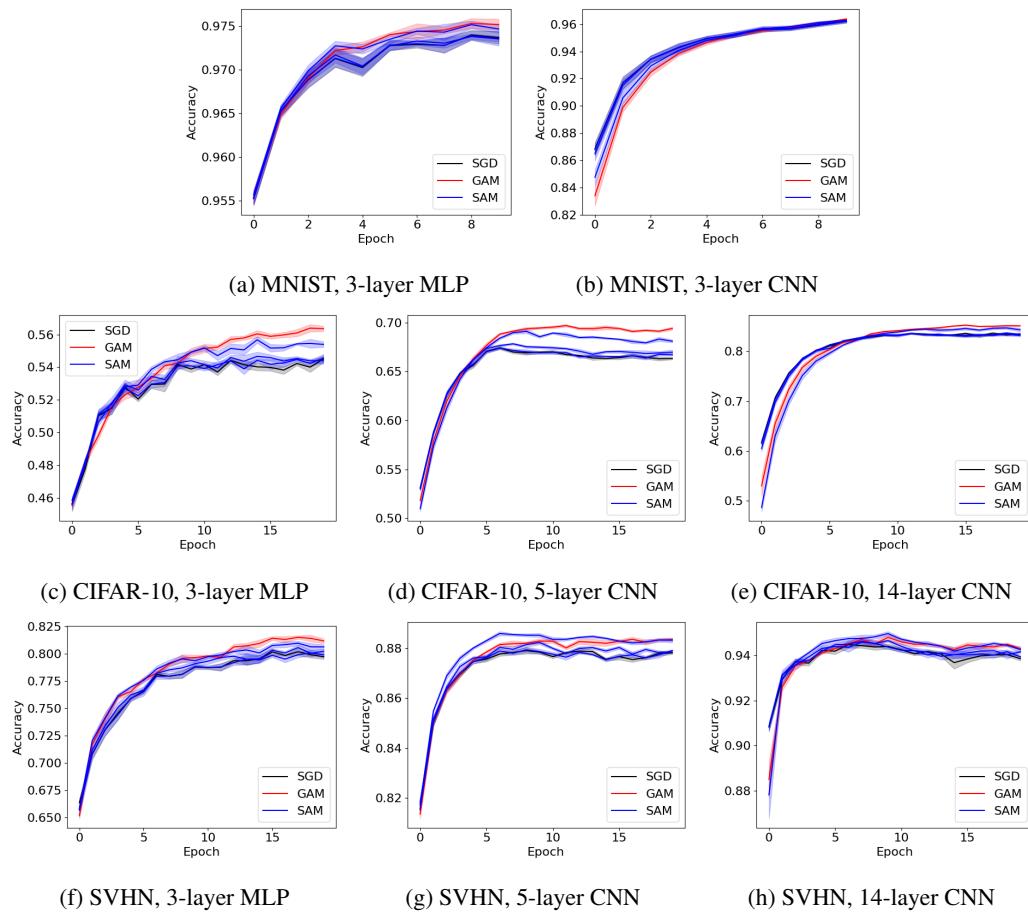
We train all methods for 20 epochs with batch size 100 using Adam optimizer (Kingma, 2015) at  
 learning rate  $10^{-3}$ . For GAM, we use  $T = 2$  perturbation steps and tune  $\gamma_s$  using Adam at learning  
 rate  $10^{-3}$ . All experiments are conducted over 5 random seeds. For GAM, we use the following  
 discrepancy function:  $\Delta(g_\theta, \bar{g}_\theta) = -g_\theta^T \bar{g}_\theta$  and set  $\epsilon = 10^{-3}$ . For CRSAM, we set step size to 0.1,  
 $\alpha = 0.1$ ,  $\beta = 0.01$ .  
 855

SVHN is made available to us for non-commercial use only.  
 856

857  
858 E.4 IMAGENET

859 We consider a softplus-activated ResNet-50 He et al. (2016) with standard settings; the only replace-  
 860 ment is ReLU with softplus.  
 861

864 We train all methods for 20 epochs with batch size 32 using Adam optimizer (Kingma, 2015) at  
865 learning rate  $10^{-3}$ . For GAM, we use  $T = 2$  perturbation steps and tune  $\gamma$ s using Adam at learning  
866 rate  $10^{-3}$ . Due to computational constraints, we run one experimental trial. For GAM, we use the  
867 following discrepancy function:  $\Delta(g_\theta, \bar{g}_\theta) = -g_\theta^T \bar{g}_\theta$  and set  $\epsilon = 10^{-3}$ .  
868  
869 ImageNet is made available to us for non-commercial use only.  
870  
871  
872  
873  
874  
875  
876  
877  
878  
879  
880  
881  
882  
883  
884  
885  
886  
887  
888  
889  
890  
891  
892  
893  
894  
895  
896  
897  
898  
899  
900  
901  
902  
903  
904  
905  
906  
907  
908  
909  
910  
911  
912  
913  
914  
915  
916  
917

918 F ADDITIONAL RESULTS  
919  
920  
921  
922  
923  
924  
925  
926  
927  
928

950  
951  
952  
953  
954  
955  
956  
957  
958  
959  
960  
961  
962  
963  
964  
965  
966  
967  
968  
969  
970  
971

Figure 5: Test error over the course of training for various networks trained on MNIST, CIFAR-10 and SVHN with different methods: stochastic gradient descent (SGD), sharpness aware minimization (SAM) with different parameter values  $\gamma_1$ , and generalization aware minimization (GAM). Margins indicate standard errors over 5 trials.

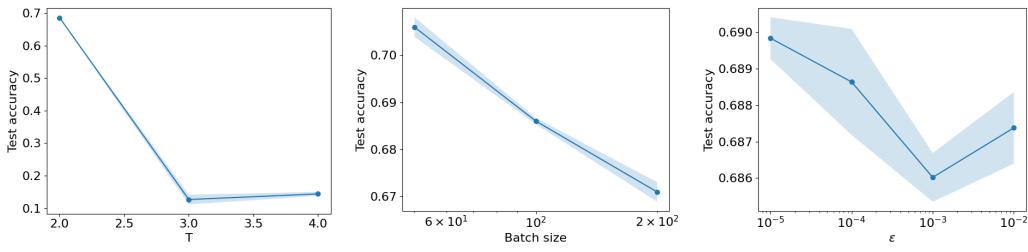


Figure 6: Test set accuracies of a CNN model trained on the CIFAR-10 dataset with generalization aware minimization (GAM) under different hyperparameter choices. By default, we use  $T = 2$ , batch size of 100, and  $\epsilon = 10^{-3}$ . Mean results and standard errors are reported over 5 trials.