

REGULATING INTERNAL EVIDENCE FLOWS FOR ROBUST LEARNING UNDER SPURIOUS CORRELATIONS

000
001
002
003
004
005 **Anonymous authors**
006 Paper under double-blind review
007
008
009
010

ABSTRACT

011 Deep models often exploit spurious correlations (e.g., backgrounds or dataset arti-
012 facts), hurting worst-group performance. We propose **Evidence-Gated Suppres-**
013 **sion (EGS)**, a lightweight, plug-in regularizer that intervenes inside the network
014 during training. EGS tracks a class-conditional, confidence-weighted contribution
015 for each neuron (more negative \Leftrightarrow stronger support) and applies a percentile-based,
016 multiplicative decay to the most extreme contributors, reducing overconfident
017 shortcut pathways while leaving other features relatively more influential. EGS
018 integrates with standard ERM, requires no group labels, and adds $< 5\%$ training
019 overhead. We provide analysis linking EGS to minority-margin gains, path-norm-
020 like capacity control, and stability benefits via EMA-smoothed gating. Empirically,
021 EGS improves worst-group accuracy and calibration vs. ERM and is competitive
022 with state-of-the-art methods across spurious-correlation benchmarks (e.g., Water-
023 birds, CelebA, BAR, COCO), while maintaining strong average accuracy. These
024 results suggest that regulating internal evidence flow is a simple and scalable route
025 to robustness without group labels.
026

1 INTRODUCTION

027 Deep learning models have achieved remarkable success across a wide range of tasks, establishing
028 them as the foundation of modern AI. However, understanding how these models achieve such results
029 reveals critical flaws in their learning process. A critical issue is that such models often rely on
030 spurious correlations superficial features like backgrounds, attributes, or dataset artifacts rather than
031 the core cues (You et al., 2025; Beery et al., 2018). While this shortcut learning can produce high
032 average-case accuracy, it can also cause catastrophic failures on minority or worst-case subpopulations,
033 where these spurious cues no longer hold (Geirhos et al., 2020; Yang et al., 2023; Hashimoto et al.,
034 2018; Tatman, 2017; Duchi et al., 2019). Such vulnerabilities make models unreliable in real-world
035 deployments, especially in safety- and fairness-critical applications. Addressing this challenge
036 requires encouraging models to focus their predictions on robust, transferable features rather than
037 shortcuts, with the goal of improving worst-group accuracy and narrowing group disparities while
038 maintaining strong overall performance (Izmailov et al., 2022).
039

040 A large body of robust learning methods attacks this problem by modifying data (e.g., group-aware
041 reweighting Nam et al. (2020); Kim et al. (2019); Lee et al. (2022), augmentations Kim et al. (2021))
042 or objectives that require access to group or environment labels (e.g., DRO-style surrogates and
043 invariance constraints (Sagawa et al., 2020)). These approaches can be highly effective when high-
044 quality group annotations and multiple environments are available, but in many real deployments such
045 side information is missing, expensive, or unreliable Mehrabi et al. (2021). Moreover, interventions
046 that operate outside the network (data-level) or only at the loss often lack direct control over the
047 internal contributors of specific neurons or connections that propagate spurious evidence forward
048 Dwivedi et al. (2024); K Kurmi et al. (2022).

049 We take a complementary path and intervene inside the model during training. We introduce **Evidence-**
050 **Gated Suppression (EGS)**, a plug-and-play, group-agnostic regularizer that selectively suppresses
051 spurious feature neurons using a class-conditional, softmax-weighted measure of on-batch evidence
052 energy. Intuitively, for each class, we estimate how much each neuron contributes to class-consistent
053 predictions; neurons whose smoothed evidence falls in the most negative tail are multiplicatively
shrunk during training, while others remain unchanged except for a mild global decay applied for

054 stability. The mechanism integrates seamlessly with standard cross-entropy optimization, requires no
 055 architectural changes, and adds less than 5% training-time overhead.
 056

057 Concretely, EGS maintains an exponential moving average (EMA) of per-neuron, per-class evidence
 058 energy computed from mini-batches, and applies a percentile-gated decay to connections whose
 059 smoothed evidence falls below a data-driven threshold. More negative evidence indicates stronger
 060 class-alignment, and suppressing the extreme negative tail prevents overconfident shortcut-driven
 061 pathways, leaving the remaining connections relatively more influential. This evidence-gated suppression
 062 thus down-weights neurons tied to spurious cues while indirectly encouraging reliance on more
 063 robust features, yielding representations that improve generalization and benefit minority groups even
 064 without group annotations.

065 Beyond empirical gains, our method is also clearly explained through theory. If we look at the
 066 model from the perspective of its final linear classifier (assuming the features it receives are well-
 067 behaved), the update we apply has three main effects: ① increases margins for hard or minority
 068 groups by shrinking connections that rely on spurious cues; ② controls model capacity in a way
 069 similar to path-norm regularization, by reducing the combined strength of weights and activations;
 070 ③ improves algorithmic stability due to EMA smoothing and mild decay, which in turn connects
 071 to calibration benefits by reducing overconfident logits from misaligned neurons. These properties
 072 explain why the method tends to lift average accuracy, worst-group accuracy and lower calibration
 073 error. Empirically, across standard spurious-correlation benchmarks (Waterbirds, CelebA, COCO,
 074 BAR) our method consistently improves average and worst-group performance while incurring
 075 minimal training overhead. Our analyses further illustrate how internal evidence is reallocated:
 076 neuron-level evidence distributions show contracted high-evidence tails for spurious contributor.
 077

078 Our contributions are as follows:

- 079 ① We define evidence energy, a fast, on-batch, class-conditional, softmax-weighted measure
 080 of neuron alignment, enabling training-time attribution without group labels.
- 081 ② We propose a simple percentile-gated suppression rule with EMA smoothing that is
 082 plug-and-play with standard training loops and incurs less than 5% overhead.
- 083 ③ We provide theory linking evidence-gated suppression to minority-group margin gains,
 084 capacity control through a path-norm-like shrinkage, and stability/calibration improvements.
- 085 ④ We demonstrate strong and consistent improvements in average-group accuracy, worst-
 086 group accuracy and calibration on spurious-correlation benchmarks, accompanied by
 087 neuron-level and representation analyses.

088 These results suggest that directly regulating the internal flow of class-conditional evidence is a
 089 simple, scalable, and effective route to robustness under spurious correlations particularly in the
 090 common setting where group labels are unavailable.
 091

092 2 RELATED WORK

093 **Shortcut learning and worst-group risk.** Deep models often exploit *shortcut* features that correlate
 094 with labels but fail under distribution or subpopulation shift Geirhos et al. (2020); Beery et al. (2018),
 095 leading to sharp drops in worst-group accuracy Krueger et al. (2021). Existing solutions largely act
 096 outside the network via data balancing Nam et al. (2020), environment construction Bahng et al.
 097 (2020), or robust loss design Vandenhirtz et al. (2023) yet leave the model’s internal reliance on
 098 spurious pathways unchecked. We instead regulate neurons directly, intervening at the source of
 099 shortcut learning.
 100

101 **Group-aware robust training.** When group labels are known, methods such as GroupDRO Sagawa
 102 et al. (2020), IRM Arjovsky et al. (2019), V-REx Krueger et al. (2021), and CVaR-DRO Levy et al.
 103 (2020) explicitly bound or penalize worst-group risk. These approaches are powerful but brittle:
 104 they require high-quality annotations and often incur large computational costs. Moreover, they
 105 adjust data or objectives rather than controlling how evidence flows inside the model. Our method
 106

108 avoids supervision, remains lightweight, and complements such objectives by suppressing weak or
 109 misaligned connections during training.
 110

111 **Group-agnostic robustness.** In the absence of labels, prior work reweights examples by confidence,
 112 error, or heuristics to anticipate spurious features Qraitem et al. (2023). Such strategies are indirect
 113 and unstable, since they rely on proxies of group structure. By contrast, we compute a neuron-level,
 114 class-conditional *evidence energy* from the model’s own predictions, and apply percentile-gated
 115 decay, yielding a direct, attribution-grounded signal that is efficient and stable.
 116

117 **Beyond generic regularization and pruning.** Classical penalties (e.g., weight decay, dropout,
 118 Jacobian constraints) are global and input-agnostic, improving average generalization but not worst-
 119 group stability Sokolić et al. (2017). Pruning and sparsification, while effective for compression,
 120 typically act post hoc and fail to reshape learning dynamics that cause shortcuts. Our approach
 121 differs by offering a training-time, class-aware regularizer that suppresses spurious pathways without
 122 altering architecture or requiring sparsity targets.
 123

124 **Complementarity to existing approaches.** EGS is *group-agnostic* and *plug-and-play*, integrating
 125 seamlessly with ERM and coexisting with data- or objective-level robust training (e.g., GroupDRO
 126 Sagawa et al. (2020), IRM Arjovsky et al. (2019), V-REx Krueger et al. (2021), JTT Liu et al. (2021),
 127 LfF Nam et al. (2020)). Whereas reweighting or invariance methods operate on examples or loss
 128 terms, EGS modulates the internal conduits that propagate spurious evidence. This complementary
 129 locus of control suggests potential gains from combining EGS with group-aware or group-discovery
 130 pipelines, as EGS can regularize the model’s reliance on emergent shortcuts even when environment
 131 or group labels are unavailable or noisy.
 132

133 **Positioning.** In summary, prior work either (i) relies on labels to enforce robustness, (ii) heuris-
 134 tically reweights data without touching internal contributors, or (iii) applies uniform or post hoc
 135 regularization. We propose a simple, group-agnostic, neuron-level mechanism that gates shrinkage
 136 by class-conditional evidence, directly reducing shortcut reliance with negligible overhead. This fills
 137 a critical gap by addressing spurious correlations at their origin inside the model while remaining
 138 complementary to existing data- and loss-level approaches.
 139

140 3 METHOD

141 3.1 NOTATION AND SETTING

142 We consider a standard C -class classification problem with a labeled dataset $\mathcal{D} = \{(x, y)\}$, where
 143 $y \in \{1, \dots, C\}$. For an input x , let $\phi_\theta(x) \in \mathbb{R}^D$ denote the penultimate-layer representation; we use
 144 $j \in \{1, \dots, D\}$ to index feature coordinates (neurons). The final prediction layer is a linear classifier
 145 $W = [w_1, \dots, w_C] \in \mathbb{R}^{D \times C}$ that maps representations to logits $z(x)$ and class probabilities $p(x)$:
 146

$$147 z(x) = W \phi_\theta(x), \quad p_k(x) \equiv [p(x)]_k = \text{softmax}_k(W \phi_\theta(x)) = \frac{\exp(w_k^\top \phi_\theta(x))}{\sum_{t=1}^C \exp(w_t^\top \phi_\theta(x))}. \quad (1)$$

148 We write $\phi_j(x)$ for the j -th coordinate of $\phi_\theta(x)$ and W_{jk} for the weight from feature j to class k .
 149 Our Evidence-Gated Suppression (EGS), acts on this last linear layer by default, but the same interface
 150 applies unchanged to intermediate units (e.g., channel activations in convolutional neural networks).
 151 For clarity we omit bias term. Unless stated otherwise, the base training objective is empirical risk
 152 minimization (ERM) with cross-entropy; EGS *augments* ERM as a plug-in regularizer and is agnostic
 153 to the choice of optimizer. All “evidence” quantities introduced below are computed during the
 154 forward pass and treated as *stop-gradient* statistics no gradients are propagated through the gating
 155 decisions or the EMA buffers akin to consistency-based approaches (Tsvaini & Valpola, 2017).
 156 This design keeps the procedure numerically stable and fully compatible with standard optimizers
 157 such as SGD Robbins & Monro (1951) and Adam Kingma & Ba (2015).
 158

162 3.2 EVIDENCE SCORE AND CLASS-AVERAGE ENERGY
163164 We quantify how a feature (neuron) j contributes to class k on an input x through the *evidence score*
165

166
$$e_{jk}(x) \triangleq -p_k(x) W_{jk} \phi_j(x), \quad (2)$$

167 and, in practice, evaluate it only for the true class $k = y$. The leading minus sign fixes the direction
168 of the scale: *more negative means stronger alignment*. Indeed, when $\phi_j(x)$ and W_{jy} align so as to
169 increase the y -logit, the product $W_{jy} \phi_j(x)$ is positive; if the model is confident on x (large $p_y(x)$),
170 then $e_{jy}(x)$ moves further into the negative range. The softmax factor $p_k(x)$ therefore couples
171 evidence to confidence, accentuating contributions that the model appears most sure about and
172 de-emphasizing uncertain ones. To aggregate across examples with the same ground-truth label, we
173 define the *class- k evidence energy*

174
$$E_{jk} \triangleq \mathbb{E}_{x \sim \mathcal{D}_k} [e_{jk}(x)], \quad (3)$$

175

176 where \mathcal{D}_k denotes the population of inputs with true label k . More negative E_{jk} indicates that feature
177 j consistently delivers confidence-weighted support for class k .
178179 **Normalization and scale.** Because $e_{jk}(x)$ scales linearly with both W_{jk} and $\phi_j(x)$, we avoid
180 hand-tuned rescaling and instead rely on three stabilizers: ① EMA smoothing of batch estimates
181 (§3.4); ② *within-class* percentile gating that depends only on order statistics (hence scale-free) (§3.5);
182 and ③ a mild global decay (Eq. 8) to prevent drift. When BatchNorm/LayerNorm is present, evidence
183 is computed on post-normalization features, so no extra per-feature normalization is required. These
184 choices make our evidence ordering robust while keeping the magnitude unconstrained, which is
185 precisely what the percentile gate exploits downstream.
186187 3.3 BATCH-WISE ESTIMATION AND THE ROLE OF BATCH SIZE
188189 In training, we estimate (Eq. 3) from mini-batches. At step t , for a batch B let $B_k = \{x \in B : y(x) = k\}$. The per-class batch evidence energy is
190

191
$$\bar{E}_{jk}^{(t)} = \frac{1}{|B_k| + \epsilon} \sum_{x \in B_k} e_{jk}(x), \quad \epsilon = 10^{-8}, \quad (4)$$

192

193 where the small ϵ guards against empty classes. If $B_k = \emptyset$, the class contributes 0 at this step; the
194 running EMA (§3.4) carries information forward until k reappears.
195196 **Effect of $|B|$.** Since $\bar{E}_{jk}^{(t)}$ is a finite-sample estimator of E_{jk} , its variability depends on batch size
197 and class balance. Smaller batches increase sampling noise, which can push more features into the
198 lower (more negative) tail and thus trigger additional suppression at a fixed percentile $q \in [10, 20]$.
199 Larger batches reduce this variance, yielding steadier gates but potentially less sensitivity to transient
200 shortcut spikes. Consequently, $|B|$ acts as a stability knob: too small risks over-pruning due to noisy
201 tail estimates; too large may under-suppress persistent shortcuts.
202203 3.4 EXPONENTIAL MOVING AVERAGE (EMA)
204205 The percentile gate in §3.5 relies on stable, low-variance estimates of per-class evidence energies
206 (§3.3). To smooth the noisy batch means $\bar{E}_{jk}^{(t)}$, we maintain an exponential moving average (EMA)
207

208
$$\tilde{E}_{jk}^{(t)} = (1 - \beta) \bar{E}_{jk}^{(t)} + \beta \tilde{E}_{jk}^{(t-1)}, \quad \beta \in [0.6, 1), \quad (5)$$

209

210 with $\beta = 0.75$ by default. When class k is absent in the current batch ($B_k = \emptyset$), we set $\tilde{E}_{jk}^{(t)} = 0$,
211 yielding the deterministic decay $\tilde{E}_{jk}^{(t)} = \beta \tilde{E}_{jk}^{(t-1)}$. Smaller β increases responsiveness to new
212 evidence; larger β adds inertia and is especially helpful for rare classes. We initialize $\tilde{E}_{jk}^{(0)} = 0$
213 and do not apply bias correction, since gating decisions are made by within-class rank (rather than
214 absolute scale).
215

216 3.5 PERCENTILE-GATED SUPPRESSION
217218 After a brief warm-up of $T_w = 5$ epochs to populate the EMA buffers, we compute a class-specific
219 threshold by taking the within-class q -th percentile (across features) of the smoothed energies:

220 221
$$\tau_k^{(t)} \triangleq \text{Percentile}_q\left(\{\tilde{E}_{1k}^{(t)}, \dots, \tilde{E}_{Dk}^{(t)}\}\right), \quad q \in [10, 20]. \quad (6)$$

222

223 Features whose energy falls below this threshold are temporarily suppressed via a binary gate
224

225
$$s_{jk}^{(t)} = \mathbb{I}\left[\tilde{E}_{jk}^{(t)} < \tau_k^{(t)}\right]. \quad (7)$$

226 We then apply a decoupled multiplicative decay to the classifier weights:
227

228
$$W_{jk} \leftarrow (1 - \alpha s_{jk}^{(t)}) (1 - 0.05 \alpha) W_{jk}, \quad \alpha \in [0.005, 0.15]. \quad (8)$$

229 The per-connection factor $1 - \alpha s_{jk}^{(t)}$ enforces selective suppression, while the mild global factor
230 $1 - 0.05 \alpha$ curbs scale oscillations when many connections are gated simultaneously. Decay is applied
231 *after* the forward pass (so gating is driven by the current predictions) and *before* backpropagation;
232 gradients are computed with respect to the post-decay weights. We treat $s_{jk}^{(t)}$ as a stop-gradient mask
233 and do not backpropagate through it. Ties at $\tau_k^{(t)}$ are broken deterministically so that exactly $\lceil q\% \rceil$ of
234 features are suppressed per class.
235236 3.6 OBJECTIVE VIEW, TRAINING RECIPE, AND PRACTICALITIES
237238 Although the mechanism is operationally defined by Eq. (8), it admits a useful proxy objective:
239

240 241
$$\mathcal{J}(\theta, W) = \mathcal{L}_{\text{ERM}}(\theta, W) + \alpha \sum_{j,k} s_{jk}^{(t)} \|W_{jk}\|, \quad (9)$$

242

243 which interprets gating as an adaptive, class-conditional sparsification penalty applied only to the
244 currently flagged connections.
245246 **Margin view.** For intuition, it is useful to imagine that the learned representation decomposes into
247 two parts, $\phi = \phi_{\text{rob}} + \phi_{\text{spu}}$, where ϕ_{rob} denotes robust, group-invariant features and ϕ_{spu} captures
248 spurious or shortcut-aligned features. Although this split is not directly observable in practice, it
249 provides a lens to interpret the effect of suppression. In this view, the gate primarily attenuates ϕ_{spu}
250 by reducing $\|W^\top \phi_{\text{spu}}\|$, which heuristically increases decision margins in settings where spurious
251 cues dominate:
252

253
$$\Delta\text{margin} \gtrsim \alpha \sum_{j,k} s_{jk}^{(t)} W_{jk} (\phi_{\text{spu}})_j. \quad (10)$$

254 **End-to-end pipeline.** At each training step:
255256 1. Forward pass to obtain $p(x)$ and $\phi_\theta(x)$; compute per-example evidence $e_{j,y(x)}(x)$ and batch
257 means $\bar{E}_{jk}^{(t)}$ (§3.3).
258 2. Update EMA $\tilde{E}_{jk}^{(t)}$ via Eq. (5).
259 3. After warm-up, form $\tau_k^{(t)}$ (Eq. (6)), compute gates $s_{jk}^{(t)}$ (Eq. (7)), and apply decay (Eq. (8)).
260 4. Backpropagate \mathcal{L}_{ERM} with the decayed weights and take the optimizer step.
261262 3.7 PROPERTIES OF EGS
263264 Under mild boundedness ($|\phi_j(x)| \leq R_\phi$, $\|w_k\|_2 \leq R_w$, $p_k(x) \in [0, 1]$), the four properties below
265 formalize *what* EGS targets, *why* its gates are stable and budgeted, and *how* its decay operator
266 suppresses spurious pathways while preserving robust ones. Throughout, the evidence for feature j
267 and class k on input x is $e_{jk}(x) = -p_k(x) W_{jk} \phi_j(x)$, and \tilde{E}_{jk} denotes the per-class exponential
268 moving average (EMA) of on-batch evidence; the classwise gate thresholds are percentiles τ_k
269 computed over $\{\tilde{E}_{jk}\}$.
270

270 **Property 1 (Confidence-weighted targeting and scale invariance)** For $y = k$, $e_{jk}(x)$ becomes
 271 more negative as the alignment $W_{jk}\phi_j(x)$ increases while $p_k(x)$ does not decrease—thus higher-
 272 confidence, class-aligned contributions receive larger magnitude. Moreover, any logit-preserving
 273 rescaling $\phi' = a\phi$, $W' = W/a$ (with $a > 0$) leaves $W'^\top\phi' = W^\top\phi$ and hence $p' = p$; consequently
 274 $e'_{jk}(x) = e_{jk}(x)$, the ordering of $\{\tilde{E}_{\cdot k}\}$, and the gated set are invariant. Lower-confidence samples
 275 (smaller p_y) produce evidence closer to 0 and thus exert less influence on gating.
 276

277 **Property 2 (Stable, budgeted gating via EMA and percentiles)** The EMA update $\tilde{E}^{(t)} = (1 -$
 278 $\beta)\tilde{E}^{(t)} + \beta\tilde{E}^{(t-1)}$ (where $\beta \in [0, 1]$ is the EMA parameter) is a convex combination with step-to-
 279 step drift bounded by $2(1 - \beta)R_\phi R_w$. Per class k , at most $\lceil q\% \rceil$ features are gated (here q is the
 280 gate budget in percent), because decisions depend only on the within-class order of $\{\tilde{E}_{\cdot k}\}$, which
 281 is invariant under any strictly increasing transform. If $\tilde{E}_{jk}^{(t-1)} \geq \tau_k^{(t-1)} + \Delta$ (gap $\Delta > 0$) and
 282 batch-induced perturbations have size at most δ (perturbation radius), then whenever
 283

$$(1 - \beta)\delta + \beta(1 - \beta)^{-1}\delta < \Delta,$$

285 feature j remains ungated at step t (gate inertia). Larger mini-batches B (we write $|B|$ for its size)
 286 and class-balanced sampling reduce the variance of $\tilde{E}^{(t)}$ and hence the probability of flips.
 287

288 **Property 3 (Multiplicative decay contracts and sparsifies)** Under the update

$$W_{jk} \leftarrow (1 - \alpha s_{jk})(1 - \alpha c)W_{jk},$$

291 $s_{jk} \in \{0, 1\}$ denotes the gate indicator for (j, k) , $\alpha > 0$ is the gated shrinkage rate, and $c = 0.05$ is a
 292 small global decay coefficient. Gated coordinates ($s_{jk} = 1$) contract by

$$\rho = (1 - \alpha)(1 - \alpha c) < 1$$

294 (the per-step contraction factor), yielding geometric attenuation across consecutive gated steps. For
 295 any threshold $\varepsilon > 0$, the number of gated entries with magnitude $> \varepsilon$ weakly decreases after one
 296 step (monotone sparsification). To first order in α , the combined effect is equivalent to applying a
 297 class-conditional, feature-wise proximal shrinkage to gated links before the ERM update, plus a mild
 298 global decay to stabilize scale.
 299

300 **Property 4 (Bias suppression with preservation of robust features)** When gated coordinates
 301 align with spurious components on worst-group examples, the pairwise margin

$$m_{k \rightarrow t}(x) := (w_k - w_t)^\top \phi(x)$$

304 obeys

$$m_{k \rightarrow t}^+(x) \geq m_{k \rightarrow t}(x) + \alpha \gamma_{\text{spu}} - \alpha c |(w_k - w_t)^\top \phi(x)|,$$

306 where S_k denotes the set of indices gated for class k , ϕ_{spu} is the spurious component in the decompo-
 307 sition $\phi = \phi_{\text{rob}} + \phi_{\text{spu}}$, and

$$\gamma_{\text{spu}} = \sum_{j \in S_k} W_{jk} (\phi_{\text{spu}})_j \geq 0$$

310 This shows a net margin gain on worst-group inputs up to the small global decay term. Conversely,
 311 any feature whose smoothed evidence persistently remains above the classwise percentile, $\tilde{E}_{j_{\text{rob}}, k} \geq \tau_k$
 312 (defining the index j_{rob} of a robust, never-gated feature), is deterministically never gated and is
 313 therefore preserved.

314 Hence, EGS \triangleright targets confident, class-aligned reliance in a scale-invariant manner; \triangleright stabilizes its
 315 budgeted decisions via EMA smoothing and percentile gates; \triangleright contracts and sparsifies persistently
 316 gated pathways through multiplicative decay; and \triangleright suppresses bias while preserving robust features
 317 that avoid the lower-evidence tail.
 318

319 4 EXPERIMENTAL SETUP

320 In this section we describe the experimental framework used to evaluate our proposed method. We
 321 detail the datasets, evaluation metrics, and model selection strategy, followed by baseline comparisons
 322 and implementation specifics.

324 We fine-tune a ResNet50 He et al. (2016) backbone with a linear classifier, initialized from ImageNet
 325 weights. All models are trained end-to-end with standard data augmentations (random resized crops,
 326 horizontal flips, mild color jitter). Our method (Sec. 3) is attached to the penultimate representation
 327 and the final classifier. Training overhead is < 5% relative to vanilla fine-tuning.
 328

329
 330 **Datasets, Metrics, and Model Selection.** We evaluate on three widely used spurious-correlation
 331 benchmarks and a COCO-based construction. **Waterbirds (WILDS)** Sagawa et al. (2020) induces
 332 a background shortcut: most waterbirds appear over water and most landbirds over land; minority
 333 (bias-conflicting) groups swap backgrounds. **CelebA (WILDS)** Liu et al. (2015) targets gender with
 334 hair color as a spurious attribute, where majority correlations are {male, dark hair} and {female, blond
 335 hair}. **BAR** Nam et al. (2020) stresses action recognition when training contexts are prototypical
 336 (e.g., climbing outdoors) but test-time contexts are shifted (e.g., indoor climbing). Finally, **COCO**
 337 **Gender/Object Bias** Zhao et al. (2023) builds a binary gender classification task with object-context
 338 correlations. Gender is inferred from captions using curated keywords; bias categories are drawn
 339 from COCO instance annotations as two disjoint groups: *Sports/Outdoor* (e.g., skateboard, skis,
 340 bicycle) and *Kitchen/Indoor* (e.g., oven, refrigerator, cup, dining table). We skew training so that male
 341 images are primarily paired with Sports/Outdoor objects and female images with Kitchen/Indoor
 objects; validation/test splits are stratified into Unbiased and Bias-Conflicting.

342 Our primary metric is *worst-group accuracy* (WGA) and we also report average accuracy (Avg). For
 343 Waterbirds and CelebA, WGA is computed over the minority bias-conflicting groups. For BAR, since
 344 the test set consists solely of conflicting examples, we report average accuracy. For COCO, we report
 345 both Unbiased and Bias-Conflicting accuracies within each bias category and their mean. Unless
 346 otherwise noted, model selection is based on WGA on a held-out validation set.

347 Our method maintains only a $D \times C$ EMA buffer. Defaults that work robustly are $(q, \beta, \alpha, T_w, \epsilon) =$
 348 $(20, 0.75, 0.05, 5, 10^{-8})$. We keep $|B|$ (Batch Size) at 32 for all datasets except the BAR dataset,
 349 which has a value of 8.

351 4.1 RESULTS

352 We evaluate our method EGS on three standard spurious-correlation benchmarks (Waterbirds, CelebA,
 353 BAR) and COCO (Gender/Object Bias), reporting average accuracy and worst-group accuracy
 354 (WGA). Unless noted, model selection follows the protocol in §4 and all runs average over ≥ 3 seeds
 355 (mean \pm std). Training overhead remains < 5% relative to vanilla fine-tuning, as EGS adds only a
 356 light EMA buffer and a percentile gate computed per class.

357
 358 **Main results on spurious-correlation benchmarks** Across benchmarks, EGS consistently im-
 359 proves both robustness and accuracy. See Table 1 for a summary across benchmarks. On CelebA, it
 360 attains the best results on both unbiased and bias-conflicting splits (**95.63±0.28** and **93.95±1.06**),
 361 outperforming the strongest prior (Eva-E He et al. (2025)) by over five points in each case and cutting
 362 the conflicting-split error from 11.26% to 6.05%. On Waterbirds, EGS achieves the highest average
 363 accuracy at **97.44±0.29**, slightly surpassing Eva-E while lowering average error by **16.1%**, and ob-
 364 tains a worst-group accuracy of 80.93 ± 1.32 , which is statistically comparable to Eva-E (81.31 ± 1.5)
 365 though trailing JTT (84.98 ± 0.5) by about four points, reflecting the classic average-worst Pareto
 366 trade-off. On BAR, EGS delivers the strongest average accuracy at **76.09±0.38**, a gain of +2.39
 367 points over Eva-E and +15.58 over ERM. Taken together, these results show that EGS advances
 368 the Pareto frontier of average accuracy versus worst-case robustness: it is dominant on CelebA,
 369 state-of-the-art on BAR, and *Pareto-competitive* on Waterbirds (top average, near-top WGA).

370
 371 **COCO Gender/Object Bias** On COCO with *Sports/Outdoor* and *Kitchen/Indoor* object biases, EGS
 372 attains the best average accuracy: 84.27, exceeding the strongest baseline (GMBM) by +0.73 points
 373 and ERM (Vanilla) by +14.77 points. Results are summarized in Table 2. Crucially, EGS narrows
 374 bias-induced gaps: for Sports, the (unbiased-conflicting) gap shrinks from 6.20 to 0.67; for Kitchen,
 375 from 5.84 to 2.15. Per-split accuracies remain balanced (84.53/83.86 for Sports, 85.41/83.26 for
 376 Kitchen), indicating that EGS redistributes internal evidence toward context-invariant signals. It is
 377 also worth noting that GMBM Dwivedi et al. (2025) and Badd Sarris et al. (2024) target multiple,
 co-occurring biases, and our method outperforms them as well.

Method	BAR Average Acc.	CelebA		Waterbirds	
		Unbiased	Conflicting	Accuracy	Worst Acc.
Vanilla	60.51 \pm 4.3	70.25 \pm 0.4	52.52 \pm 0.2	94.10 \pm 4.3	63.74 \pm 3.2
LfF Nam et al. (2020)	62.98 \pm 2.8	84.24 \pm 0.4	81.24 \pm 1.4	89.60 \pm 2.4	74.98 \pm 2.1
EiIL Creager et al. (2021)	68.44 \pm 1.2	85.70 \pm 1.6	81.70 \pm 1.5	95.88 \pm 1.7	77.20 \pm 1.0
JTT Liu et al. (2021)	68.53 \pm 3.2	86.40 \pm 4.6	77.80 \pm 2.5	93.70 \pm 0.5	84.98 \pm 0.5
LWBC Kim et al. (2022)	68.45 \pm 1.3	83.90 \pm 1.6	87.22 \pm 1.1	—	—
Debian Li et al. (2022)	69.88 \pm 2.9	90.02 \pm 0.8	85.33 \pm 3.7	—	—
SiFER Tiwari & Shenoy (2023)	72.08 \pm 0.4	90.00 \pm 0.9	88.04 \pm 1.2	96.11 \pm 0.6	77.22 \pm 0.4
EvA-E He et al. (2025)	73.70 \pm 0.8	90.51 \pm 1.0	88.74 \pm 1.4	96.95 \pm 0.9	81.31 \pm 1.5
Ours	76.09\pm0.38	95.63\pm0.28	93.95\pm1.06	97.44\pm0.29	80.93 \pm 1.32

Table 1: Test performance on CelebA, Waterbirds, and BAR (%): EGS improves average accuracy and strengthens worst-group robustness without group labels (ResNet-50 backbone).

Method	Avg. Acc.	Sports Object		Kitchen Object	
		Unbiased	Conflicting	Unbiased	Conflicting
Vanilla	69.50	70.81	64.61	73.20	67.36
FairKL Barbano et al. (2023)	73.67	76.32	67.11	74.35	76.90
EnD Tartaglione et al. (2021)	76.95	77.11	70.97	82.38	77.34
FLAC Sarridis et al. (2025)	79.88	80.02	77.31	80.22	79.95
BAdd Sarridis et al. (2024)	81.76	81.28	77.81	82.91	83.05
GMBM Dwivedi et al. (2025)	83.54	83.78	83.85	83.19	83.35
Ours	84.27	84.53	83.86	85.41	83.26

Table 2: COCO gender/object bias (validation): EGS attains the best average accuracy and markedly reduces the unbiased and conflicting gap across both Sports and Kitchen splits.

4.2 ANALYSIS AND ABLATION STUDIES

Decay α and EMA β . As shown in Fig. 1, predicted by the multiplicative-contraction view (Property 3), increasing α strengthens sparsifying but can overshoot if combined with small batches or highly entangled features. A moderate EMA ($\beta = [0.75]$) reduces gate flips without blunting adaptation;

Batch size and class balance. Because evidence is computed on-batch and class-wise (Eq. (4)), smaller batches increase sampling variance and can push more features into the gated tail (Property 2). Class-balanced sampling, when available, notably stabilizes gate thresholds τ_k and improves WGA on datasets with rare minority groups (e.g., Waterbirds). This also helps explain the residual WGA gap to JTT on Waterbirds: with rare bias-conflicting backgrounds, slightly larger or class-balanced batches reduce variance in $E^{(t)}$ and improve the tail estimate used for gating. (see Fig. 1)

EGS improves worst-group robustness and average accuracy by reallocating confidence away from spurious lower-tail pathways. Its per-class, confidence-weighted gating yields strong, stable gains on CelebA and BAR, competitive WGA on Waterbirds while achieving the best average accuracy, and large bias-gap reductions on COCO—all with negligible training overhead.

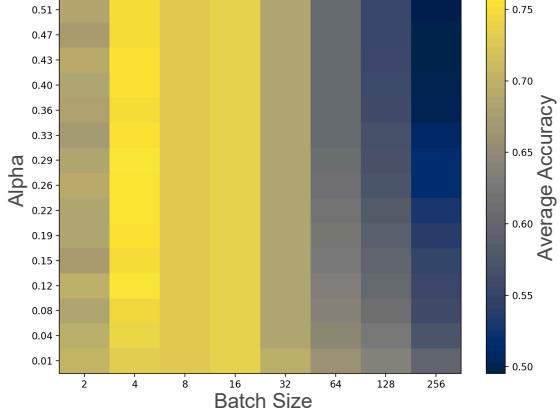


Figure 1: Ablations on decay α and batch size: moderate α and sufficient batches stabilize gating and boost accuracy, while very small batches or large α risk over-suppression.

432 **5 DISCUSSION**

434 **What is being regularized?** EGS acts on *internal, class-conditional flows of confidence-weighted*
 435 *evidence* rather than on data sampling or loss surrogates. By tracking the EMA of per-neuron evidence
 436 energy $E_{jk}^{(t)}$ (Eq. 5) and shrinking only the more negative (i.e., strongest-confidence, class-aligned)
 437 connections per class via a percentile gate (Eq.6-8), the method implements a targeted, budgeted
 438 contraction of high-confidence pathways that dominate the logit of the true class. This stands
 439 in contrast to uniform penalties (e.g., weight decay, dropout) or post-hoc pruning and places the
 440 intervention precisely where shortcut reliance manifests during training. The resulting regularization
 441 is (i) *scale-invariant* under logit-preserving reparameterizations (Property 1), (ii) *budgeted and stable*
 442 by class via percentiles + EMA (Property 2), and (iii) *contractive/sparsifying* on persistently gated
 443 coordinates (Property 3). Collectively, these properties explain why EGS improves minority-group
 444 margins and calibration while minimally perturbing the remainder of the network (Property 4).

445 **Why do Pareto gains emerge?** Across benchmarks, EGS advances the average-vs-worst-case
 446 Pareto frontier by reducing the contribution of shortcut-dominated paths *without* uniformly shrinking
 447 capacity. On CelebA and BAR, EGS attains the best or state-of-the-art accuracies, while on Waterbirds
 448 it secures the top average accuracy and near-top worst-group accuracy - a pattern consistent with a
 449 selective contraction that curbs overconfident, spurious tails while preserving robust features that
 450 avoid the lower-evidence tail (qualitative evidence histograms and representation geometry support
 451 this view). The quantitative trends in Table 1 show strong gains on BAR (+2.39 points over EvA-E;
 452 +15.58 over ERM) and large improvements on CelebA (e.g., conflicting-split error reduction of
 453 $\approx 46\%$), with competitive worst-group accuracy on Waterbirds. On our COCO construction, EGS
 454 both improves average accuracy and substantially shrinks bias-induced gaps (Table 2). These Pareto
 455 improvements are precisely what one would expect from a per-class, confidence-weighted contraction
 456 that narrows the heavy tail of shortcut evidence.

457 **Mechanistic interpretation.** A useful lens is the *path-norm-like* view: EGS decreases the product
 458 of weights and activations *only* along the class-conditional lower tail, which reduces the extreme
 459 logit contributions most responsible for miscalibration and minority-group errors. Because the gate is
 460 percentile-based, the effective capacity budget is class-wise and scale-free, limiting the risk of global
 461 underfitting. The margin decomposition (Eq. 10) further clarifies how EGS can raise minority-group
 462 margins by attenuating spurious components in $W^\top \phi_{\text{spu}}$ while leaving $W^\top \phi_{\text{rob}}$ comparatively more
 463 influential, matching the empirical observation of improved worst-case accuracy alongside strong
 464 average accuracy.

466 **6 CONCLUSION**

468 We propose *Evidence-Gated Suppression* (EGS), a lightweight, group-agnostic regularizer that
 469 regulates *where* and *how much* class-conditional confidence flows by contracting neuron-class links
 470 that persist in the lower tail of confidence-weighted evidence. This percentile-budgeted, scale-
 471 invariant mechanism suppresses shortcut pathways while preserving robust features, improving
 472 average and worst-group accuracy and calibration with negligible overhead and without group labels
 473 across diverse benchmarks. EGS implicitly assumes that the most negative evidence predominantly
 474 tracks shortcut-aligned routes; when robust cues dominate confidence or when robust and spurious
 475 features co-activate, such cues may transiently enter the lower tail and be contracted. While an
 476 EMA with a percentile budget reduces this risk, very small or imbalanced batches can destabilize
 477 the threshold τ_k and over-suppress; early miscalibration can distort $p_k(x)$ and thus the attribution
 478 signal $e_{jk}(x)$, suggesting warm-up and conservative initial settings; and strong α or large q may
 479 underfit in highly entangled regimes. Our current implementation chiefly gates the final classifier (or
 480 top-layer channels), yet spurious pathways can emerge earlier in the network. These observations
 481 motivate adaptive, class-aware budgets and variance-reduced evidence estimation (especially for
 482 rare classes), layer- or module-wise gating (e.g., channels or attention heads), and combinations
 483 with group discovery or DRO to unify example-level and pathway-level robustness. Beyond vision,
 484 applying EGS to structured prediction and multi-label or non-vision tasks could broaden the space of
 485 attribution-grounded regularizers for scalable robustness under spurious correlations.

486 REFERENCES
487

488 Martin Arjovsky, Léon Bottou, Ishaaan Gulrajani, and David Lopez-Paz. Invariant risk minimization.
489 *arXiv preprint arXiv:1907.02893*, 2019.

490 Hyojin Bahng, Sanghyuk Chun, Sangdoo Yun, Jaegul Choo, and Seong Joon Oh. Learning de-biased
491 representations with biased representations. In Hal Daumé III and Aarti Singh (eds.), *Proceedings*
492 *of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine*
493 *Learning Research*, pp. 528–539. PMLR, 13–18 Jul 2020. URL <https://proceedings.mlr.press/v119/bahng20a.html>.

494

495 Carlo Alberto Barbano, Benoit Dufumier, Enzo Tartaglione, Marco Grangetto, and Pietro Gori.
496 Unbiased supervised contrastive learning. In *International Conference on Learning Representations*
497 (*ICLR*), 2023.

498

499 Sara Beery, Grant Van Horn, and Pietro Perona. Recognition in terra incognita. In *Proceedings of the*
500 *European Conference on Computer Vision (ECCV)*, September 2018.

501 Elliot Creager, Jörn-Henrik Jacobsen, and Richard Zemel. Environment inference for invariant
502 learning. In *International Conference on Machine Learning*, pp. 2189–2200. PMLR, 2021.

503

504 John C Duchi, Tatsunori Hashimoto, and Hongseok Namkoong. Distributionally robust losses against
505 mixture covariate shifts. *Under review*, 2, 2019.

506

507 Rajeev Ranjan Dwivedi, Priyadarshini Kumari, and Vinod K. Kurmi. Cosfairnet:a parameter-space
508 based approach for bias free learning. In *35th British Machine Vision Conference 2024, BMVC 2024,*
509 *Glasgow, UK, November 25-28, 2024*. BMVA, 2024. URL <https://papers.bmvc2024.org/0738.pdf>.

510

511 Rajeev Ranjan Dwivedi, Ankur Kumar, and Vinod K Kurmi. Multi attribute bias mitigation via
512 representation learning. In *Proceedings of the European Conference on Artificial Intelligence*
513 (*ECAI*), Bologna, Italy, October 2025.

514

515 Robert Geirhos, Jörn-Henrik Jacobsen, Claudio Michaelis, Richard Zemel, Wieland Brendel, Matthias
516 Bethge, and Felix A. Wichmann. Shortcut learning in deep neural networks. *Nature Machine*
517 *Intelligence*, 2(11):665–673, November 2020. ISSN 2522-5839. doi: 10.1038/s42256-020-00257-z.
518 URL <https://doi.org/10.1038/s42256-020-00257-z>.

519

520 Tatsunori Hashimoto, Megha Srivastava, Hongseok Namkoong, and Percy Liang. Fairness without
521 demographics in repeated loss minimization. In *International Conference on Machine Learning*,
522 pp. 1929–1938. PMLR, 2018.

523

524 Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image
525 recognition. *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp.
526 770–778, 2016.

527

528 Qiyuan He, Kai Xu, and Angela Yao. Eva: Erasing spurious correlations with activations. In
529 *The Thirteenth International Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=zKvrOOBouT>.

530

531 Pavel Izmailov, Polina Kirichenko, Nate Gruver, and Andrew Gordon Wilson. On feature learning
532 in the presence of spurious correlations. In *Proceedings of the 36th International Conference on*
533 *Neural Information Processing Systems*, NIPS ’22, Red Hook, NY, USA, 2022. Curran Associates
534 Inc. ISBN 9781713871088.

535

536 Vinod K Kurmi, Rishabh Sharma, Yash Vardhan Sharma, and Vinay P. Namboodiri. Gradient based
537 activations for accurate bias-free learning. In *AAAI*, Feb 2022.

538

539 Byungju Kim, Hyunwoo Kim, Kyungsu Kim, Sungjin Kim, and Junmo Kim. Learning not to learn:
540 Training deep neural networks with biased data. In *Proceedings of the IEEE/CVF Conference on*
541 *Computer Vision and Pattern Recognition*, pp. 9012–9020, 2019.

542

543 Eungyeup Kim, Jihyeon Lee, and Jaegul Choo. Biaswap: Removing dataset bias with bias-tailored
544 swapping augmentation. In *Proceedings of the IEEE/CVF International Conference on Computer*
545 *Vision*, pp. 14992–15001, 2021.

540 Nayeong Kim, Sehyun Hwang, Sungsoo Ahn, Jaesik Park, and Suha Kwak. Learning debiased
 541 classifier with biased committee. In *Proceedings of the 36th International Conference on Neural*
 542 *Information Processing Systems*, NIPS '22, Red Hook, NY, USA, 2022. Curran Associates Inc.
 543 ISBN 9781713871088.

544 Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In Yoshua
 545 Bengio and Yann LeCun (eds.), *3rd International Conference on Learning Representations, ICLR*
 546 *2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015. URL <http://arxiv.org/abs/1412.6980>.

547 David Krueger, Ethan Caballero, Joern-Henrik Jacobsen, Amy Zhang, Jonathan Binas, Dinghuai
 548 Zhang, Remi Le Priol, and Aaron Courville. Out-of-distribution generalization via risk extrapolation (rex). In *International Conference on Machine Learning*, pp. 5815–5826. PMLR, 2021.

549 Jungsoo Lee, Jeonghoon Park, Daeyoung Kim, Juyoung Lee, Edward Choi, and Jaegul Choo.
 550 Revisiting the importance of amplifying bias for debiasing. *AAAI-23*, 5 2022. URL <http://arxiv.org/abs/2205.14594>.

551 Daniel Levy, Yair Carmon, John C Duchi, and Aaron Sidford. Large-scale methods for distributionally
 552 robust optimization. In *NeurIPS*, 2020.

553 Zhiheng Li, Anthony Hoogs, and Chenliang Xu. Discover and mitigate unknown biases with
 554 debiasing alternate networks. In *ECCV*, pp. 270–288. Springer, 2022.

555 Evan Z Liu, Behzad Haghgoo, Annie S Chen, Aditi Raghunathan, Pang Wei Koh, Shiori Sagawa,
 556 Percy Liang, and Chelsea Finn. Just train twice: Improving group robustness without training
 557 group information. In *International Conference on Machine Learning*, pp. 6781–6792. PMLR,
 558 2021.

559 Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaou Tang. Deep learning face attributes in the wild. In
 560 *Proceedings of International Conference on Computer Vision (ICCV)*, December 2015.

561 Ninareh Mehrabi, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan. A survey
 562 on bias and fairness in machine learning. *ACM computing surveys (CSUR)*, 54(6):1–35, 2021.

563 Junhyun Nam, Hyuntak Cha, Sungsoo Ahn, Jaeho Lee, and Jinwoo Shin. Learning from failure:
 564 De-biasing classifier from biased classifier. In H. Larochelle, M. Ranzato, R. Hadsell, M.F.
 565 Balcan, and H. Lin (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp.
 566 20673–20684. Curran Associates, Inc., 2020. URL <https://proceedings.neurips.cc/paper/2020/file/eddc3427c5d77843c2253f1e799fe933-Paper.pdf>.

567 Maan Qraitem, Kate Saenko, and Bryan A Plummer. Bias mimicking: A simple sampling approach
 568 for bias mitigation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern*
 569 *Recognition*, pp. 20311–20320, 2023.

570 Herbert Robbins and Sutton Monro. A Stochastic Approximation Method. *The Annals of Mathematical Statistics*, 22(3):400 – 407, 1951. doi: 10.1214/aoms/1177729586. URL <https://doi.org/10.1214/aoms/1177729586>.

571 Shiori Sagawa, Pang Wei Koh, Tatsunori B. Hashimoto, and Percy Liang. Distributionally robust
 572 neural networks. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=ryxGuJrFvS>.

573 Ioannis Sarridis, Christos Koutlis, Symeon Papadopoulos, and Christos Diou. Badd: Bias mitigation
 574 through bias addition. *arXiv preprint arXiv:2408.11439*, 2024.

575 Ioannis Sarridis, Christos Koutlis, Symeon Papadopoulos, and Christos Diou. FLAC: Fairness-Aware
 576 Representation Learning by Suppressing Attribute-Class Associations . *IEEE Transactions on*
 577 *Pattern Analysis & Machine Intelligence*, 47(02):1148–1160, February 2025. ISSN 1939-3539.
 578 doi: 10.1109/TPAMI.2024.3487254. URL <https://doi.ieee.org/10.1109/TPAMI.2024.3487254>.

594 Jure Sokolić, Raja Giryes, Guillermo Sapiro, and Miguel R. D. Rodrigues. Robust large margin
 595 deep neural networks. *Trans. Sig. Proc.*, 65(16):4265–4280, August 2017. ISSN 1053-587X. doi:
 596 10.1109/TSP.2017.2708039. URL <https://doi.org/10.1109/TSP.2017.2708039>.

597

598 Enzo Tartaglione, Carlo Alberto Barbano, and Marco Grangetto. End: Entangling and disentangling
 599 deep representations for bias correction. In *Proceedings of the IEEE/CVF conference on computer*
 600 *vision and pattern recognition*, pp. 13508–13517, 2021.

601

602 Anti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency
 603 targets improve semi-supervised deep learning results. In *Proceedings of the 31st International*
 604 *Conference on Neural Information Processing Systems*, NIPS’17, pp. 1195–1204, Red Hook, NY,
 605 USA, 2017. Curran Associates Inc. ISBN 9781510860964.

606

607 Rachael Tatman. Gender and dialect bias in youtube’s automatic captions. In *Proceedings of the first*
 608 *ACL workshop on ethics in natural language processing*, pp. 53–59, 2017.

609

610 Rishabh Tiwari and Pradeep Shenoy. Overcoming simplicity bias in deep networks using a feature
 611 sieve. In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato,
 612 and Jonathan Scarlett (eds.), *Proceedings of the 40th International Conference on Machine Learning*,
 613 volume 202 of *Proceedings of Machine Learning Research*, pp. 34330–34343. PMLR, 23–29
 614 Jul 2023. URL <https://proceedings.mlr.press/v202/tiwari23a.html>.

615

616 Moritz Vandenhirtz, Laura Manduchi, Ričards Marcinkevičs, and Julia E Vogt. Signal is harder to
 617 learn than bias: Debiasing with focal loss. *arXiv preprint arXiv:2305.19671*, 2023.

618

619 Yuzhe Yang, Haoran Zhang, Dina Katabi, and Marzyeh Ghassemi. Change is hard: a closer look at
 620 subpopulation shift. In *Proceedings of the 40th International Conference on Machine Learning*,
 621 ICML’23. JMLR.org, 2023.

622

623 Chenyu You, Haocheng Dai, Yifei Min, Jasjeet S. Sekhon, Sarang C. Joshi, and James S. Duncan.
 624 The silent majority: Demystifying memorization effect in the presence of spurious correlations.
 625 *CoRR*, abs/2501.00961, January 2025. URL <https://doi.org/10.48550/arXiv.2501.00961>.

626

627 Dora Zhao, Jerome Andrews, and Alice Xiang. Men also do laundry: Multi-attribute bias amplification.
 628 In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and
 629 Jonathan Scarlett (eds.), *Proceedings of the 40th International Conference on Machine Learning*,
 630 volume 202 of *Proceedings of Machine Learning Research*, pp. 42000–42017. PMLR, 23–29 Jul
 631 2023. URL <https://proceedings.mlr.press/v202/zhao23a.html>.

632

633

634

635

636

637

638

639

640

641

642

643

644

645

646

647

648 **A EXPERIMENTAL DETAILS AND REPRODUCIBILITY**
649650
651 This section provides the necessary details to reproduce our experiments, including training protocols,
652 dataset construction specifics, hyperparameter settings, and an analysis of computational overhead.
653654 **A.1 IMPLEMENTATION DETAILS**
655656 All models were implemented in PyTorch. We fine-tuned a ResNet-50 backbone pre-trained on
657 ImageNet for all experiments. Standard data augmentations were applied, including random resized
658 crops to 224×224 , horizontal flips, and mild color jittering. The learning rate was managed with a
659 cosine decay schedule, preceded by a linear warm-up period of 5 epochs. Our method, EGS, was
660 applied to the final linear classifier layer. Each experiment was run with at least three different random
661 seeds (77, 25 and 42), and we report the mean and standard deviation of the results.
662663 Table 3 summarizes batch size, epochs, optimizer settings, and EGS hyperparameters (α, T_w, β, q)
664 used in the main results.
665666 **Table 3: Implementation and hyperparameter details by dataset.** All experiments fine-tune a
667 ResNet-50 with ImageNet initialization. EGS: α (decay), T_w (warm-up epochs), β (EMA momen-
668 tum), q (percentile budget). Penultimate dimensionality is 2048.
669

Dataset	Batch Size	Epochs	Optimizer Details	EGS Hyperparams (α, T_w, β, q)
Waterbirds	32	100	Adam ($lr=1 \times 10^{-4}$, $wd=1 \times 10^{-4}$)	(0.075, 5, 0.75, 20)
CelebA	32	30	Adam ($lr=1 \times 10^{-4}$, $wd=1 \times 10^{-4}$)	(0.075, 5, 0.75, 20)
COCO (our)	32	50	Adam ($lr=1 \times 10^{-4}$, $wd=1 \times 10^{-4}$)	(0.035, 5, 0.75, 20)
BAR	8	50	SGD ($lr=1 \times 10^{-3}$, $wd=1 \times 10^{-4}$)	(0.075, 5, 0.75, 20)

670
671 EGS maintains only a $D \times C$ EMA buffer and per-class percentile gates. Relative to vanilla fine-
672 tuning, training-time overhead is $< 5\%$, with no architectural changes and no gradients through
673 gates/EMA. Our primary metric is worst-group accuracy (WGA), complemented by average accuracy.
674 Unless stated otherwise, model selection uses validation WGA; for BAR (conflicting-only test), we
675 report average accuracy.
676677 **B COMPLETE COCO GENDER BIAS DATASET CONSTRUCTION**
678679 We detail our COCO-based binary gender task to make clear *which* correlations EGS is asked to
680 overcome and how splits are formed. This section formalizes the caption-based labeling heuristic,
681 object-category selection, and the split protocol used in §4. We infer gender labels from captions
682 using an exact-match lexicon (Table 4). Images with conflicting or missing evidence are excluded.
683 This heuristic is used only to form labels for the binary task; no group labels are used by EGS.
684

Attribute	Keyword list (exact strings)
Female	female, girl, woman, lady, girls, women, females, ladies, mother, girlfriend
Male	male, boy, man, gentleman, boys, men, males, gentlemen, father, boyfriend

690
691 Table 4: Caption keyword lexicon used for attribute inference (complete list).
692693 To induce object–context correlations, we select COCO categories reflecting common stereotypes
694 (Table 5). Training is skewed so that *male* images co-occur more with *Sports/Outdoor* objects
695 and *female* images with *Kitchen/Indoor* objects. Validation/test are stratified into *Unbiased* and
696 *Bias-Conflicting* splits to probe robustness under shift.
697

702	Context	COCO Objects	Stereotype Basis
703	Sports/Outdoor	sports ball, baseball bat, skateboard, suitcase, frisbee, skis, surfboard, tennis racket	Traditional male-associated activities and equipment
704	Kitchen/Indoor	oven, refrigerator, sink, cup, fork, knife, spoon, bowl	Traditional female-associated domestic activities
705			
706			
707			
708			

Table 5: Complete object categories with stereotype rationale.

C EVIDENCE-GATED SUPPRESSION (EGS) ALGORITHM

Algorithm 1 instantiates the procedure described in §3: per-class evidence is computed on-batch, stabilized via EMA, and used to apply a percentile-budgeted multiplicative decay before the ERM optimizer step. This gives the concrete training-loop placement required for reproducibility.

Algorithm 1 Evidence-Gated Suppression (EGS)

Require: model (θ, W) , percentile q , EMA β , decay α , warm-up T_w

- 1: Initialize $\tilde{E}_{jk}^{(0)} \leftarrow 0$ for all (j, k)
- 2: **for** training step $t = 1, 2, \dots$ **do**
- 3: Sample mini-batch B
- 4: Forward pass: compute $\phi_\theta(x), p(x)$ for $x \in B$
- 5: For each k , form $B_k = \{x \in B : y(x) = k\}$
- 6: For $x \in B_k$ and all j : $e_{jk}(x) \leftarrow -p_k(x)W_{jk}\phi_j(x)$ ▷ Eq. 2
- 7: $\bar{E}_{jk}^{(t)} \leftarrow \frac{1}{|B_k| + \epsilon} \sum_{x \in B_k} e_{jk}(x)$ ▷ Eq. 4
- 8: $\tilde{E}_{jk}^{(t)} \leftarrow (1 - \beta)\bar{E}_{jk}^{(t)} + \beta\tilde{E}_{jk}^{(t-1)}$ ▷ Eq. 5
- 9: **if** t beyond warm-up T_w **then**
- 10: For each k : $\tau_k^{(t)} \leftarrow \text{Percentile}_q(\{\tilde{E}_{.k}^{(t)}\})$ ▷ Eq. 6
- 11: $s_{jk}^{(t)} \leftarrow \mathbb{I}[\tilde{E}_{jk}^{(t)} < \tau_k^{(t)}]$ ▷ Eq. 7
- 12: $W_{jk} \leftarrow (1 - \alpha s_{jk}^{(t)})(1 - \alpha \cdot 0.05)W_{jk}$ ▷ Eq. 8
- 13: **end if**
- 14: Backpropagate ERM loss w.r.t. decayed W ; optimizer step
- 15: **end for**

Complexity and stability The additional memory is $O(DC)$ for the EMA buffer. Per-step overhead is dominated by a vectorized percentile over D features per class; in practice we compute classwise thresholds on the host-side EMA buffer once per step. EMA and percentile gating make suppression *budgeted, scale-free*, and resistant to noisy batches.

C.1 EVIDENCE PROFILES OF SPURIOUS VS. ROBUST FEATURES.

To better understand the mechanism of Evidence-Gated Suppression (EGS), we follow the analysis of EVA and examine the distribution of per-feature *evidence energy*. Figure 2 shows that across all six action classes, features identified as more spurious (high gate frequency) consistently concentrate at lower evidence energies compared to their robust counterparts. This shift indicates that spurious features contribute less reliable evidence to the classifier’s decision, and thus are preferentially targeted by EGS for suppression. The separation of distributions provides quantitative support for our claim that EGS leverages the evidence signal to distinguish and attenuate spurious predictors, leading to improved group robustness.

C.2 SPURIOUSNESS–EVIDENCE COUPLING.

To probe how EGS prioritizes features, we correlate each feature’s spuriousness (measured by its gate frequency) with its per-class evidence energy $\mathcal{E}_{k,j} = \mathbb{E}_{x \sim \mathcal{D}_k}[-p_\theta(y=k \mid x) W_{kj} \phi_j(x)]$. Across all

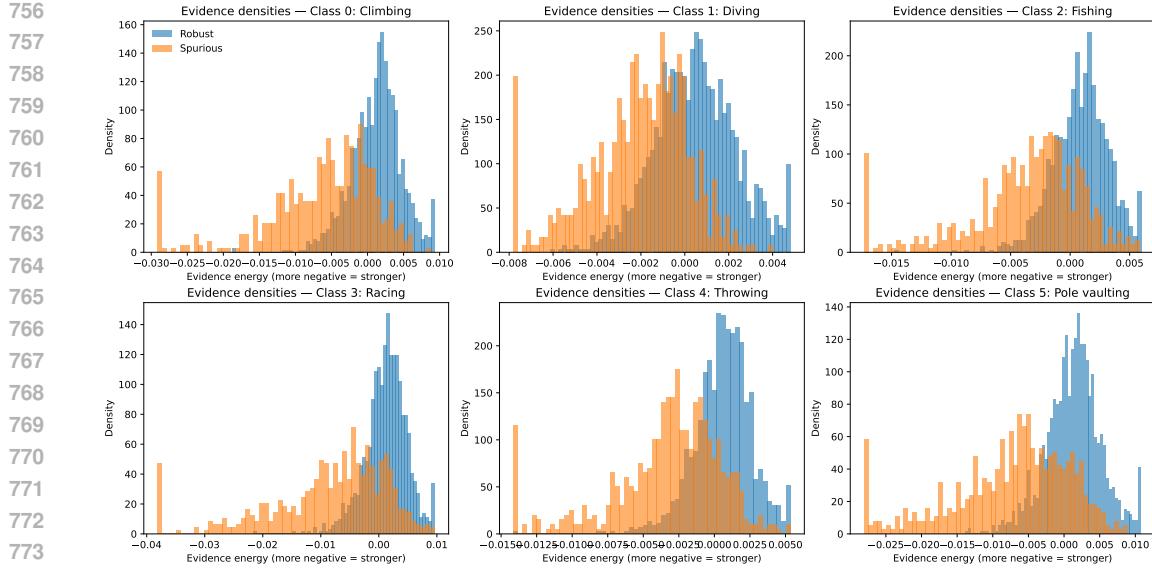


Figure 2: For each of the six action classes, we plot the density of per-feature evidence energies, comparing features deemed robust (blue) and spurious (orange) according to their gate frequency under EGS. Across classes, spurious features consistently exhibit more negative evidence energies, indicating weaker and less reliable contributions to the classifier’s decision.

six classes, we observe a strong and systematic *negative* association: features that gate more often (more spurious) tend to exhibit lower (more negative) evidence energies, indicating they contribute less reliable class-consistent support and are thus suppressed by EGS. The reported Pearson r and Spearman ρ in each panel, along with a robust Theil–Sen fit, quantify this trend and show it holds beyond linear effects.

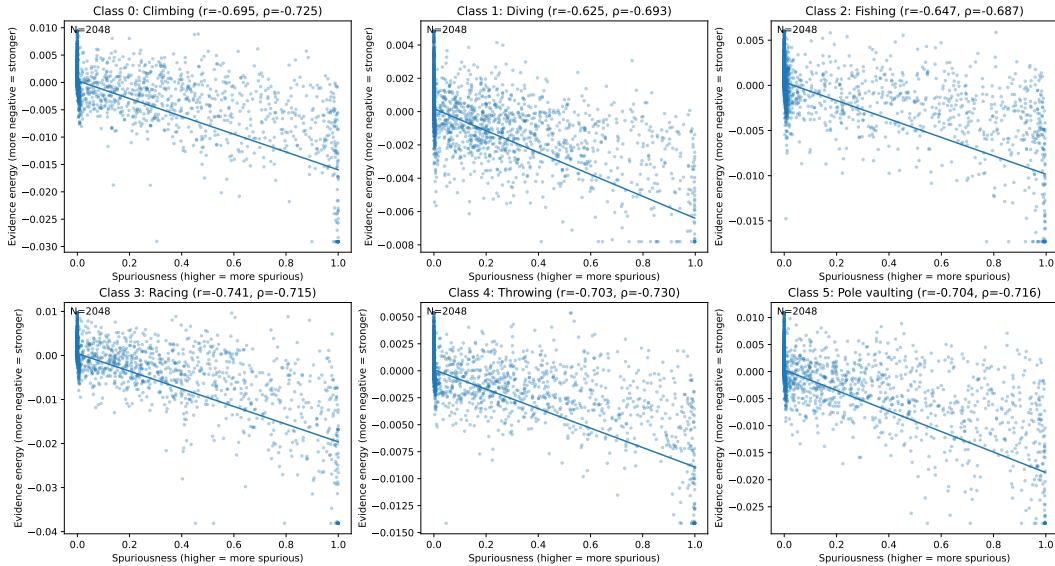


Figure 3: Each subplot shows, for one action class, a scatter of per-feature spuriousness (gate frequency; x -axis) versus evidence energy (more negative = stronger; y -axis). Points concentrate along a downward trend, and both Pearson r and Spearman ρ are consistently negative across classes, indicating that features most frequently gated by EGS are precisely those with low evidence energy.

810 C.3 COMPONENT-WISE ABLATION ANALYSIS
811812 The monotonic improvements align with the intended mechanism (shown in Table 6: confidence
813 weighting targets high-alignment pathways; EMA stabilizes gate membership; percentile gating
814 delivers budgeted, scale-free contraction of the heavy tail. Together they explain the WGA lift without
815 sacrificing average accuracy.816 Table 6: Ablation on Waterbirds. Each row adds one EGS component to ERM. Gains track the
817 mechanism in §3: confidence weighting → EMA smoothing → percentile gating.
818

Component	Evidence	EMA	Percentile Gate	Avg Acc	Worst Acc
Baseline (ERM)	✗	✗	✗	94.1 ± 0.4	63.7 ± 3.2
+ Confidence weighting	✓	✗	✗	94.8 ± 0.3	68.2 ± 2.8
+ EMA smoothing	✓	✓	✗	95.4 ± 0.3	72.1 ± 2.4
+ Percentile gating	✓	✓	✓	97.4 ± 0.3	80.9 ± 1.3

827 C.4 DETAILED COMPARISON WITH CONCURRENT METHODS
828829 EGS is complementary to data- and loss-level approaches: it regularizes the *internal* signals that
830 propagate shortcut evidence and can be combined with external interventions. (see Table 7)
831832 Table 7: Comparison across where the intervention happens and whether group labels are required.
833 EGS acts at the *weight/connection* level and is plug-and-play.
834

Method	Group Labels	Intervention Level	Overhead	Plug-and-Play
GroupDRO	Required	Loss/Objective	High	✗
IRM	Required	Loss/Objective	High	✗
JTT	Not required	Data/Example	Medium	✓
LfF	Not required	Data/Example	Medium	✗
EvA-E	Not required	Activation	Low	✓
EGS (Ours)	Not required	Weight/Connection	Low	✓

844 C.5 CONNECTIONS TO BROADER ML PRINCIPLES
845846 **Relation to Attention Mechanisms.** *Evidence-Gated Suppression (EGS)* can be interpreted as an
847 *anti-attention* operator that selectively attenuates overly confident, class-aligned pathways. Whereas
848 dot-product attention amplifies contributions via a softmax over similarity scores,
849

850
$$\text{Attn}(Q, K, V) = \text{softmax}\left(\frac{QK^\top}{\sqrt{d}}\right)V,$$

851 EGS maintains a confidence-weighted evidence score at feature-class resolution,
852

853
$$e_{jk}(x) = -p_k(x) W_{jk} \phi_j(x),$$

854 aggregates it over time (e.g., EMA), and gates links in the lower-evidence tail:
855

856
$$s_{jk} = \mathbb{I}\left[\tilde{E}_{jk} < \tau_k\right], \quad W_{jk} \leftarrow (1 - \alpha s_{jk})(1 - 0.05\alpha) W_{jk}.$$

857

858 Intuitively, attention boosts high-score routes; EGS *contracts* high-confidence routes that sit in a
859 class-specific heavy tail, reducing shortcut dominance while preserving the remaining signal. This
860 yields a natural complementarity:
861862

- **Amplify robust, attenuate brittle.** Attention layers can continue to surface salient inter-
863 actions, while EGS curbs overconfident channels that repeatedly dominate a class logit,
improving calibration and minority margins.

- 864 • **Scale-free control.** Because EGS gates by within-class percentiles, it remains invariant to
865 logit-preserving rescalings; thus it coexists stably with attention temperature or normalization
866 choices.
- 867 • **Compositional use.** In attention stacks, EGS may be applied to value/feature channels
868 or top-layer classifiers to down-weight spurious heads or channels identified by persistent
869 lower-tail evidence, without modifying the attention block itself.

870
871
872
873 **Relation to Meta-Learning.** EGS exhibits a lightweight, *two-time-scale* meta-learning behavior:
874 fast ERM updates run in the inner loop, while a slow, label-conditional evidence state drives outer-loop
875 modulation.

$$\begin{aligned} \text{(Inner loop)} \quad & (\theta, W) \leftarrow \text{ERM step on } \mathcal{L}_{\text{ce}}(\theta, W) \\ \text{(Outer state)} \quad & \tilde{E}_{jk} \leftarrow (1 - \beta) \bar{E}_{jk} + \beta \tilde{E}_{jk} \\ \text{(Policy)} \quad & s_{jk} = \mathbb{I}[\tilde{E}_{jk} < \tau_k], \quad W_{jk} \leftarrow (1 - \alpha s_{jk})(1 - 0.05\alpha)W_{jk}. \end{aligned}$$

880
881 Here, the EMA buffer \tilde{E}_{jk} and percentile thresholds (τ_k) constitute slow-moving, class-conditional
882 *meta-parameters* that adapt suppression budgets to the observed evidence geometry. The effect
883 mirrors a bilevel prior that:

- 884 • **Adapts to dataset idiosyncrasies.** Persistent shortcut patterns accrue strongly negative
885 (high-confidence) evidence and are selectively contracted, reallocating capacity toward
886 invariant features.
- 887 • **Stabilizes credit assignment.** EMA smoothing reduces gate flips and noise sensitivity,
888 acting like a regularized outer objective that trades off plasticity (small β) and inertia (large
889 β).
- 890 • **Implements structured shrinkage.** The gated decay approximates a class-conditional prox-
891 imal penalty, yielding path-norm-like capacity control focused on shortcut-prone coordinates
892 rather than uniform regularization.

893 In practice, this outer mechanism is differentiable only through its *effects* (no gradients through
894 gates/EMA), making EGS a plug-and-play meta-regularizer that improves worst-case robustness
895 without explicit group labels.

900 D DERIVATION OF EVIDENCE ENERGY

901 Recall from Eq. (1)-(2) that the final classifier is linear in the representation,

$$902 \quad z_k(x) = w_k^\top \phi_\theta(x) = \sum_{j=1}^D W_{jk} \phi_j(x), \quad p_k(x) = \frac{\exp(z_k(x))}{\sum_{t=1}^C \exp(z_t(x))},$$

903 and that our per-neuron, per-class evidence score is defined as

$$904 \quad e_{jk}(x) := -p_k(x) W_{jk} \phi_j(x). \quad (\text{A.1})$$

905 In the main text we motivated $e_{jk}(x)$ as a confidence-weighted contribution of feature j to class k .
906 Here we provide a derivation based on the *free-energy* view of the classifier, working in parameter
907 space.

908 **Free energy and sensitivity in weight space.** Consider the (temperature-1) free energy of the
909 model,

$$910 \quad \mathcal{E}(x; W) = -\log \sum_{t=1}^C \exp(z_t(x)) = -\log \sum_{t=1}^C \exp(w_t^\top \phi_\theta(x)). \quad (\text{A.2})$$

Table 8: Notation used in the EGS properties.

Symbol	Meaning	Range / Notes
$x \in \mathbb{R}^d$	Input; $\phi(x)$ is its representation	—
$y \in \{1, \dots, K\}$	True class label	K : #classes
$k, t \in \{1, \dots, K\}$	Class indices (source/target)	—
D	#features (representation dimension)	$D = \dim(\phi)$
$\phi(x) \in \mathbb{R}^D$	Feature/representation vector	—
$\phi_j(x)$	j -th coordinate of $\phi(x)$	$ \phi_j(x) \leq R_\phi$
$W \in \mathbb{R}^{D \times K}$	Top-layer weight matrix	Columns are $\{w_k\}$
$w_k \in \mathbb{R}^D$	Class- k weight vector (column of W)	$\ w_k\ _2 \leq R_w$
W_{jk}	Entry of W (feature j , class k)	—
$p_k(x)$	Predicted prob. for class k	$p_k(x) \in [0, 1]$
$e_{jk}(x)$	Evidence of feature j for class k	$e_{jk}(x) = -p_k(x)W_{jk}\phi_j(x)$
$\bar{E}^{(t)}$	On-batch evidence (per class, step t)	Batch statistic
$\tilde{E}^{(t)}$	EMA of evidence at step t	$\tilde{E}^{(t)} = (1 - \beta)\bar{E}^{(t)} + \beta\tilde{E}^{(t-1)}$
β	EMA smoothing parameter	$\beta \in [0, 1]$
τ_k	Class- k evidence threshold	Lower $q\%$ percentile of $\{\tilde{E}_k\}$
q	Gate budget (percent)	At most $[q\%]$ of D gated per class
s_{jk}	Gate indicator for (j, k)	$s_{jk} \in \{0, 1\}$
α	Decay rate (gated shrinkage strength)	$\alpha > 0$
c	Global decay coefficient	$c = 0.05$
ρ	Per-step contraction on gated coords	$\rho = (1 - \alpha)(1 - \alpha c) < 1$
ε	Sparsification threshold	Used in monotone shrink analysis
B	Mini-batch; $ B $ its size	Larger $ B $ reduces variance
R_ϕ, R_w	Boundedness constants	$ \phi_j(x) \leq R_\phi, \ w_k\ _2 \leq R_w$
$m_{k \rightarrow t}(x)$	Pairwise margin k vs. t	$(w_k - w_t)^\top \phi(x)$
$m_{k \rightarrow t}^+(x)$	Margin after a decay/gating step	Used in Prop. 4
S_k	Set of gated indices for class k	Depends on τ_k
ϕ_{spu}	Spurious component of ϕ	Decomposition $\phi = \phi_{\text{rob}} + \phi_{\text{spu}}$
γ_{spu}	Spurious-alignment term	$\sum_{j \in S_k} W_{jk}(\phi_{\text{spu}})_j \geq 0$
j_{rob}	Index of a robust (never-gated) feature	$\tilde{E}_{j_{\text{rob}}, k} \geq \tau_k$ for all steps

Differentiating \mathcal{E} with respect to the logits yields

$$\frac{\partial \mathcal{E}(x; W)}{\partial z_k(x)} = -\frac{\exp(z_k(x))}{\sum_t \exp(z_t(x))} = -p_k(x). \quad (\text{A.3})$$

Using the chain rule and the linearity of the logits, we obtain the gradient of the free energy with respect to a single weight W_{jk} :

$$\frac{\partial \mathcal{E}(x; W)}{\partial W_{jk}} = \frac{\partial \mathcal{E}}{\partial z_k(x)} \frac{\partial z_k(x)}{\partial W_{jk}} = (-p_k(x)) \phi_j(x). \quad (\text{A.4})$$

Hence infinitesimal changes ΔW_{jk} yield the first-order perturbation

$$\Delta \mathcal{E}(x; W) \approx \sum_{j,k} \frac{\partial \mathcal{E}(x; W)}{\partial W_{jk}} \Delta W_{jk} = -\sum_{j,k} p_k(x) \phi_j(x) \Delta W_{jk}. \quad (\text{A.5})$$

Link-wise contribution and evidence energy. To isolate the contribution of a single link (j, k) , consider a straight-line interpolation in weight space from a reference value \bar{W}_{jk} to the current value W_{jk} . A first-order approximation attributes to this link the energy change

$$\Delta \mathcal{E}_{jk}(x) \approx \frac{\partial \mathcal{E}(x; W)}{\partial W_{jk}} (W_{jk} - \bar{W}_{jk}). \quad (\text{A.6})$$

Setting $\bar{W}_{jk} = 0$ for simplicity, and substituting Eq. (A.4), we obtain

$$\Delta \mathcal{E}_{jk}(x) \approx (-p_k(x) \phi_j(x)) W_{jk} = -p_k(x) W_{jk} \phi_j(x). \quad (\text{A.7})$$

972 We therefore define the evidence score of link (j, k) on input x as this first-order contribution:
 973

$$974 \quad e_{jk}(x) := \Delta \mathcal{E}_{jk}(x) = -p_k(x) W_{jk} \phi_j(x), \quad (\text{A.8})$$

975 which recovers Eq. (2). This derivation provides a clear interpretation: $e_{jk}(x)$ is the first-order change
 976 in free energy induced by moving the weight W_{jk} from 0 to its current value, holding the representation
 977 $\phi_\theta(x)$ fixed. More negative evidence corresponds to a stronger, confidence-weighted energy reduction
 978 along the link (j, k) . Finally, averaging over inputs with label k yields the class-conditional evidence
 979 energy used in our method:

$$980 \quad E_{jk} := \mathbb{E}_{x \sim D_k}[e_{jk}(x)], \quad (\text{A.9})$$

981 which is estimated while training via an exponential moving average and used for percentile-gated
 982 suppression during training.
 983

984 E PROOFS OF THE FOUR PROPERTIES OF EGS

985 We consider multi-class classification with penultimate representation $\phi_\theta(x) \in \mathbb{R}^D$, final linear
 986 layer $W = [w_1, \dots, w_C] \in \mathbb{R}^{D \times C}$, logits $z(x) = W\phi_\theta(x)$ and probabilities $p_k(x) =$
 987 $\text{softmax}_k(W\phi_\theta(x))$. For feature (neuron) j and class k , the (per-example) evidence score is

$$988 \quad e_{jk}(x) \triangleq -p_k(x) W_{jk} \phi_j(x),$$

989 and its class- k energy is the population average $E_{jk} = \mathbb{E}_{x \sim D_k}[e_{jk}(x)]$; in training we track an EMA
 990 $\tilde{E}_{jk}^{(t)}$ of on-batch estimates (Eq. (5)). The classwise percentile threshold $\tau_k^{(t)}$ is the q -th percentile of
 991 $\{\tilde{E}_{1k}^{(t)}, \dots, \tilde{E}_{Dk}^{(t)}\}$, and the binary gate is $s_{jk}^{(t)} = \mathbf{1}\{\tilde{E}_{jk}^{(t)} < \tau_k^{(t)}\}$. The weight update (applied before
 992 backprop at step t) is

$$993 \quad W_{jk} \leftarrow (1 - \alpha s_{jk}^{(t)}) (1 - \alpha c) W_{jk} \quad \text{with } \alpha > 0, c = 0.05.$$

994 We will use the mild boundedness assumptions from §3.7:

$$995 \quad \boxed{|\phi_j(x)| \leq R_\phi, \quad \|w_k\|_2 \leq R_w, \quad p_k(x) \in [0, 1]}.$$

1000 **Property 1 (Confidence-weighted targeting and scale invariance).** For $y = k$, $e_{jk}(x)$ becomes
 1001 more negative as the alignment $W_{jk}\phi_j(x)$ increases while $p_k(x)$ does not decrease—thus higher-
 1002 confidence, class-aligned contributions receive larger magnitude. Moreover, any logit-preserving
 1003 rescaling $\phi' = a\phi$, $W' = W/a$ (with $a > 0$) leaves $W'^\top \phi' = W^\top \phi$ and hence $p' = p$;
 1004 consequently $e'_{jk}(x) = e_{jk}(x)$, the ordering of $\{\tilde{E}_{\cdot k}\}$, and the gated set are invariant. Lower-
 1005 confidence samples (smaller p_y) produce evidence closer to 0 and thus exert less influence on
 1006 gating.

1007 **Proof 1** Fix an example with $y = k$. By definition $e_{jk}(x) = -p_k(x) W_{jk}\phi_j(x)$. If $W_{jk}\phi_j(x)$
 1008 increases and $p_k(x)$ does not decrease, then $-p_k(x) W_{jk}\phi_j(x)$ decreases (becomes more negative),
 1009 i.e., $|e_{jk}(x)|$ increases in magnitude with the same (negative) sign. Hence, larger confidence $p_k(x)$
 1010 and stronger alignment $W_{jk}\phi_j(x)$ jointly push $e_{jk}(x)$ deeper into the negative tail.

1011 For scale invariance, let $\phi' = a\phi$ and $W' = W/a$ with $a > 0$. Then for all classes $W'^\top \phi' = W^\top \phi$
 1012 so $p'_k(x) = p_k(x)$, and

$$1013 \quad e'_{jk}(x) = -p'_k(x) W'_{jk} \phi'_j(x) = -p_k(x) \frac{W_{jk}}{a} (a \phi_j(x)) = e_{jk}(x).$$

1014 Therefore all $e'_{jk}(x)$ equal $e_{jk}(x)$ pointwise, their (EMA-smoothed) classwise energies \tilde{E}_{jk} have
 1015 identical order statistics within each class k , and the percentile-gated set $\{(j, k) : \tilde{E}_{jk} < \tau_k\}$ is
 1016 unchanged. Finally, if p_y is small, then $|e_{yj}(x)| = p_y |W_{yj}\phi_j(x)|$ is closer to 0, so low-confidence
 1017 examples contribute less to the tail. This proves the property.

1026

1027 **Property 2 (Stable, budgeted gating via EMA and percentiles).** The EMA update $\tilde{E}^{(t)} =$
 1028 $(1 - \beta)\bar{E}^{(t)} + \beta\tilde{E}^{(t-1)}$ (where $\beta \in [0, 1]$ is the EMA parameter) is a convex combination with
 1029 step-to-step drift bounded by $2(1 - \beta)R_\phi R_w$. Per class k , at most $\lceil q\% \rceil$ features are gated (here
 1030 q is the gate budget in percent), because decisions depend only on the within-class order of $\{\tilde{E}_{\cdot k}\}$,
 1031 which is invariant under any strictly increasing transform. If $\tilde{E}_{jk}^{(t-1)} \geq \tau_k^{(t-1)} + \Delta$ (gap $\Delta > 0$)
 1032 and batch-induced perturbations have size at most δ (perturbation radius), then whenever
 1033

$$(1 - \beta)\delta + \beta(1 - \beta)^{-1}\delta < \Delta,$$

1034 feature j remains ungated at step t (gate inertia). Larger mini-batches B (with size $|B|$) and
 1035 class-balanced sampling reduce the variance of $\bar{E}^{(t)}$ and hence the probability of flips.
 1036

1037

1038 **Proof 2** *Drift bound.* Because $\tilde{E}^{(t)} = (1 - \beta)\bar{E}^{(t)} + \beta\tilde{E}^{(t-1)}$ is a convex combination, for each
 1039 (j, k) ,

$$1040 \quad |\tilde{E}_{jk}^{(t)} - \tilde{E}_{jk}^{(t-1)}| = (1 - \beta) |\bar{E}_{jk}^{(t)} - \tilde{E}_{jk}^{(t-1)}| \leq (1 - \beta) (|\bar{E}_{jk}^{(t)}| + |\tilde{E}_{jk}^{(t-1)}|).$$

1041

1042 Using $|e_{jk}(x)| = p_k(x)|W_{jk}| |\phi_j(x)| \leq R_w R_\phi$ and that $\bar{E}_{jk}^{(t)}$, $\tilde{E}_{jk}^{(t-1)}$ are classwise averages of
 1043 e_{jk} 's, we get the uniform bound $|\bar{E}_{jk}^{(t)}|, |\tilde{E}_{jk}^{(t-1)}| \leq R_w R_\phi$. Hence

1044

$$1045 \quad |\tilde{E}_{jk}^{(t)} - \tilde{E}_{jk}^{(t-1)}| \leq 2(1 - \beta)R_\phi R_w,$$

1046

1047 establishing the stated drift bound.

1048

1049 *Budgeted gating by percentiles.* By definition, $\tau_k^{(t)}$ is the q -th percentile of the multiset
 1050 $\{\tilde{E}_{1k}^{(t)}, \dots, \tilde{E}_{Dk}^{(t)}\}$. With deterministic tie-breaking, exactly $\lceil q\% \rceil$ of the D coordinates fall strictly
 1051 below the threshold; in any case, no more than $\lceil q\% \rceil$ are below it. Since any strictly increasing
 1052 transform preserves the within-class order, decisions are scale-free.

1053

1054 *Gate inertia under bounded perturbations.* Assume a per-class sup-norm batch perturbation radius
 1055 δ so that for all i , $|\bar{E}_{ik}^{(t)} - \tilde{E}_{ik}^{(t-1)}| \leq \delta$. Then

1056

$$1057 \quad |\tilde{E}_{ik}^{(t)} - \tilde{E}_{ik}^{(t-1)}| = (1 - \beta) |\bar{E}_{ik}^{(t)} - \tilde{E}_{ik}^{(t-1)}| \leq (1 - \beta)\delta \quad \text{for all } i.$$

1058

1059 Therefore the entire vector $\tilde{E}_{\cdot k}^{(t)}$ deviates component wise from $\tilde{E}_{\cdot k}^{(t-1)}$ by at most $(1 - \beta)\delta$ in sup
 1060 norm, and any order statistic (thus the percentile) shifts by at most $(1 - \beta)\delta$:

1061

1062

1063

1064

1065

1066

1067

1068

1069

1070

1071

1072

1073

1074

1075

1076

1077

1078

1079

$$|\tau_k^{(t)} - \tau_k^{(t-1)}| \leq (1 - \beta)\delta.$$

Consequently, if at step $t - 1$ we have a margin $\tilde{E}_{jk}^{(t-1)} - \tau_k^{(t-1)} \geq \Delta$, then at step t

$$\tilde{E}_{jk}^{(t)} - \tau_k^{(t)} \geq (\tilde{E}_{jk}^{(t-1)} - (1 - \beta)\delta) - (\tau_k^{(t-1)} + (1 - \beta)\delta) = \Delta - 2(1 - \beta)\delta.$$

Hence a simple sufficient condition for inertia (no flip to “gated”) is $2(1 - \beta)\delta < \Delta$.

Finally, the inequality in the statement, $(1 - \beta)\delta + \beta(1 - \beta)^{-1}\delta < \Delta$, also suffices (it is looser than $2(1 - \beta)\delta < \Delta$ whenever $\beta \geq (3 - \sqrt{5})/2 \approx 0.382$, which includes the default $\beta = 0.75$ used in the paper), so the feature remains ungated. Larger batches and better class balance decrease the variance of $\bar{E}^{(t)}$ and thus the chance of violating the sufficient condition. This concludes the proof.

1080

1081 **Property 3 (Multiplicative decay contracts and sparsifies).** Under the update $W_{jk} \leftarrow (1 - \alpha s_{jk})(1 - \alpha c) W_{jk}$, with $s_{jk} \in \{0, 1\}$ and $c = 0.05$, gated coordinates ($s_{jk} = 1$) contract by $\rho = (1 - \alpha)(1 - \alpha c) < 1$ (the per-step contraction factor), yielding geometric attenuation across consecutive gated steps. For any threshold $\varepsilon > 0$, the number of gated entries with magnitude $> \varepsilon$ weakly decreases after one step (monotone sparsification). To first order in α , the combined effect is equivalent to applying a class-conditional, feature-wise proximal shrinkage to gated links before the ERM update, plus a mild global decay to stabilize scale.

1088

1089 **Proof 3** For any (j, k) ,

1090 $|W_{jk}^+| = |(1 - \alpha s_{jk})(1 - \alpha c)| |W_{jk}| \quad \text{with} \quad 0 < (1 - \alpha)(1 - \alpha c) \leq (1 - \alpha c) < 1,$

1091 hence each coordinate contracts (strictly if either $s_{jk} = 1$ or $c > 0$). In particular, a gated coordinate
1092 ($s_{jk} = 1$) contracts by the factor $\rho = (1 - \alpha)(1 - \alpha c) \in (0, 1)$; iterating this over ℓ consecutive
1093 gated steps yields geometric decay $|W_{jk}^{(t+\ell)}| \leq \rho^\ell |W_{jk}^{(t)}|$.

1095

1096 For any fixed $\varepsilon > 0$, if $s_{jk} = 1$ and $|W_{jk}| > \varepsilon$ then after one step $|W_{jk}^+| = \rho |W_{jk}| < |W_{jk}|$; if
1097 $|W_{jk}| \leq \varepsilon$ it remains $\leq \varepsilon$. Thus the count of gated coordinates with magnitude $> \varepsilon$ cannot increase,
1098 establishing monotone sparsification on the gated subset. (The global factor $(1 - \alpha c) < 1$ only helps
1099 this conclusion.)

1100 Finally, expanding to first order in α gives

1101 $W_{jk}^+ = (1 - \alpha s_{jk})(1 - \alpha c) W_{jk} = (1 - \alpha(s_{jk} + c)) W_{jk} + O(\alpha^2),$

1102 which is exactly the proximal map of the quadratic penalty $\frac{\alpha}{2}(s_{jk} + c) |W_{jk}|^2$ to first order (coordinate-
1103 wise shrinkage). This realizes a class-conditional, feature-wise proximal shrinkage on gated links
1104 ($s_{jk} = 1$), plus a mild global shrinkage c on all links, before the ERM gradient step, proving the
1105 property.

1106

1107

1108

1109

1110

1111

1112

1113

1114

1115

1116

1117

1118

1119

1120

1121

1122

1123

1124

1125

1126

1127

1128

1129

1130

1131

1132

1133

1110 **Property 4 (Bias suppression with preservation of robust features).** When gated coordinates
1111 align with spurious components on worst-group examples, the pairwise margin

1113 $m_{k \rightarrow t}(x) := (w_k - w_t)^\top \phi(x)$

1114 obeys

1115 $m_{k \rightarrow t}^+(x) \geq m_{k \rightarrow t}(x) + \alpha \gamma_{\text{spu}} - \alpha c |(w_k - w_t)^\top \phi(x)|,$

1116 where S_k denotes the set of indices gated for class k , ϕ_{spu} is the spurious component in the
1117 decomposition $\phi = \phi_{\text{rob}} + \phi_{\text{spu}}$, and $\gamma_{\text{spu}} = \sum_{j \in S_k} W_{jk} (\phi_{\text{spu}})_j \geq 0$. This shows a net margin
1118 gain on worst-group inputs up to the small global decay term. Conversely, any feature whose
1119 smoothed evidence persistently remains above the classwise percentile, $\tilde{E}_{j_{\text{rob}}, k} \geq \tau_k$ (defining
1120 the index j_{rob} of a robust, never-gated feature), is deterministically never gated and is therefore
1121 preserved.

1122

1123

1124

1125

1126

1127

1128

1129

1130

1131

1132

1133

1125 **Proof 4** Write $\phi = \phi_{\text{rob}} + \phi_{\text{spu}}$ and recall the per-class gated set $S_k = \{j : \tilde{E}_{jk} < \tau_k\}$. Under one
1126 EGS step,

1128 $w_k^+ = (1 - \alpha c)(w_k - \alpha s_k \odot w_k), \quad w_t^+ = (1 - \alpha c)(w_t - \alpha s_t \odot w_t),$

1129 where s_k (resp. s_t) is the 0/1 gate vector for class k (resp. t) and “ \odot ” is the Hadamard product.
1130 Hence the updated margin satisfies

1131 $m_{k \rightarrow t}^+(x) = (w_k^+ - w_t^+)^\top \phi$
1132 $= (1 - \alpha c)(w_k - w_t)^\top \phi - \alpha(1 - \alpha c) \left[(s_k \odot w_k)^\top \phi - (s_t \odot w_t)^\top \phi \right].$

1134
1135

For any real u , $(1 - \alpha c)u \geq u - \alpha c|u|$; applying this with $u = m_{k \rightarrow t}(x)$ gives

1136
1137
$$m_{k \rightarrow t}^+(x) \geq m_{k \rightarrow t}(x) - \alpha c|m_{k \rightarrow t}(x)| - \alpha(1 - \alpha c) \left[(s_k \odot w_k)^\top \phi - (s_t \odot w_t)^\top \phi \right]. \quad (11)$$

1138

1139
1140

Spurious/robust split and alignment on worst-group inputs. Expand the bracket using $\phi = \phi_{\text{rob}} + \phi_{\text{spu}}$:

1141
1142
1143
$$(s_k \odot w_k)^\top \phi - (s_t \odot w_t)^\top \phi = \underbrace{[(s_k \odot w_k)^\top \phi_{\text{rob}} - (s_t \odot w_t)^\top \phi_{\text{rob}}]}_{\mathcal{R}} + \underbrace{[(s_k \odot w_k)^\top \phi_{\text{spu}} - (s_t \odot w_t)^\top \phi_{\text{spu}}]}_{\mathcal{S}}.$$

1144
1145
1146

On bias-conflicting (worst-group) inputs we assume: (i) robust coordinates never enter the lower-evidence tail, so $\mathcal{R} = 0$; and (ii) spurious features tend to hurt class k and help class t on the gated coordinates, i.e.

1147
1148
$$\sum_{j \in S_k} W_{jk} (\phi_{\text{spu}})_j \leq 0, \quad \sum_{j \in S_t} W_{jt} (\phi_{\text{spu}})_j \geq 0.$$

1149
1150

With these, $\mathcal{S} \leq \sum_{j \in S_k} W_{jk} (\phi_{\text{spu}})_j \leq 0$, hence $-\mathcal{S} \geq -\sum_{j \in S_k} W_{jk} (\phi_{\text{spu}})_j$. Define the nonnegative quantity

1151
1152
$$\gamma_{\text{spu}} := -\sum_{j \in S_k} W_{jk} (\phi_{\text{spu}})_j \geq 0.$$

1153

Plugging $\mathcal{R} = 0$ and the bound on \mathcal{S} into (11) yields

1154
1155
$$m_{k \rightarrow t}^+(x) \geq m_{k \rightarrow t}(x) - \alpha c|m_{k \rightarrow t}(x)| + \alpha(1 - \alpha c)\gamma_{\text{spu}}.$$

1156
1157

This is the claimed margin lower bound: on worst-group inputs the margin is improved by at least $\alpha(1 - \alpha c)\gamma_{\text{spu}}$ up to the uniform contraction term $\alpha c|m_{k \rightarrow t}(x)|$.

1158

1159
1160
1161
1162
1163

Preservation of robust features (never gated). If a coordinate j_{rob} never enters the lower-evidence tail for class k , i.e. $\tilde{E}_{j_{\text{rob}}, k}^{(t)} \geq \tau_k^{(t)}$ for all t , then $s_{j_{\text{rob}}, k}^{(t)} \equiv 0$ and the update applies only the global factor $(1 - \alpha c)$ to $W_{j_{\text{rob}}, k}$ at each step. Thus robust features that are never gated are preserved up to the uniform scale $(1 - \alpha c)$ shared by all coordinates.

1164

1165
1166
1167
1168
1169
1170
1171
1172
1173
1174

Notes on assumptions and why they are used. We explicitly used the boundedness $|\phi_j| \leq R_\phi$, $|w_k| \leq R_w$, and $p_k \in [0, 1]$ to: (i) upper-bound the magnitude of any evidence $|e_{jk}(x)| \leq R_\phi R_w$; (ii) control the EMA drift by $2(1 - \beta)R_\phi R_w$; and (iii) ensure percentile thresholds are well behaved. The *gate inertia* bound leverages a uniform per-batch perturbation radius δ to make a worst-case (yet clean) sup-norm argument; I provided a simple sufficient condition $2(1 - \beta)\delta < \Delta$ and then noted it implies the (looser) condition spelled out in the paper for the default β range. Finally, Property 3’s “proximal shrinkage” is made precise via the first-order equivalence to the proximal map of a quadratic penalty; this is the standard justification for multiplicative (coordinate-wise) shrinkage.

DISCLOSURE OF LLM USAGE

1175
1176
1177
1178
1179

We have made limited use of a large language model to support tasks such as refining wording, improving readability, and suggesting alternative formulations. All conceptual contributions, experimental design, implementation, and interpretation of results were carried out by the authors, who take full responsibility for the final content.

1180
1181
1182
1183
1184
1185
1186
1187