

# 000 RDNAS: ROBUST DUAL-BRANCH NEURAL ARCHI- 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 RDNAS: ROBUST DUAL-BRANCH NEURAL ARCHI- TECTURE SEARCH

Anonymous authors

Paper under double-blind review

## ABSTRACT

Deep neural networks achieve impressive accuracy yet remain highly susceptible to adversarial perturbations, limiting their deployment in safety-critical domains. We propose **RDNAS**, a robust dual-branch neural architecture search framework that jointly optimizes standard (clean) accuracy and adversarial robustness. RDNAS introduces a dual-branch cell with separate “normal” and “robust” pathways, fused via a lightweight attention module to capture complementary representations without significantly enlarging the search space. To reliably score candidate operations under adversarial training, we develop **ROSE** (Robust Outlier-Aware Shapley Estimator), which stabilizes Shapley-based evaluation via median-of-means smoothing and interquartile-range filtering, reducing bias from noisy gradients. RDNAS consistently discovers architectures that outperform both hand-crafted networks and state-of-the-art robust NAS baselines across CIFAR-10, CIFAR-100, SVHN, and Tiny-ImageNet. Notably, it achieves 52.6% PGD<sup>20</sup> robustness on CIFAR-10 while maintaining strong clean accuracy. Extensive ablations validate the effectiveness of the dual-branch design, attention-based fusion, and robustness-aware search. Overall, RDNAS provides a scalable and effective framework for discovering architectures resilient to adversarial attacks.

## 1 INTRODUCTION

Deep neural networks (DNNs) have achieved strong performance across vision tasks such as classification, segmentation, and detection (Krizhevsky et al., 2012; He et al., 2015). Yet their deployment in safety-critical domains—autonomous driving (Chen & Huang, 2017), medical imaging (Taigman et al., 2014), or biometric authentication—remains limited by their vulnerability to adversarial perturbations: small, imperceptible input changes that can drastically alter predictions. This fragility motivates research into adversarial robustness, where adversarial training (Madry et al., 2019; Zhang et al., 2019) is the prevailing defense. While effective, adversarial training alone cannot fully exploit robustness potential if the network architecture itself is not designed with robustness in mind.

Neural architecture search (NAS) offers a way to automatically discover architectures with strong performance. However, most NAS approaches—including differentiable (Heuillet et al., 2024; Yang et al., 2020; Chen et al., 2019b), zero-shot (Mellor et al., 2021; Peng et al., 2024), reinforcement learning (Zoph & Le, 2017; Baker et al., 2017), or evolutionary methods (Zhou et al., 2023; 2024)—optimize almost exclusively for clean accuracy, while treating robustness as secondary. Recent robust NAS studies (Simon et al., 2022; Feng et al., 2025a) incorporate adversarial training into the search loop, but typically score candidate operations via noisy adversarial gradients, leading to unstable updates and inconsistent robustness (Dong et al., 2025b; Cheng et al., 2023). For example, Shapley-based methods (Xiao et al., 2022) can be skewed by outliers, yielding unreliable attribution.

We propose **RDNAS** (Robust Dual-Branch NAS), a robustness-aware framework that jointly optimizes clean (standard) accuracy and adversarial robustness. RDNAS builds on three ideas: **(i)** A dual-branch cell that disentangles feature representation learning into a *normal* and a *robust* pathway, fused through a lightweight attention mechanism. Unlike ensembles or mixture-of-experts, this design does not duplicate operators or widen the search combinatorics; instead, it introduces a single extra dimension (one additional cell type) while keeping the operator set unchanged, thereby maintaining search efficiency. **(ii)** Adversarial training directly embedded in the inner loop of NAS, steering the search toward architectures that are in-

054 herently robust to perturbations. **(iii) ROSE** (Robust Outlier-Aware Shapley Estimator), a  
 055 stabilized Shapley-based scoring mechanism that combines median-of-means smoothing with  
 056 interquartile-range filtering to dampen gradient noise and reward rare but decisive operation effects.  
 057 To keep the search tractable, we couple ROSE  
 058 with a few-sample evaluation strategy that re-  
 059 duces both the supernet training and scoring  
 060 costs, achieving stability without significant com-  
 061 putational overhead. Moreover, inspired by ro-  
 062 bustness studies on wide architectures (Wu et al.,  
 063 2025), RDNAS prioritizes width over depth,  
 064 achieving resilience with fewer layers and simi-  
 065 lar parameter budgets to conventional NAS cells.  
 066

067 We validate RDNAS on CIFAR-10 (Krizhevsky,  
 068 2009), CIFAR-100 (Krizhevsky, 2009), SVHN,  
 069 and Tiny-ImageNet (Deng et al., 2009) against  
 070 diverse adversarial attacks (including FGSM  
 071 (Goodfellow et al., 2015), PGD (Madry et al.,  
 072 2019), and transfer-based attacks (Papernot et al.,  
 073 2017)). As shown in Figure 1, RDNAS consis-  
 074 tently surpasses both hand-crafted and NAS base-  
 075 lines in clean and robust accuracy. Our contribu-  
 076 tions are threefold: **(i)** A dual-branch cell design  
 077 that explicitly separates and fuses clean and ro-  
 078 bust pathways, providing complementary repre-  
 079 sentations without inflating the search space. **(ii)** A robustness-aware scoring mechanism (ROSE)  
 080 that stabilizes Shapley attribution under adversarial training through principled statistical techniques.  
 081 **(iii)** Empirical evidence that RDNAS discovers architectures with superior robustness and accuracy  
 082 across datasets and attacks, while maintaining computational efficiency.

083 To ensure both rigor and reproducibility, we complement our empirical study with **theoretical anal-**  
 084 **ysis** of the proposed estimator and **release full source code** and search logs in the appendix and  
 085 supplementary material.

## 2 RELATED WORK

### 2.1 ADVERSARIAL ROBUSTNESS (ATTACKS VS. DEFENSES)

086 Deep neural networks (DNNs) have achieved remarkable performance across many tasks yet remain  
 087 vulnerable to adversarial examples, i.e., small (often human-imperceptible) perturbations that can  
 088 significantly alter model predictions. Adversarial attacks are commonly categorized into *white-box*  
 089 and *black-box* settings. In the white-box case, the attacker assumes full access to model parameters  
 090 and gradients: FGSM (Goodfellow et al., 2015) performs a fast single-step perturbation aligned with  
 091 the loss gradient, whereas PGD (Madry et al., 2019) performs multi-step projected gradient ascent  
 092 within an  $\ell_\infty$  ball of radius  $\epsilon$  and serves as a strong baseline for white-box evaluation. In the black-  
 093 box case, attacks operate under limited information, relying either on transferability (adversarial  
 094 examples crafted on a surrogate model) or on query-based methods that estimate gradients from  
 095 model outputs. AutoAttack (Croce & Hein, 2020) provides a parameter-free evaluation suite that  
 096 combines strong white-box attacks (APGD-CE/APGD-DLR, FAB) with a black-box component  
 097 (Square), and is widely used as a reliable robustness benchmark.

098 To counter adversarial threats, a variety of defenses have been proposed. Among existing strategies,  
 099 adversarial training has proven to be one of the most effective—in particular, PGD-based adversarial  
 100 training, which is formulated as a min–max optimization between the model and the adversary.  
 101 Notable extensions include TRADES (Zhang et al., 2019), which balances standard accuracy and  
 102 robustness through a regularized loss. Despite these advances, the robustness of a model remains  
 103 closely tied to its underlying architecture. Architectures not explicitly designed for robustness often  
 104 exhibit an inherent trade-off between clean accuracy and adversarial performance. This observation  
 105 underscores the need for robustness-aware NAS, where adversarial objectives are integrated directly  
 106 into the search process.

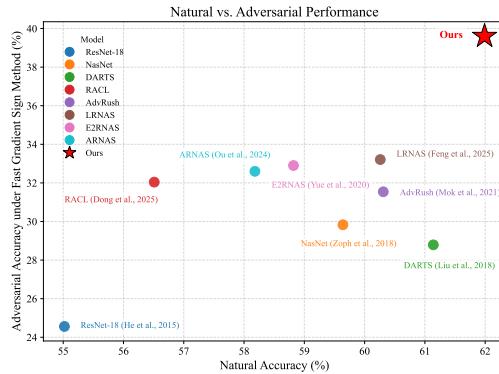


Figure 1: Clean vs. adversarial (FGSM) accuracy of various architectures on CIFAR-100. **Ours** achieves a favorable trade-off between robustness and clean performance, outperforming NAS-based and hand-crafted models.

108  
109

## 2.2 ROBUSTNESS-AWARE NEURAL ARCHITECTURE SEARCH

110  
111  
112  
113  
114  
115  
116

Neural architecture search (NAS) automates the design of DNNs, optimizing trade-offs among accuracy, efficiency, and model complexity. Early efforts considered *macro-level* search over entire networks (Zoph & Le, 2017), followed by *micro-level* approaches that learn reusable cells and stack them to build full models (Pham et al., 2018). To improve scalability, DARTS (Liu et al., 2019) introduced a continuous relaxation that forms a Supernet with differentiable operation selection; subsequent variants such as GDAS (Dong & Yang, 2019b) and PC-DARTS (Xu et al., 2019) further improved the efficiency of gradient-based search.

117  
118  
119  
120  
121  
122  
123  
124  
125  
126  
127  
128  
129

While NAS has achieved strong results in clean settings, adversarial robustness has only recently become a central focus. Early work incorporated robustness into DARTS-style frameworks via adversarial training or robustness-aware objectives—e.g., smoothness priors, certified bounds, or Jacobian regularization—to guide the search (Guo et al., 2020a; Zela et al., 2020; Mok et al., 2021; Hosseini et al., 2021). Others explored broader search spaces, hybrid designs, and topological changes to improve robustness, revealing that simple strategies such as ensembling or width expansion can enhance adversarial accuracy (Vargas et al., 2019; Yang et al., 2024b; Devaguptapu et al., 2021; Liu et al., 2023). More recent methods adopt multi-objective formulations (Geraeinejad et al., 2021) to jointly consider accuracy, robustness, and efficiency, while designing specialized search spaces (Sun et al., 2025) and robustness-aware proxies (Jung et al., 2023). Additional strategies—including knowledge distillation (Nath et al., 2024), architectural disentanglement (Yang et al., 2024a), and adaptive channel allocation (Li et al., 2024)—further refine the search process (Chen et al., 2020). Meanwhile, techniques like Fair-DARTS (Chu et al., 2020) and BatchQuant (Bai et al., 2021) study robustness under constraints such as discretization and quantization.

130  
131  
132  
133  
134  
135  
136

Collectively, these developments reflect growing momentum toward NAS (especially DARTS-based) frameworks that explicitly optimize adversarial robustness across architectures and domains. Nevertheless, many approaches still depend on predefined templates and heuristic evaluations (e.g., PGD accuracy), which can introduce statistical noise and obscure robustness-critical factors. This motivates the need for principled, robustness-aware NAS frameworks that dynamically identify sensitivity in operations and architectural choices, enabling the discovery of inherently robust models.

137  
138

## 2.3 SHAPLEY VALUES IN ARCHITECTURE SEARCH

139  
140  
141  
142

The Shapley value, rooted in cooperative game theory (Nowak & Radzik, 1994), provides a principled way to assess the contribution of individual components to a collective reward. Let  $N = \{1, \dots, n\}$  index the components and let  $v : 2^N \rightarrow \mathbb{R}$  assign a value to each subset  $S \subseteq N$ . The Shapley value of component  $i \in N$  is

143  
144  
145

$$\phi_i(v) = \frac{1}{n!} \sum_{\pi \in \mathcal{P}(N)} [v(\text{pre}_i(\pi) \cup \{i\}) - v(\text{pre}_i(\pi))], \quad (1)$$

146  
147  
148  
149  
150  
151  
152  
153  
154

where  $\mathcal{P}(N)$  is the set of all permutations of  $N$  and  $\text{pre}_i(\pi)$  denotes the set of components preceding  $i$  in permutation  $\pi$ . Shapley-based approaches have been widely used for feature attribution and, more recently, for NAS (Xiao et al., 2022) to identify influential operations. AutoShape (Fang et al., 2023) and GraphNAP++ (Oloulade et al., 2025) apply Shapley-guided pruning or architecture selection. KernelSHAP-NAS (Tran et al., 2025) adapts KernelSHAP to approximate operation influence in a supernet more robustly than earlier sampling-based methods. D-DARTS (Heuillet et al., 2023) introduces a Shapley-based loss to guide distributed differentiable NAS, promoting diversity and cell-specific optimization. Beyond architecture search, SDL (Djenouri et al., 2023) leverages Shapley values for model consensus in general-purpose vision systems, and various algorithms have been proposed to estimate Shapley values more accurately and efficiently (Chen et al., 2023).

155  
156  
157  
158  
159  
160  
161

Nevertheless, effectiveness under adversarial evaluation can be hindered by high variance in marginal contributions due to input perturbations and limited sample sizes. To address this, we develop **ROSE** with two robust statistical techniques: the interquartile range (IQR) (Rousseeuw & Hubert, 2017), which flags outlier contributions using quartiles  $Q_1$  and  $Q_3$  (with  $\text{IQR} = Q_3 - Q_1$ ), and the median-of-means (MoM) estimator (Lecué & Lerasle, 2017). These tools yield a more stable and robust estimation pipeline for operation importance under adversarial noise. This design grounds ROSE in established statistical principles while extending Shapley-based NAS to adversarial settings for the first time; we further provide detailed theoretical analysis in the appendix.

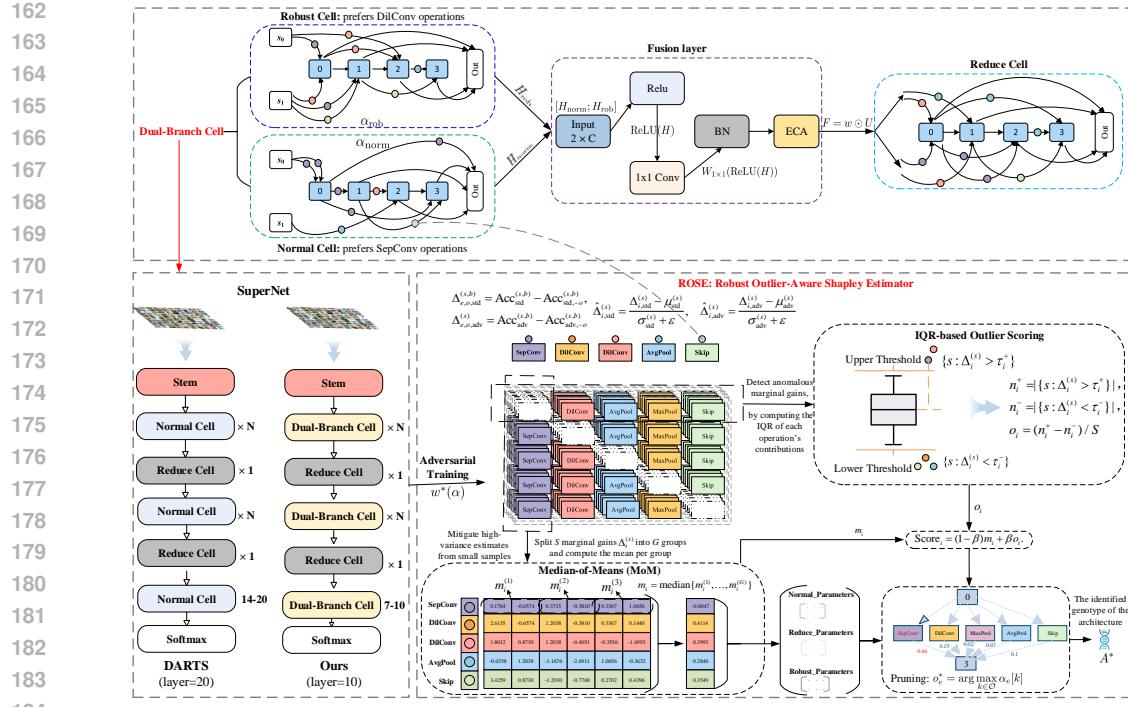


Figure 2: Overview of the RDNAS framework: (left) the dual-branch cell design; (middle) a comparison between standard DARTS (20 layers) and our dual-branch network (10 layers); (right) the ROSE estimator.

### 3 METHODOLOGY

#### 3.1 OVERVIEW OF RDNAS

RDNAS follows the DARTS paradigm in which an architecture is built by stacking cells—directed acyclic graphs (DAGs) with  $N$  intermediate nodes and a predefined operation set  $\mathcal{O}$  (e.g., depthwise convolutions, pooling, identity). For each edge  $(i, j)$ , we maintain logits  $\alpha^{(i,j)} = \{\alpha_o^{(i,j)}\}_{o \in \mathcal{O}}$  that parameterize a softmax mixture of operations:

$$\tilde{o}^{(i,j)}(x^{(i)}) = \sum_{o \in \mathcal{O}} \frac{\exp(\alpha_o^{(i,j)})}{\sum_{o' \in \mathcal{O}} \exp(\alpha_{o'}^{(i,j)})} o(x^{(i)}). \quad (2)$$

As illustrated in Figure 2, we enhance the standard DARTS backbone—typically 20 normal cells and 2 reduction cells—by replacing each normal cell with a *dual-branch* cell while halving the depth (10 dual-branch cells plus 2 reduction cells). The two branches process features in parallel to emphasize clean and robust pathways, respectively; at each fusion point an *Efficient Channel Attention* (ECA) module adaptively reweights and aggregates channel-wise responses, enabling selective retention of informative representations (see Appendix A.2 for Grad-CAM visualizations). The final representation is fed to a softmax classifier.

To support efficient search, we relax discrete operation choices into continuous architecture parameters  $\alpha$  and jointly optimize network weights  $\omega$  and  $\alpha$  via bilevel optimization:

$$\min_{\alpha} \mathcal{L}_{\text{val}}(\omega^*(\alpha), \alpha), \quad (3)$$

$$\text{s.t. } \omega^*(\alpha) = \arg \min_{\omega} \mathcal{L}_{\text{train}}(\omega, \alpha), \quad (4)$$

where the inner objective adopts a composite loss on clean and adversarial examples,

$$\mathcal{L}_{\text{train}} = \lambda \cdot \mathcal{L}_{\text{clean}} + (1 - \lambda) \cdot \mathcal{L}_{\text{adv}}. \quad (5)$$

(details in Sec. 4.2), and the outer objective is evaluated under the same robustness settings. RDNAS employs robustness-aware bilevel optimization to balance accuracy and adversarial resilience.

216 To improve stability under noisy gradients, we introduce **ROSE**, which scores operations on each  
 217 edge  $(i, j)$  and guides lightweight pruning. By combining MoM smoothing with IQR filtering,  
 218 ROSE yields robust importance estimates. After convergence, we discretize by retaining the top op-  
 219 eration(s) per edge based on ROSE scores, producing a compact model that balances clean accuracy  
 220 and adversarial robustness.

221

222 

### 3.2 DUAL-BRANCH CELL DESIGN

223

224 Conventional gradient-based NAS methods (e.g., DARTS) rely on a *single* shared cell to process  
 225 both clean and adversarial data. Under adversarial training, this shared cell must simultaneously  
 226 serve two competing objectives, and we empirically observe that the resulting architectures tend to  
 227 favor either clean accuracy or robustness, but rarely both. Although one could attempt a single cell  
 228 with a composite loss (e.g., TRADES (Zhang et al., 2019)), we find that explicitly separating the  
 229 optimization paths allows each branch to specialize, leading to superior overall performance. This  
 230 design reduces interference between objectives and provides more stable architecture search under  
 231 adversarial training.

232

233 To this end, we introduce the *Dual-Branch Cell*, which decouples learning into two parallel sub-  
 234 networks: a normal branch for clean accuracy and a robust branch for adversarial resilience. Both  
 235 branches process the same inputs  $(s_0, s_1)$  using a shared operator class  $\mathcal{C}$  but independent architec-  
 236 ture parameters  $\alpha_{\text{norm}}$  and  $\alpha_{\text{rob}}$ :

237

$$H_{\text{norm}} = \mathcal{C}(s_0, s_1; \alpha_{\text{norm}}), \quad H_{\text{rob}} = \mathcal{C}(s_0, s_1; \alpha_{\text{rob}}), \quad (6)$$

238

239 where  $H_{\text{norm}}, H_{\text{rob}} \in \mathbb{R}^{C \times H \times W}$ . Their outputs are concatenated along the channel dimension to  
 240 form a joint representation:

241

$$H = [H_{\text{norm}}; H_{\text{rob}}] \in \mathbb{R}^{2C \times H \times W}. \quad (7)$$

242

243 **Naively fusing both branches without weighting treats clean and robust features as equally impor-  
 244 tant at all times, which can limit expressiveness.** To more effectively exploit the complementary  
 245 representations, we adopt an attention-based fusion mechanism inspired by Efficient Channel At-  
 246 tention (ECA) (Wang et al., 2020). First, we reduce the channel dimension with a  $1 \times 1$  convolution  
 247 followed by ReLU activation:

248

$$U = W_{1 \times 1}(\text{ReLU}(H)) \in \mathbb{R}^{C \times H \times W}. \quad (8)$$

249

250 Next, global average pooling extracts a compact channel descriptor:

251

$$z = \text{GAP}(U) \in \mathbb{R}^C, \quad (9)$$

252

253 which is passed through a 1D convolution (with an adaptively determined kernel size as in ECA)  
 254 and a sigmoid activation to generate channel-wise attention weights:

255

$$w = \sigma(\text{Conv1D}(z)) \in (0, 1)^C. \quad (10)$$

256

257 These weights modulate the feature map via element-wise multiplication:

258

$$F = w \odot U \in \mathbb{R}^{C \times H \times W}. \quad (11)$$

259

260 The fused output  $F$  is forwarded to the next layer. **By explicitly separating optimization paths and**  
 261 **adaptively fusing them, the Dual-Branch Cell encourages each branch to specialize: the normal**  
 262 **branch can focus on fine-grained discriminative patterns, while the robust branch learns smoother,**  
 263 **perturbation-insensitive structures. This reduces conflict between clean and robust objectives and**  
 264 **stabilizes architecture search.** Our ablation in Table 4 shows that introducing ECA-based fusion  
 265 consistently improves both clean and PGD<sup>20</sup> accuracy over a dual-branch design without attention.  
 266 **Moreover, Grad-CAM visualizations in Appendix A.2 qualitatively confirm that the two branches at-**  
 267 **tend to different yet complementary regions under clean and adversarial inputs, providing empirical**  
 268 **evidence that the dual-branch design effectively captures diverse features.**

269

270 

### 3.3 ADVERSARILY ROBUST SEARCH

271

272 Following recent advances in robust NAS, we integrate adversarial training directly into the inner  
 273 optimization loop to promote robustness-aware search. **While conventional DARTS assumes subnet**

---

270 **Algorithm 1** Adversarially Robust Architecture Search

---

271 1: **Input:** Search space  $\mathcal{S} = \{o_{ij}\}$ , warm-up epochs  $N_w$ , search epochs  $N_s$ , #Shapley samples  $S$ , #MoM  
 272 groups  $G$ , IQR sensitivity  $\gamma$ , training data  $D_{\text{train}}$ , validation data  $D_{\text{val}}$   
 273 2: **Output:** Adversarially robust architecture  $A^*$   
 274 3: Initialize architecture logits  $\alpha_n \in \mathbb{R}^{E \times O}$  for each branch  $b \in \{\text{normal, reduce, robust}\}$  and weights  $\omega$ ;  
 275 4: **for** epoch = 1 **to**  $N_w + N_s$  **do**  
 276 5:   Train  $\omega$  on  $D_{\text{train}}$  by minimize the composite training loss in (5);  
 277 6:   **if** epoch >  $N_w$  **then**  
 278 7:     Compute the clean and adversarial accuracy deltas  $\Delta_{e,o,\text{std}}^{(s,b)}$  and  $\Delta_{e,o,\text{adv}}^{(s,b)}$  by (15);  
 279 8:     Compute the final  $\text{Score}_{e,o}^{(b)}$  by ROSE in (19) based on (15) to (18);  
 280 9:     Update logits:  $\alpha \leftarrow \alpha + \eta \cdot \text{Score}$ ;  
 281 10:   Row-normalize each edge of  $\alpha$  so its logits sum to 1;  
 282 11:   **end if**  
 283 12: **end for**  
 284 13: Pick the operation for each edge with highest logit to form  $A^*$ ;  
 285 14: **Return:**  $A^*$

---

286 ranking is consistent between small-search and large-inference networks, this assumption often fails  
 287 under adversarial training due to high-variance gradients. To address this, we leverage ROSE as  
 288 a statistically stable operation scorer, which mitigates the effect of noisy gradients and provides  
 289 reliable guidance for architecture updates.

290 Specifically, during each epoch, clean inputs  $x$  are transformed into adversarial examples  $x^{\text{adv}}$  using  
 291 a fast PGD- $k$  method (e.g.,  $k = 7$  steps):

$$293 \quad x^{\text{adv}} = \arg \max_{\|\delta\|_\infty \leq \epsilon} \mathcal{L}_{\text{train}}(f(x + \delta), y), \quad (12)$$

294 where  $\epsilon$  constrains the perturbation magnitude and  $\mathcal{L}_{\text{train}}$  is the composite training loss that mitigates  
 295 overfitting to clean samples and steers weight learning toward robust representations.

296 After obtaining the optimal network weights  $w^*(\alpha)$ , the outer optimization fixes  $w^*(\alpha)$  and updates  
 297 architecture parameters  $\alpha$  by minimizing the ROSE-weighted validation loss:

$$300 \quad \mathcal{L}_{\text{val}}(w^*(\alpha), \alpha) = - \sum_{b=1}^3 \sum_{e=1}^E \sum_{o=1}^O p_{e,o}^{(b)}(\alpha) \text{Score}_{e,o}^{(b)} \quad (13)$$

301 where  $b$  indexes normal, reduce, and robust branches, and  $p_{e,o}^{(b)}(\alpha) = \text{softmax}(\alpha_{e,:}^{(b)})_o$  is the proba-  
 302 bility of selecting operation  $o$  on edge  $e$  of branch  $b$ . The ROSE score  $\text{Score}_{e,o}^{(b)}$  stabilizes operation  
 303 evaluation under noisy adversarial gradients, ensuring that the outer-loop update favors consistently  
 304 robust operations.

305 Upon convergence, architecture discretization selects the most probable operation per edge:

$$306 \quad o_e^* = \arg \max_{k \in \mathcal{O}} \alpha_e[k] \quad (14)$$

307 The resulting normal and robust cells are paired into Dual-Branch Cells, interleaved with Reduce  
 308 Cells, and stacked to form a compact backbone that balances clean accuracy and adversarial robust-  
 309 ness. To reduce search cost and maintain practical efficiency, we adopt a small-sample search strat-  
 310 egy: only a fraction of the training data is used per search epoch, which, combined with lightweight  
 311 10-cell networks, drastically reduces computation while preserving ranking stability under ROSE  
 312 guidance. The overall search procedure is summarized in Algorithm 1. This setup ensures that  
 313 RDNAS discovers robust architectures efficiently, mitigating both the DARTS-related ranking in-  
 314 consistencies and the high computational overhead of adversarial training.

### 320 3.4 ROSE: ROBUST OUTLIER-AWARE SHAPLEY ESTIMATOR

321 After revisiting the Shapley foundations, we find that certain operations disproportionately influence  
 322 adversarial robustness. However, under adversarial training, architecture evaluation is inherently  
 323 noisy: gradient estimates fluctuate due to stochastic perturbations, and individual operations may

324 appear important or unimportant depending on the sampled adversarial examples. This motivates  
 325 ROSE, which refines Shapley-based scoring to capture both stable contributions and rare but critical  
 326 effects. Unlike conventional approaches that score all operations jointly, ROSE applies separate  
 327 evaluations to the normal, reduction, and robust cells, yielding three branch-specific score matrices.  
 328 By decoupling branches, ROSE also preserves branch-specific sensitivity, aligning with our dual-  
 329 branch cell design.

330 ROSE enhances the standard Monte Carlo Shapley sampler using two complementary techniques:  
 331 1) *IQR-based outlier detection*, operations whose standardized marginal gains fall outside the in-  
 332 terquartile range are highlighted, allowing rare yet decisive contributions to be recognized rather  
 333 than averaged away. This helps capture edges that are critical under adversarial conditions. 2) *MoM*  
 334 *smoothing*, by partitioning permutations into groups and aggregating their means, MoM reduces the  
 335 variance introduced by stochastic gradients, producing a stable Shapley estimate that reliably ranks  
 336 operations across noisy adversarial evaluations.

337 Let  $b \in \{\text{normal, reduce, robust}\}$  index each cell branch, and  $\mathcal{O}^{(b)}$  denote its candidate operations,  
 338 with the rest of the supernet fixed. Over  $S$  random permutations, we compute clean and adversarial  
 339 accuracy deltas for each operation  $o$  on edge  $e$  in branch  $b$ :

$$\Delta_{e,o,\text{std}}^{(s,b)} = \text{Acc}_{\text{std}}^{(s,b)} - \text{Acc}_{\text{std},-o}^{(s,b)}, \quad \Delta_{e,o,\text{adv}}^{(s,b)} = \text{Acc}_{\text{adv}}^{(s,b)} - \text{Acc}_{\text{adv},-o}^{(s,b)} \quad (15)$$

340 where  $\text{Acc}_{\text{std}}^{(s,b)}$  (or  $\text{Acc}_{\text{adv}}^{(s,b)}$ ) denotes validation accuracy with operation  $o$ , and  $\text{Acc}_{\text{std},-o}^{(s,b)}$  (or  
 341  $\text{Acc}_{\text{adv},-o}^{(s,b)}$ ) denotes the accuracy after removing  $o$ . These deltas are standardized per permutation to  
 342 reduce variance:

$$\hat{\Delta}_{i,\text{std}}^{(s)} = \frac{\Delta_{i,\text{std}}^{(s)} - \mu_{\text{std}}^{(s)}}{\sigma_{\text{std}}^{(s)} + \epsilon}, \quad \hat{\Delta}_{i,\text{adv}}^{(s)} = \frac{\Delta_{i,\text{adv}}^{(s)} - \mu_{\text{adv}}^{(s)}}{\sigma_{\text{adv}}^{(s)} + \epsilon} \quad (16)$$

343 where  $\mu^{(s)}$  and  $\sigma^{(s)}$  are empirical statistics across operations, and  $\epsilon$  ensures numerical stability.

344 To capture anomalies, we compute the IQR for each  $\{\hat{\Delta}_{e,o}^{(s,b)}\}_{s=1}^S$ , set upper and lower thresholds  
 345  $\tau_{e,o}^{(b)+}$  and  $\tau_{e,o}^{(b)-}$ , and derive an outlier score:

$$v_{e,o}^{(b)} = \frac{|\{s : \hat{\Delta}_{e,o}^{(s,b)} > \tau_{e,o}^{(b)+}\}| - |\{s : \hat{\Delta}_{e,o}^{(s,b)} < \tau_{e,o}^{(b)-}\}|}{S} \quad (17)$$

346 Here,  $v_{e,o}^{(b)} > 0$  suggests frequent robustness gains when  $o$  is dropped, while  $v_{e,o}^{(b)} < 0$  indicates  
 347 critical importance.

348 For stability, the MoM estimate is computed by partitioning samples into  $G$  folds, averaging within  
 349 each fold, and taking the median:

$$m_{e,o}^{(b)} = \text{median} \left\{ \frac{1}{|\mathcal{G}_k|} \sum_{s \in \mathcal{G}_k} \hat{\Delta}_{e,o}^{(s,b)} \right\}_{k=1}^G \quad (18)$$

350 Intuitively, MoM ensures that the ranking of operations is not dominated by extreme but noisy  
 351 gradients, while IQR ensures that rare but impactful operations are not ignored. This combination  
 352 provides a principled, statistically robust evaluation of each candidate.

353 The final ROSE score combines both components:

$$\text{Score}_{e,o}^{(b)} = (1 - \beta) m_{e,o}^{(b)} + \beta v_{e,o}^{(b)} \quad (19)$$

354 where  $\beta \in [0, 1]$  balances *steady reward* and *occasional indispensability*. Empirically,  $\beta \in [0.3, 0.5]$   
 355 yields architectures that are both reliably accurate and resilient under adversarial conditions, demon-  
 356 strating that ROSE effectively guides the search toward robust-optimal configurations.

357 We use the ROSE score to rank (and, if enabled, lightly prune) candidate operations in each cell.  
 358 This selective mechanism stabilizes operation evaluation under noisy adversarial training and high-  
 359 lights components that are consistently valuable or occasionally critical, providing a clear justifica-  
 360 tion for ROSE’s necessity. Sensitivity experiments on  $\beta$  are presented in Sec. A.3, with additional  
 361 hyperparameter studies for the entire framework provided in Appendix A.3.

378 

## 4 EXPERIMENTAL STUDIES

380 

### 4.1 BENCHMARK DATASETS AND BASELINES

383 Following established practices in robust NAS research (Guo et al., 2019; Mok et al., 2021),  
 384 we evaluate our proposed method, RDNAS, on four benchmark datasets: CIFAR-10, CIFAR-  
 385 100 (Krizhevsky, 2009), SVHN, and Tiny-ImageNet-200 (Le & Yang, 2015).

386 We compare RDNAS against a comprehensive set of peer methods, including manually designed  
 387 architectures such as ResNet-18 (He et al., 2015) and DenseNet-121 (Huang et al., 2017), standard  
 388 NAS frameworks like DARTS (Liu et al., 2019), NASNet (Zoph et al., 2018), and PDARTS (Chen  
 389 et al., 2019a), GDAS (Dong & Yang, 2019b), SETN (Dong & Yang, 2019a), and ENAS (Pham et al.,  
 390 2018); training-free NAS baselines including TTNAS (Lin et al., 2024) and MOTE-NAS (Zhang  
 391 et al., 2024); as well as robust NAS approaches such as RobNet (Guo et al., 2020b), AdvRush (Mok  
 392 et al., 2021), RACL (Dong et al., 2025a), E2RNAS (Yue et al., 2020), ARNAS (Ou et al., 2024),  
 393 and LRNAS (Feng et al., 2025a). **We ensure fair comparisons by reporting FLOPs and parameter**  
 394 **counts for all baselines, highlighting that improvements of RDNAS are achieved under comparable**  
 395 **or lower computational budgets.**

396 

### 4.2 SEARCH CONFIGURATION AND EVALUATION PROTOCOL

399 Given that RDNAS adopts a shallower architecture, we conduct the search on a 10-cell network  
 400 (instead of 8) and *adjust* the initial channel count to 32 (the common setting is 36) to maintain  
 401 training–search consistency. The search runs for 50 epochs using adversarial training based on 7-step  
 402 PGD with a step size of  $2/255$  and perturbation bound of  $8/255$ . The adversarial loss regularization  
 403 coefficient is set to 0.5. Other settings follow DARTS. Network weights ( $\omega$ ) are optimized using  
 404 SGD with momentum 0.9, learning rate  $\eta_w = 0.025$  (cosine decay), and weight decay  $3 \times 10^{-4}$ .  
 405 Architecture parameters ( $\alpha$ ) use Adam with  $\eta_\alpha = 3 \times 10^{-4}$ ,  $\beta = (0.5, 0.999)$ , and weight decay  
 406  $10^{-3}$ . **To substantially reduce search cost and accelerate convergence under adversarial training,**  
 407 **we adopt a small-sample search: on CIFAR-10, only 1,000 training samples and 500 validation**  
 408 **samples are used. ROSE is critical here**, as it stabilizes architecture ranking under small-sample  
 409 **noisy gradients, enabling reliable discovery of robust architectures in a fraction of typical search**  
 410 **time. To the best of our knowledge, this is the first attempt to combine small-sample one-shot NAS**  
 411 **with adversarial training.**

411 For final evaluation, we stack 10 cells with an initial channel count of 32. Stem / Reduction-1 /  
 412 Reduction-2 channels are set to 32 / 128 / 128. Following AdvRush (Mok et al., 2021), we adversarially  
 413 train each model for 120 epochs using 7-step PGD (step size 0.01,  $\epsilon = 8/255$ ). Optimization  
 414 uses SGD with momentum 0.9 and weight decay  $1 \times 10^{-4}$ . Learning rates are 0.1 (CIFAR-10/100)  
 415 and 0.01 (SVHN), decayed at epochs 90 and 104. Batch size is 32. All experiments are performed  
 416 on an NVIDIA GeForce RTX 4080 Super GPU.

418 Table 1: Robust accuracy on CIFAR-10. All attacks use  $\ell_\infty$  budget  $\epsilon = 8/255$ . Columns report  
 419 clean (Nat.) and white-box attacks; the last two columns show AutoAttack (AA) and GPU-days.  
 420 Best in each column is **bold**.

Model	Params	FLOPs	Evaluation ( $\ell_\infty$ )					AA	GPUday
			Nat.	FGSM	PGD <sup>20</sup>	PGD <sup>100</sup>	APGD-CE		
ResNet-18	11.2M	37.67M	84.09%	54.64%	45.86%	45.53%	44.54%	43.22%	–
DenseNet-121	7.0M	59.83M	85.95%	58.46%	50.49%	49.92%	49.11%	47.46%	–
DARTS	3.3M	547.44M	85.92%	58.96%	51.45%	49.32%	48.32%	47.73%	1.0
PDARTS	3.4M	550.75M	85.38%	59.12%	51.32%	50.91%	49.96%	48.52%	0.3
RobNet	5.6M	800.40M	85.00%	59.22%	52.09%	51.14%	50.41%	48.56%	3.3
DSRNA	2.0M	336.23M	80.93%	54.49%	49.11%	48.89%	48.67%	44.87%	0.4
LRNAS	2.2M	346.10M	84.26%	59.89%	50.20%	49.90%	49.83%	49.07%	0.4
RACL	3.6M	568.86M	85.13%	59.45%	51.89%	51.63%	51.09%	<b>50.23%</b>	0.5
AdvRush	4.2M	668.53M	86.38%	60.32%	52.29%	51.80%	51.42%	50.05%	0.7
<b>RDNAS (Ours)</b>	4.4M	1.30G	<b>86.56%</b>	<b>60.44%</b>	<b>52.62%</b>	<b>52.24%</b>	<b>52.05%</b>	49.98%	0.2

Table 2: Transfer-based black-box accuracy on CIFAR-10/100 (%).

Dataset	Params (M)	Source	Target					
			DARTS	AdvRush	LRNAS	RACL	PDARTS	Ours
CIFAR-10	3.3	DARTS	—	65.22	63.01	65.22	66.21	<b>67.54</b>
	4.2	AdvRush	65.40	—	63.52	65.45	66.29	<b>67.60</b>
	2.3	LRNAS	67.96	68.67	—	67.83	68.92	<b>69.46</b>
	3.6	RACL	65.44	65.46	62.54	—	64.97	<b>65.80</b>
	3.4	PDARTS	65.25	65.24	62.61	63.81	—	<b>65.52</b>
	4.4	Ours	<b>67.16</b>	66.80	63.85	65.79	66.32	—
CIFAR-100	3.3	DARTS	—	43.63	42.78	43.29	40.17	<b>46.37</b>
	4.2	AdvRush	42.46	—	41.91	42.83	40.01	<b>46.54</b>
	2.8	LRNAS	43.35	44.27	—	44.73	41.87	<b>47.06</b>
	3.6	RACL	42.68	41.83	42.25	—	40.39	<b>46.72</b>
	3.4	PDARTS	40.06	41.39	38.38	40.71	—	<b>45.94</b>
	4.4	Ours	47.04	<b>47.83</b>	44.61	47.65	47.18	—

Table 3: Cross-dataset transfer results (%),  $\epsilon=8/255$ , PGD<sup>20</sup> step size 2/255). Best result in each sub-table is **bold**.

CIFAR-100			SVHN			Tiny-ImageNet-200					
Model	Clean	FGSM	PGD <sup>20</sup>	Model	Clean	FGSM	PGD <sup>20</sup>	Model	Clean	FGSM	PGD <sup>20</sup>
ResNet-18	55.57	26.03	21.44	ResNet-18	92.06	88.73	69.51	ResNet-18	36.26	16.08	13.94
AdvRush	<b>60.31</b>	31.54	27.38	DenseNet-121	95.10	93.01	89.58	DenseNet-121	<b>47.56</b>	22.98	18.06
RACL	59.18	<b>34.40</b>	<b>30.41</b>	ARNAS	95.84	94.43	<b>92.02</b>	PDARTS	45.94	24.36	<b>22.74</b>
ARNAS	58.18	32.60	29.54	AdvRush	96.53	94.95	91.14	AdvRush	46.42	24.20	<b>22.89</b>
<b>RDNAS</b>	<b>61.99</b>	<b>39.60</b>	29.42	<b>RDNAS</b>	<b>97.88</b>	<b>96.37</b>	<b>95.80</b>	<b>RDNAS</b>	<b>56.84</b>	<b>24.49</b>	19.40

### 4.3 RESULTS ON WHITE-BOX ATTACKS

Table 1 shows that **RDNAS** attains the strongest overall white-box results: it is best on Clean, FGSM, PGD<sup>20</sup>, PGD<sup>100</sup>, and APGD-CE, and only 0.25 pp behind the best prior AutoAttack score (49.98% vs. 50.23% for RACL), with a moderate 4.4M parameters. Our FLOPs (1.30G) are higher than some DARTS-style baselines due to the dual-branch design with ECA fusion, but this extra compute yields consistent gains across all white-box attacks. The cost is practical: the search budget is small (0.2 GPU-days), and the halved depth (10 cells) helps latency despite per-cell compute. Moreover, the robustness *transfers*: Table 2 shows best or tied-best transfer-based black-box performance in most source→target pairs, and Table 3 confirms strong cross-dataset results (CIFAR-100, SVHN, Tiny-ImageNet-200). Overall, the modest FLOPs increase buys attack- and dataset-general robustness, making the accuracy–compute trade-off favorable.

### 4.4 RESULTS ON BLACK-BOX ATTACKS

We evaluate transfer-based black-box robustness on CIFAR-10 and CIFAR-100 using adversarial examples generated by PGD (untargeted,  $\epsilon = 8/255$ , step size 2/255, 20 steps, random starts) from five source models: DARTS, PDARTS, RACL, LRNAS, and AdvRush. Unless otherwise noted, attacks are single-source transfers (no ensemble). As summarized in Table 2, RDNAS outperforms peers in most transfer settings. Notably on CIFAR-100, it retains top robustness when attacked by adversarial inputs generated from the above sources. Despite its moderate size (4.4M parameters), RDNAS achieves 65.5%–69.5% robust accuracy on CIFAR-10 and 45.9%–47.8% on CIFAR-100, offering favorable robustness–parameter trade-offs relative to smaller models.

### 4.5 CROSS-DATASET TRANSFERABILITY

To test generalization, we transfer the RDNAS architecture searched on CIFAR-10 to CIFAR-100, SVHN, and Tiny-ImageNet-200, and retrain from scratch on each target dataset (no re-search). As shown in Table 3, RDNAS attains strong clean and adversarial accuracies. On CIFAR-100, it surpasses ResNet-18 and NASNet in robustness; on SVHN, it achieves 97.88% clean accuracy and 96.37% FGSM robustness; on Tiny-ImageNet-200, it delivers 56.84% clean accuracy and remains competitive under PGD, trailing AdvRush slightly.

486 Table 4: Ablation of ECA, adversarial-training during search, and search space. Metrics are Top-1  
 487 (%) on clean and **PGD**<sup>20</sup> ( $\epsilon = 8/255$ , step size 2/255); Params in millions.

ID	ECA	Adv. search	Search space	Params (M)	Clean (%)	PGD <sup>20</sup> (%)
A	✗	✗	Single Cell (20 cells)	4.20	84.8	50.3
B	✗	✓	Single Cell (20 cells)	4.28	85.6	51.0
C	✓	✓	Single Cell (20 cells)	4.30	<u>86.0</u>	51.5
D	✗	✓	Dual-Branch Cell (10 cells)	4.26	85.7	<u>51.8</u>
E	✓	✓	Dual-Branch Cell (10 cells)	4.32	<b>86.5</b>	<b>52.6</b>

## 497 4.6 ABLATION STUDIES

498  
 499 **Search-Space Ablation:** As shown in Table 4, a reduced RDNAS search space (10 operations)  
 500 outperforms the full DARTS space (20 operations) under FGSM and PGD, while maintaining sim-  
 501 ilar clean accuracy. **This demonstrates that our carefully designed operation set suffices for robust**  
 502 **architecture discovery and reduces search complexity.**

503 **Adversarial Training Ablation:** Clean-only search yields strong natural accuracy but poor PGD ro-  
 504 bustness. Integrating adversarial training during search significantly improves robustness, validating  
 505 its importance.

506 **Attention Module Ablation:** Removing the Efficient Channel Attention (ECA) module slightly de-  
 507 grades clean accuracy and noticeably reduces adversarial robustness. Reintroducing ECA improves  
 508 both, confirming its effectiveness. **Grad-CAM visualizations (Appendix A.2) further show that the**  
 509 **two branches focus on complementary regions, providing qualitative support for the attention mod-**  
 510 **ule’s role in robust feature integration.**

511 **ROSE Ablation:** ROSE is enabled by default in our adversarial-search setting. To assess its contri-  
 512 bution, we run an additional search variant without ROSE, using standard Shapley estimates. This  
 513 variant shows higher run-to-run variability and slightly lower PGD robustness, indicating that ROSE  
 514 stabilizes small-sample adversarial NAS and aids in discovering architectures with more consistent  
 515 robust performance.

## 517 4.7 RESULTS ON NAS-BENCH-201 (CIFAR-10)

518 Table 5 reports validation/test accuracy and  
 519 search cost on NAS-Bench-201 (CIFAR-10).  
 520 We list classic gradient-based baselines and  
 521 recent training/evaluation-free methods (e.g.,  
 522 TTNAS (Lin et al., 2024), MOTE-NAS (Zhang  
 523 et al., 2024)). Our method achieves the best trade-  
 524 off between accuracy and search time. This ex-  
 525 periment demonstrates that, even under a few-  
 526 sample search regime, the ROSE estimator re-  
 527 mains stable and reliably discovers robust archi-  
 528 tectures.

## 530 5 CONCLUSION

531 This work presents RDNAS, a robust NAS frame-  
 532 work that jointly optimizes for accuracy and ad-  
 533 versarial robustness. We propose a dual-branch cell search space and enhance Shapley-value esti-  
 534 mation using MoM and IQR to identify critical operations. Extensive experiments across multiple  
 535 benchmarks and attack settings demonstrate the effectiveness and transferability of RDNAS, with  
 536 results highlighting that heterogeneous cell placement across depths is crucial for balancing robust-  
 537 ness and accuracy.

Method	Search (s)	val	test
DARTS-V2	35781.80	$39.77 \pm 0.00$	$54.30 \pm 0.00$
GDAS	31609.80	$89.89 \pm 0.08$	$93.61 \pm 0.09$
SETN	34139.53	$84.04 \pm 0.28$	$87.64 \pm 0.00$
ENAS	14058.80	$37.51 \pm 3.19$	$53.89 \pm 0.58$
TTNAS (training-free)	1146	$91.02 \pm 0.11$	$93.94 \pm 0.38$
MOTE-NAS (K=5)	2239	$90.89 \pm 0.13$	$93.86 \pm 0.15$
<b>Ours</b>	<b>731</b>	<b><math>91.13 \pm 0.36</math></b>	<b><math>93.97 \pm 0.35</math></b>
<b>Ground Truth</b>	—	<b>91.61</b>	<b>94.37</b>

538 Table 5: Validation and test accuracy with  
 539 corresponding search cost on NAS-Bench-201  
 (CIFAR-10). RDNAS achieves the best trade-  
 off between accuracy and efficiency, demon-  
 strating stable performance even under limited  
 search budgets. The results include the average  
 and standard deviations for 3 runs.

540 REFERENCES  
541

542 Haoping Bai, Meng Cao, Ping Huang, and Jiulong Shan. Batchquant: Quantized-for-all architecture  
543 search with robust quantizer. *Advances in Neural Information Processing Systems*, 34:1074–1085,  
544 2021.

545 Bowen Baker, Otkrist Gupta, Nikhil Naik, and Ramesh Raskar. Designing neural network architec-  
546 tures using reinforcement learning, 2017. URL <https://arxiv.org/abs/1611.02167>.  
547

548 Hanlin Chen, Baochang Zhang, Song Xue, Xuan Gong, Hong Liu, Rongrong Ji, and David Do-  
549 errmann. Anti-bandit neural architecture search for model defense. In *European conference on*  
550 *computer vision*, pp. 70–85. Springer, 2020.

551 Hugh Chen, Ian C Covert, Scott M Lundberg, and Su-In Lee. Algorithms to estimate shapley value  
552 feature attributions. *Nature Machine Intelligence*, 5(6):590–601, 2023.

553

554 Xin Chen, Lingxi Xie, Jun Wu, and Qi Tian. Progressive differentiable architecture search: Bridging  
555 the depth gap between search and evaluation. In *2019 IEEE/CVF International Conference on*  
556 *Computer Vision (ICCV)*, pp. 1294–1303, 2019a. doi: 10.1109/ICCV.2019.00138.

557

558 Xin Chen, Lingxi Xie, Jun Wu, and Qi Tian. Progressive differentiable architecture search: Bridg-  
559 ing the depth gap between search and evaluation, 2019b. URL <https://arxiv.org/abs/1904.12760>.  
560

561 Zhilu Chen and Ximeng Huang. End-to-end learning for lane keeping of self-driving cars. In  
562 *2017 IEEE Intelligent Vehicles Symposium (IV)*, pp. 1856–1860, 2017. doi: 10.1109/IVS.2017.  
563 7995975.

564

565 Zhi Cheng, Yanxi Li, Minjing Dong, Xiu Su, Shan You, and Chang Xu. Neural architecture search  
566 for wide spectrum adversarial robustness. In *Proceedings of the AAAI Conference on Artificial*  
567 *Intelligence*, volume 37, pp. 442–451, 2023.

568

569 Xiangxiang Chu, Tianbao Zhou, Bo Zhang, and Jixiang Li. Fair darts: Eliminating unfair advantages  
570 in differentiable architecture search. In *European conference on computer vision*, pp. 465–480.  
Springer, 2020.

571

572 Francesco Croce and Matthias Hein. Reliable evaluation of adversarial robustness with an ensemble  
573 of diverse parameter-free attacks, 2020. URL <https://arxiv.org/abs/2003.01690>.

574

575 Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hier-  
576 archical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*,  
577 pp. 248–255, 2009. doi: 10.1109/CVPR.2009.5206848.

578

579 Chaitanya Devaguptapu, Devansh Agarwal, Gaurav Mittal, Pulkit Gopalani, and Vineeth N Balasub-  
580 ramanian. On adversarial robustness: A neural architecture search perspective. In *Proceedings of*  
*the IEEE/CVF international conference on computer vision*, pp. 152–161, 2021.

581

582 Youcef Djennouri, Ahmed Nabil Belbachir, Tomasz Michalak, and Anis Yazidi. Shapley deep learn-  
583 ing: A consensus for general-purpose vision systems. In *Proceedings of the IEEE/CVF Interna-*  
*tional Conference on Computer Vision*, pp. 1224–1233, 2023.

584

585 Minjing Dong, Yanxi Li, Yunhe Wang, and Chang Xu. Adversarially robust neural architectures.  
*IEEE Transactions on Pattern Analysis and Machine Intelligence*, 47(5):4183–4197, 2025a. doi:  
586 10.1109/TPAMI.2025.3542350.

587

588 Minjing Dong, Yanxi Li, Yunhe Wang, and Chang Xu. Adversarially robust neural architectures.  
*IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2025b.

589

590 Xuanyi Dong and Yi Yang. One-shot neural architecture search via self-evaluated template net-  
591 work. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*. IEEE, Octo-  
592 ber 2019a. doi: 10.1109/iccv.2019.00378. URL [http://dx.doi.org/10.1109/ICCV.](http://dx.doi.org/10.1109/ICCV.2019.00378)  
593 2019.00378.

594 Xuanyi Dong and Yi Yang. Searching for a robust neural architecture in four gpu hours. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 1761–1770,  
 595 2019b.

596

597 Yunfei Fang, Caihong Mu, and Yi Liu. Autoshape: Automatic design of click-through rate prediction  
 598 models using shapley value. In *Pacific Rim International Conference on Artificial Intelligence*,  
 599 pp. 29–40. Springer, 2023.

600

601 Yuqi Feng, Zeqiong Lv, Hongyang Chen, Shangce Gao, Fengping An, and Yanan Sun. Lrnas: Dif-  
 602 ferentiable searching for adversarially robust lightweight neural architecture. *IEEE Transactions*  
 603 *on Neural Networks and Learning Systems*, 36(3):5629–5643, 2025a. doi: 10.1109/TNNLS.2024.  
 604 3382724.

605

606 Yuqi Feng, Yuwei Ou, Jiahao Fan, and Yanan Sun. Zero-cost proxy for adversarial robustness  
 607 evaluation. In *The Thirteenth International Conference on Learning Representations*, 2025b.  
 608 URL <https://openreview.net/forum?id=zHf7h0feer>.

609

610 Vahid Geraeinejad, Sima Sinaei, Mehdi Modarressi, and Masoud Daneshtalab. Roco-nas: Robust  
 611 and compact neural architecture search. In *2021 International Joint Conference on Neural Net-  
 612 works (IJCNN)*, pp. 1–8. IEEE, 2021.

613

614 Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial  
 615 examples, 2015. URL <https://arxiv.org/abs/1412.6572>.

616

617 Minghao Guo, Yuzhe Yang, Rui Xu, Ziwei Liu, and Dahua Lin. When nas meets robustness: In  
 618 search of robust architectures against adversarial attacks. *arXiv preprint arXiv:1911.10695*, 2019.

619

620 Minghao Guo, Yuzhe Yang, Rui Xu, Ziwei Liu, and Dahua Lin. When nas meets robustness: In  
 621 search of robust architectures against adversarial attacks. In *Proceedings of the IEEE/CVF con-  
 622 ference on computer vision and pattern recognition*, pp. 631–640, 2020a.

623

624 Minghao Guo, Yuzhe Yang, Rui Xu, Ziwei Liu, and Dahua Lin. When nas meets robustness: In  
 625 search of robust architectures against adversarial attacks, 2020b. URL <https://arxiv.org/abs/1911.10695>.

626

627 Hyeonjeong Ha, Minseon Kim, and Sung Ju Hwang. Generalizable lightweight proxy for robust nas  
 628 against diverse perturbations, 2023. URL <https://arxiv.org/abs/2306.05031>.

629

630 Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recog-  
 631 nition, 2015. URL <https://arxiv.org/abs/1512.03385>.

632

633 Alexandre Heuillet, Hedi Tabia, Hichem Arioui, and Kamal Youcef-Toumi. D-darts: Distributed  
 634 differentiable architecture search. *Pattern Recognition Letters*, 176:42–48, 2023.

635

636 Alexandre Heuillet, Ahmad Nasser, Hichem Arioui, and Hedi Tabia. Efficient automation of neural  
 637 network design: A survey on differentiable neural architecture search. *ACM Computing Surveys*,  
 638 56(11):1–36, 2024.

639

640 Ramtin Hosseini, Xingyi Yang, and Pengtao Xie. Dsrna: Differentiable search of robust neural  
 641 architectures. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern  
 642 Recognition*, pp. 6196–6205, 2021.

643

644 Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q. Weinberger. Densely connected  
 645 convolutional networks. In *2017 IEEE Conference on Computer Vision and Pattern Recognition  
 646 (CVPR)*, pp. 2261–2269, 2017. doi: 10.1109/CVPR.2017.243.

647

648 Steffen Jung, Jovita Lukasik, and Margret Keuper. Neural architecture design and robustness: A  
 649 dataset. In *The Eleventh International Conference on Learning Representations*, 2023.

650

651 Alex Krizhevsky. Learning multiple layers of features from tiny images. 2009. URL <https://api.semanticscholar.org/CorpusID:18268744>.

648 Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep  
 649 convolutional neural networks. In F. Pereira, C.J. Burges, L. Bottou, and K.Q. Weinberger  
 650 (eds.), *Advances in Neural Information Processing Systems*, volume 25. Curran Associates, Inc.,  
 651 2012. URL [https://proceedings.neurips.cc/paper\\_files/paper/2012/file/c399862d3b9d6b76c8436e924a68c45b-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2012/file/c399862d3b9d6b76c8436e924a68c45b-Paper.pdf).

652  
 653  
 654 Ya Le and Xuan S. Yang. Tiny imagenet visual recognition challenge. 2015. URL <https://api.semanticscholar.org/CorpusID:16664790>.

655  
 656 Guillaume Lecué and Matthieu Lerasle. Robust machine learning by median-of-means : theory and  
 657 practice, 2017. URL <https://arxiv.org/abs/1711.10306>.

658  
 659 Chao Li, Jia Ning, Han Hu, and Kun He. Adaptive channel allocation for robust differentiable  
 660 architecture search. *IEEE Transactions on Neural Networks and Learning Systems*, 2024.

661  
 662 Yiqi Lin, Yuki Endo, Jinho Lee, and Shunsuke Kamijo. Neural architecture search via trainless  
 663 pruning algorithm: A bayesian evaluation of a network with multiple indicators. *Electronics*, 13  
 664 (22), 2024. ISSN 2079-9292. doi: 10.3390/electronics13224547. URL <https://www.mdpi.com/2079-9292/13/22/4547>.

665  
 666 Aishan Liu, Shiyu Tang, Siyuan Liang, Ruihao Gong, Boxi Wu, Xianglong Liu, and Dacheng Tao.  
 667 Exploring the relationship between architectural design and adversarially robust generalization.  
 668 In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp.  
 669 4096–4107, 2023.

670  
 671 Hanxiao Liu, Karen Simonyan, and Yiming Yang. Darts: Differentiable architecture search, 2019.  
 672 URL <https://arxiv.org/abs/1806.09055>.

673 Gabor Lugosi and Shahar Mendelson. Mean estimation and regression under heavy-tailed  
 674 distributions—a survey, 2019. URL <https://arxiv.org/abs/1906.04280>.

675  
 676 Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu.  
 677 Towards deep learning models resistant to adversarial attacks, 2019. URL <https://arxiv.org/abs/1706.06083>.

678  
 679 Joseph Mellor, Jack Turner, Amos Storkey, and Elliot J. Crowley. Neural architecture search without  
 680 training, 2021. URL <https://arxiv.org/abs/2006.04647>.

681  
 682 Jisoo Mok, Byunggook Na, Hyeokjun Choe, and Sungroh Yoon. Advrush: Searching for adversari-  
 683 ally robust neural architectures, 2021. URL <https://arxiv.org/abs/2108.01289>.

684  
 685 Utkarsh Nath, Yancheng Wang, Pavan Turaga, and Yingzhen Yang. Rnas-cl: Robust neural architec-  
 686 ture search by cross-layer knowledge distillation. *International journal of computer vision*, 132  
 687 (12):5698–5717, 2024.

688  
 689 Andrzej S. Nowak and Tadeusz Radzik. The shapley value for n-person games in generalized char-  
 690 acteristic function form. *Games and Economic Behavior*, 6(1):150–161, 1994. ISSN 0899-8256.  
 691 doi: <https://doi.org/10.1006/game.1994.1008>. URL <https://www.sciencedirect.com/science/article/pii/S0899825684710086>.

692  
 693 Babatounde Moctard Oloulade, Jianliang Gao, Jiamin Chen, Raeed Al-Sabri, Zhenpeng Wu, and  
 694 Monir Abdullah. Shapley-guided pruning for efficient graph neural architecture prediction in  
 695 distributed learning environments. *Information Sciences*, 695:121695, 2025.

696  
 697 Yuwei Ou, Yuqi Feng, and Yanan Sun. Towards accurate and robust architectures via neural ar-  
 698 chitecture search. In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition*  
 699 (*CVPR*), pp. 5967–5976, 2024. doi: 10.1109/CVPR52733.2024.00570.

700  
 701 Nicolas Papernot, Patrick McDaniel, Ian Goodfellow, Somesh Jha, Z. Berkay Celik, and Ananthram  
 702 Swami. Practical black-box attacks against machine learning, 2017. URL <https://arxiv.org/abs/1602.02697>.

702 Yameng Peng, Andy Song, Haytham M. Fayek, Vic Ciesielski, and Xiaojun Chang. Swap-nas:  
 703 Sample-wise activation patterns for ultra-fast nas, 2024. URL <https://arxiv.org/abs/2403.04161>.

705 Hieu Pham, Melody Y. Guan, Barret Zoph, Quoc V. Le, and Jeff Dean. Efficient neural architecture  
 706 search via parameter sharing, 2018. URL <https://arxiv.org/abs/1802.03268>.

708 Peter J. Rousseeuw and Mia Hubert. Anomaly detection by robust statistics. *WIREs Data Mining  
 709 and Knowledge Discovery*, 8(2), November 2017. ISSN 1942-4795. doi: 10.1002/widm.1236.  
 710 URL <http://dx.doi.org/10.1002/widm.1236>.

711 Ramprasaath R. Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh,  
 712 and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based lo-  
 713 calization. *International Journal of Computer Vision*, 128(2):336–359, October 2019. ISSN  
 714 1573-1405. doi: 10.1007/s11263-019-01228-7. URL <http://dx.doi.org/10.1007/s11263-019-01228-7>.

716 Christian Simon, Piotr Koniusz, Lars Petersson, Yan Han, and Mehrtash Harandi. Towards a robust  
 717 differentiable architecture search under label noise. In *Proceedings of the IEEE/CVF Winter  
 718 Conference on Applications of Computer Vision*, pp. 3256–3266, 2022.

720 Shuoyang Sun, Kaiwen Zhang, Hao Fang, Bin Chen, Jiawei Li, Enze Huo, and Shu-Tao Xia. Rob-  
 721 nas: Robust neural architecture search for point cloud adversarial defense. In *ICASSP 2025-2025  
 722 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 1–5.  
 723 IEEE, 2025.

724 Yaniv Taigman, Ming Yang, Marc’Aurelio Ranzato, and Lior Wolf. Deepface: Closing the gap to  
 725 human-level performance in face verification. In *2014 IEEE Conference on Computer Vision and  
 726 Pattern Recognition*, pp. 1701–1708, 2014. doi: 10.1109/CVPR.2014.220.

727 Thanh Hai Tran, Dac Tam Nguyen, Minh Duc Ngo, Long Doan, Ngoc Hoang Luong, and Huynh  
 728 Thi Thanh Binh. Kernelshap-nas: a shapley additive explanatory approach for characterizing  
 729 operation influences. *Neural Computing and Applications*, pp. 1–17, 2025.

731 Danilo Vasconcellos Vargas, Shashank Kotyan, and SPM IIIT-NR. Evolving robust neural architec-  
 732 tures to defend from adversarial attacks. *arXiv preprint arXiv:1906.11667*, 3, 2019.

733 Li W, Yao S, Wan B, Xiao L, Hou C, Zhong Y, and Zhou W. Tdnn achitecture with efficient channel  
 734 attention and improved residual blocks for accurate speaker recognition. *Scientific reports*, 15  
 735 (1):23484, 2025. doi: 10.1038/s41598-025-09386-0. URL <https://doi.org/10.1038/s41598-025-09386-0>.

737 Qilong Wang, Banggu Wu, Pengfei Zhu, Peihua Li, Wangmeng Zuo, and Qinghua Hu. Eca-net:  
 738 Efficient channel attention for deep convolutional neural networks. In *2020 IEEE/CVF Con-  
 739 ference on Computer Vision and Pattern Recognition (CVPR)*, pp. 11531–11539, 2020. doi:  
 740 10.1109/CVPR42600.2020.01155.

742 Sanghyun Woo, Jongchan Park, Joon-Young Lee, and In So Kweon. Cbam: Convolutional block  
 743 attention module, 2018. URL <https://arxiv.org/abs/1807.06521>.

744 Juntao Wu, Ziyu Song, Xiaoyu Zhang, Shujun Xie, Longxin Lin, and Ke Wang. Vision transfor-  
 745 mers beat wideresnets on small scale datasets adversarial robustness. In *Proceedings of the AAAI  
 746 Conference on Artificial Intelligence*, volume 39, pp. 886–894, 2025.

748 Han Xiao, Ziwei Wang, Zheng Zhu, Jie Zhou, and Jiwen Lu. Shapley-nas: Discovering opera-  
 749 tion contribution for neural architecture search. In *Proceedings of the IEEE/CVF Conference on  
 750 Computer Vision and Pattern Recognition (CVPR)*, 2022.

751 Yuhui Xu, Lingxi Xie, Xiaopeng Zhang, Xin Chen, Guo-Jun Qi, Qi Tian, and Hongkai Xiong.  
 752 Pc-darts: Partial channel connections for memory-efficient architecture search. In *International  
 753 Conference on Learning Representations*, 2019.

755 Qiancheng Yang, Yong Luo, and Bo Du. Training-free robust neural network search via pruning. In  
 2024 IEEE International Conference on Multimedia and Expo (ICME), pp. 1–6. IEEE, 2024a.

756 Shangshang Yang, Xiangkun Sun, Ke Xu, Yuanchao Liu, Ye Tian, and Xingyi Zhang. Hybrid  
 757 architecture-based evolutionary robust neural architecture search. *IEEE Transactions on Emerging*  
 758 *Topics in Computational Intelligence*, 8(4):2919–2934, 2024b.

760 Yeming Yang, Qingling Zhu, Jianping Luo, Ka-Chun Wong, Qiuzhen Lin, and Jianqiang Li. Trnas:  
 761 A training-free robust neural architecture search. In *Proceedings of the IEEE/CVF International*  
 762 *Conference on Computer Vision (ICCV)*, pp. 2336–2345, October 2025.

764 Zhaohui Yang, Yunhe Wang, Xinghao Chen, Boxin Shi, Chao Xu, Chunjing Xu, Qi Tian, and Chang  
 765 Xu. Cars: Continuous evolution for efficient neural architecture search. In *Proceedings of the*  
 766 *IEEE/CVF conference on computer vision and pattern recognition*, pp. 1829–1838, 2020.

768 Zhixiong Yue, Baijiong Lin, Xiaonan Huang, and Yu Zhang. Effective, efficient and robust neural  
 769 architecture search, 2020. URL <https://arxiv.org/abs/2011.09820>.

771 Arber Zela, Thomas Elsken, Tonmoy Saikia, Yassine Marrakchi, Thomas Brox, and Frank Hutter.  
 772 Understanding and robustifying differentiable architecture search. In *International Conference*  
 773 *on Learning Representations*, 2020.

775 Hongyang Zhang, Yaodong Yu, Jiantao Jiao, Eric P. Xing, Laurent El Ghaoui, and Michael I. Jordan.  
 776 Theoretically principled trade-off between robustness and accuracy, 2019. URL <https://arxiv.org/abs/1901.08573>.

778 Jianpeng Zhang, Xiaomin Chen, Bing Yang, Qingbiao Guan, Qi Chen, Jian Chen, Qi Wu, Yu-  
 779 tong Xie, and Yong Xia. Advances in attention mechanisms for medical image segmentation.  
 780 *Computer Science Review*, 56:100721, 2025. ISSN 1574-0137. doi: <https://doi.org/10.1016/j.cosrev.2024.100721>. URL <https://www.sciencedirect.com/science/article/pii/S1574013724001047>.

784 Yuming Zhang, Jun Wei Hsieh, Xin Li, Ming-Ching Chang, Chun-Chieh Lee, and Kuo-Chin Fan.  
 785 MOTE-NAS: Multi-objective training-based estimate for efficient neural architecture search. In  
 786 *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024. URL  
 787 <https://openreview.net/forum?id=jKLjKeZfzv>.

789 Xun Zhou, Zhenkun Wang, Liang Feng, Songbai Liu, Ka-Chun Wong, and Kay Chen Tan. Toward  
 790 evolutionary multitask convolutional neural architecture search. *IEEE Transactions on Evolutionary*  
 791 *Computation*, 28(3):682–695, 2023.

793 Xun Zhou, Songbai Liu, AK Qin, and Kay Chen Tan. Evolutionary neural architecture search for  
 794 transferable networks. *IEEE Transactions on Emerging Topics in Computational Intelligence*,  
 795 2024.

797 Barret Zoph and Quoc V. Le. Neural architecture search with reinforcement learning, 2017. URL  
 798 <https://arxiv.org/abs/1611.01578>.

800 Barret Zoph, Vijay Vasudevan, Jonathon Shlens, and Quoc V. Le. Learning transferable architectures  
 801 for scalable image recognition. In *2018 IEEE/CVF Conference on Computer Vision and Pattern*  
 802 *Recognition*, pp. 8697–8710, 2018. doi: 10.1109/CVPR.2018.00907.

## 804 A APPENDIX

### 807 A.1 THEORETICAL JUSTIFICATION OF ROSE ESTIMATOR

808 This appendix provides a concentration bound for the Robust Outlier-aware Shapley Estimator  
 809 (ROSE) under mild assumptions.

810 A.1.1 ASSUMPTIONS  
811812 **Assumption 1 (Bounded Marginal Gains).** For any branch  $b$ , edge  $e$ , operation  $o$ , and permutation  
813  $s$ ,

814 
$$|\Delta_{e,o,\text{std}}^{(s,b)}| \leq M, \quad |\Delta_{e,o,\text{adv}}^{(s,b)}| \leq M, \quad (20)$$

815 where  $M > 0$  is a finite constant. Since the Shapley value  $\phi_{e,o}^{(b)}$  is a convex combination of bounded  
816 marginal gains, we also have  $|\phi_{e,o}^{(b)}| \leq M$ .  
817818 **Assumption 2 (Permutation Independence).** The  $S$  Monte-Carlo permutations are sampled inde-  
819 pendently and partitioned into buckets for the MoM estimator; independence implies inter-bucket  
820 independence.  
821822 A.1.2 TRIANGLE INEQUALITY FOR ROSE SCORE  
823824 ROSE combines a Median-of-Means (MoM) estimator  $m_{e,o}^{(b)}$  and an outlier-aware score  $v_{e,o}^{(b)} \in$   
825  $[-1, 1]$ :

826 
$$\text{Score}_{e,o}^{(b)} = (1 - \beta) m_{e,o}^{(b)} + \beta v_{e,o}^{(b)}, \quad 0 \leq \beta < 1. \quad (21)$$

827 Subtracting the true Shapley value  $\phi_{e,o}^{(b)}$  gives  
828

829 
$$|\text{Score}_{e,o}^{(b)} - \phi_{e,o}^{(b)}| = \left| (1 - \beta)(m_{e,o}^{(b)} - \phi_{e,o}^{(b)}) + \beta(v_{e,o}^{(b)} - \phi_{e,o}^{(b)}) \right| \quad (22)$$

830 
$$\leq (1 - \beta) |m_{e,o}^{(b)} - \phi_{e,o}^{(b)}| + \beta |v_{e,o}^{(b)} - \phi_{e,o}^{(b)}|. \quad (22')$$

832 Because  $|v_{e,o}^{(b)}| \leq 1$  and  $|\phi_{e,o}^{(b)}| \leq M$ , define  $C := 1 + M$  for later use.  
833834 A.1.3 PROBABILITY RELAXATION  
835836 We upper-bound the deviation probability:  
837

838 
$$\Pr(|\text{Score}_{e,o}^{(b)} - \phi_{e,o}^{(b)}| \geq \varepsilon). \quad (23)$$

839 When  $\varepsilon > \beta C$ , inequality equation 22' implies  
840

841 
$$|m_{e,o}^{(b)} - \phi_{e,o}^{(b)}| \geq \frac{\varepsilon - \beta C}{1 - \beta}. \quad (24)$$

842 Hence  
843

844 
$$\Pr(|\text{Score}_{e,o}^{(b)} - \phi_{e,o}^{(b)}| \geq \varepsilon) \leq \Pr\left(|m_{e,o}^{(b)} - \phi_{e,o}^{(b)}| \geq \frac{\varepsilon - \beta C}{1 - \beta}\right). \quad (25)$$

845 A.1.4 CONCENTRATION OF MOM ESTIMATOR  
846847 According to the MoM concentration bound (Lugosi & Mendelson, 2019), if the bucket (group) size  
848 is at least  $n_0$ , then for any  $t > 0$ ,

849 
$$\Pr(|m_{e,o}^{(b)} - \phi_{e,o}^{(b)}| \geq t) \leq 2 \exp\left(-\frac{n_0 t^2}{2M^2}\right). \quad (26)$$

851 Plugging  $t = \frac{\varepsilon - \beta C}{1 - \beta}$  into equation 26 and using equation 25 yields  
852

853 
$$\Pr(|\text{Score}_{e,o}^{(b)} - \phi_{e,o}^{(b)}| \geq \varepsilon) \leq 2 \exp\left(-\frac{n_0(\varepsilon - \beta C)^2}{2M^2(1 - \beta)^2}\right). \quad (27)$$

855 A.1.5 FINAL RELAXATION  
856857 If  $\varepsilon > 2\beta C$  then  $(\varepsilon - \beta C)^2 \geq \varepsilon^2/4$  and  $(1 - \beta)^2 \leq 1$ , so  
858

859 
$$\Pr(|\text{Score}_{e,o}^{(b)} - \phi_{e,o}^{(b)}| \geq \varepsilon) \leq 2 \exp\left(-\frac{n_0 \varepsilon^2}{8M^2}\right). \quad (28)$$

860 
$$\varepsilon > 2\beta(1 + M).$$

861 Thus, ROSE enjoys a sub-Gaussian tail (up to a constant factor), completing the proof. Equa-  
862 tion equation 28 establishes a sub-Gaussian tail  
863

864 
$$\Pr(|\text{Score} - \phi| \geq \varepsilon) \leq 2 \exp\left(-\frac{n_0 \varepsilon^2}{8M^2}\right).$$

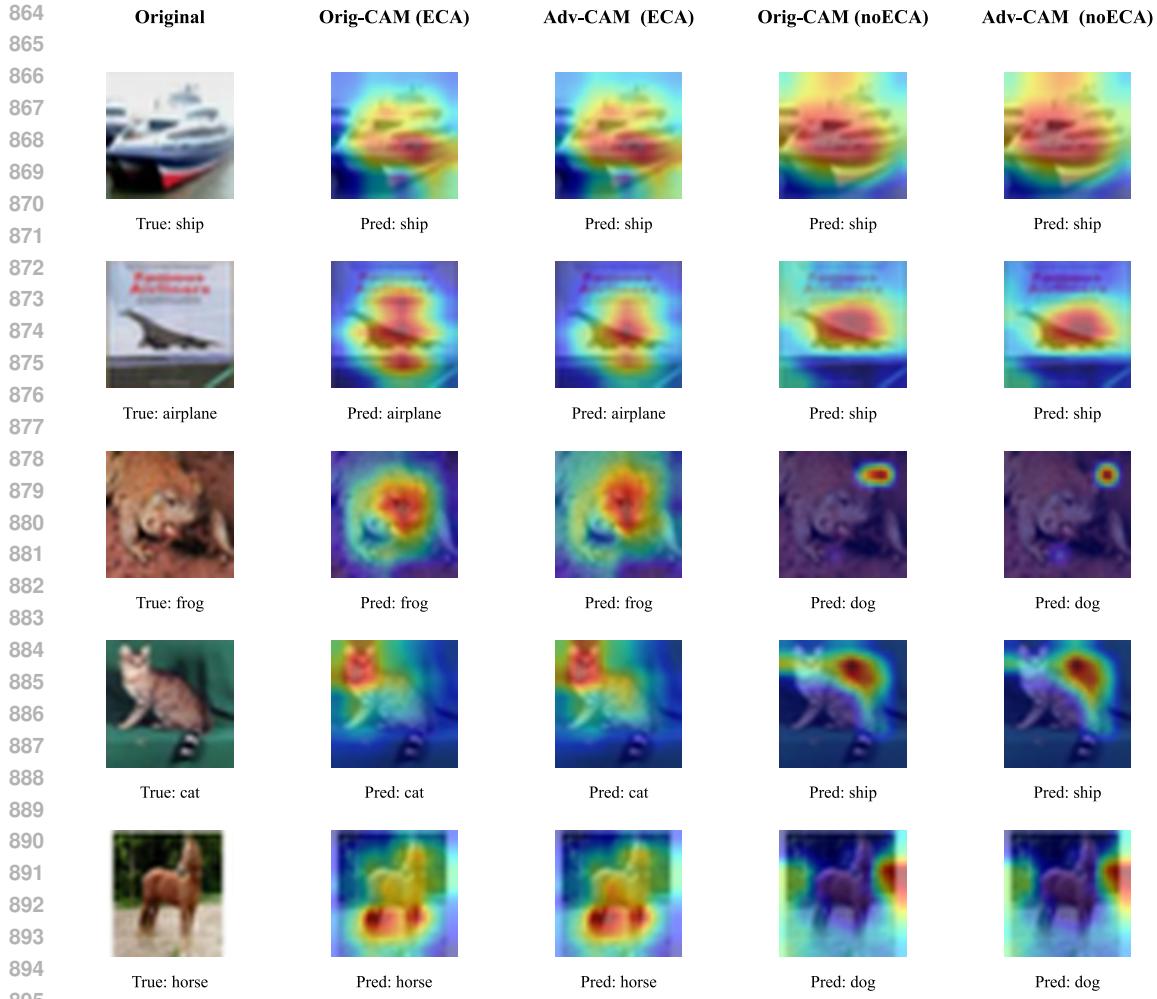


Figure 3: Grad-CAM visualizations comparing models with and without ECA under clean and adversarial inputs.

**Discussion.** The MoM mechanism counteracts heavy-tailed noise, while the outlier-aware term  $v_{e,o}$  suppresses extreme gradients, yielding robustness even when adversarial or extreme samples are present. The guarantees rely only on inter-bucket independence and thus extend naturally to multi-GPU or multi-node distributed settings. The error–probability relation also couples the bucket size  $n_0$  with any desired confidence level, allowing practitioners to minimize sampling—and hence search cost—while meeting accuracy targets. Finally, the coefficient  $\beta$  acts as a design knob between statistical reliability and outlier penalization; the theory requires  $\varepsilon > 2\beta(1 + M)$ , so exponentially decaying  $\beta$  or rescaling gains (reducing  $M$ ) keeps evaluations within the admissible region, transforming heuristic “ $\beta$  tuning” into a quantitatively grounded procedure. Although we conduct the search with only 1k/0.5k labeled samples on CIFAR-10, our theoretical guarantees concern the *estimation* sample size used by ROSE (i.e., the number of independent evaluations per MoM bucket), rather than the dataset size itself. By increasing the number of permutations, data augmentations, and resamplings, we effectively enlarge  $n_0$ . With a small  $\beta$  and bounded gains (small  $M$ ), the sub-Gaussian tail bound still applies—albeit with a looser constant—hence our small-data search remains consistent with the theory.

## A.2 VISUALIZATION OF ECA MODULE’S EFFECT

To examine how the ECA (Efficient Channel Attention) module affects robustness and interpretability, we visualize class activation maps (CAMs) using Grad-CAM (Selvaraju et al., 2019) under both

918 clean and adversarial inputs. We use the last convolutional block for Grad-CAM, and normalize  
 919 each CAM to  $[0, 1]$  (per image) before overlay. Adversarial CAMs are generated with untargeted  
 920 PGD-20 under  $\ell_\infty$  budget  $\epsilon = 8/255$  (step size 2/255, random starts). Figure 3 shows a grid of  
 921 five representative CIFAR-10 samples and their corresponding CAMs, using the same inputs across  
 922 models for a fair comparison.

- 923 • **Column 1:** Original image with ground-truth label.
- 924 • **Columns 2–3 (with ECA):** CAMs under clean and adversarial inputs.
- 925 • **Columns 4–5 (w/o ECA):** CAMs for the same architecture without ECA.

926 We observe that models with ECA tend to produce more focused, class-consistent activation maps.  
 927 For instance, in the frog example (row 3), clean and adversarial CAMs remain centered on the object  
 928 when ECA is enabled, whereas the w/o-ECA counterpart exhibits a pronounced shift that coincides  
 929 with a misclassification to “dog.” In the cat example (row 4), the ECA model preserves attention over  
 930 the cat body, while the w/o-ECA variant misclassifies it as “ship” and shifts attention to background.  
 931 Similar patterns appear in the horse sample (row 5), suggesting improved stability of attention under  
 932 perturbations.

933 Overall, these qualitative results indicate that ECA can promote both adversarial robustness and  
 934 interpretability by enforcing more discriminative and stable attention over semantically meaningful  
 935 regions.

### 936 A.3 HYPERPARAMETER SETTINGS AND SEARCH DETAILS

937 To ensure fair comparison and reproducibility, we follow standard settings widely adopted in robust  
 938 NAS for optimization and adversarial attacks. Unless otherwise noted, all experiments use:

- 939 • **Optimizer:** SGD with momentum 0.9, initial learning rate 0.025, cosine annealing;
- 940 • **Weight decay:**  $3 \times 10^{-4}$ ;
- 941 • **Search epochs:** 50;
- 942 • **Adversarial training:** PGD-7 with step size 2/255 and  $\epsilon = 8/255$  (untargeted, random  
 943 starts; evaluation protocol in Sec. 4.2).

944 We only tune the two hyperparameters introduced by the **ROSE** module:

- 945 •  $\beta$ : trade-off coefficient between the median-of-means (steady contribution) and the IQR-  
 946 based outlier score (occasional indispensability);
- 947 •  $\lambda = (\lambda_1, \lambda_2)$ : asymmetric IQR scale factors for upper/lower fences, used to detect sig-  
 948 nificant outliers in clean/adv gains, respectively, with

$$949 \tau^+ = Q_3 + \lambda_1 \text{IQR}, \quad \tau^- = Q_1 - \lambda_2 \text{IQR}.$$

950 We conduct a lightweight grid search over:

- 951 •  $\beta \in \{0, 0.2, 0.3, 0.4, 0.5, 1\}$ ;
- 952 •  $(\lambda_1, \lambda_2) \in \{1.2, 1.5, 2.0\} \times \{1.2, 1.5, 2.0\}$ .

953 Table 6 summarizes CIFAR-10 results under different  $(\lambda_1, \lambda_2)$  choices. The best-performing setting  
 954 is  $(\lambda_1=1.2, \lambda_2=2.0)$ , which yields strong clean accuracy and adversarial robustness.

### 955 A.4 CELL VISUALIZATION OF THE FINAL ARCHITECTURE

956 To enhance reproducibility and provide architectural insights, we visualize the three cell types dis-  
 957 covered by RDNAS on CIFAR-10: the *Normal* cell, the *Reduction* cell, and the *Robust* cell. Each  
 958 cell is a directed acyclic graph in which nodes denote intermediate feature maps and edges denote  
 959 selected operations. Inputs  $c_{k-2}$  and  $c_{k-1}$  are the outputs of the two preceding cells, and  $c_k$  is the  
 960 current cell’s output. The final genotype is obtained from the discretized architecture at the end of  
 961 search.

972 Table 6: CIFAR-10 performance under different  $(\lambda_1, \lambda_2)$  for asymmetric IQR thresholds. Best is  
 973 **bold**, second best is underlined. Attacks follow Sec. 4.2.  
 974

$\lambda_1$	$\lambda_2$	Clean	FGSM	PGD <sup>20</sup>	PGD <sup>100</sup>
1.5	1.2	84.69%	58.41%	50.45%	50.07%
1.5	1.5	<u>85.45%</u>	<u>59.55%</u>	<u>52.18%</u>	<u>51.73%</u>
1.2	2.0	<b>86.56%</b>	<b>60.44%</b>	<b>52.62%</b>	<b>52.24%</b>
2.0	1.5	85.05%	57.74%	49.97%	49.61%

981 Table 7: CIFAR-10 performance under different  $\beta$  values. Best is **bold**. Attacks follow Sec. 4.2.  
 982

$\beta$	Clean	FGSM	PGD <sup>20</sup>	PGD <sup>100</sup>
0	85.25%	<u>60.21%</u>	<u>52.55%</u>	<u>52.08%</u>
0.2	85.06%	58.68%	50.54%	50.20%
0.3	<b>86.56%</b>	<b>60.44%</b>	<b>52.62%</b>	<b>52.24%</b>
0.4	85.49%	59.63%	52.22%	51.67%
0.5	84.83%	58.48%	50.99%	50.47%
1	84.45%	59.47%	52.45%	52.13%

- **Normal cell.** Prioritizes expressive operations (e.g., sepconv-3x3, sepconv-5x5), forming deeper computation paths while maintaining gradient flow with occasional skipconnect.
- **Reduction cell.** Responsible for spatial downsampling, often combining pooling (e.g., maxpool-3x3) with separable convolutions to expand the receptive field efficiently.
- **Robust cell.** Tends to favor smoother or more redundant topologies (e.g., dilconv-5x5, skipconnect), which are associated with improved stability under perturbations.

1000 Together, the three cells emphasize complementary goals—expressiveness (Normal), efficiency (Reduction), and stability (Robust)—which helps explain the observed balance between clean accuracy  
 1001 and adversarial robustness.  
 1002

## 1004 A.5 FULL-SCALE VISUALIZATION OF PRIMITIVE TRENDS

1006 To provide a comprehensive view of search dynamics, we visualize the high-resolution trigger trends  
 1007 of each primitive in both the *Normal* and *Robust* cells across 50 search epochs. While the main  
 1008 paper provides compact summaries for space, the figures below reveal finer-grained patterns and  
 1009 fluctuations. For clarity, “exception triggers” (method-specific events such as IQR outliers) are  
 1010 counted per epoch; see Sec. 3.4 for the formal definition.

1011 We observe that primitives such as sepconv-3x3/5x5 and dilconv-5x5 are *often* among the  
 1012 most frequently triggered in both cells, suggesting their importance for robust architecture forma-  
 1013 tion. In contrast, skipconnect and pooling operations (e.g., avgpool-3x3, maxpool-3x3)  
 1014 are triggered less frequently, indicating a lower tendency to be prioritized in our robustness-aware  
 1015 search space.

1016 The temporal curves also exhibit periodic rises and falls, implying that the search revisits and re-  
 1017 assesses operator importance over time rather than converging prematurely. These detailed trends  
 1018 support the effectiveness of our ROSE-based exception mechanism in tracking meaningful shifts in  
 1019 primitive utility.

## 1021 A.6 ATTENTION MODULE SELECTION

1023 To assess how the fusion attention choice shapes both the searched architecture and robustness,  
 1024 we evaluate three candidates—CBAM (reduction= 16,  $k=7$ ) (Woo et al., 2018), ECAM ( $k_c=3$ ,  
 1025  $k_s=5$ ) (W et al., 2025), and ECA ( $k=3$ ) (Wang et al., 2020; Zhang et al., 2025)—by replacing the  
 1026 attention module in the search space and re-running the full NAS pipeline for each variant. Apart

1026 Table 8: Attention module comparison on CIFAR-10 (%),  $\ell_\infty$ ,  $\epsilon = 8/255$ , PGD step size 2/255).  
 1027 Single run per setting; best in each column is **bold**.

Module	Params (M)	Clean	FGSM	PGD <sup>20</sup>	PGD <sup>100</sup>	APGD <sub>CE</sub>
CBAM (reduction=16, k=7)	4.2	84.62	59.72	52.38	51.99	51.73
ECAM ( $k_c = 3$ , $k_s = 5$ )	4.5	<b>85.34</b>	<b>59.92</b>	<b>52.57</b>	<b>52.17</b>	<b>51.95</b>
ECA ( $k = 3$ )	4.4	<b>86.56</b>	<b>60.44</b>	<b>52.62</b>	<b>52.24</b>	<b>52.05</b>

1033  
 1034  
 1035  
 1036 from the attention module, all factors are held fixed: the rest of the primitives, random seeds and  
 1037 splits, search/training budgets, and optimization hyperparameters (including 7-step PGD adversarial  
 1038 training with step size 2/255,  $\epsilon=8/255$ ).

1039 As reported in Table 8, under the “re-search + retrain” setting, ECA yields the best or tied-best  
 1040 performance across all metrics. ECAM (a spatially-extended ECA) is close under strong attacks but  
 1041 slightly behind overall; CBAM drops further under stronger attacks. Because the resulting models  
 1042 have nearly identical parameter counts, these gains are unlikely to stem from model size, pointing  
 1043 instead to how attention design shapes the searched topology and robustness-oriented inductive bias.

1044  
 1045  
 1046 **Why does ECA lead the search?** We believe three factors contribute: (1) No channel bottleneck:  
 1047 ECA’s local 1D channel convolution avoids MLP reduction, preserving fine-grained discriminative  
 1048 information that is easily lost under adversarial training; (2) Local channel neighborhoods:  
 1049 small-kernel cross-channel interactions appear more stable to perturbation noise, reducing variance  
 1050 in Shapley-style scoring and gradient updates during search; (3) Lightweight and optimization-  
 1051 friendly: with minimal parameters and a short optimization path, ECA couples more benignly with  
 1052 adversarial losses, whereas spatial branches and strong reduction (CBAM/ECAM) add compute and  
 1053 can aggravate optimization instability without consistent robustness gains.

## 1054 A.7 ROBUSTNESS ACROSS PGD BUDGETS

1055 To examine how robustness evolves with the strength of the adversary, we sweep the perturbation  
 1056 budget  $\epsilon$  for untargeted  $L_\infty$  PGD on CIFAR-10 and evaluate all models under a fixed, transparent  
 1057 protocol:

- 1062 • **Attack setup.** For each  $\epsilon \in \{1, 2, 3, 4, 6, 8\}/255$  we use step size  $\alpha = \epsilon/4$ ,  $r = 5$  random  
 1063 restarts with uniform initialization in the  $\ell_\infty$  ball, and  $T$  PGD iterations. To avoid under-  
 1064 optimized attacks at small budgets,  $T$  is increased as  $\epsilon$  decreases (Table 9). No attack  
 1065 parameter is tuned per model.
- 1066 • **Evaluation.** Robust accuracy is measured on the full test set; images are clipped to  $[0, 1]$   
 1067 after each step. The same settings are applied to all compared methods.

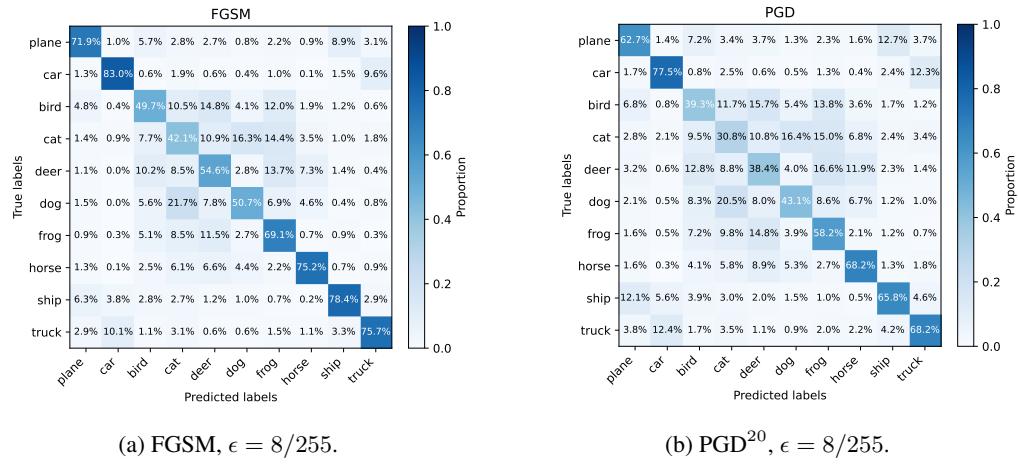
1069  
 1070  
 1071 **Results and takeaways.** Table 9 reports the robust accuracy grid. As expected, accuracy decreases  
 1072 monotonically as  $\epsilon$  grows. Across the full range of budgets, **RDNAS** consistently matches or exceeds  
 1073 baselines, with the largest margins appearing at the higher budgets (e.g.,  $\epsilon = 8/255$ ). At moderate  
 1074 budgets ( $\epsilon \in [3, 6]/255$ ), RDNAS preserves a comfortable robustness headroom while maintaining  
 1075 competitive clean accuracy (see main text), indicating that its dual-branch design and ECA fusion  
 1076 do not rely on a narrow operating point.

1077 All runs use identical preprocessing, evaluation code, and random seeds across models. We follow  
 1078 the common choice  $\alpha = \epsilon/4$  and keep restarts fixed ( $r = 5$ ) to balance attack strength and runtime.  
 1079 This grid can serve as a drop-in stress test for future methods to report budget sensitivity under a  
 1080 standardized PGD protocol.

1080 Table 9: Robust accuracy (%) under  $L_\infty$  PGD with varying budgets. Default  $\alpha = \epsilon/4$ , restarts  
 1081  $r = 5$ .

$\epsilon$	$\alpha$	$T$	$r$	RDNAS(Ours)	LRNAS	RACL	AdvRush
8/255	2/255	20	5	<b>52.62</b>	50.20	51.89	<u>52.29</u>
6/255	1.5/255	20	5	<b>62.84</b>	56.05	61.94	<u>62.09</u>
4/255	1/255	20	5	<b>72.11</b>	66.45	71.33	<u>71.94</u>
3/255	0.75/255	30	5	<b>76.08</b>	71.05	75.33	<u>75.49</u>
2/255	0.5/255	40	5	<b>79.91</b>	75.12	<u>79.87</u>	79.70
1/255	0.25/255	50	5	<b>83.34</b>	79.29	82.24	<u>83.14</u>

### A.8 ROBUSTNESS ACROSS VARIOUS SAMPLE TYPES ON CIFAR-10 UNDER $\ell_\infty$ ATTACKS



1109 Figure 7: Class-wise confusion matrices of the final RDNAS model on CIFAR-10 under white-box  
 1110  $\ell_\infty$  attacks.

1113 To better understand why a robust estimator like ROSE is needed in our adversarial NAS setting, we  
 1114 report class-wise confusion matrices of the final RDNAS model on **CIFAR-10 under white-box  $\ell_\infty$  attacks** in Figure 7.

1116 **Dataset:** CIFAR-10 test set (10,000 images, 10 classes).

1117 **Model:** The final RDNAS architecture adversarially trained

1118 **Attacks:**

- 1120 • **FGSM ( $\ell_\infty$ )** with perturbation budget  $\epsilon = 8/255$ .
- 1121 • **PGD<sup>20</sup> ( $\ell_\infty$ )** with  $\epsilon = 8/255$ , step size 2/255, 20 steps, and random starts.

1123 Figure 7 reveals a clear heterogeneity across classes: some classes (e.g., *bird*, *cat*, *deer*, *dog*) suffer  
 1124 substantial drops in diagonal accuracy and much larger mass spread over incorrect labels, whereas  
 1125 others (e.g., *plane*, *car*, *frog*, *horse*, *ship*, *truck*) remain comparatively robust under the same attack  
 1126 budget. This behavior indicates that, even at a fixed  $\ell_\infty$  budget on CIFAR-10, the induced ad-  
 1127 versarial loss distribution is **highly skewed and class-dependent**. A relatively small subset of “hard”  
 1128 examples and vulnerable classes contributes disproportionately to the gradients, leading to **heavy-  
 1129 tailed and occasionally extreme** marginal gains when estimating operation-level Shapley values in  
 1130 a weight-sharing supernet. In such a regime, naive averaging of marginal gains can be dominated  
 1131 by these outliers and yield unstable operation rankings across runs. ROSE explicitly addresses this  
 1132 issue by combining (i) a **Median-of-Means** term, which provides a robust estimate of the typical  
 1133 contribution under heavy-tailed noise, and (ii) an **IQR-based** component, which controls the influ-  
 1134 ence of rare but extreme deviations. The CIFAR-10 confusion matrices in Figure 7 thus provide  
 1135 empirical motivation for using ROSE as a robust operation-scoring mechanism in adversarial NAS.

1134 A.9 ADDITIONAL IMAGENET-1K EXPERIMENTS  
1135

1136 Following TRNAS (Yang et al., 2025), we train all architectures on ImageNet-1k with FAST-FGSM  
 1137 adversarial training for 50 epochs, perturbation radius  $\epsilon = 4/255$ , and SGD with momentum. We  
 1138 use a batch size of 512 and apply the same training protocol to DARTS (Liu et al., 2019), LR-  
 1139 NAS (Feng et al., 2025a), CRoZe (Ha et al., 2023), ZCPRob (Feng et al., 2025b), TRNAS, and  
 1140 our searched architecture. After training, we report top-1 accuracy on clean images, FGSM, and  
 1141 PGD-20 attacks.

1142 Table 10: Comparison of NAS methods on ImageNet-1k under FAST-FGSM training ( $\epsilon = 4/255$ ).  
 1143 We report top-1 accuracy (%).

Methods	Type	Clean	FGSM	PGD-20
DARTS	Clean	53.88	19.22	11.11
LRNAS	Robust	48.21	15.02	8.18
CRoZe	Zero-shot robust	49.52	16.28	9.41
ZCPRob	Zero-shot robust	52.93	18.86	10.75
TRNAS	Zero-shot robust	55.10	20.56	11.73
Ours	Robust	<b>57.09</b>	<b>22.11</b>	<b>12.68</b>

1144

1145

1146

1147

1148

1149

1150

1151

1152

1153

1154

1155

1156

1157

1158

1159

1160

1161

1162

1163

1164

1165

1166

1167

1168

1169

1170

1171

1172

1173

1174

1175

1176

1177

1178

1179

1180

1181

1182

1183

1184

1185

1186

1187

1188

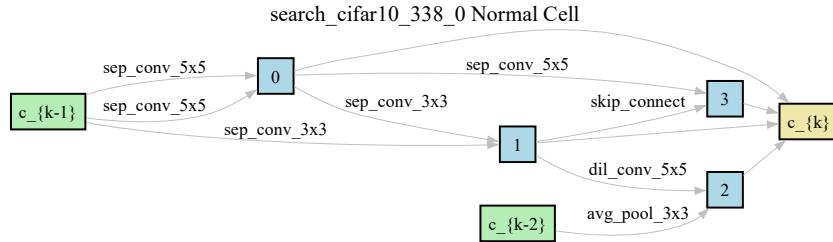
1189

1190

1191

1192

1193



(a) Normal cell discovered by RDNAS on CIFAR-10.

1204

1205

1206

1207

1208

1209

1210

1211

1212

1213

1214

1215

1216

1217

1218

1219

1220

1221

1222

1223

1224

1225

1226

1227

1228

1229

1230

1231

1232

1233

1234

1235

1236

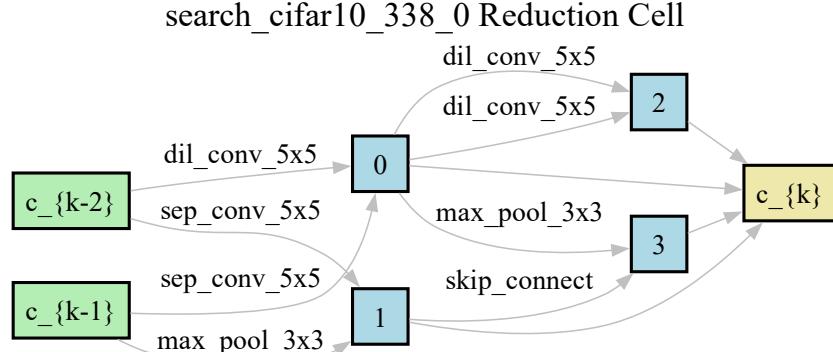
1237

1238

1239

1240

1241



(b) Reduction cell discovered by RDNAS on CIFAR-10.

1223

1224

1225

1226

1227

1228

1229

1230

1231

1232

1233

1234

1235

1236

1237

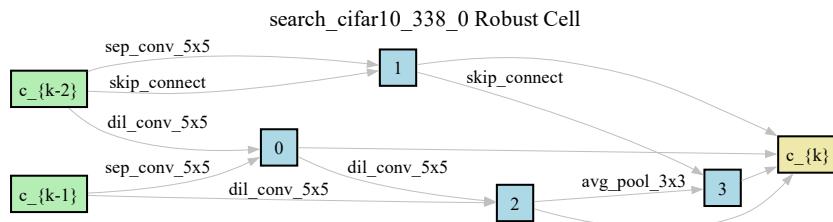
1238

1239

1240

1241

search\_cifar10\_338\_0 Robust Cell



(c) Robust cell discovered by RDNAS on CIFAR-10.

Figure 4: Cells discovered by RDNAS on CIFAR-10. From top to bottom: Normal, Reduction, and Robust cells.

1242  
1243  
1244  
1245  
1246  
1247  
1248  
1249  
1250  
1251  
1252  
1253  
1254  
1255  
1256  
1257  
1258  
1259  
1260  
1261  
1262  
1263  
1264

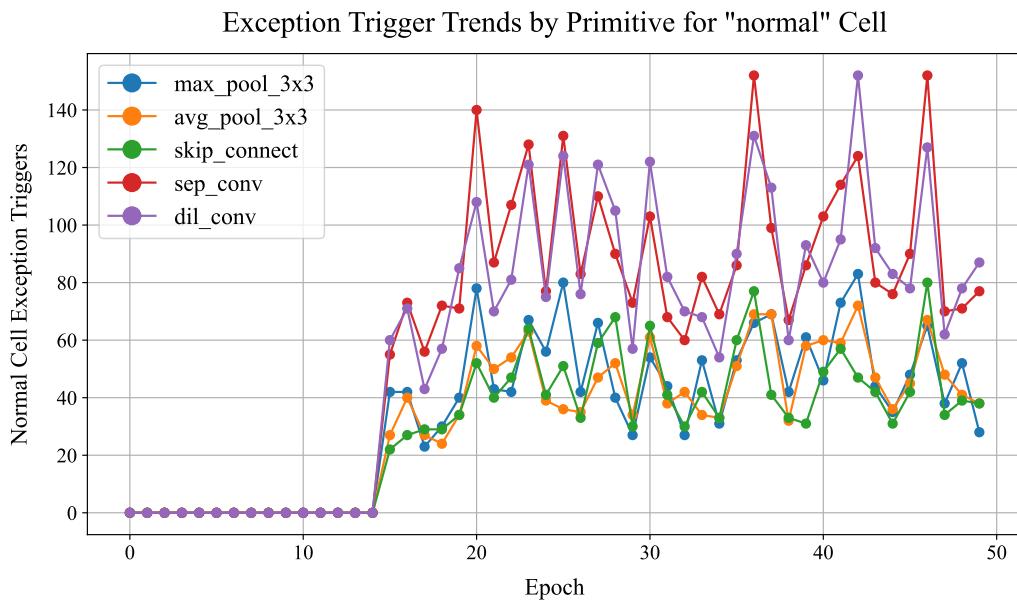


Figure 5: High-resolution trigger trends for primitives in the Normal cell across 50 search epochs.

1265  
1266  
1267  
1268  
1269  
1270  
1271  
1272  
1273  
1274  
1275  
1276  
1277  
1278  
1279  
1280  
1281  
1282  
1283  
1284  
1285  
1286  
1287  
1288  
1289  
1290  
1291  
1292  
1293  
1294

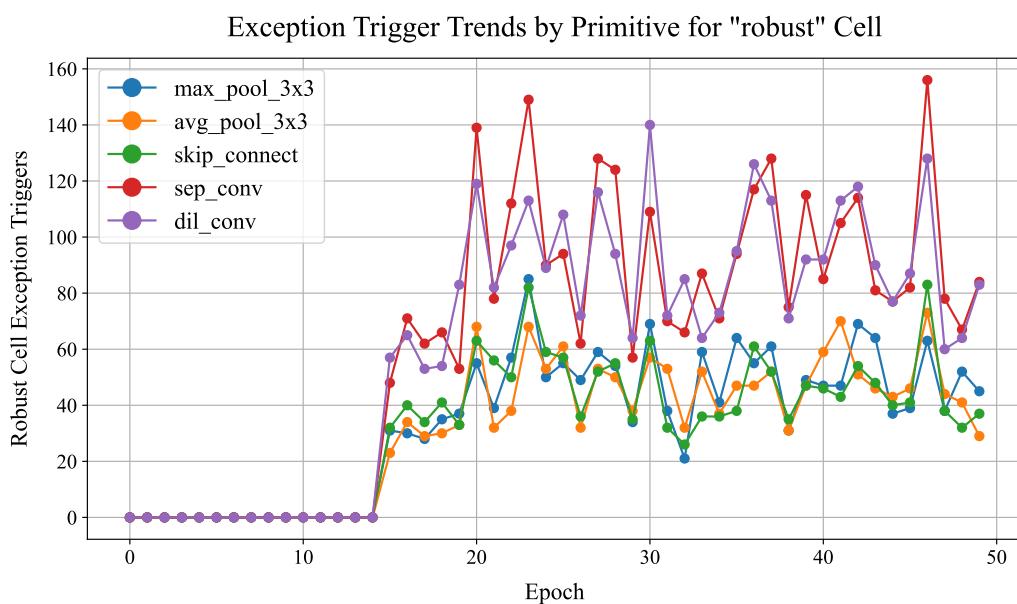


Figure 6: High-resolution trigger trends for primitives in the Robust cell across 50 search epochs.