

# 000 001 ASYMMETRIC TRAINING WITH HETEROGENEOUS 002 LOSSES: A PROBE INTO ARCHITECTURAL RESO- 003 NANCE 004 005

006 **Anonymous authors**  
007 Paper under double-blind review  
008  
009  
010  
011

## ABSTRACT

013 Is deep learning generalization necessarily rooted in optimizing a single objective?  
014 We explore an alternative view: adaptive generalization may emerge from struc-  
015 tured interactions among heterogeneous objectives. We propose an Asymmetric  
016 Training Paradigm that temporarily introduces non-competitive, per-class super-  
017 vision (Sigmoid losses) into networks optimized with competitive softmax objec-  
018 tives. This is realized through orthogonally initialized auxiliary pathways, modu-  
019 lated by a scalar coefficient  $\alpha$  and present only during training. Crucially, we em-  
020 ploy strictly controlled experiments to rule out parameter count as a confounder,  
021 identifying that simple parameter expansion yields zero gain. Our mechanistic  
022 analysis reveals two effects: (1) The proposed topology (but not mere capacity)  
023 consistently smooths the initial optimization landscape. (2) Final performance ex-  
024 hibits an architecture-dependent pattern we term Architectural Resonance, where  
025 auxiliary signals benefit models only when aligned with inductive biases. A 6-  
026 block Vision Transformer (ViT-6L) exhibits constructive gradient alignment (cosine  
027 similarity +0.19), yielding absolute accuracy gains of +9.2% on CIFAR-  
028 100. By contrast, a CNN shows destructive conflicts (cosine similarity -0.26).  
029 We further corroborate this divergence in hybrid architectures (CoAtNet), high-  
030 lighting a stage-dependent nature: transformer stages benefit from heterogeneity,  
031 while convolutional stages show limited compatibility. We validate scalability on  
032 ImageNet-1k, showing consistent top-1 gains for ViTs (up to +2.25% on ViT-  
033 B/16). Rather than functioning as a universal regularizer, our probe reveals that  
034 heterogeneous signals selectively benefit architectures with weak inductive biases  
035 (e.g., Vision Transformers), exposing a critical dependence between architectural  
036 flexibility and objective compatibility.

## 1 INTRODUCTION

038 A fundamental challenge in deep learning is understanding the complex interplay between a model’s  
039 architectural inductive biases and the training strategies it is subjected to. While auxiliary supervi-  
040 sion is a widely adopted technique for improving model performance (Szegedy et al., 2015; Lee  
041 et al., 2015; Caruana, 1997; Ruder, 2017), its application has been predominantly homogeneous,  
042 using objectives conceptually aligned with the main task. This raises a critical question we probe  
043 systematically: how do architectures intrinsically respond to fundamentally heterogeneous super-  
044 vision? Specifically, how does a system designed for “winner-takes-all” competition (via softmax)  
045 react to signals that encourage “feature coexistence” (via sigmoid)?

046 To investigate this, we propose the Asymmetric Training Paradigm (Figure 1), a framework designed  
047 as a precise scientific probe. It temporarily introduces non-competitive, sigmoid-based supervi-  
048 sion into a network through orthogonally initialized pathways, allowing us to systematically study the re-  
049 sulting internal dynamics. Our investigation reveals an architecture-dependent phenomenon that ex-  
050 tends conventional understanding of auxiliary supervision, which we term Architectural Resonance.  
051 On CIFAR-100, this paradigm produces sharply divergent outcomes: Vision Transformers achieve  
052 an accuracy gain of +9.2%, driven by constructive gradient synergy (cosine similarity +0.19), while  
053 Convolutional Neural Networks (CNNs) suffer a degradation of -8.7%, caused by persistent destruc-  
054 tive gradient conflict (cosine similarity -0.26). Crucially, strict capacity control experiments confirm

054 that these gains vanish when the auxiliary pathways are excluded from the loss, ruling out simple  
 055 parameter expansion as the cause.  
 056

057 Our work makes three key contributions:

058 (i) We present the Asymmetric Training Paradigm as a controllable framework for analyzing archi-  
 059 tecture-objective interactions at both model and stage levels.

060 (ii) Using this probe, we identify Architectural Resonance—a stage-dependent phenomenon  
 061 whereby auxiliary supervision efficacy varies with architectural inductive biases. We establish this  
 062 exists on a resilience-modulated spectrum through within-model stage differentiation (CoAtNet on  
 063 CIFAR-100) and model-level validation across architectures (ResNet/ViT on ImageNet-1k).

064 (iii) We provide quantitative characterization of the underlying mechanisms: (a) universal landscape  
 065 smoothing at initialization, and (b) architecture-specific gradient dynamics during training, revealing  
 066 how these jointly influence optimization and generalization.  
 067

## 068 2 RELATED WORK

### 071 2.1 AUXILIARY SUPERVISION AND MULTI-TASK LEARNING

073 The use of intermediate supervision is a well-established technique to combat vanishing gradients  
 074 (Szegedy et al., 2015; Lee et al., 2015). Modern applications leverage auxiliary tasks for represen-  
 075 tation learning (Gidaris et al., 2018; Chen et al., 2020) and Multi-Task Learning (MTL) (Caruana,  
 076 1997; Ruder, 2017; Kendall et al., 2018). A critical challenge in these settings is gradient conflict,  
 077 where competing objectives hinder optimization (Yu et al., 2020; Chen et al., 2018). Prior work has  
 078 largely focused on mitigating such conflicts (e.g., Gradient Surgery (Yu et al., 2020)) or enforcing  
 079 representational consistency within homogeneous objective families (Navon et al., 2022; Shamsian  
 080 et al., 2023). In contrast, our study deliberately employs heterogeneous signals (non-competitive  
 081 vs. competitive) not merely to enhance performance, but as a scientific probe to analyze how dif-  
 082 ferent architectures intrinsically respond to conflicting objectives. This perspective shifts the role  
 083 of auxiliary supervision from a Performance-driven optimization aid to a lens for understanding  
 084 architecture-specific optimization preferences.

### 085 2.2 ARCHITECTURAL INDUCTIVE BIASES

087 Our analysis is grounded in the distinct inductive biases of architectures (Goyal & Bengio, 2022).  
 088 CNNs enforce strong priors on spatial locality and translation equivariance through weight-sharing  
 089 kernels (LeCun et al., 1989; Cohen & Welling, 2016). In contrast, ViTs rely on self-attention for  
 090 global relationships (Dosovitskiy et al., 2021; Vaswani et al., 2017), but this flexibility often comes  
 091 at the cost of trainability, characterized by sharper optimization landscapes and higher sensitivity to  
 092 initialization (Xiao et al., 2021; Chen et al., 2022). While these biases are well-studied in terms of  
 093 optimization landscapes, feature geometries, and inductive bias mechanisms (Lu et al., 2022; Tuli  
 094 et al., 2021), how they govern a model’s response to heterogeneous supervisory signals remains  
 095 largely unexplored. Our work addresses this by treating the supervisory signal as a controlled vari-  
 096 able to probe these intrinsic architectural dynamics.

### 097 2.3 OPTIMIZATION LANDSCAPES AND REGULARIZATION

099 Understanding the geometry of the loss landscape is crucial for explaining generalization (Li et al.,  
 100 2018; Keskar et al., 2017). Flatter minima are generally associated with better generalization, a prin-  
 101 ciple leveraged by techniques like Sharpness-Aware Minimization (SAM) (Foret et al., 2021). Our  
 102 findings contribute to this domain by identifying a structural mechanism for landscape smoothing.  
 103 Following established methods for analyzing landscape geometry (Santurkar et al., 2019), we show  
 104 that orthogonal auxiliary branches universally reduce initial gradient norms. However, this geo-  
 105 metric benefit translates into performance gains only when auxiliary objectives resonate with archi-  
 106 tectural inductive biases. When objectives conflict with rigid spatial priors, gradient misalignment  
 107 during training overrides the initial smoothing, leading to degradation despite improved landscape  
 geometry.

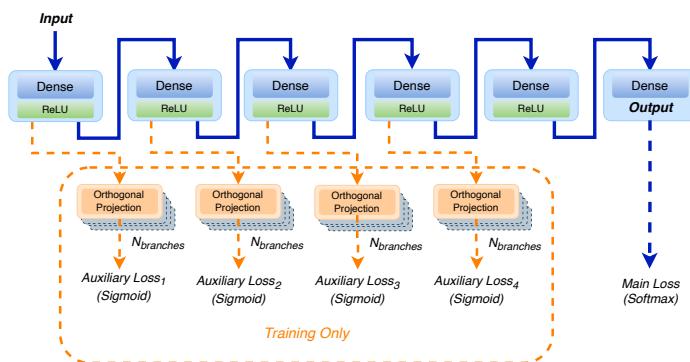


Figure 1: The Asymmetric Training Paradigm. Schematic overview using a simplified MLP architecture. Multiple auxiliary branches (orange) inject heterogeneous supervision (Sigmoid) into the backbone. These branches are removed at inference time, ensuring zero inference overhead.

### 3 METHODOLOGY

#### 3.1 THE ASYMMETRIC TRAINING PARADIGM

We introduce the Asymmetric Training Paradigm (Figure 1), a controllable framework for probing architecture-objective interactions through structured gradient modulation. Unlike standard Multi-Task Learning which seeks to optimize multiple outputs, our paradigm uses auxiliary branches strictly as training-time scaffolding to analyze how different architectures respond to heterogeneous supervision. The framework is built on three pillars: **Asymmetry** (auxiliary branches are discarded at inference), **Heterogeneity** (auxiliary objectives differ qualitatively from the primary task), and **Controlled Redundancy** (scalable orthogonal pathways enabling systematic characterization of architectural resilience).

#### 3.2 CORE HYPOTHESIS: ARCHITECTURAL RESONANCE

We propose the Architectural Resonance Hypothesis: The efficacy of heterogeneous auxiliary supervision depends on the compatibility between auxiliary signal characteristics and architectural inductive biases. This manifests as a spectrum of gradient interactions during training:

- **Constructive Interference:** When auxiliary signals align with an architecture’s inductive bias (e.g., spatially-agnostic non-competitive signals for ViTs’ global modeling capacity), they induce positive gradient alignment, enabling improved trainability and performance.
- **Destructive Interference:** When signals conflict with rigid structural priors (e.g., spatially-agnostic projections for CNNs’ locality bias), they cause persistent gradient misalignment and performance degradation.

This reveals that auxiliary supervision efficacy is architecture-dependent: the degree of resonance determines whether heterogeneous signals enhance or impair training.

#### 3.3 MECHANISM DESIGN AND CONTROLS

##### 3.3.1 ARCHITECTURAL SETUP

To isolate the impact of inductive biases, we employ three controlled lightweight backbone architectures spanning a spectrum of inductive biases: a 6-layer MLP (minimal structural bias), a 6-conv-layer CNN with spatial downsampling (rigid locality bias), and a 6-block Vision Transformer (ViT-6L, flexible attention-based bias). To establish scalability, we extend our evaluation to ImageNet-1k using standard ResNet-18/50 and ViT-Small/B-16 backbones (Section 4.1). To validate stage-level mechanisms, we further analyze the hybrid CoAtNet architecture on CIFAR-100

(Section 4.2), which combines convolutional and transformer stages within a single model to enable within-network comparisons. Full architectural details are provided in Appendix C.3.

### 3.3.2 THE GRADIENT MODULATION MECHANISM

A critical innovation in our method is separating parameter capacity from gradient flow. At each attachment point, we instantiate  $N_{\text{branches}}$  parallel branches. We distinctly categorize them into two roles to isolate the source of gains: (1) **Static Gradient Modulators** ( $k = 1 \dots N - 1$ ): These branches are initialized orthogonally and then frozen (via `stop_gradient` applied to their projection weights). Crucially, while their weights  $W_{\text{aux}}$  do not update ( $\nabla_{W_{\text{aux}}} L = 0$ ), gradients still backpropagate through these fixed projections into the shared backbone features  $h$  ( $\nabla_h L \neq 0$ ). (2) **The Active Probe** ( $k = N$ ): The final branch is fully trainable: unlike the frozen branches, its projection weights receive gradients and adapt during training. All auxiliary branches contribute Sigmoid/BCE losses, but only the last branch updates its own weights, while the earlier branches remain fixed orthogonal projections.

As a result, the backbone is optimized under a combination of rigid geometric constraints (from the static, orthogonal anchors whose gradients backpropagate through fixed projections) and adaptive error signals (from the active branch that continuously adjusts to the data). This design transforms these branches into fixed, structured “lenses” that refract the gradient flow without adding learnable capacity. By scaling  $N$ , we can systematically smooth the optimization landscape without the confounding factor of adding learnable parameters. As we empirically validate in Section 4.4 and Appendix A.1, this topology consistently reduces the initial gradient norm, creating a more favorable geometry for optimization.

### 3.3.3 CAPACITY CONTROL VARIANT

To address the concern that improvements might stem simply from “adding more parameters” (even static ones), we define a Capacity Control baseline within our methodology. In this variant, we instantiate the exact same  $N$  branches with identical initialization, but detach them from the loss graph. This ensures the model has the same parameter count and architecture, but no gradient modulation occurs. This control allows us to attribute performance gains strictly to the interaction of gradients, not the existence of parameters. Results are reported in Table 6.

### 3.3.4 VARIABLES: REDUNDANCY AND DIALOGUE STRENGTH

We manipulate two variables to map the resonance landscape:

- **Redundancy ( $N_{\text{branches}}$ ):** The number of parallel branches per attachment point, denoted as  $N \times$ . This hyperparameter modulates the intensity of initial landscape smoothing. We evaluate performance regimes up to  $20 \times$  and extend to  $300 \times$  to probe architectural limits and analyze landscape mechanics.
- **Dialogue Strength ( $\alpha$ ):** A scalar hyperparameter that balances the primary and auxiliary losses. Let  $M$  be the number of attachment points (layers) and  $N$  be the redundancy level (number of branches) at each point. The total training objective is:

$$L_{\text{total}} = L_{\text{main}} + \alpha \sum_{m=1}^M \sum_{n=1}^N L_{\text{aux}}^{(m,n)} \quad (1)$$

where  $L_{\text{aux}}^{(m,n)}$  denotes the Sigmoid loss of the  $n$ -th branch at the  $m$ -th attachment point.

### 3.3.5 HETEROGENEOUS OBJECTIVES

To induce the requisite “signal dialogue,” we enforce heterogeneity between the objectives:

- **Primary (Competitive):** Standard Softmax Cross-Entropy ( $L_{\text{main}}$ ), encouraging winner-takes-all feature discrimination.
- **Auxiliary (Coexistent):** Sigmoid Binary Cross-Entropy ( $L_{\text{aux}}$ ), computed independently per class. This encourages the model to capture non-exclusive features for each class, fundamentally differing from the Softmax dynamic. We explicitly choose this heterogeneous

216 design because, as shown in our Heterogeneity Control experiments (Section 4.4), replacing  
 217 Sigmoid with a homogeneous Softmax objective leads to performance degradation.  
 218

219 **3.4 EVALUATION PROTOCOL: BALANCING RIGOR AND SCALABILITY**  
 220

221 To ensure that our findings are both scientifically rigorous and practically applicable, we design a  
 222 distinct evaluation protocol that adapts to the scale of the problem.  
 223

224 **Mechanistic Analysis (CIFAR)** For our controlled studies, our priority is isolating the source of  
 225 gains. We employ a strict Two-Stage Tuning Protocol to rule out hyperparameter confounding: We  
 226 first exhaustively tune the Baseline to find its optimal learning rate and weight decay. We then fix  
 227 these backbone hyperparameters and search only for the dialogue strength  $\alpha$  for our method.  
 228

229 **Large-Scale Validation (ImageNet)** For ImageNet-1k, we strictly maintain fixed backbone hy-  
 230 perparameters, attaching auxiliary heads only at predefined intermediate points (Table 21). Opti-  
 231 mization follows standard protocols: linear scaling SGD for ResNet (Goyal et al., 2018) and AdamW  
 232 for ViT (adapted from (Touvron et al., 2021) using a clean baseline protocol). By tuning only  $\alpha$  (on  
 233 a held-out subset), we assess the method’s scalability in realistic, compute-constrained scenarios.  
 234

235 Crucially, across all experimental settings, including the targeted hybrid verification on CoAtNet  
 236 (Dai et al., 2021), backbone hyperparameters are kept strictly identical between the baseline and  
 237 our method. This guarantees that any observed performance difference, whether improvement or  
 238 degradation, is attributable solely to the proposed topological interaction, ruling out hyperparameter  
 239 mismatch as a confounder.

240 **4 EXPERIMENTS**

241 **4.1 MAIN RESULTS: THE DIVERGENCE OF ViT AND CNN**

242 We first evaluate the Asymmetric Training Paradigm on CIFAR-100 using three representative ar-  
 243 chitectures: MLP, CNN, and ViT-6L across 10 random seeds. As shown in Table 1, the impact of  
 244 our heterogeneous auxiliary supervision is fundamentally architecture-dependent.  
 245

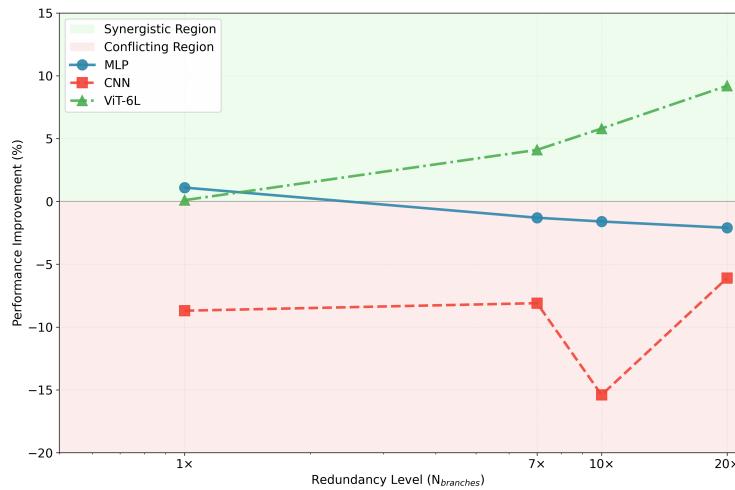
246 **Dose-Response on CIFAR-100** We observe a clear correlation between the level of redundancy  
 247 ( $N_{\text{branches}}$ ) and performance modulation (Figure 2), revealing distinct architectural preferences:  
 248

- 249 • **Vision Transformers (Constructive Synergy):** ViT-6L benefits significantly, with accu-  
 250 racy improving monotonically as redundancy increases. At  $20\times$  redundancy, it achieves a  
 251 +9.2% improvement in top-1 accuracy over the baseline. This provides evidence that ViTs  
 252 can constructively integrate dense heterogeneous signals.
- 253 • **CNNs (Destructive Conflict):** Conversely, the CNN suffers severe degradation across all  
 254 redundancy levels, experiencing up to a -15.4% decrease in accuracy. This indicates a  
 255 fundamental incompatibility between spatially-agnostic Sigmoid signals and the CNN’s  
 256 strong locality priors.
- 257 • **MLPs (Inconsistent/Noisy Interaction):** The MLP exhibits a distinct behavior. While a  
 258 single auxiliary branch ( $1\times$ ) provides mild regularization (+1.1%), increasing redundancy  
 259 leads to consistent degradation (e.g., -2.1% at  $20\times$ ). This suggests that without structural  
 260 mechanisms such as self-attention to align auxiliary signals, MLPs are destabilized by  
 261 gradient noise as redundancy scales.

262 **Scalability on ImageNet-1k** To confirm that the observed divergence is not limited to small-scale  
 263 datasets, we extend our evaluation to ImageNet-1k using ResNet-18/50 and ViT-Small/B-16 across 4  
 264 random seeds (Table 2). We strictly adhere to established training recipes (SGD for ResNet, AdamW  
 265 for ViT) while adopting a “clean baseline” setup—standard augmentation without heavy regulariza-  
 266 tion (e.g., Mixup, CutMix). The only structural difference between the Baseline and Asymmetric  
 267 models is the addition of orthogonally initialized auxiliary heads. Both share identical backbones,  
 268 data pipelines, and optimization hyperparameters, details are provided in Appendix C.4. Crucially,  
 269

270 Table 1: Architecture performance comparison across different configurations on CIFAR-100  
271

$N_{\text{branches}}$	Arch	Plain	Asymmetric	p-value
1×	MLP	$23.2 \pm 0.3$	$24.3 \pm 0.3 (+1.1; \alpha = 100.0)$	$< 0.001$
	CNN	$39.5 \pm 0.4$	$30.8 \pm 0.3 (-8.7; \alpha = 100.0)$	$< 0.001$
	ViT-6L	$35.8 \pm 1.0$	$35.9 \pm 0.5 (+0.1; \alpha = 4.642)$	0.7773
7×	MLP	$23.4 \pm 0.5$	$22.1 \pm 0.3 (-1.3; \alpha = 0.1)$	$< 0.001$
	CNN	$39.8 \pm 0.6$	$31.7 \pm 2.9 (-8.1; \alpha = 0.1)$	$< 0.001$
	ViT-6L	$35.9 \pm 0.8$	$40.0 \pm 2.1 (+4.1; \alpha = 4.642)$	$< 0.001$
10×	MLP	$23.3 \pm 0.3$	$21.7 \pm 0.2 (-1.6; \alpha = 0.1)$	$< 0.001$
	CNN	$39.7 \pm 0.6$	$24.3 \pm 2.1 (-15.4; \alpha = 1.0)$	$< 0.001$
	ViT-6L	$36.0 \pm 0.9$	$41.8 \pm 2.2 (+5.8; \alpha = 4.642)$	$< 0.001$
20×	MLP	$23.2 \pm 0.3$	$21.1 \pm 0.3 (-2.1; \alpha = 0.1)$	$< 0.001$
	CNN	$39.7 \pm 0.6$	$33.6 \pm 1.2 (-6.1; \alpha = 0.1)$	$< 0.001$
	ViT-6L	$36.1 \pm 1.0$	$45.3 \pm 1.3 (+9.2; \alpha = 1.0)$	$< 0.001$

303  
304 Figure 2: Architecture-dependent dose-response to auxiliary supervision. Each architecture shows  
305 distinct sensitivity patterns: ViT (constructive synergy), CNN (destructive conflict), MLP (incon-  
306 sistent/noisy interaction). Performance improvement plotted against branch redundancy ( $N_{\text{branches}}$ ) on  
307 CIFAR-100.  
308309 to rule out hyperparameter hacking, we tuned the dialogue strength  $\alpha$  on a 10% held-out train-  
310 ing subset, ensuring the validation set remained strictly unseen during the search. This rigorous  
311 setup guarantees that observed gains stem purely from structural resonance rather than overfitting or  
312 data-level artifacts. Consistent with CIFAR-100, ViT backbones show robust improvements (e.g.,  
313 +2.25% on ViT-B/16 at 20× redundancy,  $p < 0.001$ ), suggesting that constructive resonance scales  
314 to large-scale benchmarks. In contrast, ResNets remain neutral. This suggests that architectural  
315 resilience mechanisms (e.g., residual connections) may mitigate the gradient conflicts that caused  
316 plain CNNs to collapse, thereby reinforcing the architectural dependence of the phenomenon.317 While ViTs consistently benefit from auxiliary resonance, CNNs exhibit either destructive conflict  
318 (Plain CNN) or neutrality (ResNet). Given these distinct architectural preferences, a critical question  
319 arises: how do hybrid architectures behave? We address this in Section 4.2.  
320321 4.2 TARGETED VERIFICATION ON HYBRID ARCHITECTURE (COATNET)  
322323 To empirically validate the stage-dependent Efficacy of our Architectural Resonance hypothesis, we  
324 extend our analysis to CoAtNet (Dai et al., 2021), a hybrid architecture that integrates convolutional

324  
325  
326 Table 2: Architecture performance comparison on ImageNet-1k  
327  
328  
329  
330  
331

Architecture	Baseline	1×	10×	20×
ResNet-18	68.23±0.20	68.33±0.10	68.29±0.10	68.37±0.13
ResNet-50	73.83±0.09	73.55±0.17**	73.75±0.11	73.76±0.21
ViT-Small	69.13±0.14	70.16±0.09***	70.27±0.18**	70.35±0.29**
ViT-B/16	66.75±0.13	68.34±0.64*	68.83±0.27***	69.00±0.19***

332 *Statistical significance:* \*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$ .

333  
334  
335 stages (early) and transformer stages (late). This allows us to test whether the observed divergence  
336 is strictly stage-dependent.

337 We employ a CIFAR-100 adapted CoAtNet-0 backbone to strictly preserve the original multi-  
338 stage layout while adapting to the dataset resolution. We apply our asymmetric probe separately  
339 to the Convolutional Stage (S2) and the Transformer Stage (S3) across varying redundancy levels  
340 ( $N_{branches}$ ) to isolate their differential responses.

341 As presented in Table 3, the results reveal an observable divergence between stages within the same  
342 model: (1) **Synergy in Transformer Stages (S3):** Consistent with our main ViT results, applying  
343 auxiliary supervision to S3 yields consistent improvements. It achieves a peak accuracy of 77.08%  
344 (+1.74%), showing statistically significant resonance ( $p < 0.001$ ). (2) **Conflict in Convolutional  
345 Stages (S2):** In contrast, applying dense probes to S2 shows limited compatibility. While mild  
346 redundancy (20×) exhibits marginal response, high redundancy (300×) leads to a statistically sig-  
347 nificant tendency toward conflict (-0.80%,  $p < 0.05$ ).

348 These findings provide targeted empirical support for our Architectural Resonance Hypothesis.  
349 Comparing CoAtNet S2 with the plain CNN (Table 1) reveals the role of architectural resilience:  
350 modern components such as residual blocks and batch normalization help buffer against immediate  
351 collapse (S2: neutral at 1× vs. CNN: -8.7%), yet degradation still emerges under extreme redun-  
352 dancy (S2: -0.8% at 300×). This confirms that auxiliary supervision efficacy is not binary but stage-  
353 dependent: attention-based stages benefit consistently from heterogeneity, whereas convolutional  
354 stages exhibit limited compatibility—delayed by modern architectural features but fundamentally  
355 prone to conflict.

356  
357 Table 3: Stage-Dependent Response on CIFAR-100 adapted CoAtNet-0. Comparison of applying  
358 auxiliary supervision to CNN (S2) vs. Transformer (S3) stages. Results averaged over 5 random  
359 seeds.

Stage	1×	20×	100×	300×
S2 (CNN)	75.26±0.59 (-0.08)	75.88±0.14 (+0.54**)	75.59±0.44 (+0.25)	<b>74.54±0.48</b> <b>(-0.80*)</b>
S3 (Transformer)	75.91±0.28 (+0.57**)	76.76±0.37 (+1.42***)	<b>77.08±0.27</b> <b>(+1.74***)</b>	76.82±0.30 (+1.48***)
Baseline	$75.34 \pm 0.20$			
Phenomenon	ViT starts gaining	Strong Synergy (S3)	<b>Peak (S3)</b>	<b>Conflict (S2)</b>

370 *Statistical significance:* \*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$ .

### 374 4.3 UNDERLYING DYNAMICS: LANDSCAPE GEOMETRY AND GRADIENT FLOW

375  
376 Having confirmed the performance gains, we now investigate the mechanistic cause. We analyze the  
377 training dynamics from two perspectives: the initial geometry of the optimization landscape and the  
directional alignment of gradients during training.

378     **Initial Phase: Landscape Smoothing** As hypothesized in Section 3, the static orthogonal  
 379     branches are designed to condition the loss landscape. Following (Santurkar et al., 2019), we quan-  
 380     tify this by measuring the Initial Gradient Norm of the main loss across 5 random seeds.  
 381

382     • **Observation:** Table 4 presents the results. Increasing redundancy ( $N_{\text{branches}}$ ) drastically  
 383     reduces the gradient norm across all architectures. Notably, for MLP and CNN,  $300\times$   
 384     redundancy reduces the gradient norm by over 90%.

385     • **Implication:** This confirms that our static topology acts as a “geometric conditioner,” cre-  
 386     ating a smoother initial surface. This is particularly beneficial for ViTs, which are known  
 387     to suffer from sharp, ill-conditioned landscapes (Chen et al., 2022).

389     Table 4: Landscape Smoothing Effect  
 390

Architecture	1×		20×		300×	
	Result	p-value	Result	p-value	Result	p-value
MLP	-29.86±20.35	0.0305	-65.38±1.95	< 0.001	-90.40±0.72	< 0.001
ViT-6L	-3.40±0.47	< 0.001	-33.68±3.62	< 0.001	-68.78±8.85	< 0.001
CNN	-16.64±4.79	0.0015	-67.22±5.52	< 0.001	-90.18±2.38	< 0.001

397     398     **Training Phase: Gradient Alignment** While landscape smoothing is generally beneficial, why  
 399     400     do CNNs degrade under the same condition? To answer this, we analyze the Cosine Similarity  
 401     402     between the gradients of the main task ( $L_{\text{main}}$ ) and the auxiliary task ( $L_{\text{aux}}$ ) throughout training  
 403     (Figure 3 and Table 5).

404     • **Constructive Synergy (ViT):** ViTs exhibit positive cosine similarity (mean +0.19), indi-  
 405     cating that the auxiliary signals point in a direction compatible with the main objective.  
 406     This alignment allows the ViT to effectively exploit the smoothed landscape established in  
 407     the initial phase, translating geometric potential into performance gains.

408     • **Destructive Conflict (CNN):** Despite the smoothed landscape, CNNs show persistent neg-  
 409     ative similarity, reaching as low as  $-0.82$  (mean  $-0.26$ ). This suggests that Sigmoid sig-  
 410     nals fundamentally conflict with the CNN’s locality-biased kernels. This intense gradient  
 411     conflict overrides the benefits of smoothing, turning the auxiliary signals into destructive  
 412     interference.

413     • **Inconsistent / Noisy Interaction (MLP):** The MLP trajectory fluctuates without a con-  
 414     sistent direction, alternating between weak positive and negative values (e.g., oscillating  
 415     between  $+0.07$  and  $-0.34$ ). Lacking strong inductive biases to orient these auxiliary sig-  
 416     nals, the interaction is effectively incoherent. Consequently, as redundancy increases, these  
 417     unaligned signals accumulate as gradient noise rather than constructive guidance, explain-  
 418     ing the degradation observed in Table 1.

420     Table 5: Detailed Gradient Conflict Analysis  
 421

Architecture	Final Sim	Avg Sim	Min Sim	Max Sim	Gradient Interaction
MLP	-0.0309	-0.0801	-0.3422	0.0701	Inconsistent / Noisy Interaction
CNN	-0.1926	-0.2574	-0.8210	0.0464	Destructive Conflict
ViT	0.2631	0.1870	-0.1845	0.3654	Constructive synergy

422     423     **Visualizing the Resonance: Attention Maps** To visibly corroborate the gradient alignment find-  
 424     ings, we visualize the self-attention maps of the final CLS token (Figure 4). Compared to the Base-  
 425     426     line, the Asymmetric ViT exhibits significantly sharper attention on the semantic object foreground,  
 427     filtering out background noise. This confirms that the constructive gradient synergy translates into  
 428     429     more semantic feature extraction.

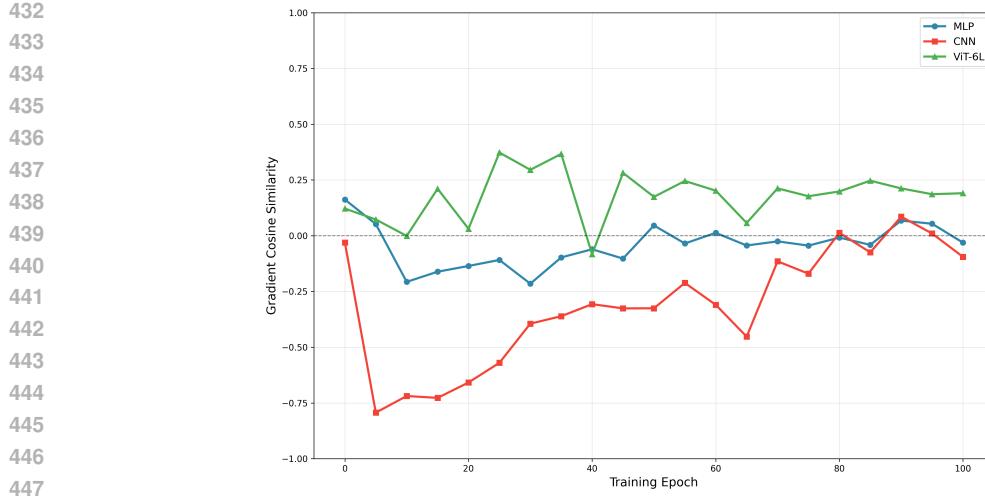


Figure 3: Gradient conflict evolution across architectures during training

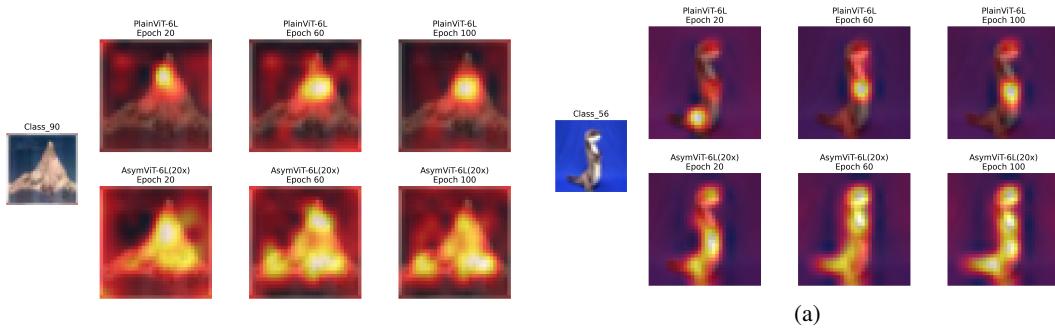


Figure 4: Visualizing Architectural Resonance. Attention maps show that Asymmetric Training induces sharper focus on the object compared to Baseline (PlainViT-6L), confirming constructive feature learning.

#### 4.4 MECHANISM VERIFICATION: RULING OUT CONFOUNDERS

A critical question is identifying the source of these gains. Is the improvement in ViTs driven by the proposed topological interaction, or simply by parameter expansion (capacity) or generic deep supervision? We address this through strictly controlled ablation studies (10 random seeds).

**Capacity Control** Is it just adding parameters? We compare against a Capacity Control baseline where the same  $N$  auxiliary branches are instantiated but detached from the loss graph (weights present but gradients blocked). As shown in Table 6 (Top), mere parameter redundancy yields zero statistical gain ( $p > 0.05$ ). For ViT-6L, increasing branches to  $20\times$  in the control group results in negligible fluctuation, whereas our active method achieves  $+9.2\%$  improvement. This rigorously rules out implicit regularization from parameter count as the cause. The gain stems from the interaction of gradients.

**Heterogeneity Control** Is it just Deep Supervision? We investigate the necessity of Heterogeneity by replacing our Sigmoid auxiliary loss with a standard Softmax auxiliary loss (Homogeneous). Table 6 (Middle) reveals that homogeneous supervision degrades performance for ViTs (e.g.,  $-5.3\%$  drop). This confirms that ViTs specifically benefit from the “non-competitive” nature of the Sigmoid signal to facilitate capacity exploitation.

**Random Init Control** Is it just initialization noise? We compare our Orthogonal Initialization strategy against standard Random Initialization for the auxiliary branches. As shown in Table 6

(Bottom), using Random Initialization fails to provide consistent gains and often leads to training instability (high variance). This suggests that the structural orthogonality is a prerequisite for effective resonance. The auxiliary branches must be geometrically aligned (via orthogonality) to probe the landscape constructively, rather than injecting unstructured noise.

490

491 Table 6: Ablation study results across different control conditions on CIFAR-100. This table investigates potential confounders. **Top:** Adding parameters without gradient flow yields no gain. **Middle:** Replacing Sigmoid with Softmax (Homogeneous) causes degradation. **Bottom:** Using Random Initialization instead of Orthogonal leads to instability (e.g., ViT -4.0%). **Contrast:** Our Asymmetric method (Sigmoid + Orthogonal) achieves **+9.2%** absolute improvement on ViT under the same  $20 \times$  condition (Table 1).

497

Control Type	$N_{branches}$	MLP	CNN	ViT-6L
Capacity Control	1×	23.2±0.3 (-0.1)	39.9±0.7 (+0.7)	36.1±1.1 (+0.1)
	10×	23.1±0.4 (-0.0)	39.8±0.8 (-0.1)	35.8±1.2 (-0.3)
	20×	23.1±0.4 (-0.0)	40.0±1.0 (+0.4)	36.1±1.1 (-0.0)
Heterogeneity Control	1×	24.1±0.3 (+0.8**)	37.4±0.4 (-2.1***)	36.1±0.5 (+0.2)
	10×	20.4±0.6 (-2.7***)	36.2±0.9 (-3.4***)	33.9±0.6 (-1.9***)
	20×	19.0±0.5 (-4.3***)	37.9±1.0 (-1.9***)	35.7±0.9 (-0.3)
Random Init. Control	1×	24.3±0.4 (+1.0***)	31.6±0.7 (-7.9***)	36.3±0.9 (+0.3)
	10×	22.8±0.3 (-0.3)	30.6±1.0 (-9.6***)	35.6±1.1 (-0.5)
	20×	23.0±0.2 (-0.3)	30.8±0.8 (-9.1***)	31.8±1.1 (-4.0***)

509 *Statistical significance:* \*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$ .

510

511

## 5 DISCUSSION AND LIMITATIONS

513

514 Our Asymmetric Training Paradigm serves as a probe for architecture–objective interactions, yielding 515 measurable gains in the tested ViT backbones and offering mechanistic insights into gradient 516 dynamics.

517

518 **Implications** Crucially, targeted verification on CoAtNet (Section 4.2) advances our understanding 519 beyond binary outcomes. It suggests a Stage-Adaptive Strategy for modern model design: heterogeneous 520 auxiliary signals may be most effective in attention-based stages, while requiring caution 521 in convolutional stages to avoid structural conflict. This challenges the assumption that “more 522 supervision is always better.”

523

524 **Limitations** Despite these findings, we acknowledge several limitations: (1) Convolutional 525 Compatibility: CNN stages in our experiments show limited compatibility with dense heterogeneous 526 probes, reflecting the rigidity of spatial inductive biases. This currently restricts applicability to 527 attention-based backbones; (2) Training Overhead: While incurring zero inference overhead, training 528 costs scale linearly with redundancy ( $N_{branches}$ ), creating memory pressure for extremely large- 529 scale pre-training; (3) Hyperparameter Search: The current approach relies on a grid search to iden- 530 tify the optimal dialogue strength  $\alpha$ , which is time-consuming and may yield sub-optimal config- 531 urations compared to dynamic schedules; (4) Theoretical Formalism: While we provide empirical 532 evidence and mechanistic analysis, a closed-form theoretical framework quantifying the relationship 533 between orthogonality, redundancy, and generalization gap remains an open challenge; (5) Scope: 534 Whether Architectural Resonance extends to other modalities (e.g., NLP) or loss types requires 535 verification.

536

537 **Future Work** We envision five key directions: (1) designing spatially-aware signals that respect 538 CNN locality; (2) developing adaptive modulation strategies for  $\alpha$  to eliminate expensive search; 539 (3) distilling the geometric smoothing effect of redundancy into computationally efficient, analyti- 540 cally equivalent formulations to reduce memory pressure; (4) mathematically formalizing resonance 541 conditions; and (5) extending the paradigm to diverse domains to test universality beyond vision.

540 REFERENCES  
541

542 Rich Caruana. Multitask learning. *Machine learning*, 28(1):41–75, 1997.

543 Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for  
544 contrastive learning of visual representations. In *International Conference on Machine Learning*,  
545 pp. 1597–1607. PMLR, 2020.

546 Xiangning Chen, Cho-Jui Hsieh, and Boqing Gong. When vision transformers outperform resnets  
547 without pre-training or strong data augmentations, 2022.

548 Zhao Chen, Vijay Badrinarayanan, Chen-Yu Lee, and Andrew Rabinovich. Gradnorm: Gradient  
549 normalization for adaptive loss balancing in deep multitask networks. In *International Conference  
550 on Machine Learning*, pp. 794–803. PMLR, 2018.

551 Taco Cohen and Max Welling. Group equivariant convolutional networks. In *International Conference  
552 on Machine Learning*, pp. 2990–2999. PMLR, 2016.

553 Zihang Dai, Hanxiao Liu, Quoc V. Le, and Mingxing Tan. Coatnet: Marrying convolution and  
554 attention for all data sizes, 2021.

555 Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas  
556 Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An im-  
557 age is worth 16x16 words: Transformers for image recognition at scale. In *International Conference  
558 on Learning Representations*, 2021.

559 Pierre Foret, Ariel Kleiner, Hossein Mobahi, and Behnam Neyshabur. Sharpness-aware minimiza-  
560 tion for efficiently improving generalization, 2021.

561 Spyros Gidaris, Praveer Singh, and Nikos Komodakis. Unsupervised representation learning by  
562 predicting image rotations. In *International Conference on Learning Representations*, 2018.

563 Anirudh Goyal and Yoshua Bengio. Inductive biases for deep learning of higher-level cognition.  
564 *Proceedings of the Royal Society A*, 478(2266):20210068, 2022.

565 Priya Goyal, Piotr Dollár, Ross Girshick, Pieter Noordhuis, Lukasz Wesolowski, Aapo Kyrola, An-  
566 drew Tulloch, Yangqing Jia, and Kaiming He. Accurate, large minibatch sgd: Training imagenet  
567 in 1 hour, 2018.

568 Alex Kendall, Yarin Gal, and Roberto Cipolla. Multi-task learning using uncertainty to weigh losses  
569 for scene geometry and semantics. In *Proceedings of the IEEE Conference on Computer Vision  
570 and Pattern Recognition*, pp. 7482–7491, 2018.

571 Nitish Shirish Keskar, Dheevatsa Mudigere, Jorge Nocedal, Mikhail Smelyanskiy, and Ping Tak Pe-  
572 ter Tang. On large-batch training for deep learning: Generalization gap and sharp minima. In  
573 *International Conference on Learning Representations*, 2017.

574 Y. LeCun, B. Boser, J. S. Denker, D. Henderson, R. E. Howard, W. Hubbard, and L. D. Jackel.  
575 Backpropagation applied to handwritten zip code recognition. *Neural Computation*, 1(4):541–  
576 551, 1989. doi: 10.1162/neco.1989.1.4.541.

577 Chen-Yu Lee, Saining Xie, Patrick Gallagher, Zhengyou Zhang, and Zhuowen Tu. Deeply-  
578 supervised nets. In *Artificial intelligence and statistics*, pp. 562–570. PMLR, 2015.

579 Hao Li, Zheng Xu, Gavin Taylor, Christoph Studer, and Tom Goldstein. Visualizing the loss land-  
580 scape of neural nets. *Advances in Neural Information Processing Systems*, 31, 2018.

581 Zhiying Lu, Hongtao Xie, Chuanbin Liu, and Yongdong Zhang. Bridging the gap between vision  
582 transformers and convolutional neural networks on small datasets, 2022.

583 Aviv Navon, Aviv Shamsian, Idan Achituv, Haggai Maron, Kenji Kawaguchi, Gal Chechik, and  
584 Ethan Fetaya. Multi-task learning as a bargaining game. In Kamalika Chaudhuri, Stefanie Jegelka,  
585 Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato (eds.), *Proceedings of the 39th Inter-  
586 national Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning  
587 Research*, pp. 16428–16446. PMLR, 17–23 Jul 2022.

594 Sebastian Ruder. An overview of multi-task learning in deep neural networks. *ArXiv*,  
595 abs/1706.05098, 2017.

596

597 Shibani Santurkar, Dimitris Tsipras, Andrew Ilyas, and Aleksander Madry. How does batch normal-  
598 ization help optimization?, 2019.

599

600 Aviv Shamsian, Aviv Navon, Neta Glazer, Kenji Kawaguchi, Gal Chechik, and Ethan Fetaya. Aux-  
601 iliary learning as an asymmetric bargaining game, 2023.

602

603 Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Du-  
604 mitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In  
605 *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1–9, 2015.

606

607 Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and  
608 Hervé Jégou. Training data-efficient image transformers & distillation through attention, 2021.

609

610 Shikhar Tuli, Ishita Dasgupta, Erin Grant, and Thomas L. Griffiths. Are convolutional neural net-  
611 works or transformers more like human vision?, 2021.

612

613 Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez,  
614 Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in Neural Infor-  
615 mation Processing Systems*, 30, 2017.

616

617 Tete Xiao, Mannat Singh, Eric Mintun, Trevor Darrell, Piotr Dollár, and Ross B. Girshick. Early  
618 convolutions help transformers see better. In *Neural Information Processing Systems*, 2021.

619

620

621

622

623

624

625

626

627

628

629

630

631

632

633

634

635

636

637

638

639

640

641

642

643

644

645

646

647

648  
649

## A APPENDIX

650  
651

## A.1 LOSS LANDSCAPE ANALYSIS

652  
653  
654  
655  
656  
657  
658  
659  
660

We conducted systematic loss landscape analysis across MLP, CNN, and ViT-6L architectures on CIFAR-10 and CIFAR-100. Following established protocols (Li et al., 2018), we visualized the loss surfaces using a  $51 \times 51$  grid centered at the initialization point, with directions determined by random Gaussian perturbations normalized to unit variance. To balance computational efficiency with statistical reliability, we randomly sampled 500 training instances for loss evaluation at each grid point. This sampling size provides sufficient statistical power while remaining computationally tractable for systematic analysis across multiple architectures and redundancy levels. The resulting visualizations reveal distinct architectural signatures in terms of loss surface smoothness and optimization landscape complexity. (Figure 5, Table 7 and 8)

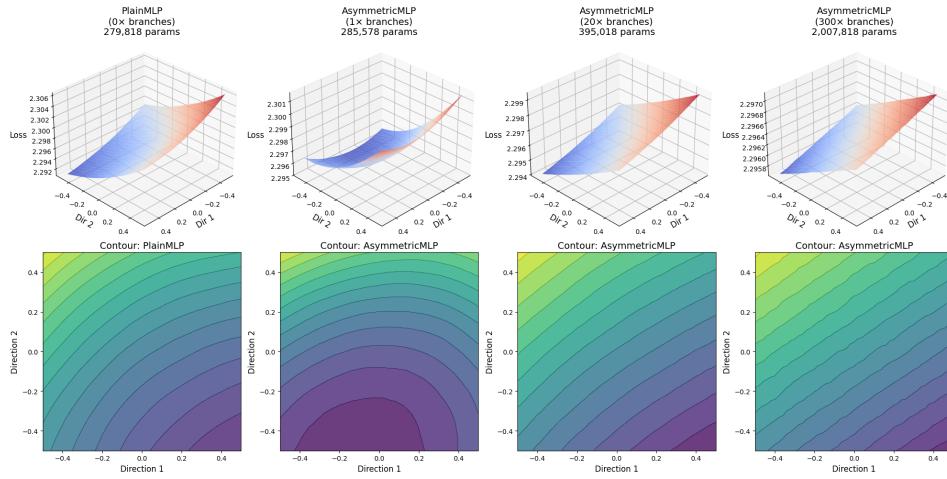
661  
662  
663  
664  
665  
666  
667  
668  
669  
670  
671  
672  
673  
674  
675  
676  
677  
678

Figure 5: Loss landscape visualization for MLP on CIFAR-10

679  
680  
681  
682  
683

Table 7: Progressive loss landscape smoothing in MLP architecture on CIFAR-10. Standard deviation (Std), range, and mean gradient magnitude all decrease systematically with increased redundancy, demonstrating that topological modifications consistently flatten the optimization surface independent of final performance outcomes. Percentages indicate relative change from baseline.

Model	Params	Std(Loss)	Range(Loss)	Mean(Grad)
Plain	0.28M	0.0032	0.0152	0.0002
Asymmetric(1x)	0.29M	0.0015 (-52.3%)	0.0064 (-57.6%)	0.0001 (-49.8%)
Asymmetric(20x)	0.40M	0.0011 (-64.4%)	0.0053 (-64.9%)	0.0001 (-65.0%)
Asymmetric(300x)	2.01M	0.0003 (-90.2%)	0.0015 (-90.3%)	0.0000 (-90.4%)

689

690

## B BOUNDARY CONDITIONS AND EXTENDED ANALYSIS

693

## B.1 SAMPLE EFFICIENCY AND GENERALIZATION IN LOW-DATA REGIMES

695  
696  
697  
698  
699

To assess the generalizability of our architectural resonance findings under data-scarce conditions, we conducted systematic few-shot learning experiments on CIFAR-10 and CIFAR-100. We hypothesize that asymmetric training benefits should be amplified in low-data regimes, where auxiliary supervision can provide crucial structural guidance when primary signals are sparse (Tables 9, 10, and 11).

700  
701

**Experimental Design.** We systematically varied the number of training samples per class from 5 to 5000, creating a comprehensive data scarcity spectrum. For each data regime, we maintained the original test set size to ensure consistent evaluation conditions. All experiments were repeated

702  
 703 Table 8: Cross-architecture comparison of loss landscape smoothing on CIFAR-100. Despite universal  
 704 landscape flattening effects (up to 90% reduction in surface roughness), architectural differences  
 705 emerge: CNN shows the most dramatic smoothing with minimal parameter increase, while ViT-  
 706 6L exhibits more modest but consistent improvements. These results demonstrate that landscape  
 707 conditioning is architecture-agnostic, yet performance benefits depend critically on architectural  
 708 resonance with auxiliary signals.

Model	Params	Std(Loss)	Range(Loss)	Mean(Grad)
<b>CNN</b>				
Plain	2.43M	0.0008	0.0040	0.0001
Asymmetric(1 $\times$ )	2.51M	0.0007 (-16.3%)	0.0033 (-17.3%)	0.0000 (-15.9%)
Asymmetric(20 $\times$ )	4.09M	0.0003 (-66.7%)	0.0013 (-68.0%)	0.0000 (-66.9%)
Asymmetric(300 $\times$ )	27.4M	0.0001 (-90.2%)	0.0004 (-90.3%)	0.0000 (-90.2%)
<b>ViT-6L</b>				
Plain	1.22M	0.0045	0.0212	0.0003
Asymmetric(1 $\times$ )	1.29M	0.0043 (-3.7%)	0.0204 (-3.7%)	0.0003 (-3.6%)
Asymmetric(20 $\times$ )	2.75M	0.0028 (-37.7%)	0.0132 (-38.0%)	0.0002 (-37.1%)
Asymmetric(300 $\times$ )	24.3M	0.0009 (-80.1%)	0.0044 (-79.3%)	0.0001 (-77.6%)

721  
 722 across 10 random seeds with stratified sampling to ensure class balance. Statistical significance was  
 723 assessed using two-tailed paired t-tests.  
 724

725 **Theoretical Motivation.** Under data scarcity, the auxiliary sigmoid branches should provide par-  
 726 ticularly valuable regularization, as the primary softmax objective becomes increasingly prone to  
 727 overfitting. This effect should be most pronounced in architectures that exhibit gradient synergy  
 728 rather than conflict.

729 Table 9: Few-shot learning performance of MLP on CIFAR-10.  
 730

Samples/Class	PlainMLP	AsymmetricMLP	Improvement	p-value
10	$22.90 \pm 1.67$	$22.99 \pm 1.71$	+0.09	0.8082
50	$28.70 \pm 1.00$	$28.97 \pm 0.82$	+0.27	0.4346
100	$31.99 \pm 1.18$	$33.30 \pm 0.93$	+1.31	0.0122
500	$38.93 \pm 1.17$	$39.44 \pm 1.05$	+0.51	0.3087
1000	$43.13 \pm 0.39$	$44.00 \pm 0.43$	+0.87	0.0043
5000	$53.57 \pm 0.39$	$53.74 \pm 0.56$	+0.17	0.4824

739 Table 10: Few-shot learning performance of MLP on CIFAR-100.  
 740

Samples/Class	PlainMLP	AsymmetricMLP	Improvement	p-value
5	$4.06 \pm 0.64$	$5.22 \pm 0.25$	+1.16	0.0002
10	$7.27 \pm 0.38$	$6.97 \pm 0.39$	-0.30	0.0619
20	$9.36 \pm 0.36$	$9.49 \pm 0.29$	+0.13	0.2965
50	$12.98 \pm 0.45$	$13.98 \pm 0.48$	+1.00	0.0041
100	$15.95 \pm 0.28$	$16.97 \pm 0.29$	+1.02	0.0000
200	$21.31 \pm 0.44$	$22.25 \pm 0.22$	+0.94	0.0004
500	$25.94 \pm 0.31$	$26.65 \pm 0.32$	+0.71	0.0007

751 B.2 GENERALIZATION STABILITY UNDER LABEL CORRUPTION  
 752

753 We evaluated model resilience under label noise by corrupting a fraction of training labels and mea-  
 754 suring performance degradation. Label noise was introduced by randomly flipping labels with prob-  
 755 abilities ranging from 10% to 90%, while maintaining the original test set for consistent evaluation  
 (Tables 13 and 14).

756 Table 11: Few-shot learning performance of ViT-6L( $20\times$ ) on CIFAR-100.  
757

Samples/Class	PlainViT-6L	AsymmetricViT-6L	Improvement	p-value
5	$5.10 \pm 0.50$	$5.59 \pm 0.33$	+0.49	0.0022
10	$7.92 \pm 0.52$	$8.79 \pm 0.49$	+0.87	0.0008
20	$9.67 \pm 0.49$	$10.66 \pm 0.68$	+0.99	0.0052
50	$15.61 \pm 0.79$	$15.58 \pm 0.51$	-0.03	0.8687
100	$20.29 \pm 0.58$	$23.16 \pm 0.78$	+2.87	0.0001
200	$28.00 \pm 0.62$	$32.23 \pm 1.79$	+4.23	0.0011
500	$43.53 \pm 1.45$	$53.28 \pm 0.65$	+9.75	0.0000

767 Table 12: Architecture Performance Comparison (CIFAR-10)  
768

Architecture	Baseline	Asymmetric( $1\times$ )	Improvement	P-Value
MLP	$49.5 \pm 0.3$	$50.6 \pm 0.4$	+1.1	0.0001
CNN	$77.0 \pm 0.6$	$77.2 \pm 1.0$	+0.2	0.5371
ViT-6L	$62.5 \pm 0.7$	$63.6 \pm 0.4$	+1.1	0.0068

## 776 C DETAILED EXPERIMENTAL CONFIGURATION

## 778 C.1 HYPERPARAMETER SETTINGS

780 All hyperparameters were determined through systematic grid search following our “Pragmatic Gold  
781 Standard” strategy to ensure fair comparison. This three-stage optimization process isolates the  
782 effect of our asymmetric training paradigm while maintaining scientific rigor.

## 784 C.1.1 OPTIMIZATION STRATEGY

786 For each architecture, we employed a principled three-stage hyperparameter search:

787 **Stage 1: Learning Rate Optimization** We fixed weight decay at  $10^{-4}$  and conducted grid search  
788 over learning rates  $\{10^{-4}, 3 \times 10^{-4}, 10^{-3}, 3 \times 10^{-3}, 10^{-2}\}$  for the baseline Plain model, training  
789 for 150 epochs and selecting the configuration yielding highest validation accuracy.

790 **Stage 2: Weight Decay Refinement** Using the optimal learning rate from Stage 1, we searched  
791 over weight decay values  $\{10^{-1}, 10^{-2}, 10^{-3}, 10^{-4}, 10^{-5}\}$  for the Plain model, again training for  
792 150 epochs.

794 **Stage 3: Auxiliary Weight Search** With optimal learning rate and weight decay fixed,  
795 we searched for the optimal auxiliary weight  $\alpha$  using logarithmic spacing:  $\alpha \in$   
796  $\{0.1, 0.215, 0.464, 1.0, 2.154, 4.642, 10.0, 21.544, 46.416, 100.0\}$  for the Asymmetric model. For  
797 CIFAR-10 MLP specifically, we employed linear spacing  $\alpha \in [0, 50]$  to accommodate its distinct  
798 optimization characteristics.

## 799 C.1.2 FINAL HYPERPARAMETER CONFIGURATIONS

801 The optimal hyperparameters (CIFAR-100) determined through our systematic search are showed  
802 in Table 15.

## 804 C.2 TRAINING CONFIGURATION

806 **Training Duration:** All final results were obtained using 200 epochs.  
807

808 **Statistical Validation:** Each configuration was evaluated across 10 independent runs with different  
809 random seeds (42-51) to ensure statistical reliability. Performance comparisons used two-tailed  
paired t-tests.

Table 13: MLP performance under label noise on CIFAR-10

Noise Level	PlainMLP	AsymmetricMLP	Improvement	P-Value
0.0%	$53.52 \pm 0.53$	$53.74 \pm 0.36$	+0.22	0.2025
10.0%	$51.82 \pm 0.47$	$52.08 \pm 0.42$	+0.26	0.1785
20.0%	$49.92 \pm 0.55$	$50.22 \pm 0.70$	+0.30	0.3667
30.0%	$47.69 \pm 0.36$	$48.42 \pm 0.45$	+0.73	0.0007
40.0%	$44.96 \pm 0.62$	$45.90 \pm 0.85$	+0.94	0.0067
50.0%	$42.13 \pm 1.09$	$42.66 \pm 0.52$	+0.53	0.2633
60.0%	$38.35 \pm 0.57$	$39.23 \pm 0.70$	+0.88	0.0530
70.0%	$32.94 \pm 0.85$	$32.05 \pm 0.93$	-0.89	0.0291
80.0%	$23.69 \pm 1.61$	$22.94 \pm 0.94$	-0.75	0.1580
90.0%	$10.43 \pm 0.53$	$10.13 \pm 0.61$	-0.30	0.1081

Table 14: MLP performance under label noise on CIFAR-100

Noise Level	PlainMLP	AsymmetricMLP	Improvement	P-Value
0.0%	$25.64 \pm 0.29$	$26.00 \pm 0.34$	+0.36	0.0040
10.0%	$24.76 \pm 0.36$	$25.36 \pm 0.47$	+0.60	0.0104
20.0%	$23.83 \pm 0.34$	$24.29 \pm 0.38$	+0.46	0.0014
30.0%	$22.54 \pm 0.25$	$23.09 \pm 0.36$	+0.55	0.0006
40.0%	$20.97 \pm 0.42$	$21.82 \pm 0.36$	+0.85	0.0023
50.0%	$19.08 \pm 0.43$	$20.15 \pm 0.28$	+1.07	0.0000
60.0%	$16.52 \pm 0.39$	$17.97 \pm 0.44$	+1.45	0.0000
70.0%	$12.76 \pm 0.63$	$14.96 \pm 0.58$	+2.20	0.0000
80.0%	$8.16 \pm 0.73$	$9.56 \pm 0.49$	+1.40	0.0005
90.0%	$3.60 \pm 0.54$	$3.49 \pm 0.43$	-0.11	0.4902

**Hardware:** All experiments were conducted on NVIDIA RTX 3090 GPUs with consistent computational environments to ensure reproducibility.

### C.3 ARCHITECTURE-SPECIFIC DETAILS

**MLP:** 6 linear layers with ReLU activations. Auxiliary branches attached after the first 4 ReLU activations.

**CNN:** 6 convolutional layers, 2 MaxPooling layers, 1 Dropout layer, and 3 linear layers. Auxiliary branches are strategically placed after ReLU activations in convolutional blocks. When a convolutional layer is immediately followed by max-pooling, the auxiliary branch is placed after the max-pooling operation to maintain spatial coherence.

**ViT-6L:** 6 Transformer blocks with 4 attention heads each and embedding dimension of 128. Auxiliary branches attached after each Transformer block output.

All auxiliary branches consist of a single linear layer with output dimension equal to the number of classes, initialized using orthogonal initialization for training stability.

### C.4 IMAGENET EXPERIMENTS

In this section, we provide documentation for ImageNet-1k. (From Table 16 to Table 23)

### C.5 COATNET EXPERIMENTS

This provides comprehensive details of the experimental setup used in our baseline experiments. All configurations follow standard practices in modern CIFAR-100 image classification research.

864  
865  
866 Table 15: Optimal hyperparameters for Table 1  
867  
868  
869  
870  
871  
872  
873  
874  
875  
876  
877  
878  
879

$N_{branches}$	Architecture	Learning Rate	Weight Decay	$\alpha$
1×	MLP	0.0001	0.001	100.0
	CNN	0.0003	0.01	100.0
	ViT-6L	0.001	0.1	4.642
7×	MLP	0.0001	0.001	0.1
	CNN	0.0003	0.001	0.1
	ViT-6L	0.001	0.01	4.642
10×	MLP	0.0001	0.001	0.1
	CNN	0.0003	0.001	1.0
	ViT-6L	0.001	0.01	4.642
20×	MLP	0.0001	0.001	0.1
	CNN	0.0003	0.001	0.1
	ViT-6L	0.001	0.01	1.0

880  
881  
882 Table 16: Optimal alpha values for Random Seed 42  
883  
884  
885  
886  
887  
888

Architecture	1X	10X	20X
ResNet-18	0.0178	0.2371	0.0237
ViT-Small	23.7137	0.0237	0.0237
ResNet-50	0.0750	0.0042	0.0178
ViT-B/16	23.7137	0.0237	0.0042

## 892 C.5.1 HYPERPARAMETER CONFIGURATION

893 Table 24 summarizes the complete set of hyperparameters used in our baseline experiments.

## 896 C.5.2 MODEL ARCHITECTURE DETAILS

898 Our baseline model uses a CIFAR-100 adapted CoAtNet backbone, structurally similar to CoAtNet-  
899 0 but scaled down for the  $32 \times 32$  input resolution. Table 25 details the stage-wise configuration.901 The architecture follows the Conv–Conv–Attention–Attention pattern proposed in the original CoAt-  
902 Net paper (Dai et al., 2021), where early stages use convolutional MBCConv blocks and later stages  
903 employ Transformer blocks. The transition from convolution to attention occurs between S2 and S3.

## 905 C.5.3 JUSTIFICATION OF CONFIGURATION CHOICES

907 **Training budget:** The 200-epoch training with batch size 128 aligns with common practice in  
908 ResNet and ViT works on CIFAR-100, providing sufficient training iterations without excessive  
909 computational cost.910 **Data augmentation:** We employ the standard CIFAR augmentation recipe: random crop with 4-  
911 pixel padding and horizontal flip, combined with per-channel normalization. Notably, we do *not* use  
912 stronger augmentations such as CutMix, Mixup, or AutoAugment, ensuring the baseline does not  
913 gain unfair advantages from advanced data augmentation techniques.

## 916 C.5.4 COATNET-CIFAR RESULTS

917 Table 26 presents the complete results across all five random seeds (42–46) for each configuration.

918 Table 17: Optimal alpha values for Random Seed 43  
919

---

Architecture	1X	10X	20X
ResNet-18	0.0042	0.5623	0.2371
ViT-Small	31.6228	0.0316	0.0056
ResNet-50	5.6234	0.0750	0.1000
ViT-B/16	23.7137	0.0316	0.0316

---

926 Table 18: Optimal alpha values for Random Seed 44  
927

---

Architecture	1X	10X	20X
ResNet-18	0.0178	0.2371	0.0178
ViT-Small	31.6228	0.0133	0.0042
ResNet-50	0.5623	0.0042	0.0316
ViT-B/16	31.6228	0.0237	0.0178

---

936 D ATTENTION PATTERN EVOLUTION ANALYSIS  
937938 D.1 DETAILED ATTENTION VISUALIZATION  
939940 Figure 6 visualizes the evolution of self-attention maps throughout the training trajectory.  
941942 D.2 QUANTITATIVE ATTENTION ANALYSIS  
943944 We measured attention pattern quality using several metrics:  
945946 Key findings:  
947

- **Peak Strength:** Asymmetric training produces more diffuse attention patterns (lower peak values)
- **Map Entropy:** Higher entropy indicates more distributed attention across spatial locations
- **Sparsity:** Lower Gini coefficient suggests more egalitarian attention distribution
- **Object Coverage:** Asymmetric models achieve near-optimal object coverage much earlier (Epoch 20 vs 100)

955 D.3 LAYER-WISE ATTENTION DEVELOPMENT  
956957 The layer-wise analysis reveals that asymmetric training guides the development of hierarchical  
958 attention patterns: - **Early layers (L1):** Both variants show similar low-level feature attention -  
959 **Middle layers (L3):** Asymmetric variant begins showing more structured patterns - **Late layers**  
960 **(L6):** Clear differentiation—asymmetric model develops coherent object-level attention while plain  
961 model remains diffuse963 E STATISTICAL VALIDATION  
964965 All reported results were validated using appropriate statistical tests. For performance comparisons,  
966 we used paired t-tests with Bonferroni correction for multiple comparisons. Effect sizes were calcu-  
967 lated using Cohen’s d, with the following interpretations: small (0.2), medium (0.5), large (0.8).968 All main results show statistical significance ( $p < 0.001$ ) with large effect sizes, confirming the  
969 reliability of our findings.971 To investigate the formation process of the final attention patterns, we visualized the evolution of  
attention across different training stages (e.g., 20, 60, 100 epochs), as shown in Figure 6. We observe

972  
973  
974  
975  
976  
977  
978  
979  
980  
981  
982  
983  
984  
985  
986  
987  
988  
989  
990  
991  
992  
993  
994  
995  
996  
997  
998  
999  
1000  
1001  
1002  
1003  
1004  
1005  
1006  
1007  
1008  
1009  
1010  
1011  
1012  
1013  
1014  
1015  
1016  
1017  
1018  
1019  
1020  
1021  
1022  
1023  
1024  
1025  
Table 19: Optimal alpha values for Random Seed 45

Architecture	1X	10X	20X
ResNet-18	0.0032	0.0422	0.1778
ViT-Small	31.6228	0.0133	0.0422
ResNet-50	17.7828	0.0042	0.0237
ViT-B/16	31.6228	0.0237	0.0237

Table 20: Training hyperparameters for ImageNet-1k experiments

Parameter	ResNet-18	ViT-Small	ResNet-50	ViT-B/16
Optimizer	SGD	AdamW	SGD	AdamW
Learning Rate	0.2	1e-3	0.2	1e-3
Momentum	0.9	-	0.9	-
Weight Decay	1e-4	0.1	1e-4	0.3
Batch Size	512	256	512	256
Total Epochs	90	150	90	150
LR Schedule	StepLR	Cosine	StepLR	Cosine
Step Epochs	30/60/80	-	30/60/80	-
Warmup	-	5%	-	5%
Min LR	-	1e-5	-	1e-5

*Note: Learning rate for ResNets uses linear batch scaling from base 0.1. No heavy regularizers (Mixup, CutMix) used. Gradient clipping and zero-initialized heads applied uniformly.*

997  
998  
999  
1000  
that the attention patterns of the Asymmetric model gradually become more holistic and comprehensive as training progresses. In contrast, the attention of the Plain model saturates earlier and consistently focuses more on local textures.

### E.1 PROCESS LEVEL: LEARNING TRAJECTORIES

1003  
1004  
1005  
1006  
1007  
1008  
1009  
1010  
1011  
1012  
1013  
1014  
1015  
1016  
1017  
1018  
1019  
1020  
1021  
1022  
1023  
1024  
1025  
This microscopic gradient behavior directly translates into dramatically different macroscopic learning dynamics, as evidenced by our analysis across Figure 8 and Tables 31-32. For CNN, the gradient conflict drives a catastrophic optimization collapse—the model converges prematurely in just 5 epochs to a inferior solution with 71% performance degradation. More tellingly, the generalization gap becomes negative by epoch 50 (-0.0242), indicating the model performs better on validation than training data—a clear symptom of learning failure. In contrast, ViT’s constructive gradient synergy guides the optimization along a more exploratory but ultimately superior trajectory, requiring 11 additional epochs but achieving both 30.6% higher validation accuracy and 47% better generalization (gap reduction from 0.5543 to 0.2934). This demonstrates that beneficial gradient alignment not only improves final performance but fundamentally enhances the learning process itself.

Table 21: Detailed Architectural Configurations for ImageNet-1k Experiments (Table 2). We summarize the backbone specifications and the exact locations of auxiliary attachment points. For ResNet, attachments are made at the feature maps of each stage prior to global pooling. For ViT, attachments are made to the CLS token at evenly spaced block intervals. The redundancy level  $N$  (number of branches per point) is a hyperparameter (e.g.,  $N = 20$ ).

Architecture	Backbone Specification	Aux. Attach Points	Feature Dim ( $d_i$ )
<b>ResNet-18</b>	4-stage CNN (BasicBlock) GAP + FC(512→1000)	End of Stages 1, 2, 3, 4 (before pooling)	64, 128, 256, 512
<b>ResNet-50</b>	4-stage CNN (Bottleneck) GAP + FC(2048→1000)	End of Stages 1, 2, 3, 4 (before pooling)	256, 512, 1024, 2048
<b>ViT-Small</b>	12 blocks, $D=384$ , 6 heads CLS head 384→1000	CLS token after Blocks 3, 6, 9	384 (at all points)
<b>ViT-B/16</b>	12 blocks, $D=768$ , 12 heads CLS head 768→1000	CLS token after Blocks 3, 6, 9	768 (at all points)

**Auxiliary Head Configuration (Shared):**

At each attachment point, we instantiate  $N$  parallel branches. The first  $N - 1$  branches are **static** (weights orthogonally initialized and frozen/non-trainable), while the  $N$ -th branch is **active** (fully trainable). All branches project features to 1000-dim logits for Sigmoid/BCE supervision.

Table 22: Two-stage hyperparameter search strategy for  $\alpha$  in Asymmetric architecture on ImageNet-1k

Stage	Search Range	Candidate Values	Strategy
Stage 1 (Coarse)	$[10^{-2}, 10^1]$ [0.01, 10.0]	$\alpha = 0.01$	
		$\alpha = 0.056$	<code>np.logspace(-2, 1, 5)</code>
		$\alpha = 0.316$	Log-uniform sampling
		$\alpha = 1.778$	5 candidates
		$\alpha = 10.0$	
Stage 2 (Fine)	$[10^{\log_{\text{best}} - 0.5}, 10^{\log_{\text{best}} + 0.5}]$	$\alpha_1$	
		$\alpha_2$	
		$\alpha_3$	<code>np.logspace(log_{\text{best}} - 0.5, log_{\text{best}} + 0.5, 9)</code>
		$\alpha_4$	
		$\alpha_{\text{best}} (\text{center})$	Local refinement
		$\alpha_6$	around best value
		$\alpha_7$	
		$\alpha_8$	9 candidates
		$\alpha_9$	

Table 23: An Example of Stage 2 candidates when Stage 1 optimal is  $\alpha = 0.316$

Index	$\alpha$ Value
1	0.100
2	0.133
3	0.178
4	0.237
5	0.316 (optimal from Stage 1)
6	0.422
7	0.562
8	0.750
9	1.000

1080

1081

1082

Table 24: Hyperparameter configuration for CoAtNet baseline experiments on CIFAR-100.

Category	Parameter	Value
Dataset	Dataset	CIFAR-100
	Split	Official train/test
Data Preprocessing	Input resolution	$32 \times 32$
	Data augmentation	Random crop (4-pixel padding) + Random horizontal flip
	Normalization	Per-channel mean/std normalization
Optimizer	Type	AdamW
	Base learning rate	$1 \times 10^{-3}$
	Weight decay	0.05
	$\beta_1, \beta_2$	0.9, 0.999
Learning Rate Schedule	Warmup epochs	5
	Schedule type	Cosine decay
	Minimum learning rate	$1 \times 10^{-6}$
Training	Total epochs	200
	Batch size	128
Regularization	Drop path rate	0.1
	Label smoothing	0.1
Model	Architecture	CoAtNet-CIFAR (fixed structure)
	Auxiliary head	None (aux_stage = None, $\alpha = 0$ )

1105

1106

1107

1108

1109

Table 25: Architecture configuration of CoAtNet-CIFAR baseline model.

Stage	Block Type	Channels	Resolution	# Blocks	Operation
S0	Conv Stem	64	$32 \times 32 \rightarrow 32 \times 32$	1	Conv $3 \times 3$
S1	MBConv	96	$32 \times 32 \rightarrow 16 \times 16$	2	+ Downsample
S2	MBConv	128	$16 \times 16 \rightarrow 8 \times 8$	2	+ Downsample
S3	Transformer	256	$8 \times 8 \rightarrow 8 \times 8$	3	Self-Attention
S4	Transformer	384	$8 \times 8 \rightarrow 4 \times 4$	2	+ Downsample

1117

1118

1119

1120

1121

1122

Table 26: Detailed test accuracy results across five random seeds for all configurations (CoAtNet-CIFAR).

Redundancy	Stage	Alpha	Seed 42	Seed 43	Seed 44	Seed 45	Seed 46	Mean	Std
1×	S2	4.642	75.63	75.34	75.94	75.18	74.21	75.26	0.66
	S3	21.544	75.94	76.14	75.64	75.60	76.21	75.91	0.28
20×	S2	0.215	75.96	75.90	75.79	76.05	75.70	75.88	0.14
	S3	4.642	76.85	76.43	77.07	76.31	77.14	76.76	0.36
100×	S2	0.1	75.34	75.38	76.32	75.54	75.39	75.59	0.41
	S3	4.642	77.41	77.32	76.95	76.74	76.96	77.08	0.26
300×	S2	0.1	74.41	73.97	74.42	75.20	74.70	74.54	0.49
	S3	0.1	76.72	77.22	76.96	76.76	76.44	76.82	0.30

1133

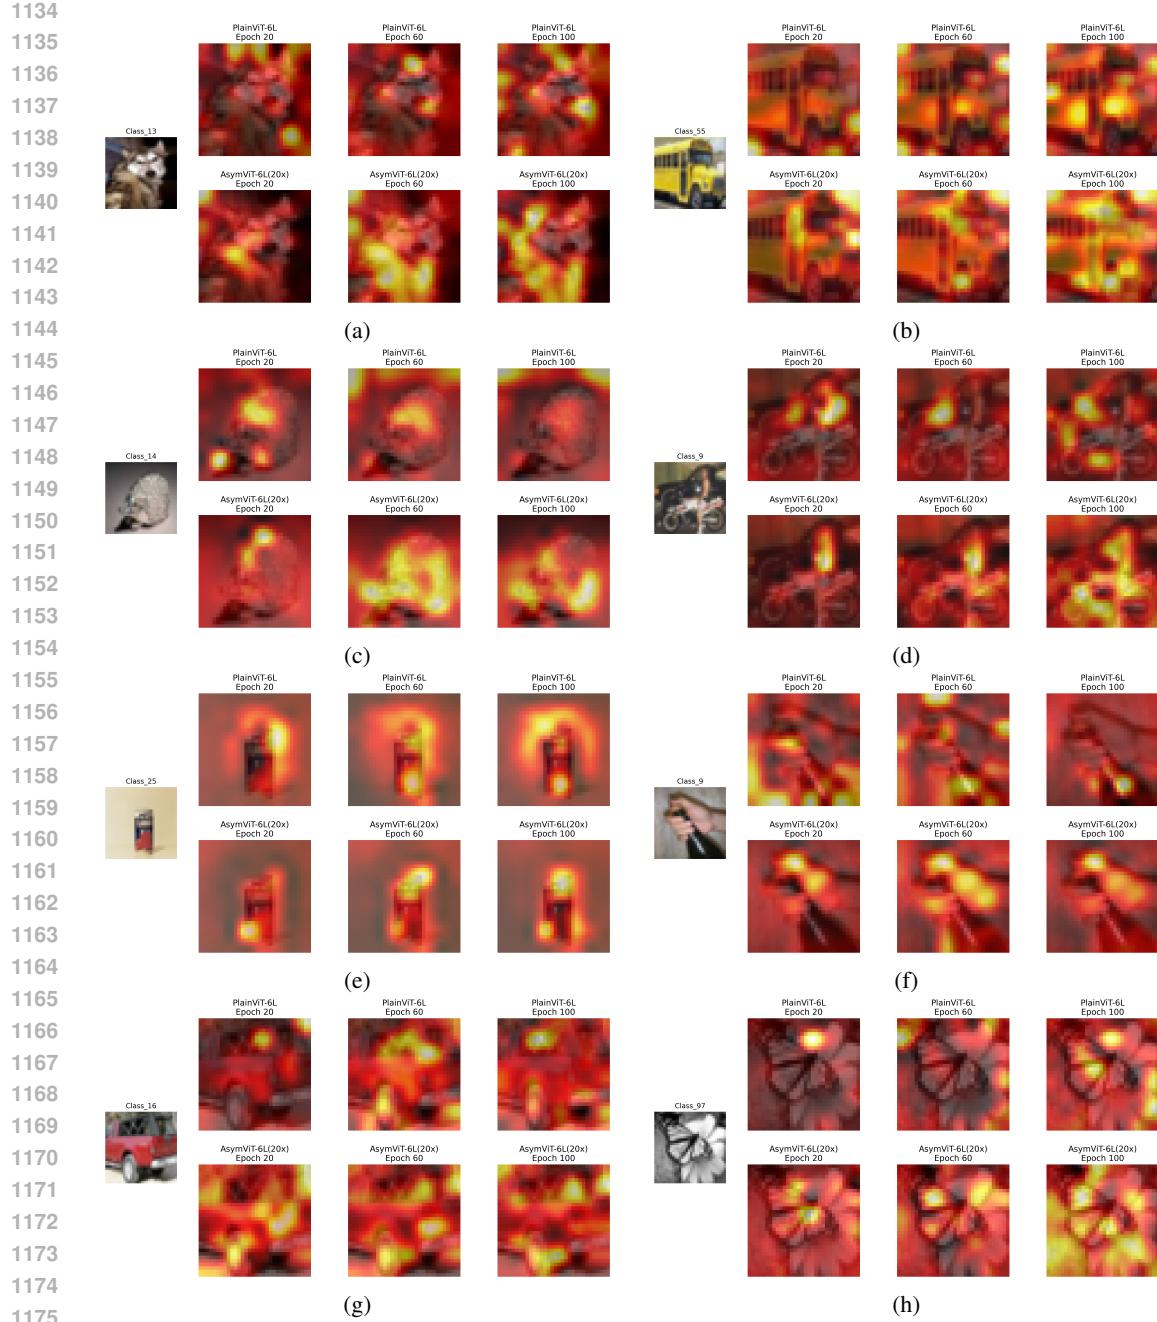


Figure 6: Comprehensive attention pattern evolution for ViT-6L on CIFAR-100.

Table 27: Attention pattern quality metrics across training epochs

Metric	Plain ViT-6L			Asymmetric ViT-6L		
	Epoch 20	Epoch 60	Epoch 100	Epoch 20	Epoch 60	Epoch 100
Peak Strength	0.092	0.063	0.051	0.053	0.049	0.045
Map Entropy	3.73	3.89	3.97	4.02	4.03	4.05
Sparsity (Gini)	0.471	0.415	0.374	0.254	0.287	0.302
Object Coverage	0.58	0.70	0.84	0.96	0.95	0.94

1188  
 1189 Table 28: Gradient Conflict Analysis. Cosine similarity analysis between main (softmax) and aux-  
 1190 illiary (sigmoid) gradients during training. ViT-6L shows consistent positive similarity (synergy),  
 1191 while CNN exhibits strong negative similarity (conflict), and MLP demonstrates near-orthogonal  
 1192 gradients with slight conflict tendency.

Architecture	Final Similarity	Average Similarity	Min Similarity	Max Similarity
MLP	-0.0309	-0.0801	-0.3422	0.0701
CNN	-0.1926	-0.2574	-0.8210	0.0464
ViT-6L	0.2631	0.1870	-0.1845	0.3654

1193  
 1194  
 1195  
 1196 Table 29: To explore the impact of different hyperparameter search strategies, we conducted a sup-  
 1197 plementary experiment for the  $N_{\text{branches}} = 3$  configuration, where hyperparameters (learning rate  
 1198 and weight decay) were independently optimized for both baseline and asymmetric models. The  
 1199 results are shown in the table, where “all active” indicates that all three auxiliary branches at each  
 1200 connection point are trainable, while “one active” means only one auxiliary branch per connection  
 1201 point is trainable. Although this “dual optimization” strategy can yield benefits in certain cases, we  
 1202 consistently adopted the “Pragmatic Gold Standard” strategy throughout the main text to isolate the  
 1203 pure effect of our paradigm.

$N_{\text{branches}}$	Architecture	Plain	Asymmetric	Improvement	p-value
3 $\times$ branches (all active)	MLP	$24.6 \pm 0.3$	$20.7 \pm 0.5$	-3.9	0.0000
	CNN	$39.9 \pm 0.6$	$36.2 \pm 0.6$	-3.7	0.0000
	ViT-6L	$36.2 \pm 0.6$	$42.1 \pm 2.7$	+5.9	0.0002
3 $\times$ branches (one active)	MLP	$24.6 \pm 0.3$	$23.7 \pm 0.6$	-0.9	0.0008
	CNN	$40.0 \pm 0.8$	$33.9 \pm 0.9$	-6.1	0.0000
	ViT-6L	$36.5 \pm 0.7$	$44.0 \pm 5.9$	+7.5	0.0045

1207  
 1208  
 1209  
 1210 Table 30: Architecture performance comparison with asymmetric training on CIFAR-10. Results  
 1211 show differential architectural responses to auxiliary supervision, with statistical significance as-  
 1212 sessed using two-tailed paired t-tests across 10 independent runs.

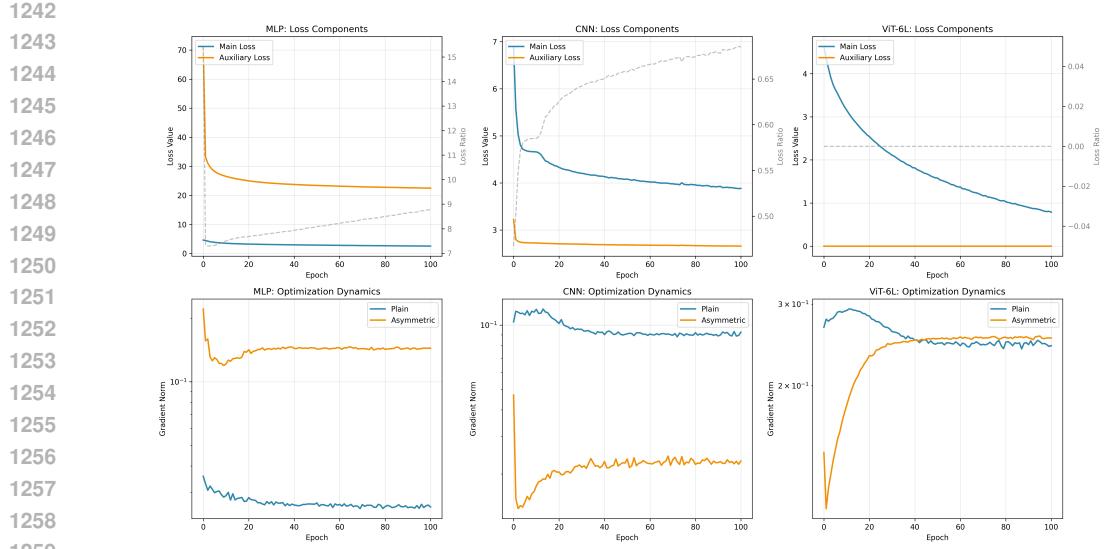
Architecture	Baseline	Asymmetric	Improvement	p-value	Significant
MLP	$49.5 \pm 0.3$	$50.6 \pm 0.4$	+1.1	0.0001	Yes
CNN	$77.0 \pm 0.6$	$77.2 \pm 1.0$	+0.2	0.5371	No
ViT-6L	$62.5 \pm 0.7$	$63.6 \pm 0.4$	+1.1	0.0068	Yes

1225  
 1226 Table 31: Architecture-dependent convergence patterns and performance outcomes.

Architecture	Convergence (Epoch)		Final Val Acc	
	Plain	Asymmetric	Plain	Asymmetric
MLP	15	21	0.2220	0.2305
CNN	20	5	0.4034	0.1171
ViT-6L	19	30	0.3498	0.4568

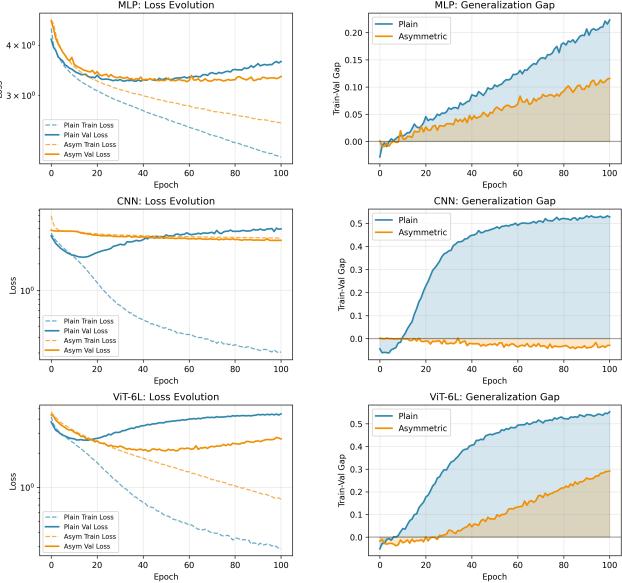
1234  
 1235 Table 32: Generalization Gap Evolution Across Training Phases

Architecture	Early(epoch10)		Middle(epoch50)		Late(epoch100)	
	Plain	Asymmetric	Plain	Asymmetric	Plain	Asymmetric
MLP	0.0098	0.0049	0.1039	0.0566	0.2234	0.1225
CNN	0.0073	0.0002	0.4665	-0.0242	0.5253	-0.0365
ViT-6L	0.0334	-0.0270	0.4717	0.0838	0.5543	0.2934



1260  
1261  
1262  
1263  
1264  
1265  
1266  
1267  
1268  
1269  
1270  
1271  
1272  
1273  
1274  
1275  
1276  
1277  
1278  
1279  
1280  
1281  
1282  
1283  
1284  
1285  
1286

Figure 7: Training dynamics comparison across architectures showing loss components and optimization trajectories. **Top row:** Evolution of main loss (softmax) and auxiliary loss (sigmoid) during training. **Bottom row:** Gradient norm dynamics for plain and asymmetric variants. Asymmetric training exhibits architecture-specific patterns: MLP shows stable auxiliary loss with reduced gradient norms, CNN demonstrates auxiliary loss divergence with increased gradient instability, while ViT-6L displays rapid auxiliary loss convergence with improved optimization stability.



1287  
1288  
1289  
1290  
1291  
1292  
1293  
1294  
1295

Figure 8: Architecture-specific learning dynamics reveal the Principle of Architectural Resonance. Comprehensive learning trajectories comparing Plain (baseline) and Asymmetric training across three architectures on CIFAR-100. Left column shows accuracy evolution (dashed: training, solid: validation); middle column displays loss curves; right column presents generalization gaps (train-val accuracy difference). CNN exhibits catastrophic degradation with massive overfitting under asymmetric training. MLP demonstrates effective regularization with reduced generalization gap but limited accuracy gains. ViT achieves substantial performance improvements with superior generalization. The divergence points (red vertical lines) mark early onset of architecture-dependent responses to heterogeneous supervision, empirically validating our core hypothesis that auxiliary signal efficacy depends fundamentally on architectural inductive biases.

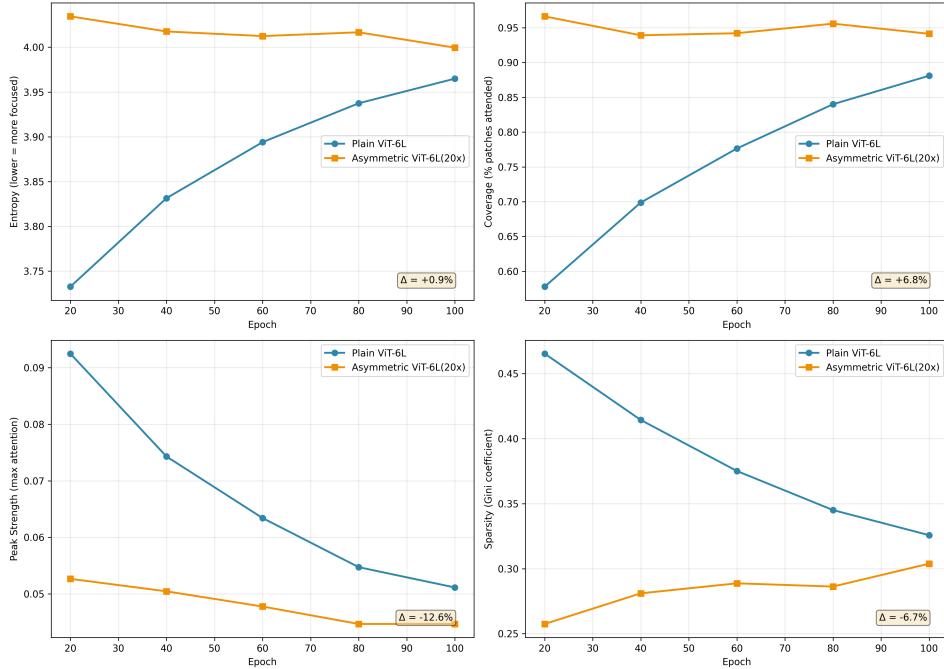


Figure 9: ViT-6L attention mechanism analysis comparing plain and asymmetric training. **Top left:** Entropy evolution showing asymmetric training maintains higher attention diversity. **Top right:** Coverage percentage demonstrating improved spatial attention coverage (+6.8%). **Bottom left:** Peak strength indicating more focused attention patterns (-12.6%). **Bottom right:** Sparsity coefficient revealing attention distribution characteristics (-6.7%). Results suggest asymmetric training promotes more comprehensive yet focused attention patterns.